Super-acceleration with cyclical step-sizes

Baptiste Goujaud, Damien Scieur, Aymeric Dieuleveut, Adrien Taylor, Fabian Pedregosa





Google Research

Preprint: https://arxiv.org/pdf/2106.09687.pdf

HeavyBall

aka gradient descent with momentum

Two parameters; step-size **h** > 0 and momentum **m** ∈ (0, 1)

$$\boldsymbol{x}_{t+1} = \boldsymbol{x}_t + m(\boldsymbol{x}_t - \boldsymbol{x}_{t-1}) - h\nabla f(\boldsymbol{x}_t)$$

Optimal among gradient-based methods on quadratics.

Stochastic variant popular in deep learning.

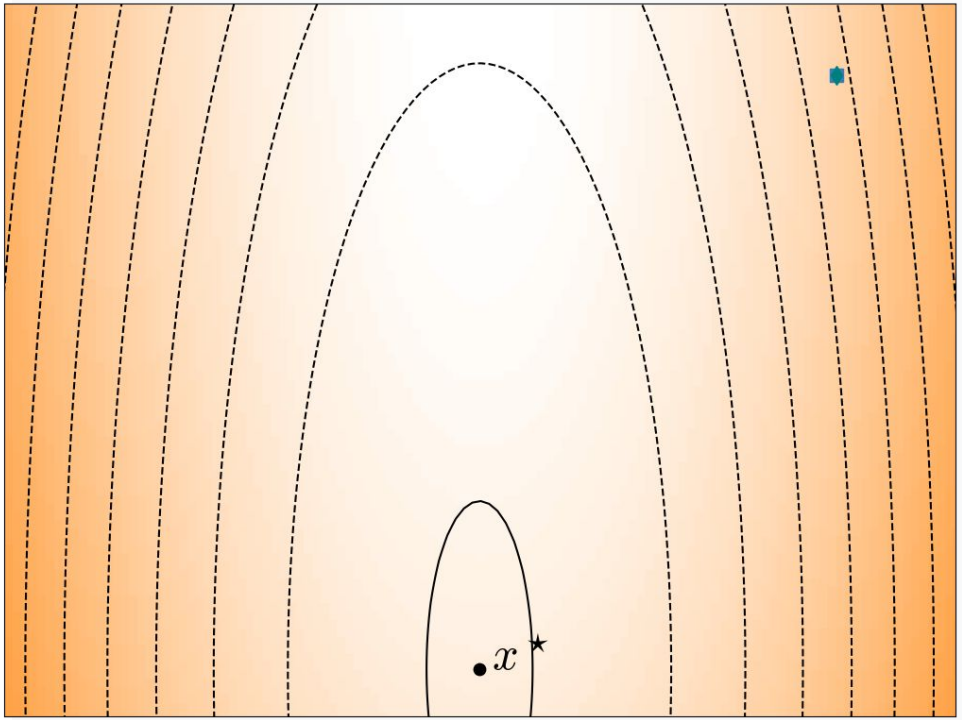

Google Research

Cyclical HeavyBall

Alternates between two step-sizes h_0 and h_1

Set $h_t = h_0$ if t is odd and $h_t = h_1$ otherwise

$$\boldsymbol{x}_{t+1} = \boldsymbol{x}_t - h_t \nabla f(\boldsymbol{x}_t) + m(\boldsymbol{x}_t - \boldsymbol{x}_{t-1})$$

Reported faster convergence (Loshchilov and Hutter, 2017; Smith, 2017)

Pervasive (TF, PyTorch, optax, etc.)

No analysis that explains why/when it works.

