

Acceleration by Random Stepsizes: Hedging, Equalization, and the Arcsine Stepsize Schedule

Jason M. Altschuler
UPenn
alts@upenn.edu

Pablo A. Parrilo
LIDS - MIT
parrilo@mit.edu

December 7, 2024

Abstract

We show that for separable convex optimization, random stepsizes fully accelerate Gradient Descent. Specifically, using inverse stepsizes i.i.d. from the Arcsine distribution improves the iteration complexity from $O(\kappa)$ to $O(\kappa^{1/2})$, where κ is the condition number. No momentum or other algorithmic modifications are required. This result is incomparable to the (deterministic) Silver Stepsize Schedule which does not require separability but only achieves partial acceleration $O(\kappa^{\log_{1+\sqrt{2}}}) \approx O(\kappa^{0.78})$. Our starting point is a conceptual connection to potential theory: the variational characterization for the distribution of stepsizes with fastest convergence rate mirrors the variational characterization for the distribution of charged particles with minimal logarithmic potential energy. The Arcsine distribution solves both variational characterizations due to a remarkable “equalization property” which in the physical context amounts to a constant potential over space, and in the optimization context amounts to an identical convergence rate over all quadratic functions. A key technical insight is that martingale arguments extend this phenomenon to all separable convex functions. We interpret this equalization as an extreme form of hedging: by using this random distribution over stepsizes, Gradient Descent converges at exactly the same rate for all functions in the function class.

1 Introduction

It is classically known that by using certain deterministic stepsize schedules $\{\alpha_t\}$, the Gradient Descent algorithm (GD)

$$x_{t+1} = x_t - \alpha_t \nabla f(x_t), \quad t = 0, 1, 2, \dots \quad (1.1)$$

solves convex optimization problems $\min_{x \in \mathbb{R}^d} f(x)$ to arbitrary precision from an arbitrary initialization x_0 . This result can be found in any textbook on convex optimization, e.g., [6, 9, 10, 30, 35, 38, 40, 45]. The present paper investigates the following question:

Is it possible to accelerate the convergence of GD by using *random* stepsizes? (1.2)

In this paper, we show that for separable convex optimization, the answer is yes. Specifically, we provide an i.i.d. stepsize schedule that makes GD converge at the optimal accelerated rate. See §1.1 for a precise statement; here we first contextualize with the relevant literature on random stepsizes.

This result first appeared in Chapter 6 of the 2018 MS Thesis [2] of the first author (advised by the second author). This was the first convergence rate for any stepsize schedule (deterministic or randomized) that improved asymptotically over the standard unaccelerated rate for constant-stepsize GD, in any setting beyond quadratics.

Note that the use of randomness for GD stepsizes is conceptually different from the standard use of randomness in Stochastic Gradient Descent (SGD). SGD has random, inexact gradient queries and typically use deterministic stepsizes. GD has exact gradient queries, and random stepsizes amounts to purposely injecting noise into the algorithm.

The special case of quadratics. This question is understood (only) for optimizing convex *quadratic* functions. In particular, the line of work [32, 47, 63] shows that for quadratic functions f that are m -strongly convex and M -smooth (i.e., $mI \leq \nabla^2 f \leq MI$), there is a unique optimal distribution for i.i.d. stepsizes that makes GD converge at a rate of

$$\left(\frac{\|x_n - x^*\|}{\|x_0 - x^*\|} \right)^{1/n} \rightarrow \frac{\sqrt{\kappa} - 1}{\sqrt{\kappa} + 1}, \quad (1.3)$$

where x^* denotes the minimizer of f , $\kappa := \frac{M}{m}$ denotes the condition number, and this convergence is in an appropriate probabilistic sense (e.g., almost-sure or L^1). This stepsize distribution is most naturally stated in terms of the inverse stepsizes $\beta_t := 1/\alpha_t$ having the Arcsine distribution on the interval (m, M) , i.e., density

$$\frac{d\mu}{d\beta} = \frac{\mathbb{1}_{\beta \in (m, M)}}{\pi \sqrt{(M - \beta)(\beta - m)}}. \quad (1.4)$$

Briefly, this Arcsine distribution arises naturally for the special case of quadratics because: (1) the convergence rate is invariant w.r.t. the stepsize order and thus depends only on the empirical measure of the stepsizes; and (2) the Arcsine distribution is the limit as $n \rightarrow \infty$ of the empirical measure of the optimal n (deterministic) stepsizes for GD [61]. See §2.1 for details.

The resulting convergence rate (1.3) requires roughly $\Theta(\sqrt{\kappa} \log \frac{1}{\varepsilon})$ iterations to produce an ε -accurate solution. Such a rate is called “accelerated” because its dependence on κ improves over the standard $\Theta(\kappa \log \frac{1}{\varepsilon})$ iteration complexity required by GD with constant stepsizes. Moreover, this convergence rate $\Theta(\sqrt{\kappa} \log \frac{1}{\varepsilon})$ is information-theoretically optimal for any first-order algorithm [39, 40], and matches the worst-case convergence rate of accelerated algorithms that change GD beyond just its stepsizes, e.g., the Conjugate Gradient method [31], Polyak’s Heavy Ball method [44], Nesterov-style momentum methods [41], and more [22]. See Table 1, left.

Beyond quadratics? For general convex optimization, momentum-based algorithms still achieve these optimally accelerated rates $\Theta(\sqrt{\kappa} \log \frac{1}{\varepsilon})$. However, existing acceleration results for random stepsizes do not extend beyond the special case of quadratic functions f —otherwise ∇f is non-linear, making GD a non-linear map, which completely breaks all previous analyses. This failure of the analyses is due to fundamental differences between quadratic optimization and convex optimization:

- The optimal deterministic stepsizes for quadratic optimization [61] (which the random stepsizes are built upon) are provably bad for convex optimization [2].
- The convergence rate is invariant w.r.t. the stepsize order for quadratic optimization (which motivates the opportunity for random stepsizes), but this is false for convex optimization [2].
- Using time-varying stepsizes is equivalent to using momentum for quadratic optimization (which enables GD to achieve accelerated rates), but this is false for convex optimization [22].

Due to these challenges, it remained unclear whether random stepsizes could lead to the accelerated rate $\Theta(\sqrt{\kappa} \log \frac{1}{\varepsilon})$ for convex optimization. In fact, it was unknown even if random stepsizes could provide an arbitrary small improvement over the $\Theta(\kappa \log \frac{1}{\varepsilon})$ rate of constant-stepsize GD in any setting beyond quadratic optimization.




FIGURE 1: The induced distribution of stepsizes α , for inverse stepsizes α^{-1} taken from the Arcsine distribution on (m, M) . The minimum stepsize is $1/M$, the maximum is $1/m$, and the median is $2/(M + m)$, which is the optimal value for constant stepsize schedules. For constant stepsize schedules, the dashed red line $\bar{\alpha} = 2/M$ is the threshold for convergence; larger stepsizes α lead to divergence, and short stepsizes α lead to slow convergence. This distribution optimally *hedges* between short and long steps. Note that the mean stepsize is $1/\sqrt{Mm}$ which is larger than the divergence threshold for constant stepsize schedules, by a factor of $\Theta(\sqrt{\kappa})$. This plot sets $\kappa = 10$, $m = 1/\kappa$, $M = 1$; the discrepancy between this distribution and standard constant stepsizes is even more dramatic for larger κ .

1.1 Contribution and discussion

This paper shows that random stepsizes provably accelerate GD beyond the setting of quadratics. Specifically, our main result (Theorem 1.2) shows that GD converges at the optimally accelerated rate for separable convex functions. This separability property is defined below and, for C^2 functions f , is equivalent to commutativity of the Hessians $\{\nabla^2 f(x)\}_{x \in \mathbb{R}^d}$, see Appendix B.1.

Definition 1.1 (Separable functions). *A function $f : \mathbb{R}^d \rightarrow \mathbb{R}$ is separable if it admits a decomposition of the form*

$$f(x) = \sum_{i=1}^d f_i([Ux]_i) \quad (1.5)$$

for some orthogonal matrix $U \in \mathbb{R}^{d \times d}$ and some functions $f_1, \dots, f_d : \mathbb{R} \rightarrow \mathbb{R}$.

Theorem 1.2 (Random stepsizes accelerate GD for separable convex optimization). *Consider any dimension d , any separable function $f : \mathbb{R}^d \rightarrow \mathbb{R}$ that is m -strongly convex and M -smooth, and any initialization point x_0 that is not equal to the minimizer x^* of f . By using i.i.d. inverse stepsizes α_t^{-1} from the Arcsine distribution (1.4), GD converges at the rate*

$$\left(\frac{\|x_n - x^*\|}{\|x_0 - x^*\|} \right)^{1/n} \rightarrow \frac{\sqrt{\kappa} - 1}{\sqrt{\kappa} + 1}, \quad (1.6)$$

where this convergence is in the almost sure and L^1 sense. Moreover, this is the unique distribution for which GD achieves this optimal rate.

We make several remarks about this result.

Optimality. The rate in Theorem 1.2 is optimal since there is an exactly matching lower bound: there is an explicit κ -conditioned quadratic function for which any (Krylov) first-order algorithm cannot contract to the optimum faster than this rate $\frac{\sqrt{\kappa}-1}{\sqrt{\kappa}+1}$ [41, Theorem 2.1.12].

Stepsize \ Function class	Quadratic	Separable	Convex
Constant	$\Theta(\kappa)$ folklore	$\Theta(\kappa)$ folklore	$\Theta(\kappa)$ folklore
Deterministic (Chebyshev/Silver)	$\Theta(\kappa^{1/2})$ [61]	$O(\kappa^{\log_1+\sqrt{2}})$ [3]	$O(\kappa^{\log_1+\sqrt{2}})$ [3]
Random (Arcsine)	$\Theta(\kappa^{1/2})$ [32, 47, 63]	$\Theta(\kappa^{1/2})$ Theorem 1.2	unknown if randomness helps

TABLE 1: Iteration complexities for minimizing κ -conditioned functions using GD with different stepsize choices. The dependence on the accuracy is omitted as it is always $\log 1/\varepsilon$. Information-theoretic lower bounds are $\Omega(\kappa^{1/2})$ in all settings [39]. Constant stepsize schedules require $\Theta(\kappa)$ iterations; this is the textbook unaccelerated rate. For quadratics (left), accelerated rates of $\Theta(\kappa^{1/2})$ can be equivalently achieved via deterministic stepsizes [61], random stepsizes [32, 47, 63], or momentum-based methods such as Heavy Ball [44] and Conjugate Gradients [31], among others [22]. Beyond quadratics (middle, right), momentum-based methods are not equivalent to varying stepsizes. Acceleration was first achieved by Nesterov’s Fast Gradient Method [41] and it has long been believed that beyond quadratics, any acceleration requires modifying GD by adding internal dynamics, e.g., momentum. [3] recently showed (conjecturally optimal) partially accelerated rates $O(\kappa^{\log_1+\sqrt{2}}) \approx O(\kappa^{0.78})$ for convex optimization with deterministic stepsizes. We prove that random stepsizes enable fully accelerated rates $\Theta(\kappa^{1/2})$ beyond quadratics. Conditional on conjectures (see §7), this suggests potential gaps between deterministic and randomized stepsizes for GD, and gaps between separable and general convex optimization for GD.

Comparison to GD with different stepsize schedules. Theorem 1.2 first appeared as Theorem 6.1 of the 2018 MS thesis [2] of the first author (advised by the second author). This was the first convergence rate for any stepsize schedule (deterministic or randomized) that improved over the textbook rate for constant-stepsize GD, in any setting beyond quadratics. Indeed, even though many stepsize schedules have been investigated in both theory and practice, none of these alternative schedules had led to an analysis that outperforms (even by an arbitrarily small amount) the “unaccelerated” contraction factor $(\frac{\|x_n - x^*\|}{\|x_0 - x^*\|})^{1/n} \leq \frac{\kappa-1}{\kappa+1}$ of constant stepsize GD, which corresponds to an iteration complexity of $\Theta(\kappa \log \frac{1}{\varepsilon})$. This comparison statement includes even adaptive stepsize schedules such as exact line search [9, 15, 42, 45], Armijo-Goldstein schedules [40], Polyak-type schedules [45], and Barzilai-Borwein-type schedules [5]. Some of these schedules fail to accelerate even on quadratics (and therefore also on separable convex problems), a classic example being exact line search [33, §3.2.2].

In the years since [2], a line of work has shown partially accelerated convergence rates for general convex optimization using deterministic stepsize schedules [2, 3, 4, 8, 12, 13, 23, 26, 27, 28, 29, 58, 62]. In particular, [3] showed the Silver Convergence Rate $O(\kappa^{\log_1+\sqrt{2}} \log \frac{1}{\varepsilon}) \approx O(\kappa^{0.78} \log \frac{1}{\varepsilon})$ for minimizing κ -conditioned functions by using a certain fractal-like deterministic stepsize schedule called the Silver Stepsize Schedule. This rate is conjecturally optimal among all possible deterministic stepsize schedules [3] and naturally extends to non-strongly convex optimization [4]. Theorem 1.2 is incomparable: it shows a fully accelerated rate $O(\kappa^{1/2} \log \frac{1}{\varepsilon})$, but under the additional assumption of separability. This incomparability raises several interesting questions about the possible benefit of randomization for GD. The future work section §7 collects several such conjectures.

Comparison to arbitrary first-order algorithms. Starting from the seminal work of Nesterov in 1983 [41], the traditional approach for obtaining accelerated convergence rates is to modify the GD algorithm beyond just changing stepsizes, for example by adding momentum, extra state, or other internal dynamics. These results hold for arbitrary convex optimization and do not require separability. See for example [11, 21, 52, 56] for some recent such algorithms, and see the recent survey [22] for a comprehensive account of this extensive line of work. The purpose of this paper is to investigate an alternative mechanism for acceleration: can GD be accelerated in its original form, by just choosing random stepsizes?

Generality. The assumptions in Theorem 1.2 can be relaxed. Separability can be generalized to radial-separability, and strong convexity and smoothness can be relaxed to less stringent assumptions on the spectrum of the Hessian of f ; see §4 for details. Note also that since GD does not depend on the choice of the basis U in the separable decomposition (1.5), Theorem 1.2 only requires the existence of some such U (which is the definition of separability). In particular, the algorithm does not need to know U . Also, Theorem 1.2 is stated for convergence in terms of distance to optimality, but extends to all standard desiderata; see §4. Theorem 1.2 also extends to settings where gradient descent is performed inexactly, e.g., due to using inexact gradients; see §6.1.

Opportunities for randomized stepsizes. Randomized algorithms have variability in their trajectories. This provides potential opportunities for running multiple times in parallel and selecting the best run, see §5.3. This best run can be significantly better than the typical run (whose rate converges almost surely to the accelerated rate, see §5.2) and the worst run (whose rate can diverge but only with exponentially small probability, see §5.1).

