

---

# Reducing the variance in online optimization by transporting past gradients

---

Anonymous Author(s)

Affiliation

Address

email

## Abstract

1 Most stochastic optimization methods use gradients once before discarding them.  
2 While variance reduction methods have shown that reusing past gradients can be  
3 beneficial when there is a finite number of datapoints, they do not easily extend  
4 to the online setting. One issue is the staleness due to using past gradients. We  
5 propose to correct this staleness using the idea of *implicit gradient transport* (IGT)  
6 which transforms gradients computed at previous iterates into gradients evaluated  
7 at the current iterate without using the Hessian explicitly. In addition to reducing  
8 the variance and bias of our updates over time, IGT can be used as a drop-in  
9 replacement for the gradient estimate in a number of well-understood methods  
10 such as heavy ball or Adam. We show experimentally that it achieves state-of-  
11 the-art results on a wide range of architectures and benchmarks. Additionally,  
12 the IGT gradient estimator yields the optimal asymptotic convergence rate for  
13 online stochastic optimization in the restricted setting where the Hessians of all  
14 component functions are equal.

## 15 1 Introduction

16 We wish to solve the following minimization problem:

$$\theta^* = \arg \min_{\theta} E_{x \sim p}[f(\theta, x)], \quad (1)$$

17 where we only have access to samples  $x$  and to a first-order oracle that gives us, for a given  $\theta$  and a  
18 given  $x$ , the derivative of  $f(\theta, x)$  with respect to  $\theta$ , i.e.  $\frac{\partial f(\theta, x)}{\partial \theta} = g(\theta, x)$ . It is known [35] that, when  
19  $f$  is smooth and strongly convex, there is a converging algorithm for Problem 1 that takes the form  
20  $\theta_{t+1} = \theta_t - \alpha_t g(\theta_t, x_t)$ , where  $x_t$  is a sample from  $p$ . This algorithm, dubbed stochastic gradient  
21 (SG), has a convergence rate of  $O(1/t)$  (see for instance [4]), within a constant factor of the minimax  
22 rate for this problem. When one has access to the true gradient  $g(\theta) = E_{x \sim p}[g(\theta, x)]$  rather than just  
23 a sample, this rate dramatically improves to  $O(e^{-\nu t})$  for some  $\nu > 0$ .

24 In addition to hurting the convergence speed, noise in the gradient makes optimization algorithms  
25 harder to tune. Indeed, while full gradient descent is convergent for constant stepsize  $\alpha$ , and also  
26 amenable to line searches to find a good value for that stepsize, the stochastic gradient method  
27 from [35] with a constant stepsize only converges to a ball around the optimum [38].<sup>1</sup> Thus,  
28 to achieve convergence, one needs to use a decreasing stepsize. While this seems like a simple  
29 modification, the precise decrease schedule can have a dramatic impact on the convergence speed.  
30 While theory prescribes  $\alpha_t = O(t^{-\alpha})$  with  $\alpha \in (1/2, 1]$  in the smooth case, practitioners often use  
31 larger stepsizes like  $\alpha_t = O(t^{-1/2})$  or even constant stepsizes.

---

<sup>1</sup>Under some conditions, it does converge linearly to the optimum [e.g., 45]

32 When the distribution  $p$  has finite support, Eq. 1 becomes a finite sum and, in that setting, it is possible  
 33 to achieve efficient variance reduction and drive the noise to zero, allowing stochastic methods to  
 34 achieve linear convergence rates [24, 17, 50, 28, 42, 5]. Unfortunately, the finite support assumption is  
 35 critical to these algorithms which, while valid in many contexts, does not have the broad applicability  
 36 of the standard SG algorithm. Several works have extended these approaches to the online setting by  
 37 applying these algorithms while increasing  $N$  [2, 14] but they need to revisit past examples multiple  
 38 times and are not truly online.

39 Another line of work reduces variance by averaging iterates [33, 22, 3, 10, 7, 6, 16]. While these  
 40 methods converge for a constant stepsize in the stochastic case<sup>2</sup>, their practical speed is heavily  
 41 dependent on the fraction of iterates kept in the averaging, a hyperparameter that is thus hard to tune,  
 42 and they are rarely used in deep learning.

43 Our work combines two existing ideas and adds a third: a) At every step, it updates the parameters  
 44 using a weighted average of past gradients, like in SAG [24, 40], albeit with a different weighting  
 45 scheme; b) It reduces the bias and variance induced by the use of these old gradients by transporting  
 46 them to “equivalent” gradients computed at the current point, similar to [11]; c) It does so implicitly  
 47 by computing the gradient at a parameter value different from the current one. The resulting gradient  
 48 estimator can then be used as a plug-in replacement of the stochastic gradient within any optimization  
 49 scheme. Experimentally, both SG using our estimator and its momentum variant outperform the most  
 50 commonly used optimizers in deep learning.

## 51 2 Momentum and other approaches to dealing with variance

52 Stochastic variance reduction methods use an average of past gradients to reduce the variance of the  
 53 gradient estimate. At first glance, it seems like their updates are similar to that of momentum [32],  
 54 also known as the heavy ball method, which performs the following updates<sup>3</sup>:

$$v_t = \gamma_t v_{t-1} + (1 - \gamma_t)g(\theta_t, x_t), \quad v_0 = g(\theta_0, x_0)$$

$$\theta_{t+1} = \theta_t - \alpha_t v_t .$$

55 When  $\gamma_t = \gamma$ , this leads to  $\theta_{t+1} = \theta_t - \alpha_t \left( \gamma^t g(\theta_0, x_0) + (1 - \gamma) \sum_{i=1}^t \gamma^{t-i} g(\theta_i, x_i) \right)$ . Hence, the  
 56 heavy ball method updates the parameters of the model using an average of past gradients, bearing  
 57 similarity with SAG [24], albeit with exponential instead of uniform weights.

58 Interestingly, while momentum is a popular method for training deep networks, its theoretical analysis  
 59 in the stochastic setting is limited [44], except in the particular setting when the noise converges  
 60 to 0 at the optimum [26]. Also surprising is that, despite the apparent similarity with stochastic  
 61 variance reduction methods, current convergence rates are slower when using  $\gamma > 0$  in the presence  
 62 of noise [39], although this might be a limitation of the analysis.

### 63 2.1 Momentum and variance

64 We propose here an analysis of how, on quadratics, using past gradients as done in momentum  
 65 does not lead to a decrease in variance. If gradients are stochastic, then  $\Delta_t = \theta_t - \theta^*$  is a random  
 66 variable. Denoting  $\epsilon_i$  the noise at timestep  $i$ , i.e.  $g(\theta_i, x_i) = g(\theta_i) + \epsilon_i$ , and writing  $\Delta_t - E[\Delta_t] =$   
 67  $\alpha \sum_{i=0}^t N_{i,t} \epsilon_i$ , with  $N_{i,t}$  the impact of the noise of the  $i$ -th datapoint on the  $t$ -th iterate, we may now  
 68 analyze the total impact of each  $\epsilon_i$  on the iterates. Figure 1 shows the impact of  $\epsilon_i$  on  $\Delta_t - E[\Delta_t]$  as  
 69 measured by  $N_{i,t}^2$  for three datapoints ( $i = 1, i = 25$  and  $i = 50$ ) as a function of  $t$  for stochastic  
 70 gradient ( $\gamma = 0$ , left) and momentum ( $\gamma = 0.9$ , right). As we can see, when using momentum, the  
 71 variance due to a given datapoint first increases as the noise influences both the next iterate (through  
 72 the parameter update) and the subsequent updates (through the velocity). Due to the weight  $1 - \gamma$   
 73 when a point is first sampled, a larger value of  $\gamma$  leads to a lower immediate impact of the noise of a  
 74 given point on the iterates. However, a larger  $\gamma$  also means that the noise of a given gradient is kept  
 75 longer, leading to little or no decrease of the total variance (dashed blue curve). Even in the case of  
 76 stochastic gradient, the noise at a given timestep carries over to subsequent timesteps, even if the old  
 77 gradients are not used for the update, as the iterate itself depends on the noise.

<sup>2</sup>Under some conditions on  $f$ .

<sup>3</sup>This is slightly different from the standard formulation but equivalent for constant  $\gamma_t$ .

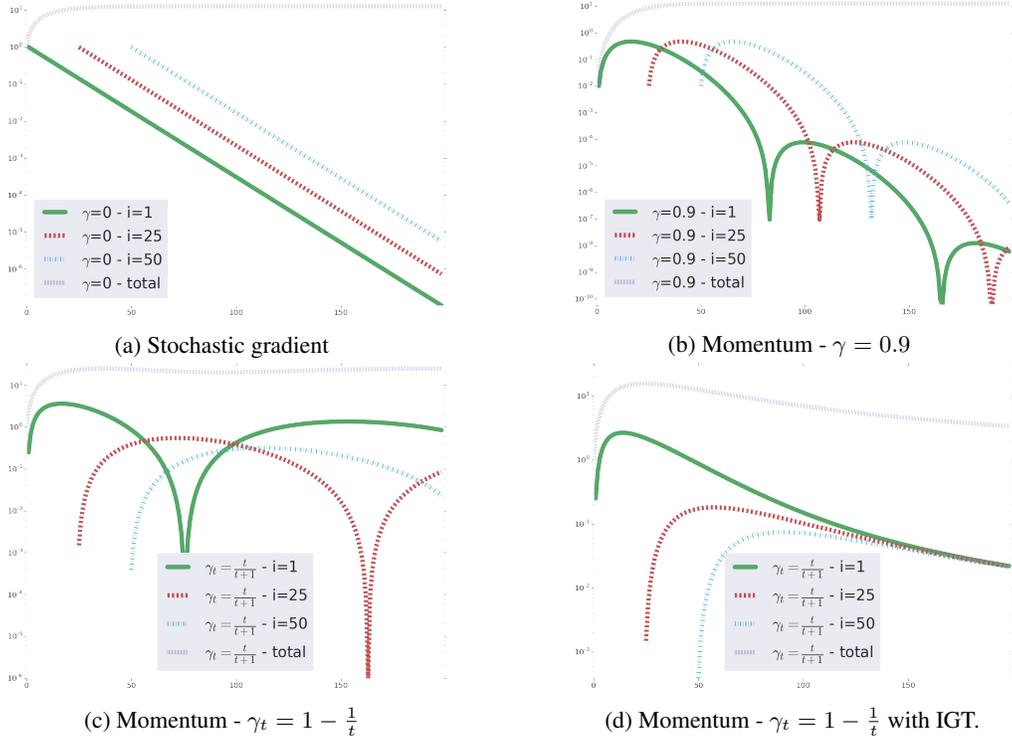


Figure 1: Variance induced over time by the noise from three different datapoints ( $i = 1$ ,  $i = 25$  and  $i = 50$ ) as well as the total variance for SG ( $\gamma = 0$ , *top left*), momentum with fixed  $\gamma = 0.9$  (*top right*), momentum with increasing  $\gamma_t = 1 - \frac{1}{t}$  without (*bottom left*) and with (*bottom right*) transport. The impact of the noise of each gradient  $\epsilon_i$  increases for a few iterations then decreases. Although a larger  $\gamma$  reduces the maximum impact of a given datapoint, the total variance does not decrease. With transport, noises are now equal and total variance decreases. The y-axis is on a log scale.

78 At every timestep, the contribution to the noise of the 1st, the 25th and the 50th points in Fig. 1 is  
 79 unequal. If we assume that the  $\epsilon_i$  are i.i.d., then the total variance would be minimal if the contribution  
 80 from each point was equal. Further, one can notice that the impact of datapoint  $i$  is only a function of  
 81  $t - i$  and not of  $t$ . This guarantees that the total noise will not decrease over time.

82 To address these two points, one can increase the momentum parameter over time. In doing so,  
 83 the noise of new datapoints will have a decreasing impact on the total variance as their gradient is  
 84 multiplied by  $1 - \gamma_t$ . Figure 1c shows the impact  $N_{i,t}^2$  of each noise  $\epsilon_i$  for an increasing momentum  
 85  $\gamma_t = 1 - \frac{1}{t}$ . The peak of noise for  $i = 25$  is indeed lower than that of  $i = 1$ . However, the variance  
 86 still does not go to 0. This is because, as the momentum parameter increases, the update is an average  
 87 of many gradients, including stale ones. Since these gradients were computed at iterates already  
 88 influenced by the noise over previous datapoints, that past noise is amplified, as testified by the higher  
 89 peak at  $i = 1$  for the increasing momentum. Ultimately, increasing momentum does not lead to a  
 90 convergent algorithm in the presence of noise when using a constant stepsize.

## 91 2.2 SAG and Hessian modelling

92 The impact of the staleness of the gradients on the convergence is not limited to momentum. In SAG,  
 93 for instance, the excess error after  $k$  updates is proportional to  $(1 - \min\{\frac{1}{16\hat{\kappa}}, \frac{1}{8N}\})^k$ , compared to  
 94 the excess error of the full gradient method which is  $(1 - \frac{1}{\kappa})^k$  where  $\kappa$  is the condition number of  
 95 the problem.<sup>4</sup> The difference between the two rates is larger when the minimum in the SAG rate is  
 96 the second term. This happens either when  $\hat{\kappa}$  is small, i.e. the problem is well conditioned and a lot

<sup>4</sup>The  $\hat{\kappa}$  in the convergence rate of SAG is generally larger than the  $\kappa$  in the full gradient algorithm.

97 of progress is made at each step, or when  $N$  is large, i.e. there are many points to the training set.  
 98 Both cases imply that a large distance has been travelled between two draws of the same datapoint.

99 Recent works showed that correcting for that staleness by modelling the Hessian [46, 11] leads to  
 100 improved convergence. As momentum uses stale gradients, the velocity is an average of current and  
 101 past gradients and thus can be seen as an estimate of the true gradient at a point which is not the  
 102 current one but rather a convex combination of past iterates. As past iterates depend on the noise  
 103 of previous gradients, this bias in the gradients amplifies the noise and leads to a non-converging  
 104 algorithm. We shall thus “transport” the old stochastic gradients  $g(\theta_i, x_i)$  to make them closer to  
 105 their corresponding value at the current iterate,  $g(\theta_t, x_i)$ . Past works did so using the Hessian or an  
 106 explicit approximation thereof, which can be expensive and difficult to compute and maintain. We  
 107 will resort to using *implicit transport*, a new method that aims at compensating the staleness of past  
 108 gradients without making explicit use of the Hessian.

### 109 3 Converging optimization through implicit gradient transport

110 Before showing how to combine the advantages of both increasing momentum and gradient transport,  
 111 we demonstrate how to transport gradients implicitly. This transport is only exact under a strong  
 112 assumption that will not hold in practice. However, this result will serve to convey the intuition behind  
 113 implicit gradient transport. We will show in Section 4 how to mitigate the effect of the unsatisfied  
 114 assumption.