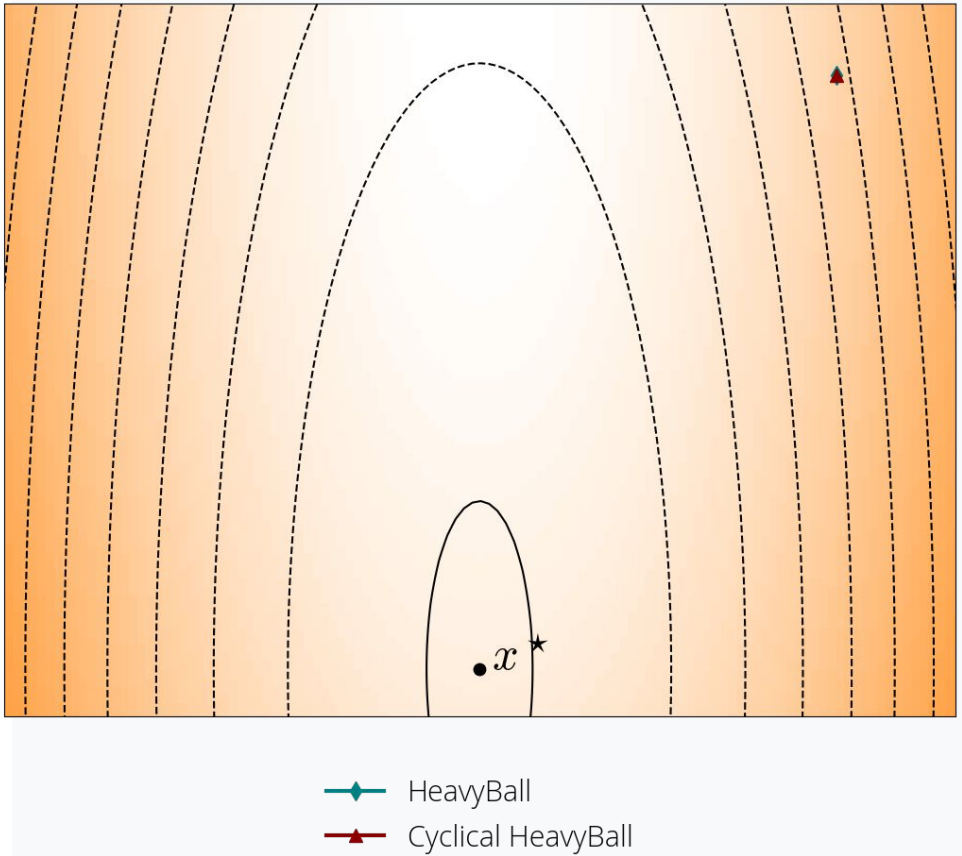

Google Research

Benchmarks

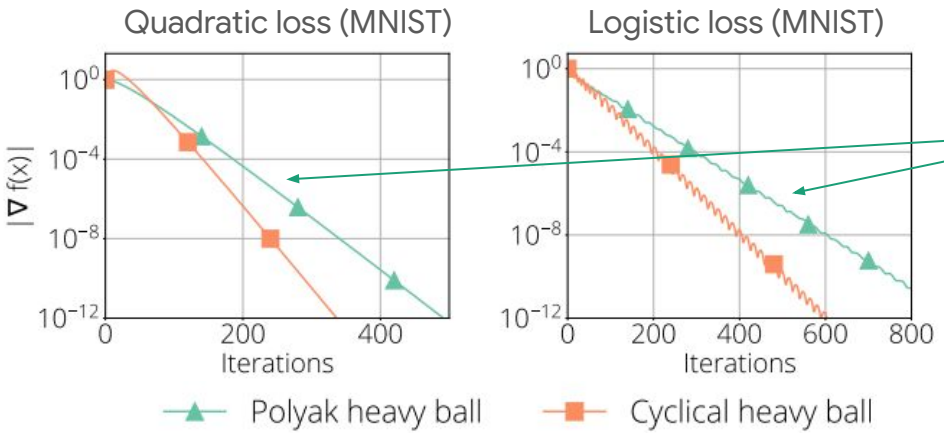

Slope = $\frac{\sqrt{L}-\sqrt{\mu}}{\sqrt{L}+\sqrt{\mu}}$

What is the slope of cyclical heavy ball?
What are the optimal parameters?