Game-theoretic connections and lower bounds. This paper considers random stepsizes in order to accelerate GD. This randomized perspective is also helpful for the dual problem of constructing lower bounds, i.e., exhibiting functions for which gradient descent converges slowly. In particular, lifting to *distributions* over functions leads to a game between an algorithm (who chooses a distribution over stepsizes) and an adversary (who chooses a distribution over functions) which is symmetric in that the Arcsine distribution is the optimal answer for both. This leads to new perspectives on classical lower bounds for GD; details in §6.2.

Techniques and connections to logarithmic potential theory. A distinguishing feature of our analysis is a connection to potential theory: the variational characterization for the optimal stepsize distribution mirrors the variational characterization for a certain equilibrium distribution which configures a unit of charged particles so as to minimize the logarithmic potential energy. The optimal distribution satisfies a remarkable “equalization property” which in the physical context amounts to a constant potential over space, and in the optimization context amounts to an identical convergence rate over all quadratic functions. A key technical insight is that martingale arguments extend this phenomenon to all separable convex functions. While links between Krylov-subspace algorithms and potential theory are classical (see the excellent survey [19]), these connections are restricted to the special case of quadratics and moreover, even for that setting, these connections have not been exploited for the analysis of GD with random stepsizes. See §2 for a conceptual overview of our analysis techniques and how we exploit this equalization property to prove Theorem 1.2.

1.2 Organization

§2 provides a conceptual overview of our analysis approach. §3 proves our main result, Theorem 1.2. §4 discusses generalizations beyond separability, strong convexity, and smoothness. §5 investigates the variability from using randomized stepsizes, in terms of notions of convergence, high probability bounds, and opportunities for parallelized optimization. §6 provides auxiliary results about stability and game-theoretic interpretations. §7 describes several directions for future work that are motivated by our results. For brevity, several proofs are deferred to the appendix.

2 Conceptual overview: acceleration by random stepsizes

Here we provide informal derivations of our main result in order to overview the key conceptual ideas. We begin by briefly recalling in §2.1 how random stepsizes can accelerate GD for the special case of quadratic optimization. Then in §2.2, we describe the challenges for acceleration beyond quadratics, and our new techniques for circumventing these challenges.

2.1 Quadratic optimization

Consider running GD on the class \mathcal{Q} of quadratic functions f that are m -strongly convex and M -smooth. What is the optimal distribution from which to draw i.i.d. stepsizes? In order to understand this, we first recall the optimal deterministic stepsizes.

Optimal deterministic stepsizes. Without loss of generality after translating,

$$f(x) = \frac{1}{2}x^T Hx \quad \text{where} \quad mI \leq H \leq MI. \quad (2.1)$$

Since the function is quadratic, its gradient is linear $\nabla f(x) = Hx$, hence GD is a linear map

$$x_{t+1} = x_t - \alpha_t \nabla f(x_t) = (I - \alpha_t H)x_t, \quad (2.2)$$

and therefore the final iterate is

$$x_n = \prod_{t=0}^{n-1} (I - \alpha_t H)x_0. \quad (2.3)$$

The convergence rate for a given stepsize schedule—in the worst case over objectives functions in \mathcal{Q} and initializations x_0 —is therefore

$$\begin{aligned} \sup_{f \in \mathcal{Q}, x_0 \neq x^*} \left(\frac{\|x_n - x^*\|}{\|x_0 - x^*\|} \right)^{1/n} &= \sup_{mI \leq H \leq MI, x_0 \neq x^*} \left(\frac{\|\prod_{t=0}^{n-1} (I - \alpha_t H)x_0\|}{\|x_0\|} \right)^{1/n} \\ &= \sup_{mI \leq H \leq MI} \left\| \prod_{t=0}^{n-1} (I - \alpha_t H) \right\|^{1/n} \end{aligned} \quad (2.4)$$

$$= \sup_{\lambda \in [m, M]} \prod_{t=0}^{n-1} |1 - \alpha_t \lambda|^{1/n}. \quad (2.5)$$

Above, the penultimate step is by definition of the spectral norm, and the final step is by diagonalizing. It is a simple and celebrated fact, dating back to Young in 1953 [61], that minimizing this convergence rate (2.5) over stepsizes is equivalent to a certain extremal polynomial problem, and that as a consequence, the optimal stepsizes minimizing (2.5) are the inverses of the roots of a (translated and scaled) Chebyshev polynomial, in any order. Said more precisely: for any horizon length n , the optimal n stepsizes are any permutation of

$$\left\{ \left(\frac{M+m}{2} + \frac{M-m}{2} \cos \left(\frac{2t+1}{2n} \pi \right) \right)^{-1} \right\}_{t=0}^{n-1} \quad (2.6)$$

and the corresponding n -step convergence rate in (2.5) is

$$\sup_{f \in \mathcal{Q}, x_0 \neq x^*} \left(\frac{\|x_n - x^*\|}{\|x_0 - x^*\|} \right)^{1/n} = \left[2 \frac{(\sqrt{\kappa} + 1)^n (\sqrt{\kappa} - 1)^n}{(\sqrt{\kappa} + 1)^{2n} + (\sqrt{\kappa} - 1)^{2n}} \right]^{1/n} \xrightarrow{n \rightarrow \infty} \frac{\sqrt{\kappa} - 1}{\sqrt{\kappa} + 1}. \quad (2.7)$$

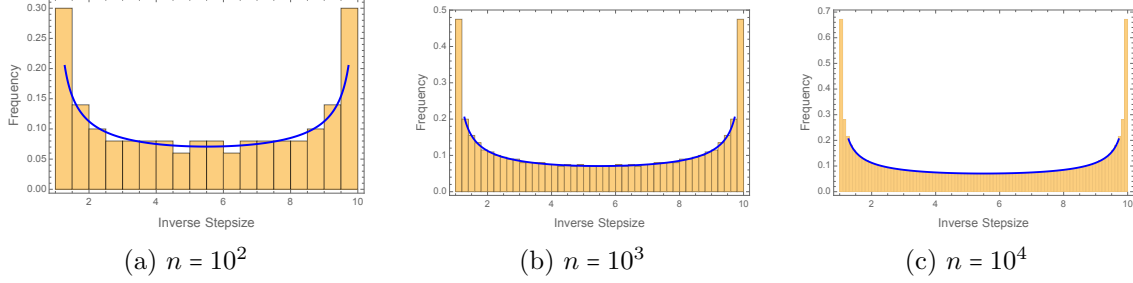


FIGURE 2: As $n \rightarrow \infty$, the empirical distribution (yellow) of the n -step inverse Chebyshev stepsize schedule (2.6) converges to the Arcsine distribution (blue). The fit is increasingly accurate as n increases from 10^2 (left) to 10^4 (right). Demonstrated here for $m = 1$ and $M = 10$.

Optimal random stepsizes. Much less well-known is the fact that this accelerated rate can also be asymptotically achieved by simply using i.i.d. stepsizes from an appropriate distribution [32, 47, 63]. The intuition is that the convergence rate (2.5) is invariant in the order of the stepsizes, instead depending only on their empirical measure, and as $n \rightarrow \infty$ one can achieve the same empirical measure (and thus the same convergence rate) by replacing the carefully chosen deterministic Chebyshev stepsize schedule (2.6) with i.i.d. stepsizes from their limiting empirical measure. See [2, Chapter 3] for a discussion of this through the lens of exchangeability and De Finetti’s Theorem.

Concretely, if the inverse stepsizes $\beta_t := 1/\alpha_t$ are i.i.d. from some distribution μ , then (forgoing measure-theoretic technicalities) the Law of Large Numbers ensures that the n -step convergence rate in (2.5) converges almost surely and in expectation to

$$\begin{aligned}
 \sup_{f \in \mathcal{Q}, x_0 \neq x^*} \left(\frac{\|x_n - x^*\|}{\|x_0 - x^*\|} \right)^{1/n} &= \sup_{\lambda \in [m, M]} \prod_{t=0}^{n-1} |1 - \alpha_t \lambda|^{1/n} \\
 &= \exp \left(\sup_{\lambda \in [m, M]} \frac{1}{n} \sum_{t=0}^{n-1} \log |1 - \beta_t^{-1} \lambda| \right) \\
 &\xrightarrow{n \rightarrow \infty} \exp \left(\sup_{\lambda \in [m, M]} \mathbb{E}_{\beta \sim \mu} \log |1 - \beta^{-1} \lambda| \right) \tag{2.8}
 \end{aligned}$$

This informal derivation suggests that the distribution μ minimizing the asymptotic convergence rate (2.8) should be the appropriate limit as $n \rightarrow \infty$ of the empirical measure for the inverse Chebyshev stepsize schedule (this limit is the Arcsine(m, M) distribution, since $a + b \cos(\Theta)$ has the Arcsine distribution on $[a - b, a + b]$ if Θ is uniformly distributed on $[0, \pi]$, see Figure 2), and moreover the optimal value should be the accelerated rate $\frac{\sqrt{\kappa-1}}{\sqrt{\kappa+1}}$. This intuition is formalized by the following lemma.

Lemma 2.1 (Extremal property of the Arcsine distribution). *The extremal problem*

$$\inf_{\mu} \sup_{\lambda \in [m, M]} \mathbb{E}_{\beta \sim \mu} \log |1 - \beta^{-1} \lambda| \tag{2.9}$$

has optimal value $\log \frac{\sqrt{\kappa-1}}{\sqrt{\kappa+1}}$, achieved by the unique optimal solution $\mu = \text{Arcsine}(m, M)$.

This lemma was proven in the optimization literature via a long, tour-de-force integral calculation [32]. In Appendix A, we present an alternative, short proof from Altschuler’s thesis [2] that argues by connecting this extremal problem for stepsize distributions to a classical extremal

problem in potential theory about electrostatic capacitors. The key feature of this connection is the following *equalization property* of the Arcsine distribution.

Lemma 2.2 (Equalization property of the Arcsine distribution). *For all $\lambda \in [m, M]$,*

$$\mathbb{E}_{\beta \sim \text{Arcsine}(m, M)} \log |1 - \beta^{-1} \lambda| = \log \frac{\sqrt{\kappa} - 1}{\sqrt{\kappa} + 1}. \quad (2.10)$$

In particular, this value does not depend on λ .

This equalization property uniquely defines the Arcsine distribution in that there is no other distribution satisfying (2.10); see Appendix A. In words, this property shows that the Arcsine distribution “equalizes” the convergence rate over all quadratic functions (parameterized by their curvatures λ). We emphasize that this is not just an algebraic identity: this equalization property also has a natural interpretation in terms of a game between an optimization algorithm (who chooses a stepsize to minimize the convergence rate) and an adversary (who chooses a quadratic function to maximize the convergence rate), see §6.2. In Appendix A we elaborate on the connections to potential theory, and leverage that connection to provide short proofs of these lemmas.

2.2 Beyond quadratic optimization

We now turn to accelerating GD beyond the special case of quadratic functions f . The central challenge is that the curvature, as measured by the Hessian $\nabla^2 f$, is no longer constant. This breaks the above quadratic analysis in several ways since it is no longer possible to summarize or re-parameterize the Hessians along the GD trajectory in terms of a single scalar $\lambda \in [m, M]$. To explain these issues, let us attempt to adapt this analysis to the general convex setting, assuming for simplicity here that f is twice differentiable (our actual analysis in §3 avoids any such assumption).

Time-varying curvature. Clearly, if f is not quadratic, then the gradient ∇f is not linear, and as a consequence the GD map is not linear. Assuming as before that $x^* = 0$ without loss of generality, one can generalize the linear GD map (2.3) to the non-quadratic setting as

$$x_n = \prod_{t=0}^{n-1} (I - \alpha_t H_t) x_0 \quad \text{where} \quad H_t := \int_0^1 \nabla^2 f(ux_t) du, \quad (2.11)$$

since by the Fundamental Theorem of Calculus, $\nabla f(x_t) = H_t x_t$ and thus $x_{t+1} = x_t - \alpha_t \nabla f(x_t) = (I - \alpha_t H_t) x_t$. The interpretation of H_t is that it is the average Hessian between $x^* = 0$ and x_t , i.e., it measures the relevant local curvature in the t -th iteration of GD. In the quadratic setting, $H_t = H$ for all t since the Hessian $\nabla^2 f \equiv H$ is constant, which recovers (2.3) from (2.11). However, for the general non-quadratic setting, the curvature H_t varies in t . This is the overarching challenge mentioned above. For the purpose of accelerating GD, this challenge manifests in two interrelated ways: the lack of *commutativity* and *predictability* of the curvature H_t along the trajectory of GD. Below we describe both challenges and how we circumvent them.

2.2.1 Issue 1: commutativity (deriving the extremal problem)

The first issue is that for convex functions, the Hessians at different points do not necessarily commute. Thus, in general, $\{H_t\}$ do not commute and therefore are not simultaneously diagonalizable. This precludes the key step (2.5) in the quadratic analysis above, which decomposes the convergence rate across eigendirections, thereby reducing the (multivariate) problem of convergence analysis to

an analytically tractable (univariate) extremal problem. This is the only reason we assume separability: it is equivalent to commuting Hessians (Lemma B.1). Concretely, separability allows us to show an analog of (2.5) in which the convergence rate of a stepsize schedule $\{\alpha_t\}$ is bounded by

$$\prod_{t=0}^{n-1} |1 - \alpha_t \lambda_t|^{1/n}, \quad (2.12)$$

the difference to the quadratic case being that there all $\lambda_t \in [m, M]$ are the same, whereas now $\lambda_t \in [m, M]$ can be time-varying and can depend on all previous stepsizes $\{\alpha_s\}_{s < t}$.

By taking an appropriate¹ limit of the n -step convergence rate (2.12), it can be shown that the (logarithm of the) asymptotic rate for i.i.d. inverse stepsizes $\beta_t := \alpha_t^{-1}$ from a distribution μ is governed by the extremal problem:

$$\inf_{\mu} \sup_{\{\mathcal{A}_t\}\text{-adapted process } \{\lambda_t \in [m, M]\}} \limsup_{n \rightarrow \infty} \frac{1}{n} \sum_{t=0}^{n-1} \mathbb{E}_{\beta_t \sim \mu} \log |1 - \beta_t^{-1} \lambda_t| \quad (2.13)$$

where $\{\mathcal{A}_t\}$ denotes the natural filtration corresponding to the inverse stepsizes $\{\beta_t\}$. Observe that (2.13) recovers the extremal problem (2.8) for the quadratic case if λ_t are constant, but in general is different because λ_t can be time-varying. This brings us to the second issue.