#### 115 3.1 Implicit gradient transport

116 Let us assume that we received samples  $x_0, \dots, x_t$  in an online fashion. We wish to approach the full  
 117 gradient  $g_t(\theta_t) = \frac{1}{t+1} \sum_{i=0}^t g(\theta_t, x_i)$  as accurately as possible. We also assume here that a) We have  
 118 a noisy estimate  $\hat{g}_{t-1}(\theta_{t-1})$  of  $g_{t-1}(\theta_{t-1})$ ; b) We can compute the gradient  $g(\theta, x_t)$  at any location  
 119  $\theta$ . We shall seek a  $\theta$  such that

$$\frac{t}{t+1} \hat{g}_{t-1}(\theta_{t-1}) + \frac{1}{t+1} g(\theta, x_t) \approx g_t(\theta_t).$$

120 To this end, we shall make the following assumption:

121 **Assumption 3.1.** *All individual functions  $f(\cdot, x)$  are quadratics with the same Hessian  $H$ .*

122 This is the same assumption as [10, Section 4.1]. Although it is unlikely to hold in practice, we shall  
 123 see that our method still performs well when that assumption is violated.

124 Under Assumption 3.1, we then have (see details in Appendix)

$$\begin{aligned} g_t(\theta_t) &= \frac{t}{t+1} g_{t-1}(\theta_t) + \frac{1}{t+1} g(\theta_t, x_t) \\ &\approx \frac{t}{t+1} \hat{g}_{t-1}(\theta_{t-1}) + \frac{1}{t+1} g(\theta_t + t(\theta_t - \theta_{t-1}), x_t). \end{aligned}$$

125 Thus, we can transport our current estimate of the gradient by computing the gradient on the new  
 126 point at a shifted location  $\theta = \theta_t + t(\theta_t - \theta_{t-1})$ . This extrapolation step is reminiscent of Nesterov’s  
 127 acceleration with the difference that the factor in front of  $\theta_t - \theta_{t-1}$ ,  $t$ , is not bounded.

#### 128 3.2 Combining increasing momentum and implicit gradient transport

129 We now describe our main algorithm, Implicit Gradient Transport (IGT). IGT uses an increasing  
 130 momentum  $\gamma_t = \frac{t}{t+1}$ . At each step, when updating the velocity, it computes the gradient of the new  
 131 point at an extrapolated location so that the velocity  $v_t$  is a good estimate of the true gradient  $g(\theta_t)$ .

132 We can rewrite the updates to eliminate the velocity  $v_t$ , leading to the update:

$$\theta_{t+1} = \frac{2t+1}{t+1} \theta_t - \frac{t}{t+1} \theta_{t-1} - \frac{\alpha}{t+1} g(\theta_t + t(\theta_t - \theta_{t-1}), x_t). \quad (\text{IGT})$$

133 We see in Fig. 1d that IGT allows a reduction in the total variance, thus leading to convergence with a  
 134 constant stepsize. This is captured by the following proposition:

135 **Proposition 3.1.** *If  $f$  is a quadratic function with positive definite Hessian  $H$  with largest eigenvalue*  
136  *$L$  and condition number  $\kappa$  and if the stochastic gradients satisfy:  $g(\theta, x) = g(\theta) + \epsilon$  with  $\epsilon$  a random*  
137 *i.i.d. noise with covariance bounded by  $BI$ , then Eq. IGT with stepsize  $\alpha = 1/L$  leads to iterates  $\theta_t$*   
138 *satisfying*

$$E[\|\theta_t - \theta^*\|^2] \leq \left(1 - \frac{1}{\kappa}\right)^{2t} \|\theta_0 - \theta^*\|^2 + \frac{d\alpha^2 B \bar{v}_0^2}{t},$$

139 with  $\nu = (2 + 2 \log \kappa)\kappa$  for every  $t > 2\kappa$ .

140 The proof of Prop. 3.1 is provided in the appendix.

141 Despite this theoretical result, two limitations remain: First, Prop. 3.1 shows that IGT does not  
142 improve the dependency on the conditioning of the problem; Second, the assumption of equal  
143 Hessians is unlikely to be true in practice, leading to an underestimation of the bias. We address the  
144 conditioning issue in the next section and the assumption on the Hessians in Section 4.

### 145 3.3 IGT as a plug-in gradient estimator

146 We demonstrated that the IGT estimator has lower variance than the stochastic gradient estimator for  
147 quadratic objectives. IGT can also be used as a drop-in replacement for the stochastic gradient in  
148 an existing, popular first order method: the heavy ball (HB). This is captured by the following two  
149 propositions:

150 **Proposition 3.2** (Non-stochastic). *In the non-stochastic case, where  $B = 0$ , variance is equal to 0*  
151 *and Heavyball-IGT achieves the accelerated linear rate  $O\left(\left(\frac{\sqrt{\kappa}-1}{\sqrt{\kappa}+1}\right)^t\right)$  using the known, optimal*

152 *heavy ball tuning,  $\mu = \left(\frac{\sqrt{\kappa}-1}{\sqrt{\kappa}+1}\right)^2$ ,  $\alpha = (1 + \sqrt{\mu})^2/L$ .*

153 **Proposition 3.3** (Online, stochastic). *When  $B > 0$ , there exist constant hyperparameters  $\alpha > 0$ ,*  
154  *$\mu > 0$  such that  $\|E[\theta_t - \theta^*]\|^2$  converges to zero linearly, and the variance is  $\tilde{O}(1/t)$ .*

155 The pseudo-code can be found in Algorithm 1.

---

#### Algorithm 1 Heavyball-IGT

---

```

1: procedure HEAVYBALL-IGT(Stepsize  $\alpha$ , Momentum  $\mu$ , Initial parameters  $\theta_0$ )
2:    $v_0 \leftarrow g(\theta_0, x_0)$  ,  $w_0 \leftarrow -\alpha v_0$  ,  $\theta_1 \leftarrow \theta_0 + w_0$ 
3:   for  $t = 1, \dots, T - 1$  do
4:      $\gamma_t \leftarrow \frac{t}{t+1}$ 
5:      $v_t \leftarrow \gamma_t v_{t-1} + (1 - \gamma_t)g\left(\theta_t + \frac{\gamma_t}{1-\gamma_t}(\theta_t - \theta_{t-1}), x_t\right)$ 
6:      $w_t \leftarrow \mu w_{t-1} - \alpha v_t$ 
7:      $\theta_{t+1} \leftarrow \theta_t + w_t$ 
8:   end for
9:   return  $\theta_T$ 
10: end procedure

```

---

## 156 4 IGT and Anytime Tail Averaging

157 So far, IGT weighs all gradients equally. This is because, with equal Hessians, one can perfectly  
158 transport these gradients irrespective of the distance travelled since they were computed. In practice,  
159 the individual Hessians are not equal and might change over time. In that setting, the transport induces  
160 an error which grows with the distance travelled. We wish to average a linearly increasing number of  
161 gradients, to maintain the  $O(1/t)$  rate on the variance, while forgetting about the oldest gradients to  
162 decrease the bias. To this end, we shall use *anytime tail averaging* [23], named in reference to the tail  
163 averaging technique used in optimization [16].

164 Tail averaging is an online averaging technique where only the last points, usually a constant  
165 fraction  $c$  of the total number of points seen, is kept. Maintaining the exact average at every  
166 timestep is memory inefficient and anytime tail averaging performs an approximate averaging using

167  $\gamma_t = \frac{c(t-1)}{1+c(t-1)} \left(1 - \frac{1}{c} \sqrt{\frac{1-c}{t(t-1)}}\right)$ . We refer the reader to [23] for additional details.

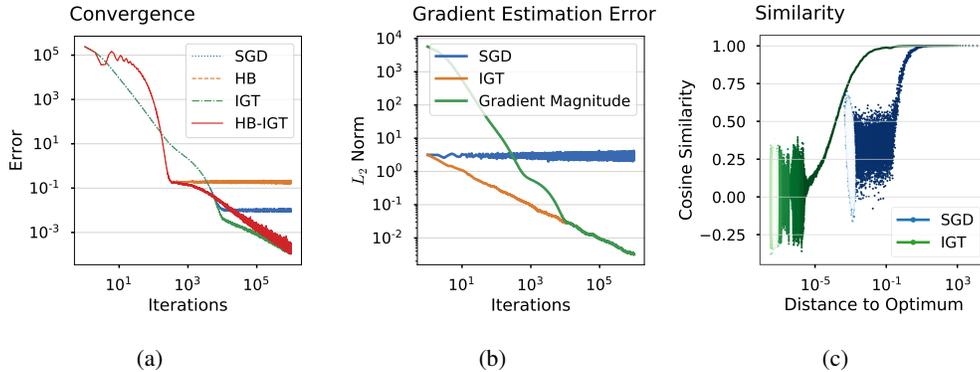


Figure 2: Analysis of IGT on quadratic loss functions. (a) Comparison of convergence curves for multiple algorithms. As expected, the IGT family of algorithms converges to the solution while stochastic gradient algorithms can not. (b) The blue and orange curves show the norm of the noise component in the SGD and IGT gradient estimates, respectively. The noise component of SGD remains constant, while it decreases at a rate  $1/\sqrt{t}$  for IGT. The green curve shows the norm of the IGT gradient estimate. (c) Cosine similarity between the full gradient and the SGD/IGT estimates.

## 168 5 Impact of IGT on bias and variance in the ideal case

169 To understand the behaviour of IGT when Assumption 3.1 is verified, we minimize a strongly convex  
 170 quadratic function with Hessian  $Q \in \mathbb{R}^{100 \times 100}$  with condition number 1000, and we have access to  
 171 the gradient corrupted by noise  $\epsilon_t$ , where  $\epsilon_t \sim N(0, 0.3 \cdot I_{100})$ . In that scenario where all Hessians are  
 172 equal and implicit gradient transport is exact, Fig. 2a confirms the  $O(1/t)$  rate of IGT with constant  
 173 stepsize while SGD and HB only converge to a ball around the optimum.

174 To further understand the impact of IGT, we study the quality of the gradient estimate. Standard  
 175 stochastic methods control the variance of the parameter update by scaling it with a decreasing  
 176 stepsize, which slows the optimization down. With IGT, we hope to have a low variance while  
 177 maintaining a norm of the update comparable to that obtained with gradient descent. To validate the  
 178 quality of our estimator, we optimized a quadratic function using IGT, collecting iterates  $\theta_t$ . For each  
 179 iterate, we computed the squared error between the true gradient and either the stochastic or the IGT  
 180 gradient. In this case where both estimators are unbiased, this is the trace of the noise covariance of  
 181 our estimators. The results in Figure 2b show that, as expected, this noise decreases linearly for IGT  
 182 and is constant for SGD.

183 We also analyse the direction and magnitude of the gradient of IGT on the same quadratic setup.  
 184 Figure 2c displays the cosine similarity between the true gradient and either the stochastic or the IGT  
 185 gradient, as a function of the distance to the optimum. We see that, for the same distance, the IGT  
 186 gradient is much more aligned with the true gradient than the stochastic gradient is, confirming that  
 187 variance reduction happens without the need for scaling the estimate.

## 188 6 Experiments

189 While Section 5 confirms the performance of IGT in the ideal case, the assumption of identical  
 190 Hessians almost never holds in practice. In this section, we present results on more realistic and  
 191 larger scale machine learning settings. All experiments are extensively described in the Appendix A  
 192 and additional baselines compared in Appendix B.

### 193 6.1 Supervised learning

194 **CIFAR10 image classification** We first consider the task of training a ResNet-56 model [12] on  
 195 the CIFAR-10 image classification dataset [19]. We use TF official models code and setup [1],  
 196 varying only the optimizer: SGD, HB, Adam and our algorithm with anytime tail averaging both on  
 197 its own (ITA) and combined with Heavy Ball (HB-ITA). We tuned the step size for each algorithm  
 198 by running experiments using a logarithmic grid. To factor in ease of tuning [48], we used Adam's

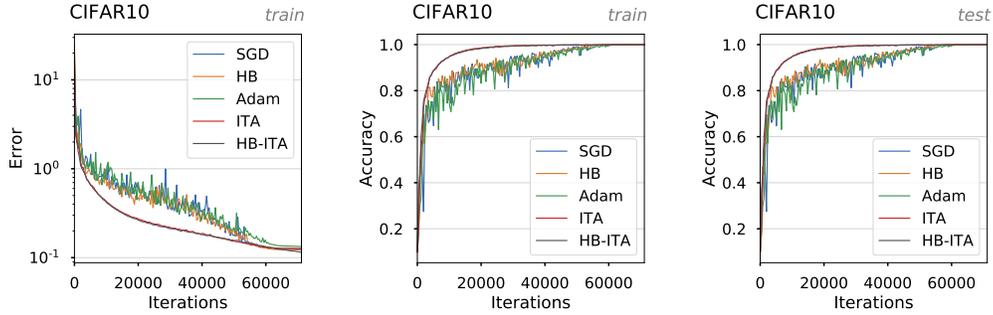


Figure 3: Resnet-56 on CIFAR10. **Left:** Train loss. **Center:** Train accuracy. **Right:** Test accuracy.

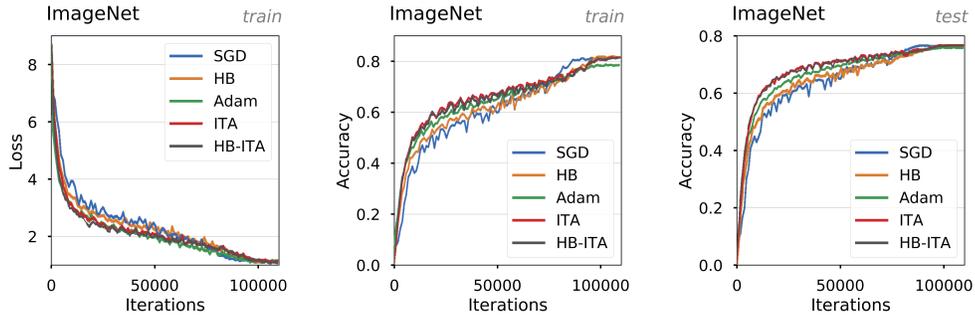


Figure 4: ResNet-50 on ImageNet. **Left:** Train loss. **Center:** Train accuracy. **Right:** Test accuracy.

199 default parameter values and a value of 0.9 for HB’s parameter. We used a linearly decreasing  
 200 stepsize as it was shown to be simple and perform well [43]. For each optimizer we selected the  
 201 hyperparameter combination that is fastest to reach a consistently attainable target train loss [43].  
 202 Selecting the hyperparameter combination reaching the lowest training loss yields qualitatively  
 203 identical curves. Figure 3 presents the results, showing that IGT with the exponential anytime tail  
 204 average performs favourably, both on its own and combined with Heavy Ball: the learning curves  
 205 show faster improvement and are much less noisy.

206 **ImageNet image classification** We also consider the task of training a ResNet-50 model[12] on  
 207 the larger ImageNet dataset [36]. The setup is similar to the one used for CIFAR10 with the difference  
 208 that we trained using larger minibatches (1024 instead of 128). In Figure 4, one can see that IGT is as  
 209 fast as Adam for the train loss, faster for the train accuracy and reaches the same final performance,  
 210 which Adam does not. We do not see the noise reduction we observed with CIFAR10, which could  
 211 be explained by the larger batch size (see Appendix A.1).