Today's topic

Google Research

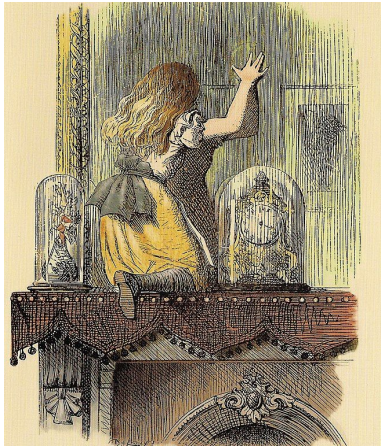

Optimization and Polynomials

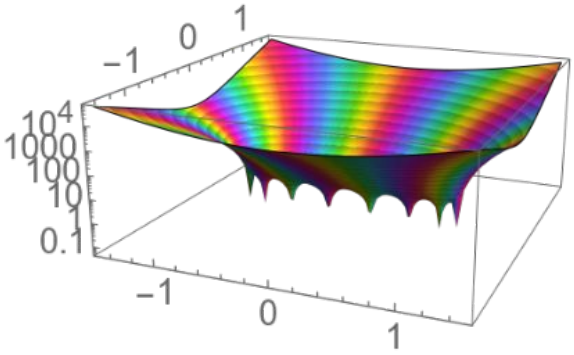

Cyclical HeavyBall

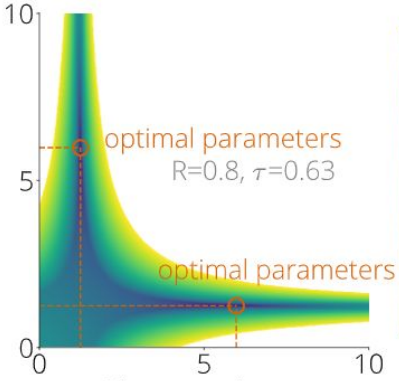

Simulations & Open problems

Google Research

Polynomials and optimization

Some problems can be posed in the space of polynomials.

Exploited in early numerical analysis [Hestenes and Stiefel (1952), Rutishauser (1959)]

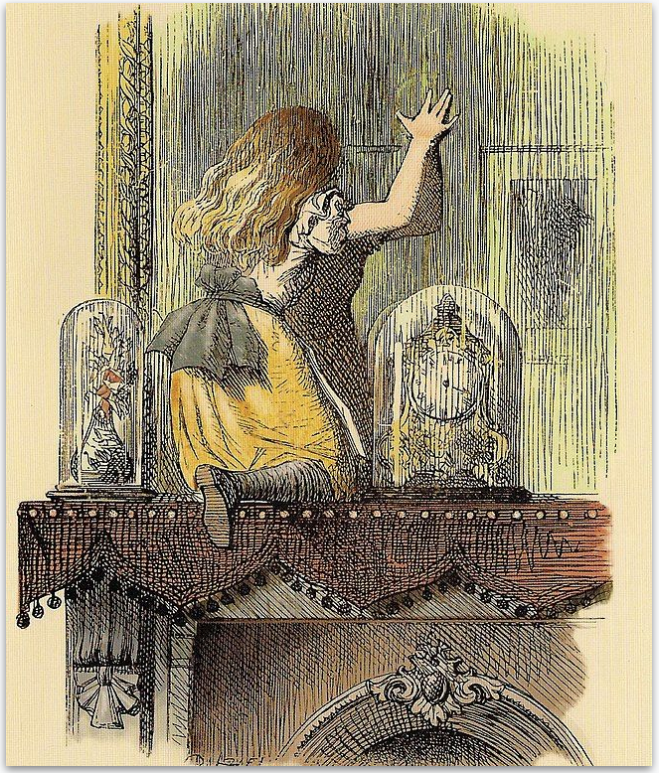

Google Research

Polynomials and Optimization

Consider Gradient Descent on

$$f(x) = \tfrac{1}{2}(x - x^\star)H(x - x_\star)$$

Then at iteration t we have

$$\begin{aligned} x_{t+1} - x_\star &= x_t - x_\star - \gamma H(x_t - x_\star) \\ &= (I - \gamma H)(x_t - x_\star) \\ &= \dots \\ &= (I - \gamma H)^{t+1}(x_0 - x_\star) \end{aligned}$$