2.2.2 Issue 2: unpredictability (solving the extremal problem)

The second issue is that the curvature is time-varying (as measured by λ_t), thus less predictable, which makes it difficult to choose stepsizes that lead to fast convergence (as measured by the extremal problem (2.13)). The key upshot of our random stepsizes is that the randomness of β_t prevents an adversarial choice of $\lambda_t \in [m, M]$ since λ_t must be “chosen” before β_t is drawn. More precisely, we exploit the equalization property of the Arcsine distribution (Lemma 2.2) to argue that for *any choice of* λ_t , the distribution of β_t ensures that

$$\mathbb{E}_{\beta_t \sim \mu} \log |1 - \beta_t^{-1} \lambda_t| = \log \frac{\sqrt{\kappa} - 1}{\sqrt{\kappa} + 1}. \quad (2.14)$$

Intuitively, this can be interpreted as using random stepsizes to *hedge* between the possible curvatures λ_t : the expected “progress” from an iteration is the same regardless of the curvature. Quantitatively, this equalization property (2.14) lets us argue an accelerated convergence rate using a martingale version of the Law of Large Numbers. In particular, this implies that the Arcsine distribution—which was optimal for the extremal problem (2.9) in the quadratic setting—achieves the same value for the extremal problem (2.13) in the convex setting, and therefore is also optimal for this setting (as it is more general).

Remark 2.3 (Overcoming unpredictability via random hedging). *Observe that in the extremal problem (2.13), we allow $\{\lambda_t\}$ to vary arbitrarily in t , and in particular we do not enforce these curvatures to be consistent with a single convex function. This consistency is the defining feature*

¹A subtle but important technical issue is that one cannot simply use the Law of Large Numbers to derive the asymptotic convergence rate from the n -step rate (c.f., (2.8) for how that worked in the quadratic case). Indeed, in the convex case here, the relevant empirical sum $\frac{1}{n} \sum_{t=0}^{n-1} \log |1 - \beta_t^{-1} \lambda_t|$ has non-i.i.d. summands: they are neither identically distributed (since λ_t is time-varying) nor independent (since λ_t can depend on previous stepsizes). Briefly, we overcome the identically distributed issue by using the equalization property of the Arcsine distribution (Lemma 2.2) to show that these summands still have the desired expectation, and we overcome the independence issue by using a martingale-based argument to show that the fresh randomness in α_t makes these summands “independent enough” to prove a specially tailored version of Kolmogorov’s Strong Law of Large Numbers for this setting. Details in §3.1.

of Performance Estimation Problems [20, 53] and is essential to the line of work on accelerating gradient descent using deterministic stepsize schedules [2, 3, 4, 8, 12, 13, 23, 26, 27, 28, 29, 58, 62]. Remarkably, random stepsizes obtain the fully accelerated rate without needing to enforce any consistency conditions. This can be interpreted through the lens of “hedging” between worst-case functions, as described in [2]. On one hand, the opportunity for faster convergence via deterministic time-varying stepsizes arises due to the rigidity of convex functions: a function being worst-case for the stepsize in one iteration may constrain its curvature from being worst-case in the next iteration. On the other hand, the opportunity for faster convergence via random stepsizes arises because the worst-case function does not “know” the realization of the random stepsizes. It is an interesting question if these two opportunities can be combined, e.g., can randomized analyses also exploit consistency conditions on $\{\lambda_t\}$ in order to generalize beyond separability? See §7.

3 Proof of main result

Here we prove Theorem 1.2. We build up to the full result by first proving it for univariate objectives, and then extending to the general (separable) multivariate setting. See §2 for a high-level overview of the analysis techniques and the two key conceptual challenges regarding unpredictability of the Hessian (dealt with in the univariate proof via the equalization property of the Arcsine distribution) and commutativity of the Hessian (dealt with in the multivariate proof via separability). For shorthand, we denote the accelerated rate by

$$R_{\text{acc}} := \frac{\sqrt{\kappa} - 1}{\sqrt{\kappa} + 1}. \quad (3.1)$$

3.1 Univariate setting

We begin by remarking that unforeseen complications can occur even in the univariate setting. For example, Polyak’s Heavy Ball method fails to converge for such functions [37, §4.6].

We first state a simple helper lemma that lets us avoid assumptions of twice-differentiability when expanding the convergence rate of GD in the form (2.12) described in the overview §2.2.

Lemma 3.1. *Suppose that $f : \mathbb{R} \rightarrow \mathbb{R}$ is m -strongly convex and M -smooth, and let x^* denote its minimizer. Then for any $x \neq x^*$, it holds that*

$$m \leq \frac{f'(x)}{x - x^*} \leq M.$$

For intuition, if f is twice-differentiable, then this lemma follows immediately from the Intermediate Value Theorem since $f'(x) = f''(z)(x - x^*)$ for some z between x and x^* , and the Hessian $f''(z) \in [m, M]$ by m -strong convexity and M -smoothness. The proof is not much more difficult without the assumption of twice-differentiability.

Proof. For the upper bound, M -smoothness of f implies $|f'(x)| = |f'(x) - f'(x^*)| \leq M|x - x^*|$. Note that $\text{sign}(f'(x) - f'(x^*)) = \text{sign}(x - x^*)$ since f is univariate. Thus, by re-arranging, we obtain the desired inequality $f'(x)/(x - x^*) \leq M$. The lower bound follows from an analogous argument, since m -strong convexity of f implies $|f'(x) - f'(x^*)| \geq m|x - x^*|$. \square

Proof of Theorem 1.2 for univariate f . Let β_t denote the inverse stepsizes; these are i.i.d. Arcsine(m, M). Defining $\lambda_t := \frac{f'(x_t)}{x_t - x^*}$, we have by definition of GD that

$$x_{t+1} - x^* = x_t - x^* - \beta_t^{-1} f'(x_t) = (1 - \beta_t^{-1} \lambda_t)(x_t - x^*).$$

Therefore for any $n \in \mathbb{N}$, the average rate after n steps is

$$R_n := \left| \frac{x_n - x^*}{x_0 - x^*} \right|^{1/n} = \left| \prod_{t=0}^{n-1} (1 - \beta_t^{-1} \lambda_t) \right|^{1/n}. \quad (3.2)$$

Note that this sequence of random variables is uniformly integrable. Indeed for each $n \in \mathbb{N}$, a crude bound gives $|\prod_{t=0}^{n-1} (1 - \beta_t^{-1} \lambda_t)|^{1/n} \leq \max_{\beta, \lambda \in [m, M]} |1 - \beta^{-1} \lambda| = \kappa - 1 < \infty$. Therefore, by the Dominated Convergence Theorem, L_1 convergence of R_n to R_{acc} will follow immediately from almost-sure convergence. To show the latter, note that since almost-sure convergence is preserved under continuous mappings (in particular the $\exp(\cdot)$ function), it suffices to show that

$$\frac{1}{n} \sum_{t=0}^{n-1} Z_t \xrightarrow{\text{a.s.}} \log R_{\text{acc}},$$

where we define

$$Z_t := \log |1 - \beta_t^{-1} \lambda_t|.$$

As mentioned in the overview section 2.2, the issue with applying the Law of Large Numbers is that $\{Z_t\}$ are not i.i.d. (For example, even though we show below that these random variables have the desired first moments—which allows us to adapt the Law of Large Numbers—one can check that $\{Z_t\}$ are neither independent nor identically distributed.)

Now observe that λ_t depends on the realizations of only the *previous* $\beta_0, \dots, \beta_{t-1}$. Moreover, λ_t lies inside $[m, M]$ by Lemma 3.1. Thus by the equalization property of the Arcsine distribution (Lemma 2.2), $\mathbb{E}[Z_t] = \log R_{\text{acc}}$ for all t , even if λ_t is chosen adversarially as a function of the previous $\beta_0, \dots, \beta_{t-1}$. Indeed, by the independence of $\{\beta_t\}$ and the Law of Iterated Expectations:

$$\begin{aligned} \mathbb{E}[Z_t] &= \mathbb{E}_{\beta_0, \dots, \beta_t} \left[\log |1 - \beta_t^{-1} \lambda_t| \right] \\ &= \mathbb{E}_{\beta_0, \dots, \beta_{t-1}} \left[\mathbb{E}_{\beta_t} \left[\log |1 - \beta_t^{-1} \lambda_t| \mid \beta_0, \dots, \beta_{t-1} \right] \right] \\ &= \mathbb{E}_{\beta_1, \dots, \beta_{t-1}} \left[\log R_{\text{acc}} \right] \\ &= \log R_{\text{acc}}. \end{aligned} \quad (3.3)$$

Therefore it suffices to show

$$\frac{1}{n} \sum_{t=0}^{n-1} (Z_t - \mathbb{E}[Z_t]) \xrightarrow{\text{a.s.}} 0. \quad (3.4)$$

We show this by adapting the standard proof of Kolmogorov's Strong Law of Large Numbers, see e.g. [59] for the classical proof. We first argue that

$$M_n := \sum_{t=0}^{n-1} \frac{Z_t - \mathbb{E}[Z_t]}{t+1}$$

forms a martingale w.r.t. its own filtration $\{F_n\}$, as stated in the following lemma. Intuitively this follows from the calculation (3.3); a formal proof is deferred to Appendix B.2.

Lemma 3.2 (Martingale helper lemma). $\{M_n\}$ is an L_1 -bounded martingale w.r.t. $\{F_n\}$.

Thus we may apply the Martingale Convergence Theorem (see e.g., [59, §11.5]) to establish that $M_\infty := \lim_{n \rightarrow \infty} M_n$ exists and is finite almost surely. Therefore by Kronecker's Lemma (recalled below), we conclude that $\frac{1}{n} \sum_{t=0}^{n-1} (Z_t - \mathbb{E}[Z_t]) \xrightarrow{\text{a.s.}} 0$, proving (3.4) as desired. \square

Lemma 3.3 (Kronecker’s Lemma). *Let $\{z_n\}$ and $\{b_n\}$ be two real-valued sequences such that each b_n is strictly positive, the b_n are strictly increasing to ∞ , and $\sum_{t=0}^{\infty} \frac{z_t}{b_t}$ converges to a finite limit. Then $\lim_{n \rightarrow \infty} \frac{1}{b_n} \sum_{t=0}^{n-1} z_t = 0$.*

3.2 Multivariate setting

Here we prove Theorem 1.2 in the general multivariate setting. The proof follows quickly from the univariate setting proved above and the definition of separability.

Proof of Theorem 1.2. First, note that it suffices to prove the theorem for functions which are separable in the standard basis, i.e., functions f of the form

$$f(x) = \sum_{i=1}^d f_i([x]_i). \quad (3.5)$$

Indeed, consider a general separable function $f(x) = \sum_{i=1}^d f_i([Ux]_i)$ where $U \in \mathbb{R}^{d \times d}$ is orthogonal. Then by the Chain Rule, $\nabla f = U^T \nabla h$ where $h(y) := \sum_{i=1}^d f_i([y]_i)$. Hence running GD $x_{t+1} = x_t - \alpha_t \nabla f(x_t)$ on f is equivalent to running GD $y_{t+1} = y_t - \alpha_t \nabla h(y_t)$ on h , in the sense that $y_t = Ux_t$ for all t . Thus, modulo an isometric change of coordinates which does not affect the convergence rate, we may assume henceforth that the objective function is separable of the form (3.5).

Since $f(x) = \sum_{i=1}^d f_i([x]_i)$, the i -th coordinate of $\nabla f(x)$ is $f'_i([x]_i)$. Thus the GD update $x_{t+1} = x_t - \beta_t^{-1} \nabla f(x_t)$ can be re-written coordinate-wise as

$$[x_{t+1}]_i = [x_t]_i - \beta_t^{-1} f'_i([x_t]_i), \quad \forall i \in [d]. \quad (3.6)$$

Now we make two observations. First, (3.6) is the update that GD would perform to minimize f_i from initialization $[x_0]_i$. Second, each f_i is m -strongly convex and M -smooth (because strong convexity and smoothness are preserved under restrictions). Thus we may apply the univariate setting of Theorem 1.2 (proved above) to argue that the asymptotic convergence rate is

$$\lim_{n \rightarrow \infty} \left(\frac{|[x_n]_i - [x^*]_i|}{|[x_0]_i - [x^*]_i|} \right)^{1/n} \stackrel{\text{a.s.}}{=} R_{\text{acc}}$$

for every coordinate $i \in S := \{i \in [d] : [x_0]_i \neq [x^*]_i\}$ for which the initialization is not already optimal. Note that $x_0 \neq x^*$ ensures S is non-empty. Since $[x^*]_i$ is the unique minimizer of f_i , and since also the number of coordinates $i \in S$ is finite, it follows that

$$\lim_{n \rightarrow \infty} \left(\frac{\|x_n - x^*\|}{\|x_0 - x^*\|} \right)^{1/n} \stackrel{\text{a.s.}}{=} \lim_{n \rightarrow \infty} \max_{i \in S} \left(\frac{|[x_n]_i - [x^*]_i|}{|[x_0]_i - [x^*]_i|} \right)^{1/n} \stackrel{\text{a.s.}}{=} R_{\text{acc}}.$$

This establishes the desired almost-sure convergence. L_1 convergence then follows from the Dominated Convergence Theorem, since we have uniform integrability by a similar argument as in the univariate setting above. \square

4 Generalizations

The purpose of this paper is to show that random stepsizes can accelerate GD on classes of strongly convex and smooth functions that are more general than quadratics. For concreteness, Theorem 1.2 focuses on functions that are separable, strongly convex, and smooth. But this result extends to

other classes of functions. We describe two extensions below, thematically splitting relaxations of the separability assumption (§4.1) from the strong convexity and smoothness assumptions (§4.2). These two types of relaxations are complementary in the sense that they individually deal with the two key issues discussed in §2.2: commutativity and unpredictability of the Hessian, respectively.

Remark 4.1 (Performance metrics). *We also mention that throughout, we measure convergence via the contraction factor $(\frac{\|x_n - x^*\|}{\|x_0 - x^*\|})^{1/n}$. The results are unchanged if the progress measure is changed from distance to function suboptimality or gradient norm, at either the initialization x_0 , termination x_n , or both. This is because for κ -conditioned functions, all these progress measures $\|x - x^*\|^2$, $f(x) - f(x^*)$, and $\|\nabla f(x)\|^2$ are equivalent up to a multiplicative factor c that is independent of n (it depends only on the strong convexity and smoothness parameters), and therefore in the limit $n \rightarrow \infty$, the convergence rate is unaffected by changing these performance metrics since $\lim_{n \rightarrow \infty} c^{1/n} = 1$.*

4.1 Beyond separability

Recall from Definition 1.1 that separability amounts to decomposability into univariate functions. Theorem 1.2 holds more generally for functions which are “radially-separable” in that they decompose into radial functions. This is formally defined as follows.

Definition 4.2 (Radial functions). *A function $r : \mathbb{R}^d \rightarrow \mathbb{R}$ is radial if it admits a decomposition of the form*

$$r(x) = h(\|x - x^*\|) \tag{4.1}$$

for some function $h : \mathbb{R} \rightarrow \mathbb{R}$ and point $x^* \in \mathbb{R}^d$.