212 **IMDb sentiment analysis** We train a bi-directional LSTM on the IMDb Large Movie Review  
 213 Dataset for 200 epochs. [27] We observe that while the training convergence is comparable to HB,  
 214 HB-ITA performs better in terms of validation and test accuracy. In addition to the baseline and  
 215 IGT methods, we also train a variant of Adam using the ITA gradients, dubbed **Adam-ITA**, which  
 216 performs similarly to Adam.

## 217 6.2 Reinforcement learning

218 **Linear-quadratic regulator** We cast the classical linear-quadratic regulator (LQR) [21] as a policy  
 219 learning problem to be optimized via gradient descent. This setting is extensively described in  
 220 Appendix A. Note that despite their simple linear dynamics and a quadratic cost functional, LQR  
 221 systems are notoriously difficult to optimize due to the non-convexity of the loss landscape. [8]

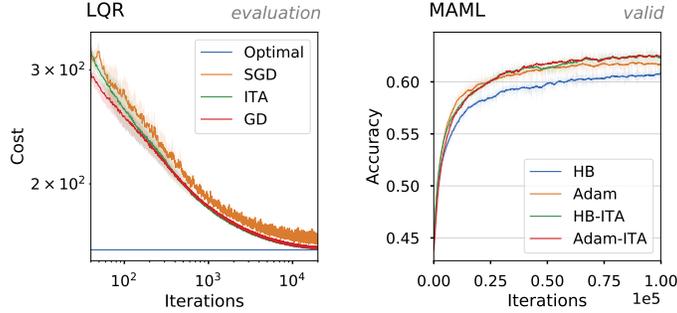


Figure 5: Validation curves for different large-scale machine learning settings. Shading indicates one standard deviation computed over three random seeds. **Left:** Reinforcement learning via policy gradient on a LQR system. **Right:** Meta-learning using MAML on Mini-Imagenet.

222 The left chart in Figure 5 displays the evaluation cost computed along training and averaged over three  
 223 random seeds. The first method (**Optimal**) indicates the cost attained when solving the algebraic  
 224 Riccati equation of the LQR – this is the optimal solution of the problem. **SGD** minimizes the costs  
 225 using the REINFORCE [47] gradient estimator, averaged over 600 trajectories. **ITA** is similar to  
 226 SGD but uses the ITA gradient computed from the REINFORCE estimates. Finally, **GD** uses the  
 227 analytical gradient by taking the expectation over the policy.

228 We make two observations from the above chart. First, ITA initially suffers from the stochastic  
 229 gradient estimate but rapidly matches the performance of GD. Notably, both of them converge to  
 230 a solution significantly better than SGD, demonstrating the effectiveness of the variance reduction  
 231 mechanism. Second, the convergence curve is smoother for ITA than for SGD, indicating that the  
 232 ITA iterates are more likely to induce similar policies from one iteration to the next. This property  
 233 is particularly desirable in reinforcement learning as demonstrated by the popularity of trust-region  
 234 methods in large-scale applications. [41, 29]

### 235 6.3 Meta-learning

236 **Model-agnostic meta-learning** We now investigate the use of IGT in the *model-agnostic meta-*  
 237 *learning* (MAML) setting. [9] We replicate the 5 ways classification setup with 5 adaptation steps on  
 238 tasks from the Mini-Imagenet dataset [34]. This setting is interesting because of the many sources  
 239 contributing to noise in the gradient estimates: the stochastic meta-gradient depends on the product  
 240 of 5 stochastic Hessians computed over only 10 data samples, and is averaged over only 4 tasks. We  
 241 substitute the meta-optimizer with each method, select the stepsize that maximizes the validation  
 242 accuracy after 10K iterations, and use it to train the model for 100K iterations.

243 The right graph of Figure 5 compares validation accuracies for three random seeds. We observe that  
 244 methods from the IGT family significantly outperform their stochastic meta-gradient counter-part,  
 245 both in terms of convergence rate and final accuracy. Those results are also reflected in the final test  
 246 accuracies where Adam-ITA (65.16%) performs best, followed by HB-ITA (64.57%), then Adam  
 247 (63.70%), and finally HB (63.08%).

## 248 7 Conclusion and open questions

249 We proposed a simple optimizer which, by reusing past gradients and transporting them, offers  
 250 excellent performance on a variety of problems. While it adds an additional parameter, the ratio of  
 251 examples to be kept in the tail averaging, it remains competitive across a wide range of such values.  
 252 Further, by providing a higher quality gradient estimate that can be plugged in any existing optimizer,  
 253 we expect it to be applicable to a wide range of problems. As the IGT is similar to momentum, this  
 254 further raises the question on the links between variance reduction and curvature adaptation. Whether  
 255 there is a way to combine the two without using momentum on top of IGT remains to be seen.

## References

- 256 [1] The TensorFlow Authors. Tensorflow official resnet model. 2018.
- 257 [2] Reza Babanezhad, Mohamed Osama Ahmed, Alim Virani, Mark Schmidt, Jakub Konečný, and  
258 Scott Sallinen. Stop wasting my gradients: Practical SVRG. In *Advances in Neural Information*  
259 *Processing Systems*, 2015.
- 260 [3] Francis Bach and Eric Moulines. Non-strongly-convex smooth stochastic approximation with  
261 convergence rate  $o(1/n)$ . In *Advances in Neural Information Processing Systems*, pages 773–781,  
262 2013.
- 263 [4] Sébastien Bubeck. Convex optimization: Algorithms and complexity. *Foundations and Trends®*  
264 *in Machine Learning*, 8(3-4):231–357, 2015.
- 265 [5] Aaron Defazio, Francis Bach, and Simon Lacoste-Julien. Saga: A fast incremental gradient  
266 method with support for non-strongly convex composite objectives. In *Advances in Neural*  
267 *Information Processing Systems*, pages 1646–1654, 2014.
- 268 [6] Aymeric Dieuleveut, Alain Durmus, and Francis Bach. Bridging the gap between constant step  
269 size stochastic gradient descent and markov chains. *arXiv preprint arXiv:1707.06386*, 2017.
- 270 [7] Aymeric Dieuleveut, Nicolas Flammarion, and Francis Bach. Harder, better, faster, stronger  
271 convergence rates for least-squares regression. *The Journal of Machine Learning Research*,  
272 18(1):3520–3570, 2017.
- 273 [8] Maryam Fazel, Rong Ge, Sham M. Kakade, and Mehran Mesbahi. Global convergence of  
274 policy gradient methods for linearized control problems. *CoRR*, abs/1801.05039, 2018.
- 275 [9] Chelsea Finn, Pieter Abbeel, and Sergey Levine. Model-agnostic meta-learning for fast adap-  
276 tation of deep networks. In *Proceedings of the 34th International Conference on Machine*  
277 *Learning-Volume 70*, pages 1126–1135. JMLR. org, 2017.
- 278 [10] Nicolas Flammarion and Francis Bach. From averaging to acceleration, there is only a step-size.  
279 In *Conference on Learning Theory*, pages 658–695, 2015.
- 280 [11] Robert M Gower, Nicolas Le Roux, and Francis Bach. Tracking the gradients using the hessian:  
281 A new look at variance reducing stochastic methods. *arXiv preprint arXiv:1710.07462*, 2017.
- 282 [12] Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. Deep residual learning for image  
283 recognition. *CoRR*, abs/1512.03385, 2015.
- 284 [13] Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. Identity mappings in deep residual  
285 networks. In *European conference on computer vision*, pages 630–645. Springer, 2016.
- 286 [14] Thomas Hofmann, Aurelien Lucchi, Simon Lacoste-Julien, and Brian McWilliams. Variance re-  
287 duced stochastic gradient descent with neighbors. In *Advances in Neural Information Processing*  
288 *Systems*, pages 2305–2313, 2015.
- 289 [15] Prateek Jain, Sham M Kakade, Rahul Kidambi, Praneeth Netrapalli, and Aaron Sidford. Ac-  
290 celerating stochastic gradient descent for least squares regression. In *Conference On Learning*  
291 *Theory*, pages 545–604, 2018.
- 292 [16] Prateek Jain, Sham M. Kakade, Rahul Kidambi, Praneeth Netrapalli, and Aaron Sidford.  
293 Parallelizing stochastic gradient descent for least squares regression: Mini-batching, averaging,  
294 and model misspecification. *Journal of Machine Learning Research*, 18(223):1–42, 2018.
- 295 [17] Rie Johnson and Tong Zhang. Accelerating stochastic gradient descent using predictive variance  
296 reduction. In *Advances in Neural Information Processing Systems*, pages 315–323, 2013.
- 297 [18] Diederik Kingma and Jimmy Ba. Adam: A method for stochastic optimization. *arXiv preprint*  
298 *arXiv:1412.6980*, 2014.
- 299 [19] Alex Krizhevsky. Learning multiple layers of features from tiny images. Technical report, 2009.
- 300

- 301 [20] Alex Krizhevsky. Learning multiple layers of features from tiny images. Technical report,  
302 Citeseer, 2009.
- 303 [21] Huibert Kwakernaak. *Linear optimal control systems*, volume 1.
- 304 [22] Simon Lacoste-Julien, Mark Schmidt, and Francis Bach. A simpler approach to obtaining  
305 an  $o(1/t)$  convergence rate for the projected stochastic subgradient method. *arXiv preprint*  
306 *arXiv:1212.2002*, 2012.
- 307 [23] Nicolas Le Roux. Anytime tail averaging. *arXiv preprint arXiv:1902.05083*, 2019.
- 308 [24] Nicolas Le Roux, Mark Schmidt, and Francis Bach. A stochastic gradient method with an  
309 exponential convergence rate for finite training sets. In *Advances in Neural Information*  
310 *Processing Systems*, pages 2663–2671, 2012.
- 311 [25] Yann LeCun and Corinna Cortes. Mnist handwritten digit database. *AT&T Labs [Online]*.  
312 Available: <http://yann.lecun.com/exdb/mnist>, 2010.
- 313 [26] Nicolas Loizou and Peter Richtárik. Momentum and stochastic momentum for stochastic gradi-  
314 ent, newton, proximal point and subspace descent methods. *arXiv preprint arXiv:1712.09677*,  
315 2017.
- 316 [27] Andrew L. Maas, Raymond E. Daly, Peter T. Pham, Dan Huang, Andrew Y. Ng, and Christopher  
317 Potts. Learning word vectors for sentiment analysis. In *Proceedings of the 49th Annual Meeting*  
318 *of the Association for Computational Linguistics: Human Language Technologies*, pages  
319 142–150, Portland, Oregon, USA, June 2011. Association for Computational Linguistics.
- 320 [28] Julien Mairal. Optimization with first-order surrogate functions. In *Proceedings of The 30th*  
321 *International Conference on Machine Learning*, pages 783–791, 2013.
- 322 [29] OpenAI, Marcin Andrychowicz, Bowen Baker, Maciek Chociej, Rafal Józefowicz, Bob McGrew,  
323 Jakub W. Pachocki, Jakub Pachocki, Arthur Petron, Matthias Plappert, Glenn Powell, Alex  
324 Ray, Jonas Schneider, Szymon Sidor, Josh Tobin, Peter Welinder, Lilian Weng, and Wojciech  
325 Zaremba. Learning dexterous in-hand manipulation. *CoRR*, abs/1808.00177, 2018.
- 326 [30] Brendan O’Donoghue and Emmanuel Candes. Adaptive restart for accelerated gradient schemes.  
327 *Foundations of computational mathematics*, 15(3):715–732, 2015.
- 328 [31] Jeffrey Pennington, Richard Socher, and Christopher Manning. Glove: Global vectors for  
329 word representation. In *Proceedings of the 2014 conference on empirical methods in natural*  
330 *language processing (EMNLP)*, pages 1532–1543, 2014.
- 331 [32] Boris T Polyak. Some methods of speeding up the convergence of iteration methods. *USSR*  
332 *Computational Mathematics and Mathematical Physics*, 4(5):1–17, 1964.
- 333 [33] Boris T Polyak and Anatoli B Juditsky. Acceleration of stochastic approximation by averaging.  
334 *SIAM Journal on Control and Optimization*, 30(4):838–855, 1992.
- 335 [34] Sachin Ravi and Hugo Larochelle. Optimization as a model for few-shot learning. 2016.
- 336 [35] Herbert Robbins and Sutton Monro. A stochastic approximation method. *Annals of Mathemati-*  
337 *cal Statistics*, 22(3):400–407, 1951.
- 338 [36] Olga Russakovsky, Jia Deng, Hao Su, Jonathan Krause, Sanjeev Satheesh, Sean Ma, Zhiheng  
339 Huang, Andrej Karpathy, Aditya Khosla, Michael Bernstein, Alexander C. Berg, and Li Fei-Fei.  
340 ImageNet Large Scale Visual Recognition Challenge. *International Journal of Computer Vision*  
341 (*IJCV*), 115(3):211–252, 2015.
- 342 [37] Mark Sandler, Andrew Howard, Menglong Zhu, Andrey Zhmoginov, and Liang-Chieh Chen.  
343 Mobilenetv2: Inverted residuals and linear bottlenecks. In *Proceedings of the IEEE Conference*  
344 *on Computer Vision and Pattern Recognition*, pages 4510–4520, 2018.
- 345 [38] Mark Schmidt. Convergence rate of stochastic gradient with constant step size. 2014.

- 346 [39] Mark Schmidt, Nicolas Le Roux, and Francis Bach. Convergence rates of inexact proximal-  
347 gradient methods for convex optimization. In *Advances in Neural Information Processing*  
348 *Systems 24*, 2011.
- 349 [40] Mark Schmidt, Nicolas Le Roux, and Francis Bach. Minimizing finite sums with the stochastic  
350 average gradient. *Mathematical Programming*, 162(1-2):83–112, 2017.
- 351 [41] John Schulman, Sergey Levine, Philipp Moritz, Michael I. Jordan, and Pieter Abbeel. Trust  
352 region policy optimization. *CoRR*, abs/1502.05477, 2015.
- 353 [42] Shai Shalev-Shwartz and Tong Zhang. Stochastic dual coordinate ascent methods for regularized  
354 loss. *Journal of Machine Learning Research*, 14(1):567–599, February 2013.
- 355 [43] Christopher J Shallue, Jaehoon Lee, Joe Antognini, Jascha Sohl-Dickstein, Roy Frostig, and  
356 George E Dahl. Measuring the effects of data parallelism on neural network training. *arXiv*  
357 *preprint arXiv:1811.03600*, 2018.
- 358 [44] Ilya Sutskever, James Martens, George Dahl, and Geoffrey Hinton. On the importance of  
359 initialization and momentum in deep learning. In *International conference on machine learning*,  
360 pages 1139–1147, 2013.
- 361 [45] Sharan Vaswani, Francis Bach, and Mark Schmidt. Fast and faster convergence of sgd for over-  
362 parameterized models and an accelerated perceptron. In *Proceedings of the 22nd International*  
363 *Conference on Artificial Intelligence and Statistics*, 2019.
- 364 [46] Hoi-To Wai, Wei Shi, Angelia Nedic, and Anna Scaglione. Curvature-aided incremental  
365 aggregated gradient method. *arXiv preprint arXiv:1710.08936*, 2017.
- 366 [47] Ronald J Williams. Simple statistical gradient-following algorithms for connectionist reinforce-  
367 ment learning. *Machine learning*, 8(3-4):229–256, 1992.
- 368 [48] Ashia C Wilson, Rebecca Roelofs, Mitchell Stern, Nathan Srebro, and Benjamin Recht. The  
369 marginal value of adaptive gradient methods in machine learning. *arXiv preprint*  
370 *arXiv:1705.08292*, 2017.
- 371 [49] Jian Zhang and Ioannis Mitliagkas. Yellowfin and the art of momentum tuning. In *SysML*, 2019.
- 372 [50] Lijun Zhang, Mehrdad Mahdavi, and Rong Jin. Linear convergence with condition number  
373 independent access of full gradients. In *Advances in Neural Information Processing Systems*,  
374 pages 980–988, 2013.