Polynomial in H

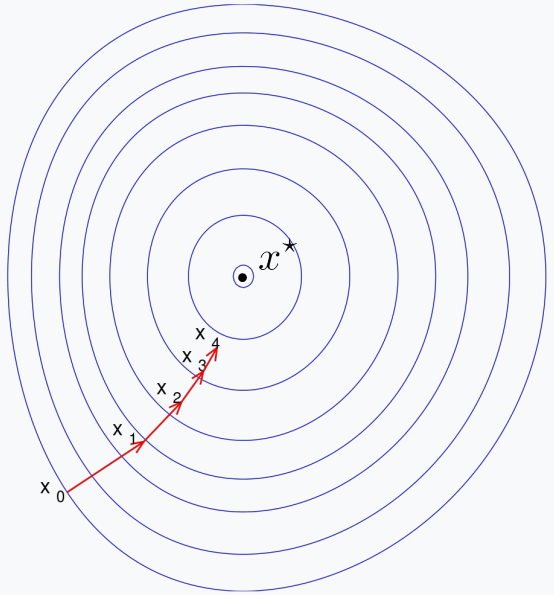

Google Research

Real-valued polynomials

Taking norms on the previous expression

Cauchy-Schwarz

$$\|\boldsymbol{x}_{t+1} - \boldsymbol{x}^\star\|_2 \leq \|(\boldsymbol{I} - \tfrac{2}{L+\mu}\boldsymbol{H})^{t+1}\|_2 \|\boldsymbol{x}_0 - \boldsymbol{x}^\star\|_2$$

Matrix 2-norm

$$\leq \max_{\lambda \in [\mu, L]} |(1 - \tfrac{2}{L+\mu}\lambda)^{t+1}| \|\boldsymbol{x}_0 - \boldsymbol{x}^\star\|_2$$

Convergence rate $\left(\frac{L-\mu}{L+\mu}\right)^t$

⚠ L, μ = largest and smallest eigenvalue of H

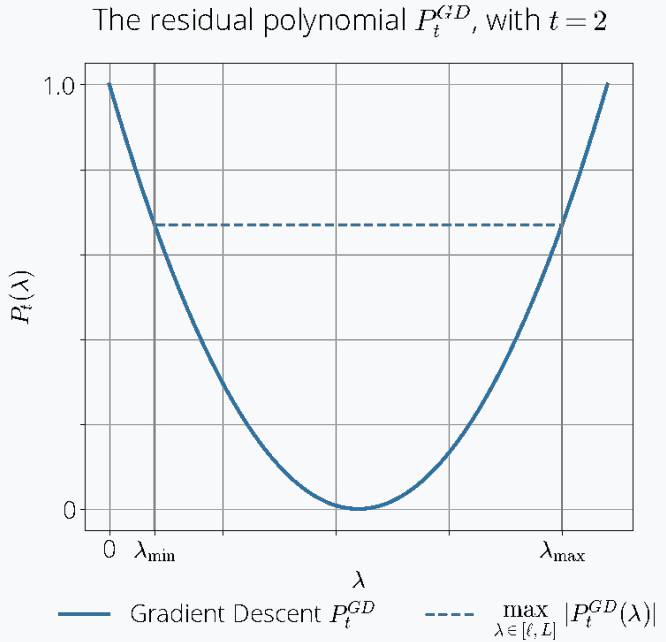

Google Research

Gradient-based Methods and Polynomials

Corollary (Convergence rate) *Let μ and L be the smallest and largest eigenvalue of $\boldsymbol{H}$ respectively. Then for any gradient-based method with residual polynomial P_t, we have*