Definition 4.3 (Radially-separable functions). *A function $f : \mathbb{R}^d \rightarrow \mathbb{R}$ is radially-separable if it admits a decomposition of the form*

$$f(x) = \sum_{i=1}^b r_i([Ux]_{S_i}), \tag{4.2}$$

for some orthogonal matrix $U \in \mathbb{R}^{d \times d}$, partition S_1, \dots, S_b of $[d]$, and radial functions r_1, \dots, r_b . (Above, $[Ux]_{S_i}$ is shorthand for the subset of coordinates of Ux in set S_i .)

Clearly, radial-separability generalizes both separability (consider the partition $S_i = \{i\}$) and radially (consider the trivial partition $S_1 = [d]$).

Theorem 4.4 (Random stepsizes accelerate GD on radially-separable functions). *Theorem 1.2 is unchanged if the assumption of separability is replaced with radial-separability.*

Proof. The proof is mostly similar to the proof of Theorem 1.2. By the same argument (see §3.2), it suffices to prove the accelerated convergence rate for a block of the decomposition. So, without loss of generality we may assume that f is radial. In words, the key idea behind the rest of the proof is that although radial functions do not have commutative Hessians—which we exploited in our main result to reduce the (multivariate) problem of convergence analysis to an analytically tractable (univariate) extremal problem, see the overview §2.2—radial functions are univariate in nature (in particular all gradients in a GD trajectory are colinear), and this enables a similar reduction.

Formally, note that the center x^* of a convex radial function is a minimizer. (Indeed, for any $x \in \mathbb{R}^d$, we have by radially and then Jensen’s inequality that $f(x) = \mathbb{E}_z f(z) \geq f(x^*)$, where the

expectation is over z drawn uniformly at random from the sphere of radius $\|x - x^*\|$ around x^* .) Thus after a possible translation, we may assume without loss of generality that $x^* = 0$, so that

$$f(x) = h(\|x\|^2)$$

for some univariate function $h : \mathbb{R} \rightarrow \mathbb{R}$. The chain rule implies $\nabla f(x) = 2h'(\|x\|^2)x$, hence the GD update $x_{t+1} = x_t - \alpha_t \nabla f(x_t)$ expands to

$$x_n = \prod_{t=0}^{n-1} (1 - \alpha_t \lambda_t) x_0 \quad \text{where} \quad \lambda_t := 2h'(\|x_t\|^2). \quad (4.3)$$

It is straightforward to show that $\lambda_t \in [m, M]$, e.g., this follows from Lemma 3.1 after restricting f to the line between x^* and x_t . Thus this gives an alternative approach for reducing the (multivariate) convergence rate to the same (univariate) product form (2.12), at which point the rest of our analysis goes through unchanged. \square

4.2 Beyond strong convexity and smoothness

Recall that for twice-differentiable functions f , the assumptions of m -strong convexity and M -smoothness are equivalent to the assumption that the spectrum of $\nabla^2 f(x)$ lies in the interval $[m, M]$ for all points x . The algorithmic techniques developed in this paper extend to functions f whose Hessians have more structured spectra, lying within some set $E \subset \mathbb{R}_{\geq 0}$ that is not necessarily an interval. Note that (strong²) convexity of f ensures a (strictly) positive real spectrum. Such extensions to more structured Hessians have led to improved results for other algorithms in quadratic settings and beyond, see e.g., [24, 34, 43].

This extension is rooted in connections to logarithmic potential theory, and proceeds as follows. Recall from the overview §2 that the asymptotic convergence rate of GD with inverse stepsizes i.i.d. from a distribution μ , is dictated by the maximum logarithmic potential of μ over $[m, M]$, i.e.,

$$\sup_{\lambda \in [m, M]} \mathbb{E}_{\beta \sim \mu} \log |1 - \beta^{-1} \lambda|, \quad (4.4)$$

and that the unique distribution minimizing this quantity is the Arcsine distribution (Lemma 2.1). This extends to spectral sets $E \subset \mathbb{R}_{\geq 0}$ beyond intervals with little change. Indeed, by an identical derivation, the convergence rate is dictated by the analogous extremal problem

$$\sup_{\lambda \in E} \mathbb{E}_{\beta \sim \mu} \log |1 - \beta^{-1} \lambda|, \quad (4.5)$$

and the unique distribution minimizing this quantity is the unique optimal stepsize distribution. What is this optimal distribution? This is a classical topic in potential theory and there is a well-developed machinery (typically based on computing conformal mappings, Green's functions, or Fekete sets) for determining this distribution explicitly in closed-form when possible (typically this only occurs for very simple sets) or otherwise approximating it; we refer the interested reader to the surveys and textbooks [19, 48, 49, 50].

²In the case of (non-strongly) convex, smooth optimization, we have $m = 0$ and $E = [0, M]$. For that case, much of the ideas in this paper still apply at a high level, but the singularity at $\lambda = 0$ makes the problem degenerate in ways that must be addressed with refined versions of these techniques. This will be discussed in a later paper.

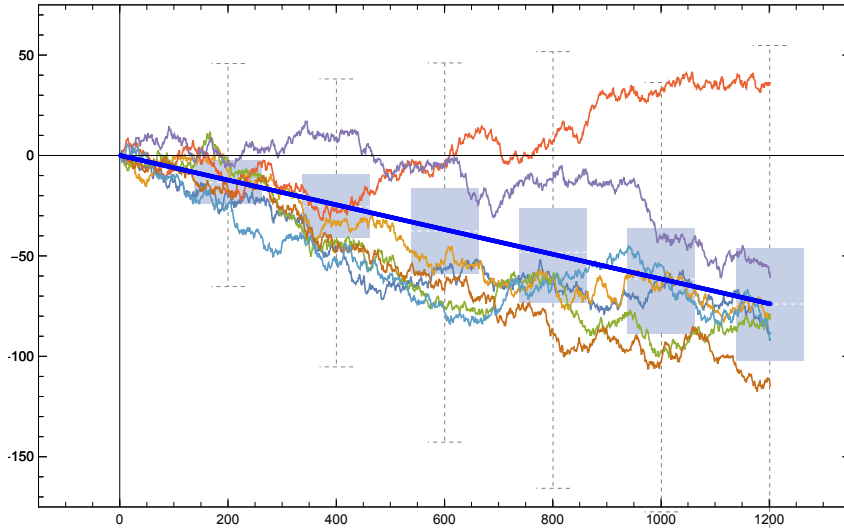


FIGURE 3: The randomness of the proposed stepsizes leads to variability between runs. Illustrated here with a boxplot for 500 runs (7 full trajectories shown) from the same initialization on a univariate quadratic function, with $\kappa = 200$; similar phenomena occur for other problem instances. The horizontal axis is the number of iterations n ; the vertical axis is $\log_{10} \frac{\|x_n - x^*\|}{\|x_0 - x^*\|}$. The worst run can diverge exponentially, but this occurs with exponentially low probability (details in §5.1). The median run converges exponentially at the optimal accelerated rate R_{acc} (Theorem 1.2). The fit is extremely precise, as shown by the bolded blue line. The best run can be substantially better than the median run, leading to the possibility of faster parallel optimization (details in §5.3).

5 Opportunities and perils of random stepsizes

Here we investigate the implications of the variability arising from random stepsizes, see Figure 3. In §5.1, we show that trajectories can diverge but only with exponentially low probability. In §5.2, we investigate typical trajectories, showing that with overwhelming probability, their rate converges to the accelerated rate. In §5.3, we discuss running multiple trajectories in parallel and selecting the best.

5.1 Divergent runs and notion of convergence

When considering randomized optimization algorithms, there are several notions for measuring the convergence rate. Two natural ones are:

$$\mathbb{E} \left[\left(\frac{\|x_n - x^*\|}{\|x_0 - x^*\|} \right)^{1/n} \right] \quad \text{vs} \quad \mathbb{E} \left[\frac{\|x_n - x^*\|}{\|x_0 - x^*\|} \right]^{1/n}. \quad (5.1)$$

These are different. Indeed, our main result shows that by running GD with random stepsizes, the former quantity is asymptotically equal to the accelerated rate $R_{\text{acc}} = \frac{\sqrt{\kappa}-1}{\sqrt{\kappa}+1} < 1$. Whereas the latter quantity can be larger than 1 (details in Appendix B.3). Why the discrepancy between these two notions of convergence rate? Which is relevant for optimization?

The discrepancy arises because the error $\|x_n - x^*\|$ is exponentially small with overwhelming probability (leading to our convergence result for the first notion), but can be exponentially large with exponentially small probability (leading to the instability for the second notion). *The key point is that for optimization algorithms, it is very desirable to succeed with overwhelming probability,*

even if performance is poor in the unlikely complement event. Simply re-run the algorithm if that happens. This why we study the quantity on the left hand side of (5.1). For intuition, we conclude this section with an illustrative toy example that explains this discrepancy in its simplest form.

Example 5.1 (Measuring convergence when exponentially rare trajectories can be exponentially large). Consider i.i.d. random variables Z_t which are each equally likely to be $1/9$ or 3 . Then $(\prod_{t=1}^n Z_t)^{1/n} = 3^{\frac{1}{n} \sum_{t=1}^n \log_3 Z_t} \rightarrow 3^{\mathbb{E}[\log_3 Z_t]} = 3^{-1/2} < 1$ by the Law of Large Numbers, hence

$$\lim_{n \rightarrow \infty} \mathbb{E} \left[\left(\prod_{t=1}^n Z_t \right)^{1/n} \right] = \frac{1}{\sqrt{3}} < 1. \quad (5.2)$$

Whereas

$$\lim_{n \rightarrow \infty} \mathbb{E} \left[\prod_{t=1}^n Z_t \right]^{1/n} = \mathbb{E}[Z_1] = \frac{14}{9} > 1. \quad (5.3)$$

In words, this discrepancy arises because $\prod_{t=1}^n Z_t$ can be exponentially large in n , but the probability of this is exponentially small. For example, $Z_1 = \dots = Z_n = 3$ occurs with probability 2^{-n} and thus by itself, this event contributes $\lim_{n \rightarrow \infty} (3^n 2^{-n})^{1/n} = 3/2 > 1$ to the expectation in (5.3). Whereas the $1/n$ normalization inside the expectation in (5.2) dampens the effect of exponentially rare blow-ups, and thus agrees with the exponentially fast convergence that occurs with overwhelming probability.

5.2 Typical runs and high probability bounds

Theorem 1.2 establishes that random stepsizes lead to an accelerated rate in the limit as the number of iterations $n \rightarrow \infty$. The same analysis arguments can also establish non-asymptotic convergence rates that apply for any finite n and hold with high probability. Indeed, recall from (3.2) that our argument in §3 relates the convergence rate (for any finite n) to an empirical sum, namely

$$\left| \frac{x_n - x^*}{x_0 - x^*} \right|^{1/n} = \exp \left(\frac{1}{n} \sum_{t=0}^{n-1} \log |1 - \beta_t^{-1} \lambda_t| \right). \quad (5.4)$$

Now, instead of using the Law of Large Numbers to argue that this empirical sum converges to its expectation (giving the accelerated rate), one can use concentration-of-measure to argue about the deviations from its expectation.

Below we state such a result. This is established by bounding the variance of the aforementioned empirical sum via a martingale-type argument and applying Chebyshev's inequality; and then extending this convergence rate to general multivariate objectives f by a union bound. Full proof details are deferred to Appendix B.4. We note that one can show more refined bounds by bounding higher-order moments or the moment generating function (see, e.g., the textbooks [17, 55, 57]), but that requires computing involved logarithmic integrals and thus for simplicity of exposition, here we state only this simple illustrative result.

Theorem 5.2 (Non-asymptotic version of Theorem 1.2). Consider the setup of Theorem 1.2. For any deviation $\varepsilon > 0$ and error probability $\delta > 0$,

$$\mathbb{P} \left(\left(\frac{\|x_n - x^*\|}{\|x_0 - x^*\|} \right)^{1/n} \geq e^\varepsilon R_{\text{acc}} \right) \leq \delta \quad (5.5)$$

if $n \geq (\pi^2 + o_\kappa(1)) \frac{d}{\delta \varepsilon^2}$.

This implies that for a typical run, the convergence rate is close to the accelerated rate; see Figure 3 for an illustration.

5.3 Best runs and parallelization

The randomness in the Arcsine Stepsize Schedule leads to a spread of convergence rates. Often the best run is substantially better than the median, which in turn is substantially better than the worst run; see Figure 3. This suggests a natural approach for parallelized optimization: run this randomized algorithm on p different processors with i.i.d. stepsizes, and take the best final iterate. One can also perform a more refined parallelized scheme: synchronize every k iterations by restarting all p processors from the best current iterate.

How much of a speedup does this parallelization give? There are two sources of tension. First, the convergence rate R_{acc} achieved by Theorem 1.2 for $p = 1$ processor is already minimax-optimal, but parallelization helps with instance-optimality which is important in practical settings, i.e., beyond worst-case settings. The second source of tension is that the results in §5.2 show that each of the p i.i.d. runs has convergence rate $R_n := (\|x_n - x^*\|/\|x_0 - x^*\|)^{1/n}$ concentrating around R_{acc} for n large, which means that R_n will be asymptotically unaffected by parallelization in the regime of fixed p and $n \rightarrow \infty$. However, for regimes where p and n scale together, there are deviations from this concentration, and the maximal such deviation in the p runs is precisely the speedup of this parallelization. See below for a calculation that illustrates this tension in a simple setting. Characterizing the tradeoff between the number of processors p , the synchronization rate k , the number of iterations n , the condition number κ , and the precision ε is an interesting direction for future work but a full understanding is out of the scope of this paper.³

Simple illustrative calculation. Here, we provide a short back-of-the-envelope calculation in order to demonstrate these tensions and the benefit of parallelization in the simple setting of univariate quadratic functions. For simplicity consider a single synchronization (i.e., $k = n$); the general case of multiple synchronizations ($k < n$) follows by replacing n by k in this calculation, and then repeating this convergence rate n/k times. For shorthand, let $x_t^{(j)}$ denote the t -th iterate on the j -th machine, let $\alpha_t^{(j)}$ denote the corresponding stepsize, and let $f(x) = \frac{\lambda}{2}(x - x^*)^2$ for $\lambda \in [m, M]$ (all univariate κ -conditioned quadratics are of this form).

Step 1: single machine. By the same analysis as in our main results (see §2.1 or §3.1) and then a Central Limit Theorem approximation, the convergence rate on the j -th machine is

$$\left(\frac{|x_n^{(j)} - x^*|}{|x_0^{(j)} - x^*|} \right)^{1/n} = \exp\left(\frac{1}{n} \sum_{t=0}^{n-1} \log |1 - \alpha_t^{(j)} \lambda| \right) \approx \exp\left(\log R_{\text{acc}} + \frac{\sigma}{\sqrt{n}} \xi_j \right) = R_{\text{acc}} \cdot \exp\left(\frac{\sigma}{\sqrt{n}} \xi_j \right),$$

where $\xi_j \sim \mathcal{N}(0, 1)$ and $\sigma = \text{Var}_\alpha(\log |1 - \alpha \lambda|)$. Note that the error of this approximation can be tightly controlled using standard techniques such as the Berry-Esseen CLT [57, Theorem 2.1.13].