## 375 A Experimental Details

376 This section provides additional information regarding the experiments included in the main text.

377 For each experimental setting we strive to follow the *reproducibility checklist*, and provide:

- 378 • a description and citation of the dataset,
- 379 • a description of pre-processing steps,
- 380 • training / validation / testing splits,
- 381 • a description of the hyper-parameter search process and chosen values for each method,
- 382 • the exact number of evaluation runs,
- 383 • a clear definition of statistics reported, and
- 384 • a description of the computing infrastructure.

### 385 A.1 CIFAR10 image classification

386 **Dataset** The CIFAR10 dataset [20] consists 50k training and 10k testing images, partitioned over 10  
387 classes. We download and pre-process the images using the TensorFlow `models` package, available  
388 at the following URL: <https://github.com/tensorflow/models>

389 **Model** We use a residual convolutional network [12] with 56 layers as defined in the `models`  
390 package. Specifically, we use the second version whose blocks are built as a batch normalization,  
391 then a ReLU activation, and then a convolutional layer. [13]

392 **Hyper-parameters** We use the exact setup from <https://github.com/tensorflow/models/officials/resnet>. As such, training is carried out with minibatches of 128 examples for 182  
393 epochs and the training data is augmented with random crops and horizontal flips. Also note this  
394 setup multiplies the step size by the size of the minibatch. One deviation from the setup is our use  
395 of a linearly decaying learning rate instead of an explicit schedule. The linearly decaying learning  
396 rate schedule is simple and was shown to perform well [43]. This schedule is specified using two  
397 parameters: the decay rate, a multiplier specifying the final step size (0.1 or 0.01), and the decay step,  
398 specifying the step at which the fully decayed rate is reached (always set to 90% of the training steps).  
399 To factor in ease of tuning[48] we used Adam’s default parameter values and a value of 0.9 for HB’s  
400 parameter. We used IGT with the exponential Anytime Tail Averaging approach. For the tail fraction,  
401 we tried two values: the number of epochs and a tenth of that number (180 and 18). We ran using the  
402 following learning rate: (1e0, 3e-1, 1e-1, 3e-2, 1e-2) for SGD, HB and the IGT variants and (1e-2,  
403 3e-3, 1e-3, 3e-4, 1e-4) for Adam. We ran a grid search over the base learning rate and its decay rate  
404 with a single run per combination. For each optimizer we selected the hyperparameter combination  
405 that is fastest to reach a consistently attainable target train loss of 0.2 [43]. Note that selecting the  
406 hyperparameter combination reaching the lowest training loss yields qualitatively identical curves.  
407

408 The resulting hyper-parameters are:

- 409 • SGD stepsize 0.3, decay 0.01
- 410 • HB stepsize 0.03, decay 0.01
- 411 • Adam stepsize 0.001, decay 0.01
- 412 • ITA stepsize 0.3, decay 0.01, tail fraction 18
- 413 • HB-ITA stepsize 0.03, decay 0.1, tail fraction 18

414 **Infrastructure and Runs** The experiments were run using P100 GPUs (single GPU).

415 **Additional Results** We provide all learning curves for the methods comparison presented in the  
416 main manuscript in figure 6.

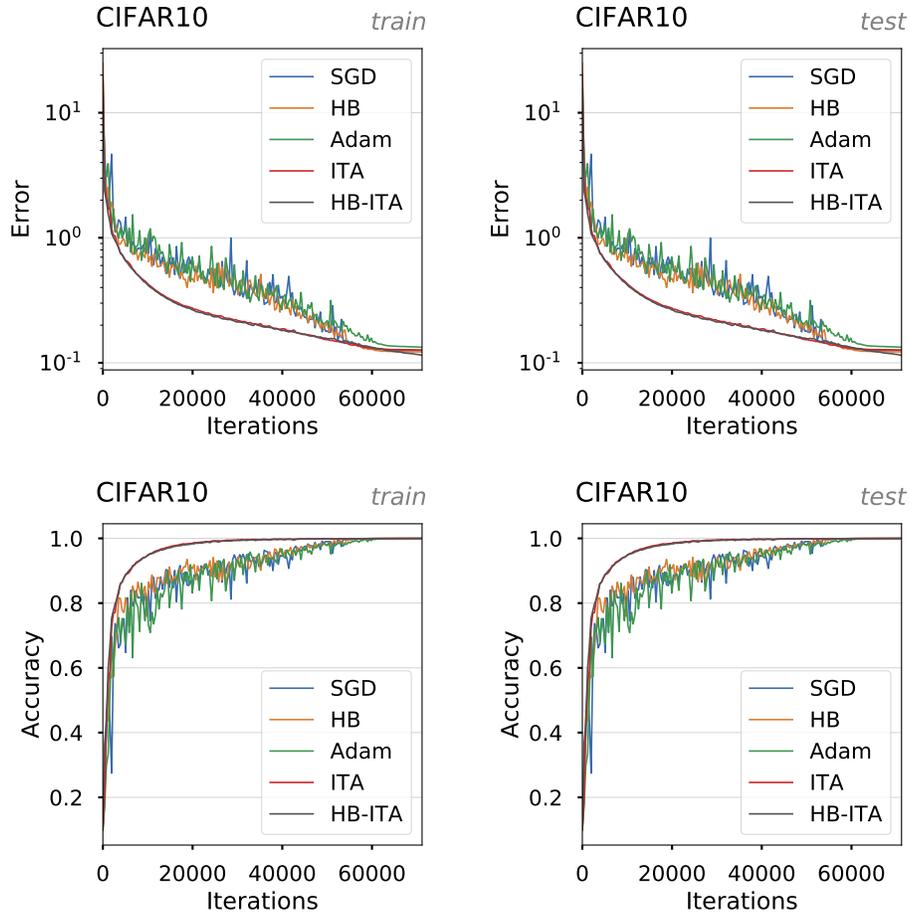


Figure 6: Convergence and accuracy curves along training for the CIFAR10 experiments comparing baseline methods to ours. **Left:** Training. **Right:** Testing.

417 **Ablation study: importance of IGT correction** We confirm the importance of the implicit gradi-  
 418 ent transport correction by running an experiment in which an increasing momentum is used without  
 419 transport. The results appear in figure 7.

420 The resulting hyper-parameters are:

421 • ATA stepsize 0.3, decay 0.01, tail fraction 18

422 • HB-ATA stepsize 0.03, decay 0.01, tail fraction 18

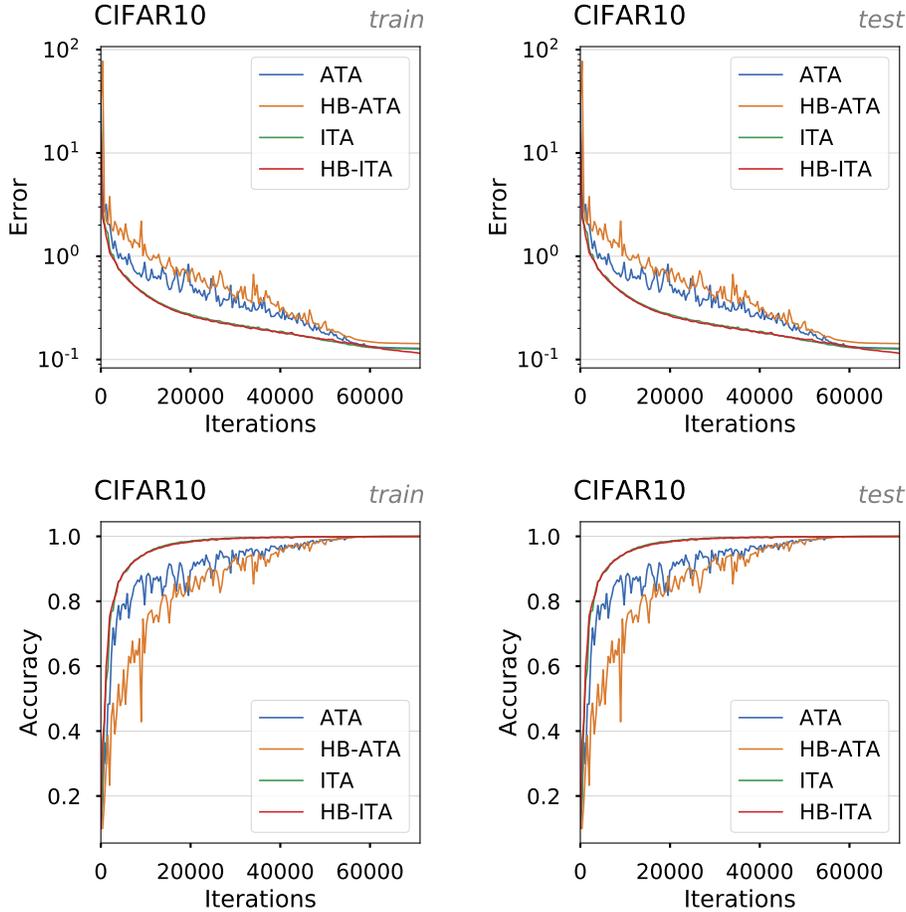


Figure 7: Convergence and accuracy curves along training for the CIFAR10 experiments comparing the use of ATA combined with our proposed implicit transport mechanism. **Left:** Training. **Right:** Testing.

423 **Effect of the batch size** We look into the effect of the batch size. To do so, we plot the number of  
 424 steps required to reach a reliably attainable training loss of 0.4 as a function of the batch size. We ran  
 425 using the following mini-batch sizes: 32, 128, 512 and 2048. Running with larger minibatches led to  
 426 out of memory errors on our single GPU setup. The results presented in figure 8 show the benefit of  
 427 IGT lowers as the batch size increases. Note that Adam’s ability to keep benefiting from larger batch  
 428 sizes is consistent with previous observations.

429 **A.2 ImageNet image classification**

430 **Dataset** We use the 2015 edition of the *ImageNet Large-Scale Visual Recognition Challenge*  
 431 (ImageNet) [36] dataset. This dataset consists of 1.2M images partitioned into 1’000 classes. We use  
 432 the pre-processing and loading utilities of the TensorFlow models package, available at the following  
 433 URL: <https://github.com/tensorflow/models>

434 **Model** Our model is again a large residual network, consisting of 50 layers. Similar to our CIFAR10  
 435 experiments above, we use the implementation described in [13].

436 **Hyper-parameters** We used the same setup and approach as for the CIFAR-10 experiments. The  
 437 setup trains for 90 epochs using mini-batches of 1024 examples. We used a larger grid for the learning  
 438 rate schedule: decay (0.1, 0.01, 0.001) and decay step fraction (0.7, 0.8, 0.9).

439 The resulting hyper-parameters are:

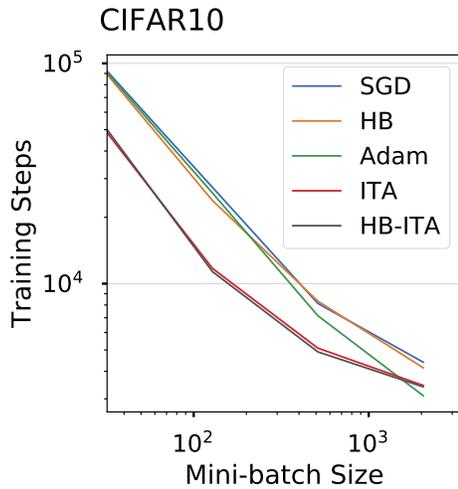


Figure 8: Effect of mini-batch size on the number of steps to reach a target training loss.

- 440      • SGD stepsize 0.3, decay 0.01, decay step 0.8
- 441      • HB stepsize 0.03, decay 0.001, decay step 0.9
- 442      • Adam stepsize 0.0001, decay 0.01, decay step 0.9
- 443      • ITA stepsize 0.3, decay 0.01, tail fraction 90, decay step 0.9
- 444      • HB-ITA stepsize 0.03, decay 0.01, tail fraction 90, decay step 0.9

445 **Infrastructure and Runs** We ran these experiments using Google TPUv2.

446 **Additional Results** We include error and accuracy curves for training and testing in Figure 9.

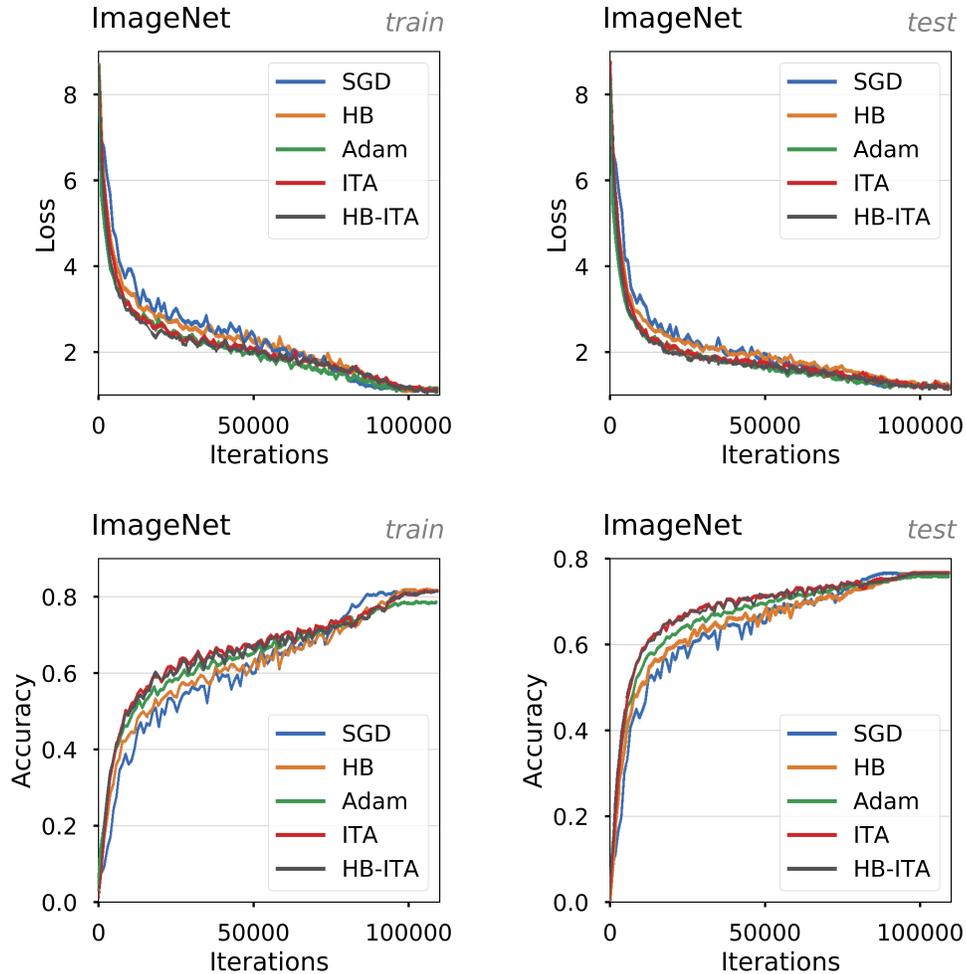


Figure 9: Convergence and accuracy curves along training for the ImageNet experiments. **Left:** Training. **Right:** Testing.