$$\|\boldsymbol{x}_t - \boldsymbol{x}^\star\| \leq \underbrace{\max_{\lambda \in [\mu, L]}}_{\text{conditioning}} \underbrace{|P_t(\lambda)|}_{\text{algorithm}} \underbrace{\|\boldsymbol{x}_0 - \boldsymbol{x}^\star\|}_{\text{initialization}} . \qquad (17)$$

- Problem difficulty enters through [μ,L], interval that contains Hessian eigenvalues.
- Algorithm enters through polynomial P_t. This polynomial verifies $P_t(0)=1$

Google Research

HeavyBall

The HeavyBall update

$$\boldsymbol{x}_{t+1} = \boldsymbol{x}_t + m(\boldsymbol{x}_t - \boldsymbol{x}_{t-1}) - h\nabla f(\boldsymbol{x}_t)$$

Gives the residual polynomial

$$P_t(\lambda) = m^{t/2}\left(\frac{2m}{1+m}\underbrace{T_t(\sigma(\lambda))}_{\text{Chebyshev 1st kind}} - \frac{m-1}{1+m}\underbrace{U_t(\sigma(\lambda))}_{\text{Chebyshev 2nd kind}}\right)$$

with $\sigma(\lambda) = \frac{1}{2\sqrt{m}}(1+m-h\,\lambda)$

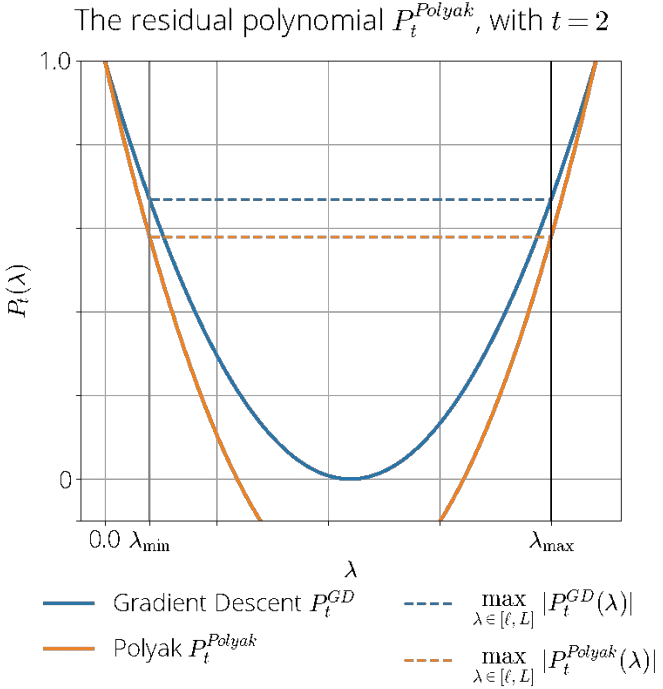

Google Research

The two faces of Chebyshev polynomials

In the [-1, 1] interval, Chebyshev polynomials are linearly bounded.

$$|T_t(\xi)| \leq 1 \quad \text{and} \quad |U_t(\xi)| \leq t+1$$

Outside, they grow exponentially.

$$T_t(\xi) = \frac{1}{2}\left(\xi - \sqrt{\xi^2-1}\right)^t + \frac{1}{2}\left(\xi + \sqrt{\xi^2-1}\right)^t$$

$$U_t(\xi) = \frac{(\xi + \sqrt{\xi^2-1})^{t+1} - (\xi - \sqrt{\xi^2-1})^{t+1}}{2\sqrt{\xi^2-1}}.$$