Step 2: best machine. Taking the minimum over all p processors gives the convergence rate of the parallelized scheme which takes the best final iterate:

$$\min_{j \in [p]} \left(\frac{|x_n^{(j)} - x^*|}{|x_0^{(j)} - x^*|} \right)^{1/n} \approx R_{\text{acc}} \cdot \exp\left(\frac{\sigma}{\sqrt{n}} \min_{j \in [p]} \xi_j \right) \approx R_{\text{acc}} \cdot \exp\left(-\sigma \sqrt{\frac{2 \log p}{n}} \right).$$

³Interesting phenomena arise when jointly optimizing k and p . For example, if $k = 1$, then moderate parallelism ($p > 1$) improves the rate in both practice (see Figure 3) and theory (see below), but $p \rightarrow \infty$ recovers exact line search which provably does not accelerate, see e.g., [9, §9.3]. This tradeoff warns of a potential drawback of over-parallelizing. On the other hand, when $k > 1$, the $p \rightarrow \infty$ parallelization is a familiarly good algorithm: it produces the optimal point in the k -dimensional Krylov subspace, which yields minimax-optimal worst-case convergence rates and is also instance-adaptive; in fact it is equivalent to Conjugate Gradients in exact arithmetic for quadratics.

Above, the final step is due to the well-known fact that the minimum of p i.i.d. standard Gaussians concentrates tightly around $-\sqrt{2\log p}$, see, e.g., [57, Chapter 2.5]. This factor $\exp(-\sigma\sqrt{(2\log p)/n})$ quantifies the speedup of parallelization since it is the size of the variation in Figure 3. By a Taylor approximation of $\log R_{\text{acc}}$, this parallelization provides asymptotic speedups when $\sigma\sqrt{(2\log p)/n} \gtrsim |\log R_{\text{acc}}| = |\log \frac{\sqrt{\kappa-1}}{\sqrt{\kappa+1}}| \approx \frac{1}{\sqrt{\kappa}}$, or equivalently when the number of processors $p \gtrsim \exp(\frac{n}{2\kappa\sigma^2})$.

6 Discussion and related results

6.1 Stability

Throughout, we have assumed for simplicity that all calculations are performed with exact gradients and in exact arithmetic. It is of practical importance to understand how inexactness affects convergence, and there is a long line of work on such questions for first-order algorithms, see, e.g., [1, 14, 16, 18, 36, 46, 51] and the references within. There are many models of inexact computation. Here, as an illustrative example, we study one common model: the *multiplicative error model for gradients* [16, 46]. Specifically, suppose that all calculations are still performed in exact arithmetic, but GD performs updates using an estimate \tilde{g}_t of the true gradient $\nabla f(x_t)$, where

$$\|\tilde{g}_t - \nabla f(x_t)\| \leq \varepsilon \|\nabla f(x_t)\|. \quad (6.1)$$

Theorem 6.1 (Inexact-gradient version of Theorem 1.2). *Consider the setup of Theorem 1.2. Suppose that GD updates in each iteration using \tilde{g}_t rather than $\nabla f(x_t)$, i.e.,*

$$x_{t+1} = x_t - \alpha_t \tilde{g}_t, \quad (6.2)$$

where \tilde{g}_t are inexact gradients satisfying the multiplicative-error assumption (6.1). Then the convergence rate of GD with i.i.d. Arcsine stepsizes is almost surely bounded by

$$\lim_{n \rightarrow \infty} \left(\frac{\|x_n - x^*\|}{\|x_0 - x^*\|} \right)^{1/n} \leq \underbrace{\frac{1}{1 + \tilde{\varepsilon} - \sqrt{(1 + \tilde{\varepsilon})^2 - 1}}}_{\text{slowdown from inexact gradients}} \cdot \underbrace{\frac{\sqrt{\kappa} - 1}{\sqrt{\kappa} + 1}}_{\text{accelerated rate } R_{\text{acc}}} \quad \text{where} \quad \tilde{\varepsilon} := \frac{2M}{M - m} \varepsilon. \quad (6.3)$$

Moreover this bound is tight, in that it holds with equality for some choice of f and inexact gradients.

The proof is deferred to Appendix B.5, since it is similar to that of Theorem 1.2. Conceptually, the only difference is that the average Hessians $\lambda_t \in [m, M]$ are now perturbed, and thus the convergence rate is governed by $\exp(\mathbb{E}_{\beta \sim \text{Arcsine}(m, M)} \log |1 - \beta^{-1} \tilde{\lambda}|)$ where $\tilde{\lambda}$ can now be slightly outside the interval $[m, M]$.

We remark that for small error ε , the slowdown factor in Theorem 6.1 is of the order

$$\frac{1}{1 + \tilde{\varepsilon} - \sqrt{(1 + \tilde{\varepsilon})^2 - 1}} = 1 + \sqrt{2\tilde{\varepsilon}} + \Theta(\tilde{\varepsilon}) = 1 + \sqrt{\frac{4M\varepsilon}{M - m}} + \Theta(\varepsilon).$$

In other words, for large condition number $\kappa \gg 1$, the number of iterations required by GD increases by a multiplicative factor of roughly $1 + 2\sqrt{\varepsilon}$ when using inexact gradients. Note that this slowdown is more than the $(1 + \Theta(\varepsilon))$ slowdown for constant-stepsize schedules [16, Theorem 5.3]. This is a manifestation of the general phenomenon that accelerated algorithms are often less stable to model mis-specification, see e.g., [18].

We also remark that by a nearly identical argument, Theorem 6.1 can be extended to a stochastic error models, i.e., where $\tilde{g}_t - \nabla f(x_t)$ are random. The slowdown factor then becomes an averaged version of the quantity in Theorem 6.1. However, this is more complicated to state, hence for simplicity the above discussion focuses just on the deterministic error model.

6.2 Game-theoretic connections and lower bounds

This paper considers random stepsizes for the algorithmic purpose of accelerating GD. It turns out that this randomized perspective is also helpful for the dual problem of constructing algorithmic lower bounds. We describe this duality here.

Recall from §2 that the optimal distribution μ over i.i.d. inverse stepsizes $\beta = \alpha^{-1}$ has the following extremal characterization:

$$\inf_{\mu \in \mathcal{M}(E)} \sup_{\lambda \in E} \mathbb{E}_{\beta \sim \mu} \log |1 - \beta^{-1} \lambda|. \quad (6.4)$$

Above, μ is constrained to $E := [m, M]$ without loss of generality (since this holds at optimality), $\lambda \in E$ should be interpreted as the curvature of a “worst-case” univariate quadratic function $f(x) = \frac{\lambda}{2}x^2$, and $\mathbb{E}_{\beta \sim \nu} \log |1 - \beta^{-1} \lambda|$ is the logarithm of the asymptotic convergence rate.

To bridge this problem of accelerating GD with the problem of constructing lower bounds, we consider a standard two-player zero-sum game where players choose their actions simultaneously and independently. In our case, one player chooses inverse stepsizes β^{-1} and the other player chooses a quadratic function (parameterized by its curvature λ). By lifting to mixed (i.e., randomized) strategies, one player chooses a distribution μ over stepsizes, the other chooses a distribution ν over quadratic functions, with payoff given by $\mathcal{R}(\mu, \nu) := \mathbb{E}_{\beta \sim \mu, \lambda \sim \nu} \log |1 - \beta^{-1} \lambda|$.

The key observation is that due to the specific payoff function, the game is essentially symmetric between the two players (modulo inversion), and as a consequence, so are the optimal mixed strategies for both players. Below, let $\text{FlippedArcsine}(m, M)$ denote the law of λ when $\lambda^{-1} \sim \text{Arcsine}(M^{-1}, m^{-1})$.

Lemma 6.2 (Optimal distribution for worst-case functions). *The (minimax) value of this game is $\log R_{\text{acc}}$, i.e., we have the saddle point condition*

$$\sup_{\nu} \mathcal{R}(\mu^*, \nu) = \mathcal{R}(\mu^*, \nu^*) = \inf_{\mu} \mathcal{R}(\mu, \nu^*)$$

with

$$\mathcal{R}(\mu^*, \nu^*) = \mathbb{E}_{\beta \sim \mu^*, \lambda \sim \nu^*} \log |1 - \beta^{-1} \lambda| = \log R_{\text{acc}},$$

and the unique optimal mixed strategies are $\mu^* = \text{Arcsine}(m, M)$ and $\nu^* = \text{FlippedArcsine}(m, M)$.

Remark 6.3 (Game-theoretic interpretation of equalization). *Notice that both optimal distributions have full support. As a consequence, at equilibrium a player is indifferent among all of their pure actions, i.e., the equalization property (Lemma 2.2) must hold.*

Proof sketch for Lemma 6.2. Lemma 2.1 implies that $\mu^* = \text{Arcsine}(m, M)$ is the unique optimal solution for the stepsize player, and that the value of the game is $\log R_{\text{acc}}$. See [2, Chapter 3.3.3] for a rigorous proof that the unique optimal solution for the other player is $\nu^* = \text{FlippedArcsine}(m, M)$; here we just provide intuition in terms of symmetry of the game. The underlying idea is to consider a “flipped game” in which strategies correspond to the inverses of strategies in the original game, i.e., one player chooses stepsizes $\tilde{\beta} = \beta^{-1} \in \tilde{E}$, and the other player chooses a quadratic function parameterized by its inverse curvature $\tilde{\lambda} = \lambda^{-1} \in \tilde{E}$. Here $\tilde{E} := [\tilde{m}, \tilde{M}]$ where $\tilde{m} := M^{-1}$ and $\tilde{M} := m^{-1}$. By lifting to mixed strategies, one player chooses a distribution $\tilde{\mu}$ over stepsizes $\tilde{\beta} \in \tilde{E}$, and the other player chooses a distribution $\tilde{\nu}$ over quadratic functions $\tilde{\lambda} \in \tilde{E}$, with payoff given by

$$\tilde{\mathcal{R}}(\tilde{\mu}, \tilde{\nu}) := \mathbb{E}_{\tilde{\beta} \sim \tilde{\mu}, \tilde{\lambda} \sim \tilde{\nu}} \log |1 - \tilde{\lambda}^{-1} \tilde{\beta}|.$$

This flipped game is equivalent to the original game due the invariance of both the payoff function $\tilde{\mathcal{R}}(\tilde{\mu}, \tilde{\nu}) = \mathcal{R}(\mu, \nu)$ and the condition number $\frac{\tilde{M}}{\tilde{m}} = \frac{M}{m}$ with respect to this inversion of the strategies. Because Arcsine(m, M) is the unique optimal solution μ^* in the original game, FlippedArcsine(\tilde{m} , \tilde{M}) is the unique optimal solution $\tilde{\mu}^*$ in the flipped game; in order to show that FlippedArcsine(m, M) is the unique optimal solution ν^* in the original game, it suffices to show that Arcsine(\tilde{m} , \tilde{M}) is the unique optimal solution $\tilde{\nu}^*$ in the flipped game. To this end, it suffices to show that an optimal solution $\tilde{\nu}^*$ must satisfy the equalization property in the flipped game, i.e.,

$$\mathcal{R}(\delta_{\tilde{\beta}}, \tilde{\nu}^*) = \mathbb{E}_{\tilde{\lambda} \sim \tilde{\nu}^*} \log |1 - \tilde{\lambda}^{-1} \tilde{\beta}|$$

is independent of $\tilde{\beta} \in \tilde{E}$. This suffices since the equalization property uniquely defines the Arcsine distribution by classical results from logarithmic potential theory, see Appendix A.

Why must an optimal solution $\tilde{\nu}^*$ for the flipped game satisfy this equalization property? An informal argument is as follows. By definition of $\tilde{\mathcal{R}}$, the symmetry between $\tilde{\mathcal{R}}$ and \mathcal{R} , and then the equalization property of the Arcsine distribution (Lemma 2.2) applied to μ^* ,

$$\mathbb{E}_{\tilde{\beta} \sim \tilde{\mu}} \mathcal{R}(\delta_{\tilde{\beta}}, \tilde{\nu}^*) = \tilde{\mathcal{R}}(\tilde{\mu}^*, \tilde{\nu}^*) = \mathcal{R}(\mu^*, \nu^*) = \log R_{\text{acc}}.$$

Thus, forgoing measure-theoretic technicalities, if $\tilde{\nu}^*$ did not satisfy the equalization property, then by the probabilistic method, there would exist some $\tilde{\beta}$ for which $\tilde{\beta} \mapsto \mathcal{R}(\delta_{\tilde{\beta}}, \tilde{\nu}^*)$ is strictly less than its average $\mathbb{E}_{\tilde{\beta} \sim \tilde{\mu}} \mathcal{R}(\delta_{\tilde{\beta}}, \tilde{\nu}^*) = \log R_{\text{acc}}$. But this means that the value of the flipped game is less than $\log R_{\text{acc}}$, hence also for the original game (since their values are identical), contradiction. \square

As a corollary, Lemma 6.2 gives a new proof that $\frac{\sqrt{\kappa}-1}{\sqrt{\kappa}+1}$ is the fastest possible convergence rate for GD. This result is classically due to [39], who proved this by constructing a single, infinite-dimensional “worst-case function” for which any Krylov-subspace algorithm converges slowly. Here we construct a “worst-case distribution” over univariate functions for which any GD stepsize schedule converges slowly in expectation. By the probabilistic method, for any specific GD stepsize schedule, we can then conclude the existence of a single univariate function for which convergence is slow. For simplicity, we state this lower bound here for i.i.d. stepsizes; this can be extended to arbitrary GD stepsize schedules and also to non-asymptotic lower bounds, see [2, Chapter 5].