### 447 A.3 IMDb sentiment analysis

448 **Dataset** The *Internet Movie Database* (IMDb) [27] consists of 25’000 training and 25’000 test  
 449 movie reviews. The objective is binary sentiment classification based on the review’s text. We  
 450 randomly split the training set in two folds of 17’536 and 7’552 reviews, the former being used for  
 451 training and the latter for testing. The data is downloaded, splitted, and pre-processed with `torchtext`  
 452 package, available at the following URL: <https://github.com/pytorch/text> More specifically,  
 453 we tokenize the text at the word-level using the `spaCy` package, and embed the tokens using the  
 454 100-dimensional GloVe 6B [31] distributed representations.

455 **Model** The model consists of an embedding lookup-table, followed by a bi-directional LSTM with  
 456 dropout, and then by a fully-connected layer. The LSTM uses 256 hidden units and the dropout rate is  
 457 set to 0.5. The whole model consists of 3.2M trainable parameters, with the embedding lookup-table  
 458 initialized with the GloVe vectors. The model is trained to minimize the `BCEWithLogitsLoss` with a  
 459 mini-batch size of 64.

460 **Hyper-parameters** For each method, we used a grid-search to find the stepsize minimizing vali-  
 461 dation error after 15 epochs. The grid starts at 0.00025 and doubles until reaching 0.1024, so as to  
 462 ensure that no chosen value lies on its boundaries. When applicable, the momentum factor is jointly

463 optimized over values 0.1 to 0.95. The final hyper-parameters are displayed in the following table for  
464 each method.

Table 1: Hyperparameters for IMDb experiments.

	HB	Adam	ASGD	HB-IGT	HB-ITA
$\alpha$	0.032	0.0005	0.064	0.128	0.064
$\mu$	0.95	0.95	n/a	0.9	0.9
$\xi$	n/a	n/a	100	n/a	n/a
$\kappa$	n/a	n/a	$10^5$	n/a	n/a

465 **Infrastructure and Runs** All IMDb experiments use a single NVIDIA GTX 1080, with PyTorch  
466 v0.3.1.post2, CUDA 8.0, and cuDNN v7.0.5. We run each final configurations with 5 different random  
467 seeds and always report the mean tendency  $\pm$  one standard deviation. Each run lasts approximately  
468 three hours and thirty minutes.

469 **Additional Results** In addition to the results reported in the main text, we include training, valida-  
470 tion, and testing curves for each method in Figure 10. Shading indicates the one standard deviation  
471 interval. Note that our focus is explicitly on optimization: in the specific case of IMDb, training for  
472 200 epochs is completely unnecessary from a generalization standpoint as performance degrades  
473 rapidly after 15-20 epochs.

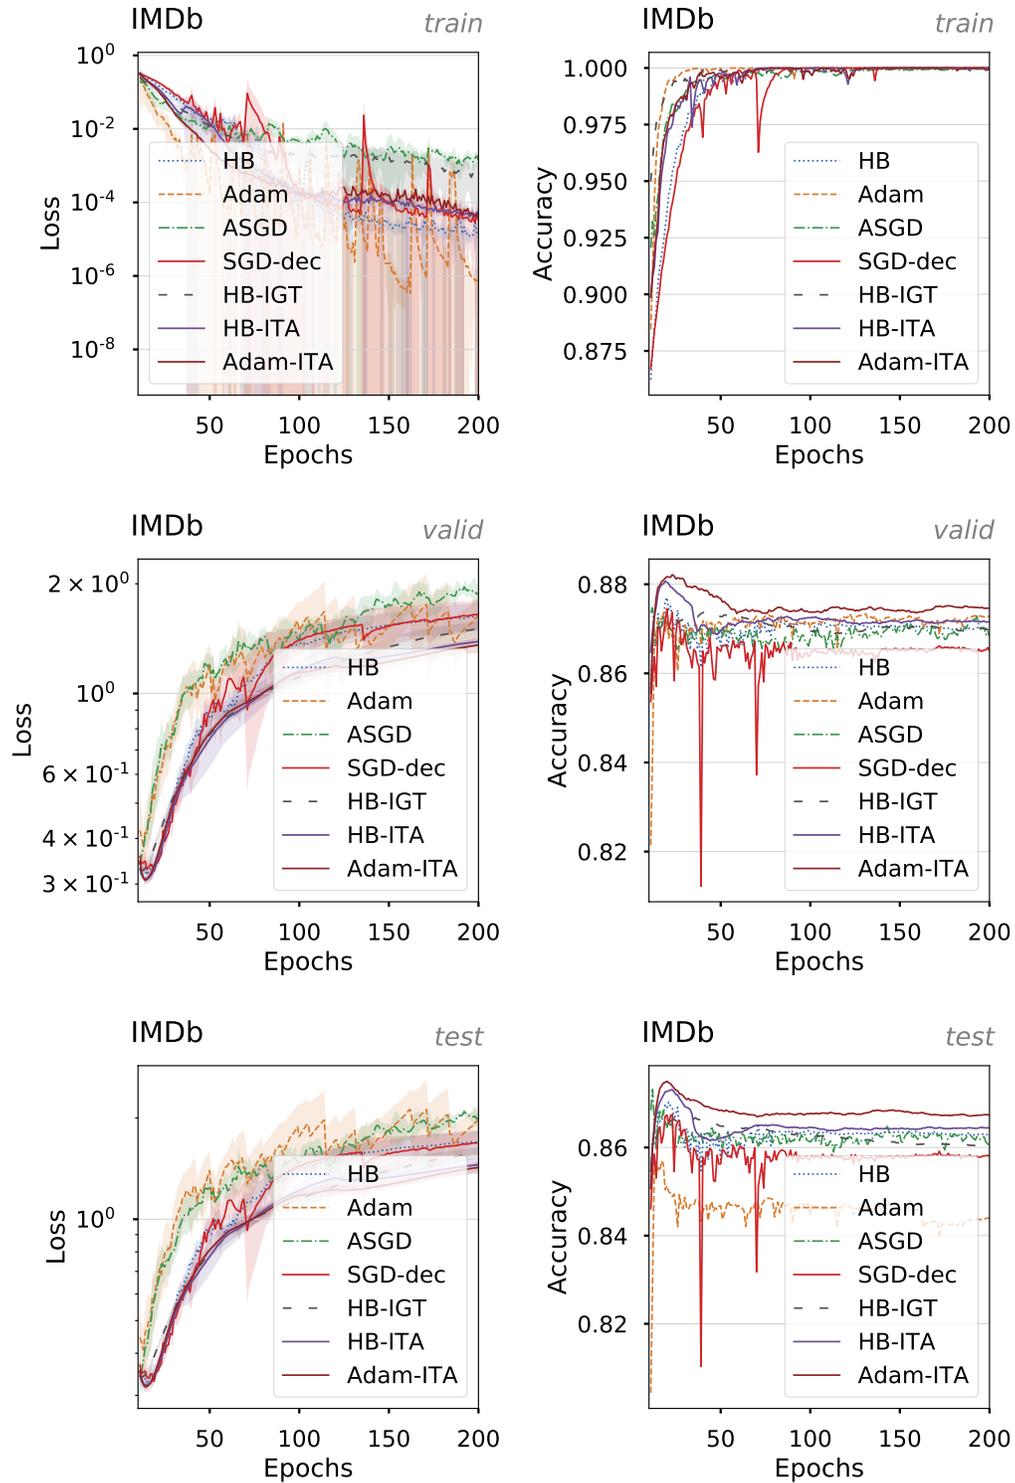


Figure 10: Convergence and accuracy curves along training for the IMDB experiments. **Left:** Convergence. **Right:** Accuracy.

474 **A.4 Linear-quadratic regulator**

475 **Setup** Our linear-quadratic regulator [21] implements the following equations. The cost functional  
 476 is evaluated at every timestep  $h$  and is given by

$$C(s_h, a_h) = s_h^\top Q s_h + a_h^\top R a_h, \quad (2)$$

477 for random symmetric positive definite matrices  $Q \in \mathbb{R}^{20 \times 20}$  and  $R \in \mathbb{R}^{12 \times 12}$  each with condition  
 478 number 3. The initial state  $s_0 \sim \mathcal{N}(0, 3 \cdot I_{20})$  is sampled around the origin, and the subsequent states  
 479 evolve according to

$$s_{h+1} = A s_h + B a_h, \quad (3)$$

480 where entries of  $A \in \mathbb{R}^{20 \times 20}$ ,  $B \in \mathbb{R}^{20 \times 12}$  are independently sampled from a Normal distribution  
 481 and then scaled such that the matrix has unit Frobenius norm. The actions are given by the linear  
 482 stochastic policy  $a_h = K s_h + \epsilon_h^a$ , where  $\epsilon_h^a \sim \mathcal{N}(0, I)$  and  $K$  are the parameters to be optimized.

483 Gradient methods in this manuscript optimize the sum of costs using the REINFORCE estimate [47]  
 484 given by

$$\nabla_K \mathbb{E} \sum_h^{10} C(s_h, a_h) = \mathbb{E} \left( \sum_h^{10} \nabla_K \log \pi_K(a_h | s_h) \right) \left( \sum_h^{10} C(s_h, a_h) \right). \quad (4)$$

485 In our experiments, the above expectation is approximated by the average of 600 trajectory rollouts.  
 486 Due to the noisy dynamics of the system, it is possible for the gradient norm to explode leading to  
 487 numerical instabilities – especially when using larger stepsizes. To remedy this issue, we simply  
 488 discard such problematic trajectories from the gradient estimator.

489 For each training iteration, we first gather 600 trajectories used for learning and then 600 more used  
 490 to report evaluation metrics.

491 **Hyper-parameters** Due to the simplicity of the considered methods, the only hyper-parameter is  
 492 the stepsize. For each method, we choose the stepsize from a logarithmically-spaced grid so as to  
 493 minimize the evaluation cost after 600 iterations on a single seed. Incidentally, the optimal stepsize  
 494 for GD, SGD, and ITA is 0.0002.

495 **Infrastructure and Runs** We use an Intel Core i7-5820K CPU to run the LQR experiments. All  
 496 methods are implemented using numpy v1.15.4. We present results averaged over 3 random seeds,  
 497 and also report the standard deviation. For stochastic gradient methods (SGD, ITA) training for 20K  
 498 iterations takes about 3h, for full-gradient method (GD) about 10h, and computing the solution of the  
 499 Riccati equation takes less than 5 seconds.

500 **Additional Results** In addition to the evaluation cost reported in the main text, we also include the  
 501 cost witnessed during training (and used for optimization) in Figure 11.

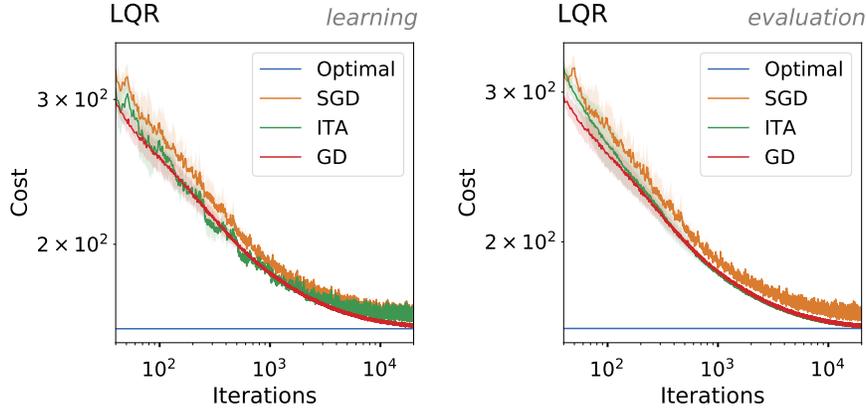


Figure 11: LQR costs along training iterations. **Left:** Costs used for learning. **Right:** Costs used for evaluation.

502 We notice that the training cost curve of ITA is not as smooth as the evaluation one. Similarly, the  
 503 observed learning costs never reach as good a minima as the evaluation ones. This phenomena  
 504 is easily clarified: during learning, ITA estimates the gradient using the shifted parameters  $K_t +$   
 505  $\frac{\gamma_t}{1-\gamma_t}(K_t - K_{t-1})$  as opposed to the true parameters  $K_t$ . Those shifted parameters are not subject to a  
 506 reduced variance, hence explaining the observed noise in the cost as well as deteriorated convergence.

#### 507 A.5 Model-agnostic meta-learning

508 **Dataset** We use the Mini-Imagenet dataset [34] in our model-agnostic meta-learning (MAML) [9]  
 509 experiments. This dataset comprises 64 training, 12 validation, and 24 test classes. For each of train,  
 510 validation, and test sets, tasks are constructed by sampling 5 classes from their respective split, and  
 511 further sampling 5 images per class. Images are downsampled to 84x84x3 tensors of RGB values.  
 512 For more details, please refer to the official code repository – which we carefully replicated – at the  
 513 following URL: <https://github.com/cbfinn/maml>

514 Our implementation departs in two ways from the reference. First, we train our models for 100k  
 515 iterations as opposed to 60k and only use 5 image samples to compute a meta-gradient whereas the  
 516 reference implementation uses 15. Second, we only use 5 adaptation steps at evaluation time, when  
 517 the reference uses 10.

518 **Model** The model closely replicates the convolutional neural network of MAML [9]. It consists  
 519 of 4 layers, each with 32 3x3 kernels, followed by batch normalization and ReLU activations. For  
 520 specific implementation details, we refer the reader to the above reference implementation.

521 **Hyper-parameters** We only tune the meta-stepsize for the MAML experiment. We set the mo-  
 522 mentum constant to 0.9, the adaptation-stepsize to 0.01, and average the meta-gradient of 4 tasks per  
 523 iterations. Due to the reduced variance in the gradients, we found it necessary to increase the  $\epsilon$  of  
 524 Adam-ITA to 0.01.