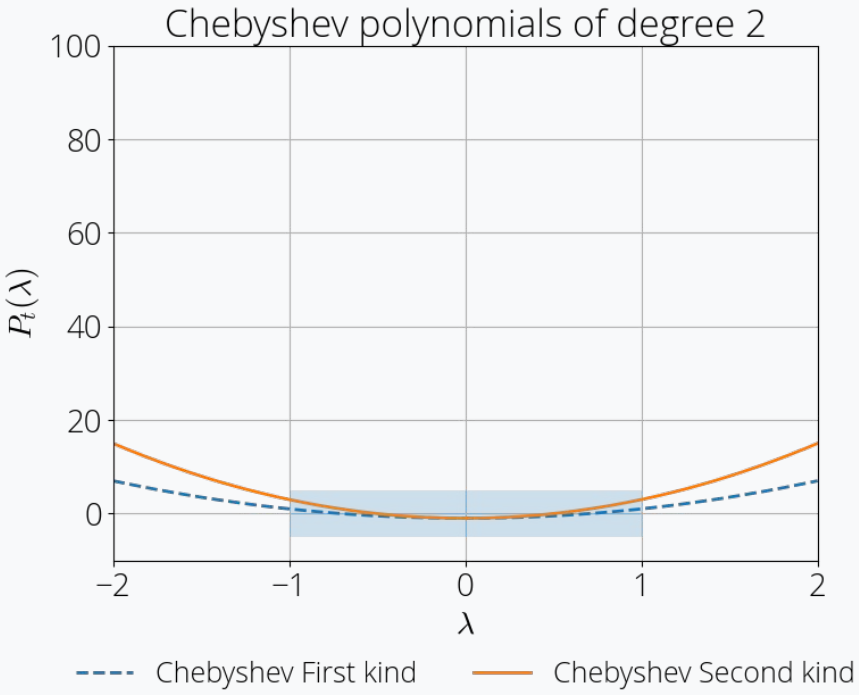

Google Research

Link function

$$\sigma(\lambda) = \frac{1}{2\sqrt{m}}(1 + m - h\,\lambda)$$

Pre-image is also an interval:

$$\sigma^{-1}([-1,1]) = \left[\frac{(1-\sqrt{m})^2}{h}, \frac{(1+\sqrt{m})^2}{h}\right]$$

Robust region: Parameters for which

$$[\mu, L] \subseteq \sigma^{-1}([-1,1])$$

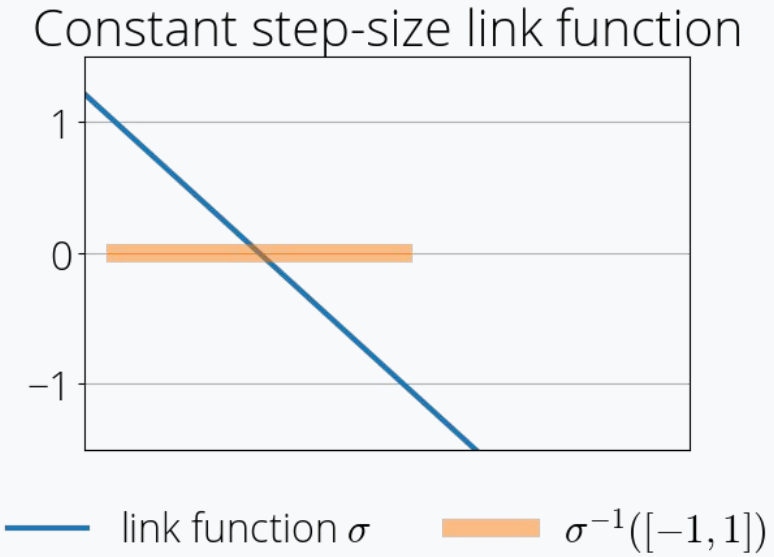

Google Research

The asymptotic rate in the robust region is $\sqrt{m}$.