Corollary 6.4. *For any distribution of i.i.d. GD stepsizes, there exists a univariate quadratic function that is m -strongly convex and M -smooth such that the asymptotic convergence rate*

$$\lim_{n \rightarrow \infty} \mathbb{E} \left[\left(\frac{\|x_n - x^*\|}{\|x_0 - x^*\|} \right)^{1/n} \right] \geq \frac{\sqrt{\kappa} - 1}{\sqrt{\kappa} + 1}.$$

Proof. Lemma 6.2 shows that in expectation over $\lambda \sim \text{FlippedArcsine}(m, M)$, the asymptotic convergence rate is R_{acc} when run on the quadratic function $f_\lambda(x) = \frac{\lambda}{2}x^2$ from initialization $x_0 = 1$. By the probabilistic method, there exists a λ^* for which the asymptotic convergence rate on $f_{\lambda^*}(x)$ is at least this average value R_{acc} . \square

Remark 6.5. *The lower bounds in this section can be interpreted as an instantiation of Yao’s minimax principle [60]: in order to lower bound the performance of a randomized algorithm (e.g., GD with the Arcsine Stepsize Schedule), we exhibit a distribution over difficult inputs (namely quadratics with curvature λ according to the FlippedArcsine distribution) for which any deterministic algorithm cannot perform well. This is implicit along the way in this section’s development.*

7 Conclusion and future work

This paper shows that randomizing stepsizes provides an alternative mechanism for accelerated convex optimization beyond the special case of quadratics. This raises several natural questions about the possible benefit of randomness for GD—or said equivalently, whether the Arcsine Stepsize Schedule can be derandomized. The current results differ on two fronts: Theorem 1.2 shows that random stepsizes achieve the fully accelerated rate $O(\kappa^{1/2} \log \frac{1}{\varepsilon})$ on separable functions, whereas [3] shows that deterministic stepsize achieves the partially accelerated “Silver” rate $O(\kappa^{\log_1 + \sqrt{2}} \log \frac{1}{\varepsilon}) \approx O(\kappa^{0.78} \log \frac{1}{\varepsilon})$ on general convex functions. These two sources of differences motivate the questions:

- Is it possible to de-randomize Theorem 1.2, i.e., construct a deterministic stepsize schedule which achieves the fully accelerated rate for separable functions? If so, a natural conjecture would be some ordering of the Chebyshev stepsize schedule.
- Does the Arcsine Stepsize Schedule fully accelerate on general (non-separable) convex functions? If not, can any randomized stepsize schedule improve over the Silver Convergence Rate for general convex functions?

There are also many questions about extensions to other optimization settings of interest. For other settings, what is the optimal stepsize distribution if not Arcsine? In any setting, is there a provable benefit of using non-i.i.d. stepsizes? Can random stepsizes be alternatively understood through discretization analyses of continuous-time gradient flows? While other mechanisms for acceleration (e.g., momentum) are now well understood [22], little is known about acceleration via randomness.

A Connections to logarithmic potential theory

Recall from the overview section §2 that the optimal distribution μ for inverse stepsizes β is given by the extremal problem

$$\inf_{\mu \in \mathcal{M}(E)} \sup_{\lambda \in E} \mathbb{E}_{\beta \sim \mu} \log |1 - \beta^{-1} \lambda|, \quad (\text{A.1})$$

where $E := [m, M]$ and $\mathcal{M}(E)$ denotes the probability distributions supported on E .⁴ Indeed, for quadratic optimization, this extremal problem exactly characterizes the asymptotic convergence rate (see §2.1); and for the more general setting of separable convex optimization, even though the relevant extremal problem is somewhat different, the solution is identical (see §2.2).

In this appendix section, we connect this extremal problem over stepsize distributions with the extremal problem for electrostatic capacitors that is classically studied in potential theory. As we explain below, although these two problems seek probability distributions minimizing different objectives, the two problems are intimately linked because they have the same optimality condition: the optimal distribution μ must satisfy the equalization property from Lemma 2.2 (i.e., the objective $\mathbb{E}_{\beta \sim \mu} \log |1 - \beta^{-1} \lambda|$ has equal value for all $\lambda \in [m, M]$). Because this equalization property uniquely characterizes the distribution (a classical fact from potential theory, see Lemma A.9), this means that the two problems have the same unique optimal solution. In this way, the equalization property plays the central role connecting these two problems.

This appendix section is organized as follows. In §A.1 we recall the electrostatics problem and highlight the syntactic similarities and differences to the stepsize problem. In §A.2 we explain the equalization property in the context of both problems and how it arises as the optimality condition. This is the key link between the two problems. In §A.3 we exploit this connection concretely by using the machinery of potential theory to solve the stepsize problem. This provides alternative proofs of Lemmas 2.1 and 2.2 that are significantly shorter than the integral calculations of [32, 47].

A.1 Extremal problem for electrostatic capacitors

The extremal problem central to logarithmic potential theory is

$$\inf_{\mu \in \mathcal{M}(E)} \iint \log \frac{1}{|\lambda - \omega|} d\mu(\omega) d\mu(\lambda), \quad (\text{A.2})$$

given a set $E \subset \mathbb{C}$. This problem is sometimes called the electrostatics problem for a capacitor due to its physical interpretation: how should you place a distribution of repelling particles on E so as to minimize the total repulsive energy? Here the repulsive energy is measured via

$$I(\mu) := \iint \log \frac{1}{|\lambda - \omega|} d\mu(\omega) d\mu(\lambda) = \mathbb{E}_{\lambda \sim \mu} \phi_\mu(\lambda),$$

which is the μ -average of the logarithmic potential energy at λ , denoted

$$\phi_\mu(\lambda) := \mathbb{E}_{\omega \sim \mu} \log \frac{1}{|\lambda - \omega|}.$$

Standard arguments based on lower semi-continuity and strict convexity of the objective ensure that an optimal distribution exists and is unique. For certain sets E , this optimal distribution can be computed in closed form. For example it is well-known that for the interval $E = [m, M]$, the unique optimal distribution is the Arcsine(m, M) distribution, see e.g., [49, Example 1.11].

⁴Although the stepsize problem was not originally written with this constraint in §2, it is easily seen that this constraint is satisfied at optimality and therefore does not change the extremal problem. See Observation A.4.

Syntactic similarities and differences. Let $E = [m, M]$. Then the stepsize extremal problem (A.1) clearly bears similarities to the electrostatics extremal problem (A.2), namely

$$\inf_{\mu \in \mathcal{M}(E)} \sup_{\lambda \in E} \phi_\mu(0) - \phi_\mu(\lambda) \quad \text{versus} \quad \inf_{\mu \in \mathcal{M}(E)} \mathbb{E}_{\lambda \sim \mu} \phi_\mu(\lambda). \quad (\text{A.3})$$

Yet there are two differences. First, the stepsize problem evaluates the logarithmic potential $\phi_\mu(\lambda)$ at a specific point $\lambda \in [m, M]$, whereas the electrostatics problem averages this potential over $\lambda \sim \mu$. Second, the stepsize problem has an additional term of $\phi_\mu(0)$.

In general, changing the objective or constraints in an extremal problem can significantly alter the optimal solution. Yet, remarkably, the optimal solution μ is identical for both problems. This is because both problems have the same optimality condition, described next.

A.2 Equalization property

Here we describe how the equalization property arises as the optimality condition in both problems.

A.2.1 Equalized potential for the electrostatics problem

A celebrated fact—sometimes called the Fundamental Theorem of Potential Theory [48, Chapter 3.3]—is that the optimal distribution μ to the electrostatics problem (A.2) satisfies an *equalization property*: the potential $\lambda \mapsto \phi_\mu(\lambda)$ is constant on E . (Note that this is equivalent to the constancy on E of the objective for the stepsize problem, namely $\lambda \mapsto \mathbb{E}_\beta \log |1 - \beta^{-1}\lambda| = \phi_\mu(0) - \phi_\mu(\lambda)$.)

This equalization property has an intuitive physical interpretation: if the potential is not constant on E , then particles will flow to points of E with lower potential, so the current state is not at equilibrium and therefore cannot be optimal. This is why the optimal solution μ is often called the equilibrium distribution. This result is attributed to Frostman; here we state a simple version of it that suffices for our setting (namely a compact interval $E = [m, M]$).

Theorem A.1 (Frostman’s Theorem). *Let $E \subset \mathbb{C}$ be compact and simply connected, and suppose that the optimal value of (A.2) is not $-\infty$. Then there is a unique optimal distribution μ^* for (A.2). Moreover it satisfies*

- $\phi_{\mu^*}(\lambda) = I(\mu^*)$ for all $\lambda \in E$.
- $\phi_{\mu^*}(\lambda) \leq I(\mu^*)$ for all $\lambda \in \mathbb{C}$.

A proof can be found in any textbook on logarithmic potential theory, e.g., [48, 50, 54]. Below we use calculus of variations to give an informal back-of-the-envelope calculation that provides intuition for how this equalization property arises as the optimality condition. Let

$$F(\mu) := \iint \log \frac{1}{|x-y|} d\mu(x)d\mu(y) + \gamma \left(\int d\mu(x) - 1 \right)$$

denote the original objective plus a Lagrangian term for the normalization constraint on μ . Then a quick calculation shows that the first variation is

$$\delta F(\mu, h) = \frac{d}{d\varepsilon} F(\mu + \varepsilon h) \Big|_{\varepsilon=0} = \int \left[\gamma + \int \log \frac{1}{|x-y|} d\mu(x) \right] dh(y).$$

In order for this first variation to vanish along any direction h , it must hold that

$$\gamma + \int \log \frac{1}{|x-y|} d\mu(x) = 0, \quad \forall y \in E.$$

This is the equalization property.

A.2.2 Equalized convergence rate for the stepsize problem

As discussed in §2.1 and formally stated in Lemma 2.1, the stepsize problem has the same equalization property at optimality. This is rigorously proven in §A.3.2 using the machinery of potential theory; to provide intuition for how this arises as the optimality condition, here we give an informal back-of-the-envelope argument based on the game-theoretic perspective developed in §6.2.

Suppose that μ is an solution to the stepsize problem (A.1). By the calculation in §6.2, there is a distribution ν (the Flipped Arcsine distribution) for which

$$\mathbb{E}_{\beta \sim \mu, \lambda \sim \nu} \log |1 - \beta^{-1} \lambda| = \log R_{\text{acc}} . \quad (\text{A.4})$$

Observe that this is the ν -average of $P_\mu(\lambda) := \mathbb{E}_{\beta \sim \mu} \log |1 - \beta^{-1} \lambda| = \phi_\mu(0) - \phi_\mu(\lambda)$. Thus, if μ does not equalize, then (forgoing measure-theoretic technicalities) there exists $\lambda \in \text{supp}(\nu) = [m, M]$ for which $P_\mu(\lambda)$ is strictly larger than the ν -average. That is, $\max_{\lambda \in [m, M]} P_\mu(\lambda) > \log R_{\text{acc}}$. But this contradicts the optimality of μ since the value of the stepsize problem (A.1) is $\log R_{\text{acc}}$. Therefore the equalization property must be satisfied by the optimal distribution μ .

A.3 Solution to the extremal problem over stepsizes

So far we have focused on developing conceptual connections between the stepsize problem and electrostatics problem. We now exploit these connections concretely by using the machinery of potential theory to solve the stepsize problem. This provides short proofs of Lemmas 2.1 and 2.2.

A.3.1 Proof of Lemma 2.2

Here we prove the equalization property. In the optimization literature, this was proven by computing the integral directly using trigonometric identities [47], or by decomposing the logarithmic integrand into the Chebyshev polynomial basis [32]. Here we mention an alternative approach by appealing to the following well-known identity from potential theory, which gives the potential for the equilibrium distribution on $[-1, 1]$ at any point z ; see e.g., [50, Example 1.11].

Fact A.2. *Let γ denote the Arcsine distribution on $[-1, 1]$. Then for any $z \in \mathbb{C}$,*

$$\mathbb{E}_{t \sim \gamma} \log \frac{1}{|z - t|} = \log 2 - \log \left| z - \sqrt{z^2 - 1} \right| \cdot \mathbf{1}_{z \notin [-1, 1]} .$$

Proof of Lemma 2.2. Let μ denote the Arcsine distribution on $[m, M]$. Note that μ is the push-forward of γ under the linear map $L(z) := \frac{M+m}{2} + \frac{M-m}{2}z$ that sends $[-1, 1]$ to $[m, M]$. Thus by a change-of-measure, the μ -potential at any $\omega \in \mathbb{C}$ is

$$\phi_\mu(\omega) = \mathbb{E}_{\beta \sim \mu} \log \frac{1}{|\beta - \omega|} = \mathbb{E}_{t \sim \gamma} \log \frac{1}{|L(t) - L(z)|} = \phi_\gamma(z) + \log \frac{2}{M - m} ,$$

where $z := L^{-1}(\omega)$. This gives the desired identity:

$$\mathbb{E}_{\beta \sim \mu} \log |1 - \beta^{-1} \lambda| = \mathbb{E}_{\beta \sim \mu} \log \frac{1}{|\beta|} - \mathbb{E}_{\beta \sim \mu} \log \frac{1}{|\lambda - \beta|} = \phi_\mu(0) - \phi_\mu(\lambda) = -\log \left| z - \sqrt{z^2 - 1} \right| = \log \frac{\sqrt{\kappa} - 1}{\sqrt{\kappa + 1}} .$$

Above the first step is by expanding the logarithm, the second step is by definition of the logarithmic potential ϕ_μ , the third step is by Fact A.2 and the above formula for ϕ_μ in terms of ϕ_γ , and the last step is by plugging in the value of $z = L^{-1}(0) = -\frac{\kappa+1}{\kappa-1}$ and simplifying. \square

Remark A.3 (Connection to Green’s function). *This expression $\log \frac{\sqrt{\kappa-1}}{\sqrt{\kappa+1}} = \log R_{\text{acc}}$ is the evaluation at 0 of Green’s function for the interval $[m, M]$. This connection essentially arises because: 1) for quadratic optimization, polynomial approximation dictates the convergence rate of GD (see §2.1), and 2) the Bernstein-Walsh Theorem characterizes the optimal rate of convergence of polynomial approximations in terms of Green’s function (see e.g., [19, Theorem 1]).*

A.3.2 Proof of Lemma 2.1

Here we prove that the Arcsine distribution is the unique optimal solution for the stepsize problem (A.1) by mirroring the standard argument from potential theory that the Arcsine distribution is the unique optimal solution to the electrostatic problem (A.2), see e.g., [49, Example 2.9]. Specifically, our approach is to show that any optimal distribution μ satisfies the equalization property (see §A.2.2 for a heuristic argument of this). Indeed, it then follows that the optimal distribution μ is unique and is the Arcsine distribution by the converse of Frostman’s Theorem (stated below), which ensures that the Arcsine distribution is the unique distribution satisfying equalization.

We begin with a simple observation.

Observation A.4 (Support of the optimal stepsize distribution). *The constraint $\mu \in \mathcal{M}(E)$ does not change the value of the stepsize problem (A.1).*

Proof. It suffices to argue that for any distribution μ (with arbitrary support), there exists some distribution $\tilde{\mu} \in \mathcal{M}(E)$ whose objective value is at least as good, i.e., $\max_{\lambda \in [m, M]} \mathbb{E}_{\tilde{\beta} \sim \tilde{\mu}} \log |1 - \tilde{\beta}^{-1} \lambda| \geq \max_{\lambda \in [m, M]} \mathbb{E}_{\beta \sim \mu} \log |1 - \beta^{-1} \lambda|$. To this end, it suffices to observe the elementary inequality

$$|1 - \lambda \tilde{\alpha}| \leq |1 - \lambda \alpha|, \quad \forall \lambda \in [m, M],$$

where $\alpha := \beta^{-1} = a + bi \in \mathbb{C}$ is arbitrary and $\tilde{\alpha} := \max(\min(a, M^{-1}), m^{-1})$ is the projection of α onto the interval $[M^{-1}, m^{-1}]$. Indeed, this pointwise inequality implies that we can project μ onto $\mathcal{M}(E)$ in this way without worsening the objective value of μ . \square

We now show that any optimal distribution, constrained to this interval, must satisfy the equalization property. For this it is helpful to recall several classical facts from potential theory. These can be found, e.g., in page 171 and Theorems 3.2, 2.8, 2.2, and 3.1, respectively, of [49].