525 For each method, we search over stepsize values on a logarithmically-spaced grid and select those  
 526 values that maximize validation accuracy after 10k meta-iterations. These values are reported in  
 527 Table 2.

	HB	Adam	HB-ITA	Adam-ITA
$\alpha$	0.008	0.001	0.008	0.0005

Table 2: Stepsizes for MAML experiments.

528 **Infrastructure and Runs** Each MAML experiment is run on a single NVIDIA GTX TITAN X,  
 529 with PyTorch v1.1.0, CUDA 8.0, and cuDNN v7.0.5. We run each configuration with 3 different  
 530 random seeds and report the mean tendency  $\pm$  one standard deviation. Each run takes approximately  
 531 36 hours, and we evaluate the validation and testing accuracy every 100 iteration.

532 **Additional Results** We complete the MAML validation curves from the main manuscript with  
 533 training and testing accuracy curves in Figure 12. Moreover, we recall the final test accuracies for  
 534 each method: Adam-ITA reaches 65.16%, HB-ITA 64.57%, Adam 63.70%, and HB 63.08%.

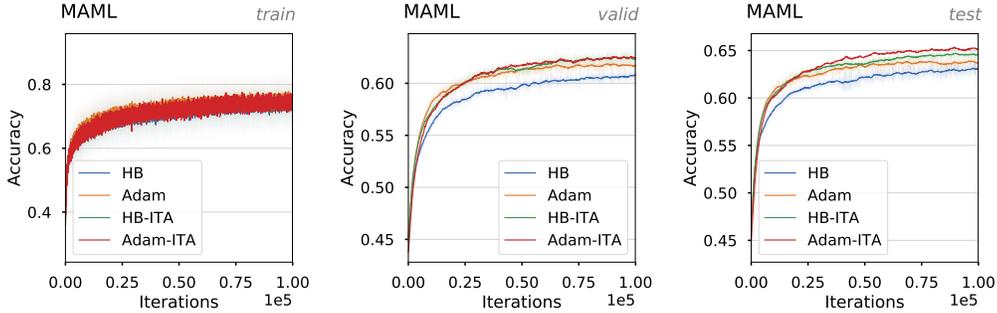


Figure 12: Training, validation, and testing accuracies for the MAML experiments along training. Shading indicates the 1 standard deviation interval. **Left:** Training. **Center:** Validation. **Right:** Testing.

## 535 B Additional Experiments

536 This section presents additional experiments to the ones included in the main text.

### 537 B.1 Baselines comparisons

538 While experiments in Section 5 highlighted properties of IGT and HB-IGT when the assumption of  
 539 identical, constant Hessians was verified, we now turn to more realistic scenarios where individual  
 540 functions are neither quadratic nor have the same Hessian to compare our proposed methods against  
 541 popular baselines for the online stochastic optimization setting. We target optimization benchmarks  
 542 and focus on training loss minimization.

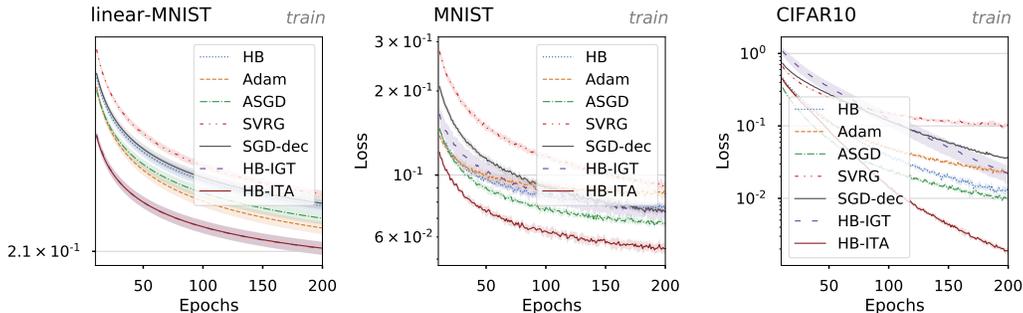


Figure 13: Training loss curves for different optimization algorithms on several popular benchmarks. For each method, the hyper-parameters are tuned to minimize the training error after 15 epochs. Algorithms using the IGT gradient estimates tend to outperform their stochastic gradient counter-parts. **Left:** Logistic regression on MNIST. **Center:** LeNet5 on MNIST. **Right:** MobileNetv2 on CIFAR10.

543 We investigate three different scenarios: (a) **linear-mnist**: a logistic regression model on MNIST,  
 544 (b) **mnist**: a modified version of LeNet5 [25] on MNIST and (c) **cifar-small**: the MobileNet2

	Linear-MNIST	MNIST	CIFAR10	IMDb
Heavyball	$92.52 \pm 0.04$	$99.08 \pm 0.07$	$91.55 \pm 0.25$	$86.90 \pm 0.67$
Adam	$92.57 \pm 0.10$	$98.99 \pm 0.05$	$89.36 \pm 0.75$	$85.62 \pm 0.63$
ASGD	$92.47 \pm 0.08$	$99.15 \pm 0.07$	$91.45 \pm 0.20$	$87.31 \pm 0.21$
SVRG	$92.51 \pm 0.04$	$99.06 \pm 0.08$	$86.84 \pm 0.17$	n/a
SGD-dec	$92.52 \pm 0.06$	$99.11 \pm 0.06$	$87.53 \pm 0.23$	$86.73 \pm 0.34$
Heavyball-IGT	$92.47 \pm 0.10$	$99.00 \pm 0.05$	$12.05 \pm 0.21$	$86.61 \pm 0.23$
Heavyball-ITA	$92.50 \pm 0.10$	$99.19 \pm 0.02$	$90.37 \pm 0.31$	$87.26 \pm 0.24$

Table 3: Test accuracies from the best validation epoch.

545 architecture [37] consisting of 19 convolutional layers on CIFAR10. All models are trained with a  
546 mini-batch size of 64, while the remaining hyper-parameters are available in Tables 4, 5, and 6.

547 For each of the above tasks, models are trained for 200 epochs. We compare the following methods:

- 548 • **HB**: the heavy ball method [33],
- 549 • **Adam** [18],
- 550 • **ASGD** [15],
- 551 • **SVRG** [17],
- 552 • **SGD-dec**: stochastic gradient method with an exponential learning rate schedule and  
553 exponential constant 0.999,
- 554 • **HB-IGT**: the heavy ball using the IGT as a plug-in estimator, and
- 555 • **HB-ITA**: same as HB-IGT but using the anytime tail averaging to forget the oldest gradients.

556 The hyperparameters of each method, and in particular the stepsize, are tuned independently according  
557 to a logarithmic grid so as to minimize the mean training error after epoch 15 on one seed. We  
558 then use those parameters on 5 random seeds and report the mean and standard deviation of the  
559 performance.

560 Figure 13 shows the training curves for the five algorithms in the three settings.

561 First, we observe that, for the logistic regression, HB-IGT performs on par with HB-ITA and far  
562 better than all the other methods, even though the assumption on the Hessians is violated. When  
563 using a ConvNet, however, we see that HB-IGT is not competitive with state-of-the-art methods such  
564 as Adam or ASGD. HB-ITA, on the other hand, with its smaller reliance on the assumption, once  
565 again performs much better than all other methods. In fact, HB-ITA not only converges to a lower  
566 final train error but also has a faster initial rate.

567 While our focus is on optimization, we also report generalization metrics in Table 3. For each  
568 algorithm, we computed the best mean accuracy after each epoch on the test set and report this value  
569 together with its standard deviation. The importance of the Anytime Tail-Averaging mechanism is  
570 again apparent: without it, Heavyball-IGT is unable to improve for more than a few epochs on the  
571 CIFAR10 validation set, regardless of the stepsize choice. On the other hand, it is evident from those  
572 results that the solutions found by Heavyball-ITA are competitive with the ones discovered by other  
573 optimization algorithms.

	HB	Adam	ASGD	HB-IGT	HB-ITA
$\alpha$	0.0128	0.0002	0.0032	0.0032	0.0016
$\mu$	0.1	0.95	n/a	0.9	0.1
$\xi$	n/a	n/a	10	n/a	n/a
$\kappa$	n/a	n/a	$10^4$	n/a	n/a

Table 4: Hyperparameters for linear-MNIST experiments.

	HB	Adam	ASGD	HB-IGT	HB-ITA
$\alpha$	0.0064	0.0016	0.0128	0.0032	0.0032
$\mu$	0.9	0.95	n/a	0.95	0.95
$\xi$	n/a	n/a	10	n/a	n/a
$\kappa$	n/a	n/a	$10^4$	n/a	n/a

Table 5: Hyperparameters for MNIST experiments.

	HB	Adam	ASGD	HB-IGT	HB-ITA
$\alpha$	0.0512	0.0512	0.1024	0.0128	0.0512
$\mu$	0.95	0.9	n/a	0.9	0.1
$\xi$	n/a	n/a	100	n/a	n/a
$\kappa$	n/a	n/a	$10^5$	n/a	n/a

Table 6: Hyperparameters for MobileNetV2 on CIFAR10 experiments.

## 574 C Proofs

### 575 C.1 Transport formula

$$\begin{aligned}
g_t(\theta_t) &= \frac{t}{t+1}g_{t-1}(\theta_t) + \frac{1}{t+1}g(\theta_t, x_t) \\
&= \frac{t}{t+1}(g_{t-1}(\theta_{t-1}) + H(\theta_t - \theta_{t-1})) + \frac{1}{t+1}g(\theta_t, x_t) && \text{(Quadratic } f) \\
&= \frac{t}{t+1}g_{t-1}(\theta_{t-1}) + \frac{1}{t+1}(g(\theta_t, x_t) + tH(\theta_t - \theta_{t-1})) \\
&= \frac{t}{t+1}g_{t-1}(\theta_{t-1}) + \frac{1}{t+1}g(\theta_t + t(\theta_t - \theta_{t-1}), x_t) && \text{(Identical Hessians)} \\
&\approx \frac{t}{t+1}\widehat{g}_{t-1}(\theta_{t-1}) + \frac{1}{t+1}g(\theta_t + t(\theta_t - \theta_{t-1}), x_t). && (\widehat{g}_{t-1} \text{ is an approximation})
\end{aligned}$$

## 576 D Proof of Prop. 3.1

577 In this proof, we assume that  $g$  is a strongly-convex quadratic function with hessian  $H$ .

578 At timestep  $t$ , we have access to a stochastic gradient  $g(\theta, x_t) = g(\theta_t) + \epsilon_t$  where the  $\epsilon_t$  are i.i.d.  
579 with variance  $C \preceq \sigma^2 H$ .

580 We first prove a simple lemma:

581 **Lemma D.1.** *If  $v_0 = g(\theta_0) + \epsilon_0$  and, for  $t > 0$ , we have*

$$v_t = \frac{t}{t+1}v_{t-1} + \frac{1}{t+1}g(\theta_t + t(\theta_t - \theta_{t-1})) + \frac{1}{t+1}\epsilon_t,$$

582 *then*

$$v_t = g(\theta_t) + \frac{1}{t+1} \sum_{i=0}^t \epsilon_i.$$

583 *Proof.* Per our assumption, this is true for  $t = 0$ . Now let us prove the result by induction. Assume  
 584 this is true for  $t - 1$ . Then we have:

$$\begin{aligned}
 v_t &= \frac{t}{t+1}v_{t-1} + \frac{1}{t+1}g(\theta_t + t(\theta_t - \theta_{t-1})) + \frac{1}{t+1}\epsilon_t \\
 &= \frac{t}{t+1}g(\theta_{t-1}) + \frac{1}{t+1}\sum_{i=0}^{t-1}\epsilon_i \\
 &\quad + \frac{1}{t+1}g(\theta_t + t(\theta_t - \theta_{t-1})) + \frac{1}{t+1}\epsilon_t \quad (\text{recurrence assumption}) \\
 &= \frac{t}{t+1}g(\theta_{t-1}) + \frac{1}{t+1}\sum_{i=0}^{t-1}\epsilon_i \\
 &\quad + g(\theta_t) - \frac{t}{t+1}g(\theta_{t-1}) + \frac{1}{t+1}\epsilon_t \quad (g \text{ is quadratic}) \\
 &= g(\theta_t) + \frac{1}{t+1}\sum_{i=0}^t\epsilon_i .
 \end{aligned}$$

585 This concludes the proof. □

586 **Lemma D.2.** Let us assume we perform the following iterative updates:

$$\begin{aligned}
 v_t &= \frac{t}{t+1}v_{t-1} + \frac{1}{t+1}g(\theta_t + t(\theta_t - \theta_{t-1})) + \frac{1}{t+1}\epsilon_t \\
 \theta_{t+1} &= \theta_t - \alpha v_t ,
 \end{aligned}$$

587 starting from  $v_0 = g(\theta_0) + \epsilon_0$ . Then, denoting  $\Delta_t = \theta_t - \theta^*$ , we have

$$\Delta_t = (I - \alpha H)^t \Delta_0 - \alpha \sum_{i=0}^{t-1} N_{i,t} \epsilon_i$$

588 with

$$\begin{aligned}
 N_{i,0} &= 0 \\
 N_{i,t} &= (I - \alpha H)N_{i,t-1} + 1_{i < t} \frac{1}{t} I .
 \end{aligned}$$

589 *Proof.* The result is true for  $t = 0$ . We now prove the result for all  $t$  by induction. Let us assume this  
 590 is true for  $t - 1$ . Using Lemma D.1, we have

$$v_{t-1} = g(\theta_{t-1}) + \frac{1}{t} \sum_{i=0}^{t-1} \epsilon_i$$

591 and thus, using  $g(\theta_{t-1}) = H\Delta_{t-1}$ ,

$$\begin{aligned}
 \Delta_t &= \Delta_{t-1} - \alpha v_{t-1} \\
 &= \Delta_{t-1} - \alpha H \Delta_{t-1} - \frac{\alpha}{t} \sum_{i=0}^{t-1} \epsilon_i \\
 &= (I - \alpha H) \Delta_{t-1} - \frac{\alpha}{t} \sum_{i=0}^{t-1} \epsilon_i \\
 &= (I - \alpha H)^t \Delta_0 - \alpha \sum_{i=0}^{t-2} (I - \alpha H) N_{i,t-1} \epsilon_i - \frac{\alpha}{t} \sum_{i=0}^{t-1} \epsilon_i \quad (\text{recurrence assumption}) \\
 &= (I - \alpha H)^t \Delta_0 - \alpha \sum_{i=0}^{t-1} N_{i,t} \epsilon_i
 \end{aligned}$$

592 with

$$N_{i,t} = (I - \alpha H)N_{i,t-1} + 1_{i < t} \frac{1}{t} I.$$

593 This concludes the proof.  $\square$

594 For the following lemma, we will assume that the Hessian is diagonal and will focus on one dimension  
 595 with eigenvalue  $h$ . Indeed, we know that there are no interactions between the eigenspaces and that  
 596 we can analyze each of them independently [30].