$$\equiv \|x_t - x_\star\| = \mathcal{O}(\sqrt{m}^t)$$

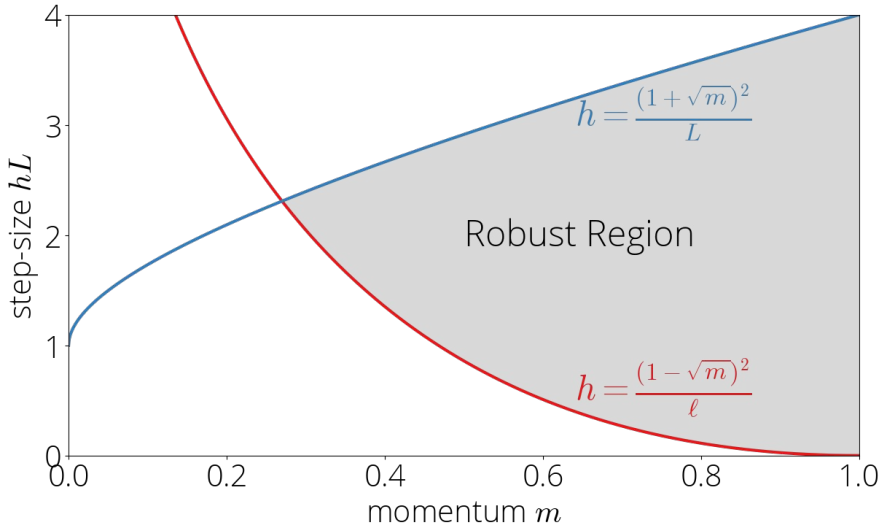

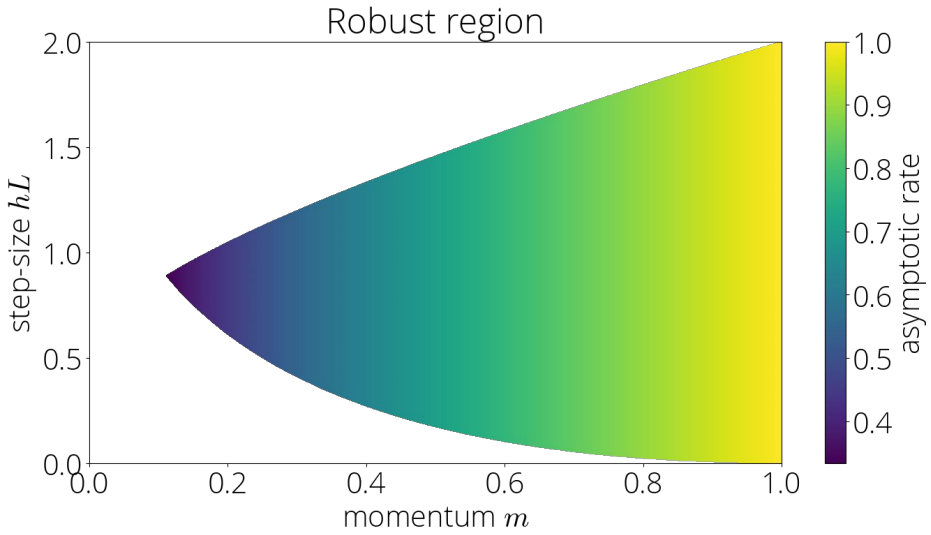

A Hitchhiker's Guide to Momentum, http://fa.bianp.net/blog/2021/hitchhiker/

Google Research

Optimal parameters (aka Polyak HeavyBall)

Minimizing ***m*** in the robust region results in (worst-case) optimal params

$$m = \left(\frac{\sqrt{L} - \sqrt{\mu}}{\sqrt{L} + \sqrt{\mu}}\right)^2 \qquad h = \left(\frac{2}{\sqrt{L} + \sqrt{\mu}}\right)^2$$