Lemma A.5 (Basic properties of logarithmic potentials). *If $\mu \in \mathcal{M}(\mathbb{C})$, then ϕ_μ is lower semicontinuous on \mathbb{C} and harmonic outside the support of μ .*

Lemma A.6 (Continuity of the equilibrium potential). *Consider the setting of Theorem A.1. Then the potential ϕ_{μ^*} is continuous on \mathbb{C} .*

Lemma A.7 (Principle of domination). *Suppose $\mu, \nu \in \mathcal{M}(\mathbb{C})$ are compactly supported and that $I(\mu) < \infty$. If for some constant c , $\phi_\mu(z) - \phi_\nu(z) \leq c$ holds μ -a.e., then it holds for all $z \in \mathbb{C}$.*

Lemma A.8 (Maximum principle). *Suppose h is a harmonic function on a domain $\Omega \subset \mathbb{C}$. If h attains a local maximum in Ω , then h is constant on Ω .*

Lemma A.9 (Converse to Frostman’s Theorem). *Suppose $\mu \in \mathcal{M}(E)$ satisfies $I(\mu) < \infty$. If $\phi_\mu(z) = c$ on E for some constant c , then μ is the equilibrium distribution for E .*

Proof of Lemma 2.1. For shorthand, let μ denote the Arcsine(m, M) distribution and note that $I(\mu) < \infty$ by Lemma 2.2. Suppose $\nu \in \mathcal{M}([m, M])$ has objective value for the stepsize problem (2.9) that is at least as good as μ . Then by the equalization property for μ (Lemma 2.2),

$$\phi_\nu(0) - \phi_\nu(\lambda) = \mathbb{E}_{\beta \sim \nu} \log |1 - \beta^{-1}\lambda| \leq \mathbb{E}_{\beta \sim \mu} \log |1 - \beta^{-1}\lambda| = \phi_\mu(0) - \phi_\mu(\lambda), \quad \forall \lambda \in [m, M].$$

Re-arranging gives

$$f(\lambda) \leq c, \quad \forall \lambda \in [m, M], \tag{A.5}$$

where $f := \phi_\mu - \phi_\nu$ and $c := f(0)$. By the principle of domination (Lemma A.7), this inequality (A.5) holds for all $\lambda \in \mathbb{C}$. Since $0 \in \Omega := \mathbb{C} \setminus [m, M]$, we conclude that f achieves its maximum on Ω at 0. But because f is harmonic on Ω (Lemma A.5), the Maximum Principle (Lemma A.8) implies that f is constant on Ω , i.e.,

$$f(\lambda) = c, \quad \forall \lambda \in \mathbb{C} \setminus [m, M]. \tag{A.6}$$

Fix any $\lambda \in [m, M]$ and let $\{z_n\}_{n \in \mathbb{N}}$ be any sequence in $\mathbb{C} \setminus [m, M]$ converging to λ . Note that $\phi_\nu(\lambda) \leq \liminf_{n \rightarrow \infty} \phi_\nu(z_n)$ by lower semicontinuity (Lemma A.5), and $\phi_\mu(\lambda) = \lim_{n \rightarrow \infty} \phi_\mu(z_n)$ by continuity of the equilibrium potential (Lemma A.6). Thus by (A.6),

$$f(\lambda) \geq \limsup_{n \rightarrow \infty} f(z_n) = c, \quad \forall \lambda \in [m, M].$$

Combining this with (A.5) shows that

$$f(\lambda) = c, \quad \forall \lambda \in [m, M].$$

Thus ν has constant logarithmic potential $\phi_\nu(\lambda) = \phi_\mu(\lambda) + c$ on $[m, M]$, and therefore by the converse to Frostman's Theorem (Lemma A.9), it must be the equilibrium distribution, a.k.a., $\nu = \text{Arcsine}(m, M)$. We conclude that the Arcsine distribution is the unique optimal solution for the stepsize problem (A.1). The optimal value then follows from Lemma 2.1. \square

B Deferred proofs

B.1 Equivalence of separability and commutative Hessians

Here we prove an alternative characterization of separability. We remark that this provides a simple way to certify that a function is not separable: it suffices to exhibit points x and y for which the Hessians $\nabla^2 f(x)$ and $\nabla^2 f(y)$ do not commute.

Lemma B.1 (Equivalence of separability and commutative Hessians). *Let $f \in C^2(\mathbb{R}^d)$. Then the following are equivalent:*

- (i) $f(x) = \sum_{i=1}^d f_i([y]_i)$ where $y = Ux$ for some orthogonal U and functions $f_i : \mathbb{R} \rightarrow \mathbb{R}$.
- (ii) $\{\nabla^2 f(x)\}_{x \in \mathbb{R}^d}$ commute.

Proof of (i) \implies (ii). The function $h(y) := f(U^T y) = \sum_{i=1}^d f_i([y]_i)$ is separable, hence its Hessian $\nabla^2 h$ is diagonal. The chain rule implies $\nabla^2 f = U^T \nabla^2 h U$, thus the Hessians of f are simultaneously diagonalizable in the basis U , hence they commute.

Proof of (ii) \implies (i). Since the Hessians commute, they are simultaneously diagonalizable, i.e., there exists an orthogonal matrix U such that $\nabla^2 f(x) = U^T D(x) U$ for all x , where $D(x)$ is diagonal. Defining $h(y) := f(U^T y)$, it follows by the chain rule that $\nabla^2 h(y) = D(U^T y)$ is diagonal. Thus all mixed partial derivatives vanish, hence $\frac{\partial}{\partial [y]_i} h(y)$ is a function only of $[y]_i$. Integrating implies $h(y) = \sum_{i=1}^d f_i([y]_i)$ for some functions f_i . Thus $f(x) = h(Ux) = \sum_{i=1}^d f_i([Ux]_i)$ as desired.

B.2 Proof of Lemma 3.2

We first isolate a useful identity: the random variables $Z_t = \log|1 - \beta_t^{-1}\lambda_t|$ defined in §3.1 are uncorrelated—despite not being independent. Intuitively, this is due to the equalization property of the Arcsine distribution (from which β_t^{-1} are independently drawn) since it ensures that each Z_t has the same (conditional) expectation regardless of the dependencies injected through the non-stationary process λ_t .

Lemma B.2. $\text{Cov}(Z_s, Z_t) = 0$ for all $s \neq t$.

Proof. Suppose $s < t$ without loss of generality. Then $\text{Cov}(Z_s, Z_t) = \mathbb{E}_{\beta_0, \dots, \beta_t}[(Z_s - \log R_{\text{acc}})(Z_t - \log R_{\text{acc}})] = \mathbb{E}_{\beta_0, \dots, \beta_{t-1}}[\mathbb{E}_{\beta_t}[(Z_s - \log R_{\text{acc}})(Z_t - \log R_{\text{acc}}) | \beta_0, \dots, \beta_{t-1}]] = 0$, where the first step is by the definition of covariance and (3.3), the second step is by the Law of Iterated Expectations, and the third step is because $\mathbb{E}_{\beta_t}[Z_t | \beta_0, \dots, \beta_{t-1}] = \log R_{\text{acc}}$ by the equalization property of the Arcsine distribution (Lemma 2.2). \square

Proof of Lemma 3.2. First we check that $\{M_n\}_{n \in \mathbb{N}}$ is indeed a martingale:

$$\mathbb{E}[M_n | F_{n-1}] = M_{n-1} + \frac{\mathbb{E}[Z_n - \mathbb{E}[Z_n] | F_{n-1}]}{n+1} = M_{n-1}.$$

The final step uses $\mathbb{E}[Z_n | F_{n-1}] = \mathbb{E}[Z_n]$, which follows by an identical argument as in (3.3) above.

Next, we show $\{M_n\}_{n \in \mathbb{N}}$ is L_1 -bounded. By the Cauchy-Schwarz inequality,

$$\mathbb{E}|M_n| \leq \sqrt{\mathbb{E}[M_n^2]} = \sqrt{\text{Var}(M_n)}$$

Thus it suffices to show that $\text{Var}(M_n)$ is uniformly bounded away from ∞ . To this end, since $\text{Cov}(Z_s, Z_t) = 0$ for any $s \neq t$ by Lemma B.2, it follows that

$$\text{Var}(M_n) = \text{Var}\left(\sum_{t=0}^{n-1} \frac{Z_t}{t+1}\right) = \sum_{t=0}^{n-1} \text{Var}\left(\frac{Z_t}{t+1}\right) \leq V \sum_{t=1}^n \frac{1}{t^2} \leq V \sum_{t=1}^{\infty} \frac{1}{t^2} = V \frac{\pi^2}{6}$$

where $V := \sup_{\lambda \in [m, M]} \text{Var}_{\beta \sim \text{Arcsine}(m, M)}(\log|1 - \beta^{-1}\lambda|)$ is a uniform upper bound on the variance of Z_t . Note that $V < \infty$ because the logarithmic potential is an integrable singularity (see also Appendix 5 of [32] for quantitative bounds on V). Thus M_n has finite variance, as desired. \square

B.3 Details for §5.1

Here we observe that random stepsizes can make GD unstable in alternative notions of convergence. Instability in the following lemma occurs for $\kappa > 4$ since then $\sqrt{\kappa} - 1 > 1$. See §5.1 for a discussion of this phenomena and why the notion of convergence studied in the rest of this paper is more relevant to optimization practice.

Lemma B.3 (Instability in alternative notions of convergence). *Consider GD with i.i.d. inverse stepsizes from $\text{Arcsine}(m, M)$. There is an m -strongly convex, M -smooth, quadratic function f such that for any initialization $x_0 \neq x^*$ and any number of iterations n ,*

$$\mathbb{E}\left[\frac{\|x_n - x^*\|}{\|x_0 - x^*\|}\right]^{1/n} \geq \sqrt{\kappa} - 1. \tag{B.1}$$

To prove this, we make use of the following integral representation of the geometric mean of two positive scalars $a, b > 0$ in terms of the weighted harmonic mean $H_t(a, b) := (ta^{-1} + (1-t)b^{-1})^{-1}$. This identity is well-known (e.g., it follows from [7, Exercise V.4.20], or directly from [25, §3.121-2] by taking $q = 1/b$ and $p = 1/b - 1/a$).

Lemma B.4 (Arcsine integral representation of the geometric mean). *For any $a, b > 0$,*

$$\mathbb{E}_{t \sim \text{Arcsine}(0,1)}[H_t(a, b)] = \sqrt{ab}.$$

This lets us compute the expectation of the proposed random stepsizes.

Lemma B.5 (Expectation of proposed random stepsizes). *For any $0 < m < M < \infty$,*

$$\mathbb{E}_{\beta \sim \text{Arcsine}(m, M)}[\beta^{-1}] = \frac{1}{\sqrt{Mm}}.$$

Proof. Re-parameterize $\beta = tM + (1-t)m$ so that $\beta^{-1} = H_t(M^{-1}, m^{-1})$. Since $\beta \sim \text{Arcsine}(m, M)$ implies $t \sim \text{Arcsine}(0, 1)$, applying Lemma B.4 provides the desired identity

$$\mathbb{E}_{\beta \sim \text{Arcsine}(m, M)}[\beta^{-1}] = \mathbb{E}_{t \sim \text{Arcsine}(0,1)}[H_t(M^{-1}, m^{-1})] = \frac{1}{\sqrt{Mm}}.$$

□

We are now ready to prove Lemma B.3.

Proof of Lemma B.3. Let β_t denote the inverse stepsizes; recall that these are i.i.d. $\text{Arcsine}(m, M)$. By expanding the definition of GD for the function $f(x) = \frac{M}{2}x^2$,

$$x_n = \prod_{t=0}^{n-1} (1 - \beta_t^{-1}M)x_0.$$

By taking expectations, using independence of the stepsizes, and applying Lemma B.5,

$$\mathbb{E}[x_n] = (1 - \sqrt{\kappa})^n x_0.$$

Thus by Jensen's inequality and then the fact that $\kappa > 1$, we conclude the desired inequality

$$\mathbb{E}\left[\frac{\|x_n - x^*\|}{\|x_0 - x^*\|}\right]^{1/n} \geq \left(\frac{\|\mathbb{E}x_n - x^*\|}{\|x_0 - x^*\|}\right)^{1/n} = \sqrt{\kappa} - 1.$$

□

B.4 Proof of Theorem 5.2

For shorthand, let $V := \sup_{\lambda \in [m, M]} \text{Var}_{\beta \sim \text{Arcsine}(m, M)}(\log|1 - \beta^{-1}\lambda|)$ denote a uniform upper bound on the variance of any Z_t . By the integral calculations in [32, Appendix 5],

$$V \leq \pi^2 + o_\kappa(1).$$

Thus because $\text{Cov}(Z_s, Z_t) = 0$ for $s \neq t$ by Lemma B.2,

$$\text{Var}\left(\frac{1}{n} \sum_{t=0}^{n-1} Z_t\right) = \frac{1}{n^2} \sum_{t=0}^{n-1} \text{Var}(Z_t) \leq \frac{V}{n} \leq \frac{\pi^2 + o_\kappa(1)}{n}.$$

Thus by Chebyshev's inequality and the equalization property of the Arcsine distribution (Lemma 2.2),

$$\mathbb{P}\left(\frac{1}{n}\sum_{t=0}^{n-1}\log|1-\beta_t^{-1}\lambda_t|\geq\log R_{\text{acc}}+\varepsilon\right)\leq\frac{\pi^2+o_\kappa(1)}{n\varepsilon^2}, \quad (\text{B.2})$$

where β_t are i.i.d. Arcsine(m, M) and $\lambda_t \in [m, M]$ is any process adapted to the filtration $\sigma(\{\beta_s\}_{s<t})$. Thus by (3.2), the average convergence rate $R_{n,i} := \left(\frac{|[x_n]_i - [x^*]_i|}{|[x_0]_i - [x^*]_i|}\right)^{1/n}$ for the function f_i satisfies

$$\mathbb{P}(R_{n,i} \geq e^\varepsilon R_{\text{acc}}) \leq \frac{\pi^2 + o_\kappa(1)}{n\varepsilon^2} \quad (\text{B.3})$$

for any coordinate $i \in S := \{i \in [d] : [x_0]_i \neq [x^*]_i\}$, i.e., for any i for which the initialization is not already optimal. Thus the overall convergence rate $R_n := \left(\frac{\|x_n - x^*\|}{\|x_0 - x^*\|}\right)^{1/n}$ for f satisfies

$$\mathbb{P}(R_n \geq e^\varepsilon R_{\text{acc}}) \leq \mathbb{P}\left(\max_{i \in S} R_{n,i} \geq e^\varepsilon R_{\text{acc}}\right) \leq \frac{d(\pi^2 + o_\kappa(1))}{n\varepsilon^2}. \quad (\text{B.4})$$

Above the first step is because $R_n \leq \max_{i \in S} R_{n,i}$ (since $R_n^2 = \frac{\sum_{i \in S} |[x_n]_i - [x^*]_i|^2}{\sum_{i \in S} |[x_0]_i - [x^*]_i|^2} \leq \max_{i \in S} \frac{|[x_n]_i - [x^*]_i|^2}{|[x_0]_i - [x^*]_i|^2} = \max_{i \in S} R_{n,i}^2$). The second step is by a union bound and the previous display. The proof is then complete by setting δ equal to the right hand side of the above display.