597 **Lemma D.3.** Denote  $r_h = 1 - \alpha h$ . We assume  $\alpha \leq \frac{1}{L}$ . Then, for any  $i$  and any  $t$ , we have

$$N_{i,t} \geq 0 \quad (\text{Positivity})$$

$$N_{i,t} = 0 \quad \text{if } t \leq i \quad (\text{Zero-start})$$

$$N_{i,t} \leq \log \left( \frac{2}{i(1-r_h)} \right) \quad \text{if } i < t \leq \frac{2}{1-r_h} \quad (\text{Constant bound})$$

$$N_{i,t} \leq \frac{\max \left\{ 1 + r_h, 2 \log \left( \frac{2}{i(1-r_h)} \right) \right\}}{t(1-r_h)} \quad \text{if } \frac{2}{1-r_h} \leq t. \quad (\text{Decreasing bound})$$

598 *Proof.* The Zero-start case  $i \geq t$  is immediate from the recursion of Lemma D.2. The Positivity  
 599 property of  $N_{i,t}$  is also immediate from the recursion since the stepsize  $\alpha$  is such that  $r_h = 1 - \alpha h$  is  
 600 positive.

601 We now turn to the Constant bound property. We have, for  $t > i$ ,

$$\begin{aligned} N_{i,t} &= r_h N_{i,t-1} + \frac{1}{t} \\ &\leq N_{i,t-1} + \frac{1}{t}. \end{aligned}$$

602 Thus,  $N_{i,t} - N_{i,t-1} \leq \frac{1}{t}$ . Summing these inequalities, we get a telescopic sum and, finally:

$$\begin{aligned} N_{i,t} &\leq \sum_{j=i+1}^t \frac{1}{j} \\ &\leq \int_{x=i}^t \frac{dx}{x} \\ &= \log \left( \frac{t}{i} \right). \end{aligned}$$

603 This bound is trivial in the case  $i = 0$ . In that case, we keep the first term in the sum separate and get

$$N_{0,t} \leq 1 + \log t.$$

604 In the remainder, we shall keep the  $\log \left( \frac{t}{i} \right)$  bound for simplicity. The upper bound on the right-hand  
 605 size is increasing with  $t$  and its value for  $t = \frac{2}{1-r_h}$  is thus an upper bound for all smaller values of  $t$ .

606 Replacing  $t$  with  $\frac{2}{1-r_h}$  leads to

$$\begin{aligned} N_{i, \frac{2}{1-r_h}} &\leq \log \left( \frac{\frac{2}{1-r_h}}{i} \right) \\ &= \log \left( \frac{2}{i(1-r_h)} \right). \end{aligned}$$

607 This proves the third inequality.

608 We shall now prove the Decreasing bound by induction. This bound states that, for  $t$  large enough,  
 609 each  $N_{i,t}$  decreases as  $O(1/t)$ . Using the second and third inequalities, we have

$$\begin{aligned} N_{i, \frac{2}{1-r_h}} &\leq \log \left( \frac{2}{i(1-r_h)} \right) \frac{\frac{2}{1-r_h}}{\frac{2}{1-r_h}} \\ &= \frac{\log \left( \frac{2}{i(1-r_h)} \right)}{1-r_h} \frac{2}{\frac{2}{1-r_h}} \\ &\leq \frac{\max \left\{ 1+r_h, 2 \log \left( \frac{2}{i(1-r_h)} \right) \right\}}{\frac{2}{1-r_h}(1-r_h)}. \end{aligned}$$

610 The maximum will help us prove the last property. Thus, for  $t = \frac{2}{1-r_h}$ , we have

$$\begin{aligned} N_{i,t} &\leq \frac{\max \left\{ 1+r_h, 2 \log \left( \frac{1}{i(1-r_h)} \right) \right\}}{t(1-r_h)} \\ &\leq \frac{\nu_i}{t}, \end{aligned}$$

611 with  $\nu_i = \frac{\max \left\{ 1+r_h, 2 \log \left( \frac{1}{i(1-r_h)} \right) \right\}}{(1-r_h)}$ . The Decreasing bound is verified for  $t = \frac{2}{1-r_h}$ .

612 We now show that if, for any  $t > \frac{2}{1-r_h}$ , we have  $N_{i,t-1} \leq \frac{\nu_i}{t-1}$ , then  $N_{i,t} \leq \frac{\nu_i}{t}$ . Assume that there is  
 613 such at  $t$ . Then

$$\begin{aligned} N_{i,t} &= r_h N_{i,t-1} + \frac{1}{t} \\ &\leq \frac{r_h \nu_i}{t-1} + \frac{1}{t} \\ &= \frac{r_h t \nu_i + t - 1}{t(t-1)} \\ &= \frac{(t-1)\nu_i + (r_h - 1)t\nu_i + \nu_i + t - 1}{t(t-1)} \\ &= \frac{\nu_i}{t} + \frac{(r_h - 1)t\nu_i + \nu_i + t - 1}{t(t-1)}. \end{aligned}$$

614 We now shall prove that  $(r_h - 1)t\nu_i + \nu_i + t - 1 = [(r_h - 1)\nu_i + 1]t + \nu_i - 1$  is negative. First, we  
 615 have that

$$\begin{aligned} (r_h - 1)\nu_i + 1 &= 1 - \max \left\{ 1+r_h, 2 \log \left( \frac{1}{i(1-r_h)} \right) \right\} \\ &\leq 0. \end{aligned}$$

616 Then,

$$[(r_h - 1)\nu_i + 1]t + \nu_i - 1 \leq 0 \iff t \geq \frac{\nu_i - 1}{(1-r_h)\nu_i - 1}$$

617 since  $(r_h - 1)\nu_i + 1 \leq 0$ . Thus, the property is true for every  $t \geq \frac{\nu_i - 1}{(1-r_h)\nu_i - 1}$ . In addition, we have

$$\begin{aligned} \nu_i &\geq \frac{1+r_h}{1-r_h} \\ \nu_i(1-r_h) &\geq 1+r_h \\ 2\nu_i(1-r_h) - 2 &\geq \nu_i(1-r_h) - 1 + r_h \\ \frac{2}{1-r_h} &\geq \frac{\nu_i - 1}{\nu_i(1-r_h) - 1}, \end{aligned}$$

618 and the property is also true for every  $t \geq \frac{2}{1-r_h}$ . This concludes the proof.  $\square$

619 Finally, we can prove the Proposition 3.1:

620 *Proof.* The expectation of  $\Delta_t$  is immediate using Lemma D.2 and the fact that the  $\epsilon_i$  are independent,  
 621 zero-mean noises. The variance is equal to  $V[\Delta_t] = \alpha^2 B \sum_{i=0}^t N_{i,t}^2$ . While our analysis was  
 622 only along one eigenspace of the Hessian with associated eigenvalue  $h$ , we must now sum over all  
 623 dimensions. We will thus define

$$\bar{\nu}_i = \frac{\max \left\{ 2 - \alpha\mu, 2 \log \left( \frac{1}{i\alpha\mu} \right) \right\}}{\alpha\mu} \quad \text{for } i > 0$$

$$\bar{\nu}_0 = \frac{2 + 2 \log \left( \frac{1}{\alpha\mu} \right)}{\alpha\mu},$$

624 which is, for every  $i$ , the maximum  $\nu_i$  across all dimensions. We get

$$\begin{aligned} V[\Delta_t] &\leq d\alpha^2 B \sum_{i=0}^t \frac{\bar{\nu}_i^2}{t^2} \\ &\leq d\alpha^2 B \sum_{i=0}^t \frac{\bar{\nu}_0^2}{t^2} \quad \text{since } \nu_i \geq \nu_{i+1} \quad \forall i \\ &\leq \frac{d\alpha^2 B \bar{\nu}_0^2}{t}. \end{aligned}$$

625 Since we have

$$E[\theta_t - \theta^*] = (I - \alpha H)^t (\theta_0 - \theta^*),$$

626 we get

$$\begin{aligned} E[\|\theta_t - \theta^*\|^2] &= \|E[\theta_t - \theta^*]\|^2 + V[\Delta_t] \\ &\leq (\theta_0 - \theta^*)^\top (I - \alpha H)^{2t} (\theta_0 - \theta^*) + \frac{d\alpha^2 B \bar{\nu}_0^2}{t} \\ &\leq \left(1 - \frac{1}{\kappa}\right)^{2t} \|\theta_0 - \theta^*\|^2 + \frac{d\alpha^2 B \bar{\nu}_0^2}{t}. \end{aligned}$$

627 This concludes the proof. □

## 628 E Proof of Proposition 3.2 and Proposition 3.3

629 In this section we list and prove all lemmas used in the proofs of Proposition 3.2 and Proposition 3.3;  
 630 all lemmas are stated in the same conditions as the proposition.

631 We start the following proposition:

632 **Proposition E.1.** *Let  $f$  be a quadratic function with positive definite Hessian  $H$  with largest eigen-*  
 633 *value  $L$  and condition number  $\kappa$  and if the stochastic gradients satisfy  $g(\theta, x) = g(\theta) + \epsilon$  with  $\epsilon$  a*  
 634 *random uncorrelated noise with covariance bounded by  $BI$ .*

635 *Then, Algorithm 1 leads to iterates  $\theta_t$  satisfying*

$$E[\theta_t - \theta^*] = \begin{pmatrix} I \\ 0 \end{pmatrix} A^t \begin{pmatrix} E[\theta_1 - \theta^*] \\ E[\theta_0 - \theta^*] \end{pmatrix} \quad (5)$$

636 *where*

$$A = \begin{pmatrix} I - \alpha H + \mu I & -\mu I \\ I & 0 \end{pmatrix} \quad (6)$$

637 *governs the dynamics of this bias. In particular, when its spectral radius,  $\rho(A)$  is less than 1, the*  
 638 *iterates converge linearly to  $\theta^*$ .*

639 In a similar fashion, the variance dynamics of Heavyball-IGT are governed by the matrix

$$D_i = \begin{pmatrix} (1 - \alpha h_i + \mu)^2 + 2\alpha^2 h_i^2 & \mu^2 & -2\mu(1 - \alpha h_i + \mu)^2 \\ 1 & 0 & 0 \\ 1 - \alpha h_i + \mu & 0 & -\mu \end{pmatrix}$$

640 If the spectral radius of  $D_i$ ,  $\rho(D_i)$ , is strictly less than 1 for all  $i$ , then there exist constants  $t_0 > 0$   
641 and  $C > 0$  for which

$$\text{Var}(\theta_t) \leq 2\alpha^2 dBC \frac{\log(t)}{t}, \quad \text{for } t > t_0$$

642 where  $B$  is a bound on the variance of noise variables  $\epsilon_i$ .

**Lemma E.2** (IGT estimator as true gradient plus noise average). *If  $v_0 = g(\theta_0) + \epsilon_0$  and for  $t > 0$  we have*

$$v_t = \frac{t}{t+1}v_{t-1} + \frac{1}{t+1}g(\theta_t + t(\theta_t - \theta_{t-1})) + \frac{1}{t+1}\epsilon_t,$$

then

$$v_t = g(\theta_t) + \frac{1}{t+1} \sum_{i=0}^t \epsilon_i.$$

643 This lemma is already proved in the previous section for the IGT estimator (Lemma D.1) and is just  
644 repeated here for completeness. We will use this result in the next few lemmas.

**Lemma E.3** (The IGT gradient estimator is unbiased on quadratics). *For the IGT gradient estimator,  $v_t$ , corresponding to parameters  $\theta_t$  we have*

$$\mathbb{E}[v_t] = g(\mathbb{E}\theta_t),$$

645 where the expectation is over all gradient noise vectors  $\epsilon_0, \epsilon_1, \dots, \epsilon_t$ .

646 *Proof.* The proof proceeds by induction. The base case holds as we have

$$\mathbb{E}[v_0] = \mathbb{E}[g_0 + \epsilon_0] = g(\theta_0).$$

647 For the inductive case, we can write

$$\begin{aligned} \mathbb{E}[v_t] &= \mathbb{E}\left[\frac{t}{t+1}v_{t-1} + \frac{1}{t+1}\hat{g}(\theta_t + t(\theta_t - \theta_{t-1}))\right] \\ &= \mathbb{E}\left[\frac{t}{t+1}v_{t-1} + \frac{1}{t+1}g_t + \frac{t}{t+1}g_t - \frac{t}{t+1}g_{t-1} + \frac{1}{t+1}\epsilon_t\right] \\ &= \frac{t}{t+1}\mathbb{E}[v_{t-1} - g_{t-1}] + \mathbb{E}[g_t] + \frac{t}{t+1}\mathbb{E}[\epsilon_t] \\ &= \mathbb{E}[g_t] = g(\mathbb{E}[\theta_t]). \end{aligned}$$

648 Where, in the third equality,  $\mathbb{E}[v_{t-1} - g_{t-1}] = 0$  by the inductive assumption, and the last equality  
649 because the gradient of a quadratic function is linear.  $\square$

**Lemma E.4** (Bounding the IGT gradient variance). *Let  $v_t$  be the IGT gradient estimator. Then*

$$\text{Var}[v_t] \leq 2h^2\text{Var}[\theta_t - \theta^*] + \frac{2B}{t},$$

650 where  $B$  is the variance of the homoscedastic noise  $\epsilon_t$ .

*Proof.*

$$\begin{aligned}
\text{Var}[v_t] &= \text{Var}\left[g_t + \frac{1}{t+1} \sum_{i=0}^t \epsilon_i\right] \\
&= \text{Var}[h\theta_t] + \text{Var}\left[\frac{1}{t+1} \sum_{i=0}^t \epsilon_i\right] \\
&\quad + 2\text{Cov}\left[h\theta_t, \frac{1}{t+1} \sum_{i=0}^t \epsilon_i\right] \\
&\leq 2\text{Var}[h\theta_t] + 2\text{Var}\left[\frac{1}{t+1} \sum_{i=0}^t \epsilon_i\right] \\
&= 2h^2\text{Var}[\theta_t - \theta^*] + 2\frac{B}{t}
\end{aligned}$$

651

□

652 Now that we have these basic results on the IGT estimator, we can analyze the evolution of the bias  
653 and variance of Heavyball-IGT. We use the quadratic assumption to decouple the vector dynamics  
654 of Heavyball-IGT into independent scalar dynamics. If the Hessian,  $H$ , has eigenvalues  $L \geq h_1 \geq$   
655  $h_2 \geq \dots \geq h_n = L/\kappa$ , then we can assume without loss of generality that  $H$  is diagonal with  
656  $H_{ii} = h_i$ .

657 **Lemma E.5** (Evolution of bias for scalar quadratic). *Assume that the Hessian, second derivative, is*  
658  *$h$ .*

*Starting with  $v_0 = g(\theta_0) + \epsilon_0$  and  $w_0 = 0$ , performing the following iterative updates (Heavyball-IGT, Algorithm 1):*

$$\begin{aligned}
v_t &= \frac{t}{t+1}v_{t-1} + \frac{1}{t+1}g(\theta + t(\theta_t - \theta_{t-1})) + \frac{1}{t+1}\epsilon_t, \\
w_{t+1} &= \mu w_t + \alpha v_t, \quad \theta_{t+1} = \theta_t - w_{t+1}
\end{aligned}$$

*results in*

$$\Delta_t = A^t \Delta_0 - \alpha \sum_{i=0}^{t-1} N_{i,t} \begin{bmatrix} \epsilon_i \\ 0 \end{bmatrix}$$

659 *where  $N_{j,0} = 0_{2 \times 2}$ ,  $N_{i,t} = AN_{i,t-1} + \mathbf{1}_{i < t} \frac{1}{t} I$ ,*

660  $\Delta_t = \begin{bmatrix} \theta_t - \theta^* \\ \theta_{t-1} - \theta^* \end{bmatrix}$  *and  $A = \begin{pmatrix} 1 - \alpha h + \mu & -\mu \\ 1 & 0 \end{pmatrix}$ .*

661 *Proof.* The proof proceeds by induction. First notice that for  $t = 0$  the equality naturally holds. We  
662 make the inductive assumption that it holds for  $t - 1$ , and start by using Lemma E.2:

$$\begin{aligned}
\Delta_t &= A\Delta_{t-1} - \frac{\alpha}{t} \sum_{i=0}^{t-1} \begin{bmatrix} \epsilon_i \\ 0 \end{bmatrix} \\
&= A(A^{t-1}\Delta_0 - \alpha \sum_{i=0}^{t-2} N_{i,t} \begin{bmatrix} \epsilon_i \\ 0 \end{bmatrix}) - \frac{\alpha}{t} \sum_{i=0}^{t-1} \begin{bmatrix} \epsilon_i \\ 0 \end{bmatrix} && \text{(Inductive assumption)} \\
&= A^t \Delta_0 - \alpha \left( \sum_{i=0}^{t-2} AN_{i,t} \begin{bmatrix} \epsilon_i \\ 0 \end{bmatrix} + \frac{1}{t} \sum_{i=0}^{t-1} \begin{bmatrix} \epsilon_i \\ 0 \end{bmatrix} \right) \\
&= A^t \Delta_0 - \alpha \sum_{i=0}^{t-1} N_{i,t} \begin{bmatrix} \epsilon_i \\ 0 \end{bmatrix} && \text{(Def. of } N_{i,t} \text{)}
\end{aligned}$$

663

□

664 **Lemma E.6** (Evolution of variance). Let  $U_t = \text{Var}[\theta_t]$  and  $V_t = \text{Cov}[\theta_t, \theta_{t-1}]$ , where  $\theta_t$  is the  $t$ -th  
665 iterate of Heavyball-IGT on a 1-dimensional quadratic function with curvature  $h$ . The following  
666 matrix describes the variance dynamics of Heavyball-IGT.

$$D = \begin{pmatrix} (1 - \alpha h + \mu)^2 + 2\alpha^2 h^2 & \mu^2 & -2\mu(1 - \alpha h + \mu)^2 \\ 1 & 0 & 0 \\ 1 - \alpha h + \mu & 0 & -\mu \end{pmatrix} \quad (7)$$

667 If the spectral radius of  $D$ ,  $\rho(D)$ , is strictly less than 1, then there exist constants  $t_0 > 0$  and  $C > 0$   
668 for which

$$\text{Var}(\theta_t) \leq 2\alpha^2 BC \frac{\log(t)}{t}$$

669 , where  $B$  is a bound on the variance of the noise.

670 *Proof.* The proof (and lemma) is similar to the proof of Lemma 9 in [49]. We start by expanding  
671  $U_{t+1}$  as follows.

$$\begin{aligned} U_{t+1} &= \mathbb{E}[(\theta_{t+1} - \bar{\theta}_{t+1})^2] \\ &= \mathbb{E}[(\theta_t - \alpha v_t + \mu(\theta_t - \theta_{t-1}) - \bar{\theta}_t + \alpha g_t - \mu(\bar{\theta}_t - \bar{\theta}_{t-1}))^2] \\ &= \mathbb{E}[(\theta_t - \alpha g_t + \mu(\theta_t - \theta_{t-1}) - \bar{\theta}_t + \alpha g_t \\ &\quad - \mu(\bar{\theta}_t - \bar{\theta}_{t-1}) + \alpha(g_t - v_t))^2] \\ &= \mathbb{E}[((1 - \alpha h + \mu)(\theta_t - \bar{\theta}_t) - \mu(\theta_{t-1} - \bar{\theta}_{t-1}))^2 \\ &\quad + \alpha^2 \mathbb{E}[(g_t - v_t)^2]] \\ &\leq \mathbb{E}[((1 - \alpha h + \mu)(\theta_t - \bar{\theta}_t) - \mu(\theta_{t-1} - \bar{\theta}_{t-1}))^2] \\ &\quad + \alpha^2 \left( 2h^2 \mathbb{E}[(\theta_t - \bar{\theta}_t)^2] + \frac{2B}{t+1} \right) \\ &\leq [(1 - \alpha h + \mu)^2 + 2\alpha^2 \mu^2] \mathbb{E}[(\theta_t - \bar{\theta}_t)^2] \\ &\quad - 2\mu(1 - \alpha h + \mu) \mathbb{E}[(\theta_t - \bar{\theta}_t)(\theta_{t-1} - \bar{\theta}_{t-1})] \\ &\quad + \mu^2 \mathbb{E}[(\theta_{t-1} - \bar{\theta}_{t-1})^2] + \alpha^2 \frac{2B}{t+1}. \end{aligned}$$

672 Where the fourth equality is obtained since we know that the IGT gradient estimator is unbiased, i.e.  
673  $\mathbb{E}[g_t - v_t] = 0$ . The first inequality stems from Lemma E.4. We similarly expand  $V_t$ .

$$\begin{aligned} V_t &= \mathbb{E}[(\theta_t - \bar{\theta}_t)(\theta_{t-1} - \bar{\theta}_{t-1})] \\ &= \mathbb{E}[(\mu(\theta_{t-1} - \bar{\theta}_{t-1}) - \mu(\theta_{t-2} - \bar{\theta}_{t-2}) + \alpha(g_t - v_t)) \\ &\quad (\theta_{t-1} - \bar{\theta}_{t-1})] \\ &= (1 - \alpha h + \mu) \mathbb{E}[(\theta_{t-1} - \bar{\theta}_{t-1})^2] \\ &\quad - \mu \mathbb{E}[(\theta_{t-1} - \bar{\theta}_{t-1})(\theta_{t-2} - \bar{\theta}_{t-2})] \end{aligned}$$

674 From the above expressions, we obtain

$$\begin{aligned}
\begin{pmatrix} U_{t+1} \\ U_t \\ V_{t+1} \end{pmatrix} &\leq D \begin{pmatrix} U_t \\ U_{t-1} \\ V_t \end{pmatrix} + \begin{pmatrix} \alpha^2 \frac{2B}{t+1} \\ 0 \\ 0 \end{pmatrix} \\
&\leq 2\alpha^2 B \sum_{i=0}^t D^i \begin{pmatrix} \frac{1}{t+1-i} \\ 0 \\ 0 \end{pmatrix} \\
&\leq 2\alpha^2 B \left( \sum_{i=0}^{s-1} D^i \begin{pmatrix} \frac{1}{t+1-i} \\ 0 \\ 0 \end{pmatrix} + \sum_{i=s}^t D^i \begin{pmatrix} \frac{1}{t+1-i} \\ 0 \\ 0 \end{pmatrix} \right)
\end{aligned}$$

675 where an inequality of vectors implies the corresponding elementwise inequalities.

676 If the spectral radius of  $D$ ,  $\rho(D)$  is strictly less than 1, then there exists constant  $C' > 0$  such that

$$\begin{aligned}
\begin{pmatrix} 1 \\ 0 \\ 0 \end{pmatrix}^T \sum_{i=0}^{s-1} D^i \begin{pmatrix} \frac{1}{t+1-i} \\ 0 \\ 0 \end{pmatrix} &\leq C' \sum_{i=0}^{s-1} \frac{1}{t+1-i} \\
&\leq \frac{C' s}{t+2-s}
\end{aligned}$$

677 If the spectral radius of  $D$ ,  $\rho(D)$ , is strictly less than 1, then there exists constant  $\zeta > 0$  and constant  
678  $C''(\zeta) > 0$  such that,  $\rho(D) + \zeta < 1$  and

$$\begin{aligned}
\begin{pmatrix} 1 \\ 0 \\ 0 \end{pmatrix}^T \sum_{i=s}^t D^i \begin{pmatrix} \frac{1}{t+1-i} \\ 0 \\ 0 \end{pmatrix} &\leq \begin{pmatrix} 1 \\ 0 \\ 0 \end{pmatrix}^T \sum_{i=s}^t D^i \begin{pmatrix} 1 \\ 0 \\ 0 \end{pmatrix} \\
&\leq C'' \sum_{i=s}^t (\rho(D) + \zeta)^s \\
&= C''(t-s+1)(\rho(D) + \zeta)^s
\end{aligned}$$

679 Let  $\rho' = \rho(D) + \zeta$  and  $s = \lceil 2 \log_{1/\rho'} t \rceil$ . Then  $(\rho(D) + \zeta)^s = 1/t^2$ , and putting the above two  
680 bounds together,

$$\begin{aligned}
U_{t+1} &\leq 2\alpha^2 B \left( \frac{2C' \log_{1/\rho'} t}{t+2-2\log_{1/\rho'} t} + C'' \frac{t-2\log_{1/\rho'} t+1}{t^2} \right) \\
&\leq 2\alpha^2 B C \frac{\log(t+1)}{t+1}
\end{aligned}$$

681 where the last inequality holds for  $t > t_0$  for some  $t_0$  and some constant  $C > 0$ .

682 □

683 We can now prove Proposition E.1.

684 *Proof of Proposition E.1.* The bias statement of the proposition follows directly from taking an  
685 expectation on the bound of Lemma B.4, and the variance statement from summing up the  $d$  different  
686 variance terms given for each scalar component by Lemma B.5. □

### 687 E.1 Proof of Proposition 3.2

688 This Proposition follows from the observation that, in the noiseless case,  $\epsilon_t = 0$  in our model. In that  
689 case, Lemma E.3 shows that Heavyball-IGT reduces to the heavy ball, and the rest follows from the  
690 optimal tuning of the heavy ball [49].

691 **E.2 Proof of Proposition 3.3**

692 *Proof.* Like we did in previous proofs, we can assume without loss of generality that the Hessian,  $H$ ,  
 693 is diagonal with elements  $h_i$ . For a diagonal  $H$ , matrix  $A$  can be permuted to be block diagonal with  
 694 blocks

$$A_i = \begin{pmatrix} 1 - \alpha h_i + \mu & \mu \\ 1 & 0 \end{pmatrix}.$$

695 To prove that  $\rho(A) < 1$  it suffices to prove that  $\rho(A_i) < 1$  for all  $i$ . For the rest of the proof we will  
 696 focus on the dynamics along a single eigendirection with curvature  $h_i$ . The rest of this proof used  $D$   
 697 to denote  $D_i$ ,  $A$  to denote  $A_i$  and  $h$  to denote  $h_i$ .

698 To make explicit the dependence of matrices  $A$  and  $D$  on hyperparameters and curvature, we write  
 699  $A(\alpha, \mu, h)$  and  $D(\alpha, \mu, h)$ . Let  $0 < \alpha < 2/(3h)$  and  $\mu_0 = 0$ . Using hyperparameters  $(\alpha, \mu_0)$  one  
 700 gets the results for gradient descent without momentum. In particular  $\rho(A(\alpha, \mu_0, h)) = |1 - \alpha h| < 1$ ,  
 701 and the spectral radius of  $D$  is  $\rho(D(\alpha, \mu_0, h)) = |(1 - \alpha h)^2 + 2\alpha^2 h^2| < 1$ .

702 We will argue that there exists  $\mu > 0$ , such that  $\rho(A(\alpha, \mu, h)) < 1$ , and the spectral radius of  $D$  is  
 703  $\rho(D(\alpha, \mu, h)) < 1$ . Then the previous lemma implies that bias converges linearly, and variance is  
 704  $O(\log(t)/t)$ .

705 To argue the existence of  $\mu > 0$ , we will perform eigenvalue perturbation analysis using the Bauer-  
 706 Fike theorem. Note that  $A(\alpha, \mu, h) = A(\alpha, \mu_0, h) + \mu \Delta_A$  where

$$\Delta_A = \begin{pmatrix} 1 & -1 \\ 0 & 0 \end{pmatrix}.$$

707 Similarly,  $D(\alpha, \mu, h) \approx D(\alpha, \mu_0, h) + \mu \Delta_D$  where

$$\Delta_D = \begin{pmatrix} 2(1 - \alpha h) & 0 & -2(1 - \alpha h) \\ 0 & 0 & 0 \\ 1 & 0 & -1 \end{pmatrix}.$$

708 This last approximate inequality is a first-order approximation, in the sense that we are working with  
 709 arbitrarily small, positive values of  $\mu$ , and we have kept terms linear in  $\mu$  but ignored higher order  
 710 terms, like  $\mu^2$ .

711 We will apply the Bauer-Fike theorem to bound the eigenvalues of  $D(\alpha, \mu, h)$ . Consider the eigende-  
 712 composition  $D(\alpha, \mu_0, h) = V \Lambda V^{-1}$ . We can compute

$$V = \begin{pmatrix} 0 & 0 & \frac{1-2\alpha h+3\alpha^2 h^2}{1-\alpha h} \\ 0 & 1 & \frac{1}{1-\alpha h} \\ 1 & 0 & 1 \end{pmatrix}$$

713 and

$$V^{-1} = \begin{pmatrix} \frac{1-\alpha h}{1-2\alpha h+3\alpha^2 h^2} & -\frac{1}{1-2\alpha h+3\alpha^2 h^2} & 0 \\ 0 & 1 & 0 \\ 1 & 0 & 0 \end{pmatrix}.$$

714 Note that because we assume  $\alpha < 2/(3h)$  we get  $1 - \alpha h > 0$ . Also,  $1 - 2\alpha h + 3\alpha^2 h^2 > 0$  regardless  
 715 of the choice of hyperparameters. This means that matrices  $V$  and  $V^{-1}$  are singular and of finite  
 716 norm. The norm of  $\Delta_D$  is also finite. The Bauer-Fike theorem states that, if  $\nu$  is an eigenvalue of  
 717  $D(\alpha, \mu_0, h)$ , then there exists an eigenvalue  $\lambda$  of  $D(\alpha, \mu, h)$  such that

$$|\lambda - \nu| \leq \|V\|_p \|V^{-1}\|_p \|\mu \Delta_D\|_p,$$

718 for any  $p$ -norm. Since by construction  $|\nu| \leq \rho(D(\alpha, \mu_0, h)) < 1$ , the above means that there exists a  
 719 sufficiently small, but strictly positive value of  $\mu$ , such that  $\lambda < 1$ . By repeating this argument for  
 720 all pairs of eigenvalues, we get the stated result. The same argument can be repeated to prove the  
 721 existence of a strictly positive  $\mu$  such that  $\rho(A(\alpha, \mu, h)) < 1$ .

722 □