B.5 Details for §6.1

Proof of Theorem 6.1. The proof is similar to that of Theorem 1.2, so we just highlight the differences. By the same argument as in §3.2, it suffices to show the result for univariate objectives f , so we henceforth restrict to this case. For shorthand, denote $E_t := (\tilde{g}_t - f'(x_t))/(x_t - x^*)$ and note that the multiplicative error assumption (6.1) implies $|E_t| \leq \varepsilon' := M\varepsilon$ by smoothness. By an identical calculation as in §3.1,

$$x_{t+1} - x^* = x_t - \alpha_t \tilde{g}_t - x^* = (1 - \alpha_t(\lambda_t + E_t))(x_t - x^*),$$

and thus

$$x_n - x^* = \prod_{t=0}^{n-1} (1 - \alpha_t(\lambda_t + E_t))(x_0 - x^*),$$

and so in particular

$$\log\left(\frac{|x_n - x^*|}{|x_0 - x^*|}\right)^{1/n} = \frac{1}{n} \sum_{t=0}^{n-1} \log|1 - \alpha_t(\lambda_t + E_t)|.$$

By a similar martingale-based Law of Large Numbers argument as in §3.1, it suffices to compute

$$\max_{\tilde{\lambda} \in [m - \varepsilon', M + \varepsilon']} \mathbb{E}_{\beta \sim \text{Arcsine}(m, M)} \log|1 - \beta^{-1}\tilde{\lambda}| \quad (\text{B.5})$$

and in particular to show that this equal to the logarithm of the upper bound in (6.3). Above, $\tilde{\lambda}$ plays the role of $\lambda + E$ where $\lambda \in [m, M]$ and $E \in [-\varepsilon', \varepsilon']$. By the same calculation as in Appendix A.3.1, this expectation simplifies to

$$\mathbb{E}_{\beta \sim \text{Arcsine}(m, M)} \log|1 - \beta^{-1}\tilde{\lambda}| = \log R_{\text{acc}} - \log(|z| - \sqrt{z^2 - 1}) \cdot \mathbb{1}_{z \notin [-1, 1]}, \quad (\text{B.6})$$

where $z := \frac{2}{M-m}(\tilde{\lambda} - \frac{M+m}{2})$. This quantity is increasing in $|z|$, thus (B.5) is maximized at both extremes $\tilde{\lambda} \in \{m - \varepsilon', M + \varepsilon'\}$ which lead to the same maximal value of $|z| = 1 + \frac{2\varepsilon'}{M-m} = 1 + \tilde{\varepsilon}$. Plugging this into (B.5) yields the desired bound.

Finally, for tightness, note that the analysis is tight at the aforementioned extremes $\lambda = m, E = -\varepsilon'$ and $\lambda = M, E = \varepsilon'$. These two extremes respectively correspond to functions $f(x) = \frac{mx^2}{2}$ and $f(x) = \frac{Mx^2}{2}$, with estimates \tilde{g}_t that maximally underestimate or overestimate the true gradient. \square

References

- [1] Naman Agarwal, Surbhi Goel, and Cyril Zhang. Acceleration via fractal learning rate schedules. In *International Conference on Machine Learning*, pages 87–99. PMLR, 2021.
- [2] Jason M. Altschuler. Greed, hedging, and acceleration in convex optimization. Master’s thesis, Massachusetts Institute of Technology, 2018.
- [3] Jason M Altschuler and Pablo A Parrilo. Acceleration by Step Size Hedging I: Multi-Step Descent and the Silver Step Size Schedule. *arXiv preprint arXiv:2309.07879*, 2023.
- [4] Jason M Altschuler and Pablo A Parrilo. Acceleration by Step Size Hedging: Silver Step Size Schedule for Smooth Convex Optimization. *Mathematical Programming*, pages 1–14, 2024.
- [5] Jonathan Barzilai and Jonathan M Borwein. Two-point step size gradient methods. *IMA Journal of Numerical Analysis*, 8(1):141–148, 1988.
- [6] Dimitri P. Bertsekas. *Nonlinear programming*. Athena Scientific, 1999.
- [7] Rajendra Bhatia. *Matrix analysis*, volume 169. Springer Science & Business Media, 2013.
- [8] Jinho Bok and Jason M. Altschuler. Accelerating proximal gradient descent via silver step sizes. *arXiv preprint*, 2024.
- [9] Stephen Boyd and Lieven Vandenberghe. *Convex optimization*. Cambridge University Press, 2004.
- [10] Sébastien Bubeck. Convex optimization: Algorithms and complexity. *Foundations and Trends in Machine Learning*, 8(3-4):231–357, 2015.
- [11] Sébastien Bubeck, Yin Tat Lee, and Mohit Singh. A geometric alternative to Nesterov’s accelerated gradient descent. *arXiv preprint arXiv:1506.08187*, 2015.
- [12] Antoine Daccache. *Performance estimation of the gradient method with fixed arbitrary step sizes*. PhD thesis, Master’s thesis, Université Catholique de Louvain, 2019.
- [13] Shuvomoy Das Gupta, Bart PG Van Parys, and Ernest K Ryu. Branch-and-bound performance estimation programming: a unified methodology for constructing optimal optimization methods. *Mathematical Programming*, pages 1–73, 2023.
- [14] Alexandre d’Aspremont. Smooth optimization with approximate gradient. *SIAM Journal on Optimization*, 19(3):1171–1183, 2008.
- [15] Etienne De Klerk, François Glineur, and Adrien B Taylor. On the worst-case complexity of the gradient method with exact line search for smooth strongly convex functions. *Optimization Letters*, 11:1185–1199, 2017.
- [16] Etienne De Klerk, Francois Glineur, and Adrien B Taylor. Worst-case convergence analysis of inexact gradient and Newton methods through semidefinite programming performance estimation. *SIAM Journal on Optimization*, 30(3):2053–2082, 2020.
- [17] Amir Dembo. *Large deviations techniques and applications*. Springer, 2009.

- [18] Olivier Devolder, François Glineur, and Yurii Nesterov. First-order methods of smooth convex optimization with inexact oracle. *Mathematical Programming*, 146:37–75, 2014.
- [19] Tobin A Driscoll, Kim-Chuan Toh, and Lloyd N Trefethen. From potential theory to matrix iterations in six steps. *SIAM Review*, 40(3):547–578, 1998.
- [20] Yoel Drori and Marc Teboulle. Performance of first-order methods for smooth convex minimization: a novel approach. *Mathematical Programming*, 145(1-2):451–482, 2014.
- [21] Dmitriy Drusvyatskiy, Maryam Fazel, and Scott Roy. An optimal first order method based on optimal quadratic averaging. *SIAM Journal on Optimization*, 28(1):251–271, 2018.
- [22] Alexandre d’Aspremont, Damien Scieur, and Adrien Taylor. Acceleration methods. *Foundations and Trends® in Optimization*, 5(1-2):1–245, 2021.
- [23] Diego Eloi. *Worst-case functions for the gradient method with fixed variable step sizes*. PhD thesis, Master’s thesis, Université Catholique de Louvain, 2022.
- [24] Baptiste Goujaud, Damien Scieur, Aymeric Dieuleveut, Adrien B Taylor, and Fabian Pedregosa. Super-acceleration with cyclical step-sizes. In *International Conference on Artificial Intelligence and Statistics*, pages 3028–3065. PMLR, 2022.
- [25] Izrail Solomonovich Gradshteyn and Iosif Moiseevich Ryzhik. *Table of integrals, series, and products*. Academic Press, 2014.
- [26] Benjamin Grimmer. Provably faster gradient descent via long steps. *SIAM Journal on Optimization*, 34(3):2588–2608, 2024.
- [27] Benjamin Grimmer, Kevin Shu, and Alex L. Wang. Accelerated gradient descent via long steps. *arXiv preprint arXiv:2309.09961*, 2023.
- [28] Benjamin Grimmer, Kevin Shu, and Alex Wang. Accelerated objective gap and gradient norm convergence for gradient descent via long steps. *arXiv preprint arXiv:2403.14045*, 2024.
- [29] Benjamin Grimmer, Kevin Shu, and Alex L Wang. Composing optimized stepsize schedules for gradient descent. *arXiv preprint arXiv:2410.16249*, 2024.
- [30] Elad Hazan. Introduction to online convex optimization. *Foundations and Trends in Optimization*, 2(3-4):157–325, 2016.
- [31] Magnus R Hestenes, Eduard Stiefel, et al. Methods of conjugate gradients for solving linear systems. *Journal of Research of the National Bureau of Standards*, 49(6):409–436, 1952.
- [32] Zdeněk Kalousek. Steepest descent method with random step lengths. *Foundations of Computational Mathematics*, 17(2):359–422, 2017.
- [33] Carl T Kelley. *Iterative methods for optimization*. SIAM, 1999.
- [34] Jonathan Kelner, Annie Marsden, Vatsal Sharan, Aaron Sidford, Gregory Valiant, and Honglin Yuan. Big-step-little-step: Efficient gradient methods for objectives with multiple scales. In *Conference on Learning Theory*, pages 2431–2540. PMLR, 2022.
- [35] Guanghui Lan. *First-order and stochastic optimization methods for machine learning*, volume 1. Springer, 2020.
- [36] VI Lebedev and SA Finogenov. Ordering of the iterative parameters in the cyclical Chebyshev iterative method. *USSR Computational Mathematics and Mathematical Physics*, 11(2):155–170, 1971.
- [37] Laurent Lessard, Benjamin Recht, and Andrew Packard. Analysis and design of optimization algorithms via integral quadratic constraints. *SIAM Journal on Optimization*, 26(1):57–95, 2016.

- [38] David G Luenberger and Yinyu Ye. *Linear and nonlinear programming*, volume 2. Springer, 1984.
- [39] Arkadii Nemirovskii and David Borisovich Yudin. *Problem complexity and method efficiency in optimization*. Wiley, 1983.
- [40] Yurii Nesterov. *Introductory lectures on convex optimization: A basic course*, volume 87. Springer Science & Business Media, 1998.
- [41] Yurii E. Nesterov. A method of solving a convex programming problem with convergence rate $O(1/k^2)$. *Soviet Math. Dokl.*, 27(2):372–376, 1983.
- [42] J. Nocedal and S. J. Wright. *Numerical Optimization*. Springer, 1999.
- [43] Samet Oymak. Provable super-convergence with a large cyclical learning rate. *IEEE Signal Processing Letters*, 28:1645–1649, 2021.
- [44] Boris T Polyak. Some methods of speeding up the convergence of iteration methods. *USSR Computational Mathematics and Mathematical Physics*, 4(5):1–17, 1964.
- [45] Boris T. Polyak. *Introduction to optimization*. Optimization Software, Inc., 1987.
- [46] Boris Teodorovich Polyak. Convergence of methods of feasible directions in extremal problems. *USSR Computational Mathematics and Mathematical Physics*, 11(4):53–70, 1971.
- [47] Luc Pronzato and Anatoly Zhigljavsky. Gradient algorithms for quadratic optimization with fast convergence rates. *Computational Optimization and Applications*, 50(3):597–617, 2011.
- [48] Thomas Ransford. *Potential theory in the complex plane*, volume 28. Cambridge University Press, 1995.
- [49] Edward B Saff. Logarithmic potential theory with applications to approximation theory. *Surveys in Approximation Theory*, 5:165–200, 2010.
- [50] Edward B Saff and Vilmos Totik. *Logarithmic potentials with external fields*, volume 316. Springer Science & Business Media, 2013.
- [51] Mark Schmidt, Nicolas Roux, and Francis Bach. Convergence rates of inexact proximal-gradient methods for convex optimization. *Advances in neural information processing systems*, 24, 2011.
- [52] Adrien Taylor and Yoel Drori. An optimal gradient method for smooth strongly convex minimization. *Mathematical Programming*, 199(1-2):557–594, 2023.
- [53] Adrien B Taylor, Julien M Hendrickx, and François Glineur. Smooth strongly convex interpolation and exact worst-case performance of first-order methods. *Mathematical Programming*, 161(1-2):307–345, 2017.
- [54] Masatsugu Tsuji. *Potential theory in modern function theory*. Maruzen, 1959.
- [55] Ramon Van Handel. Probability in high dimension. *Lecture Notes (Princeton University)*, 2(3):2–3, 2014.
- [56] Bryan Van Scoy, Randy A Freeman, and Kevin M Lynch. The fastest known globally convergent first-order method for minimizing strongly convex functions. *IEEE Control Systems Letters*, 2(1):49–54, 2018.
- [57] Roman Vershynin. *High-dimensional probability: An introduction with applications in data science*, volume 47. Cambridge University Press, 2018.
- [58] Bofan Wang, Shiqian Ma, Junfeng Yang, and Danqing Zhou. Relaxed proximal point algorithm: Tight complexity bounds and acceleration without momentum. *arXiv preprint arXiv:2410.08890*, 2024.

- [59] David Williams. *Probability with martingales*. Cambridge University Press, 1991.
- [60] Andrew Chi-Chin Yao. Probabilistic computations: Toward a unified measure of complexity. In *Symposium on Foundations of Computer Science (FOCS)*, pages 222–227. IEEE, 1977.
- [61] David Young. On Richardson’s method for solving linear systems with positive definite matrices. *Journal of Mathematics and Physics*, 32(1-4):243–255, 1953.
- [62] Zehao Zhang and Rujun Jiang. Accelerated gradient descent by concatenation of stepsize schedules. *arXiv preprint arXiv:2410.12395*, 2024.
- [63] Anatoly Zhigljavsky, Luc Pronzato, and Elena Bukina. An asymptotically optimal gradient algorithm for quadratic optimization with low computational cost. *Optimization Letters*, 7(6):1047–1059, 2013.