Asymptotic convergence rate: $\sqrt{m} = \dfrac{\sqrt{L} - \sqrt{\mu}}{\sqrt{L} + \sqrt{\mu}}$

Google Research

2. Cyclical HeavyBall

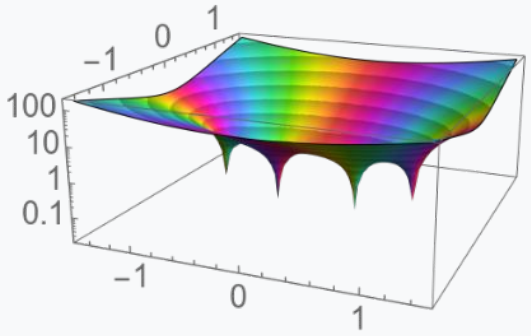

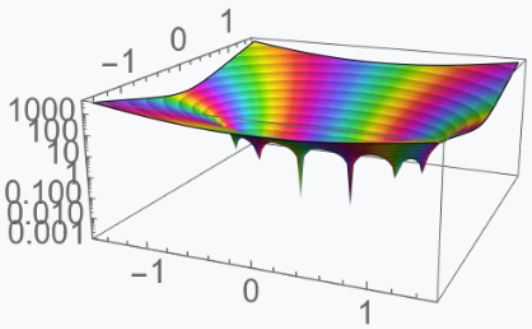

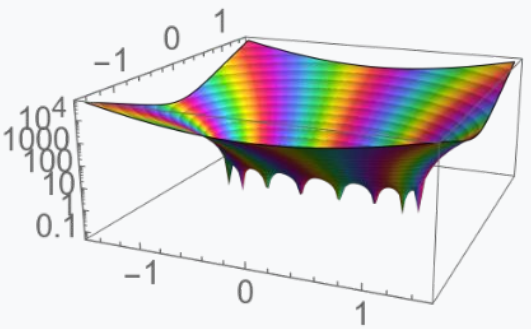

Google Research

Cyclical HeavyBall

Alternates between 2 step-sizes

Set $h_t = h_0$ if t is odd and $h_t = h_1$ otherwise

$$\boldsymbol{x}_{t+1} = \boldsymbol{x}_t - h_t \nabla f(\boldsymbol{x}_t) + m(\boldsymbol{x}_t - \boldsymbol{x}_{t-1})$$

Google Research

Analysis of Cyclical HeavyBall

- Coefficients in recurrence now depends on t

$$P_{2t+1}(\lambda) = (1 + m + h_1\lambda)P_{2t}(\lambda) - mP_{2t-1}(\lambda)$$
$$P_{2t+2}(\lambda) = (1 + m + h_0\lambda)P_{2t+1}(\lambda) - mP_{2t}(\lambda)$$

- Known in the OP field as "orthogonal polynomials with varying coefficients" [Chihara (1968), Van Assche (1985)]

- Analyzed by chaining iterations:

$$P_{2t+2}(\lambda) = ((1 + m + h_0\lambda)(1 + m + h_1\lambda) - 2m)\,P_{2t}(\lambda) - m^2 P_{2t-2}(\lambda)$$

TS Chihara

Google Research

Cyclical step-sizes

The residual polynomial for the cyclical HeavyBall method at even iterations is

$$P_{2t}(\lambda) = m^t \left(\frac{2m}{1+m} T_{2t}(\zeta(\lambda)) + \frac{1-m}{1+m} U_{2t}(\zeta(\lambda)) \right) , \qquad (6)$$

$$\text{with } \zeta(\lambda) = \frac{1+m}{2\sqrt{m}} \sqrt{(1 - \frac{h_0}{1+m}\lambda)(1 - \frac{h_1}{1+m}\lambda)}.$$

Same than HeavyBall except for link function ζ

Google Research

Complex Chebyshev polynomials

Image of link function can now be real or imaginary

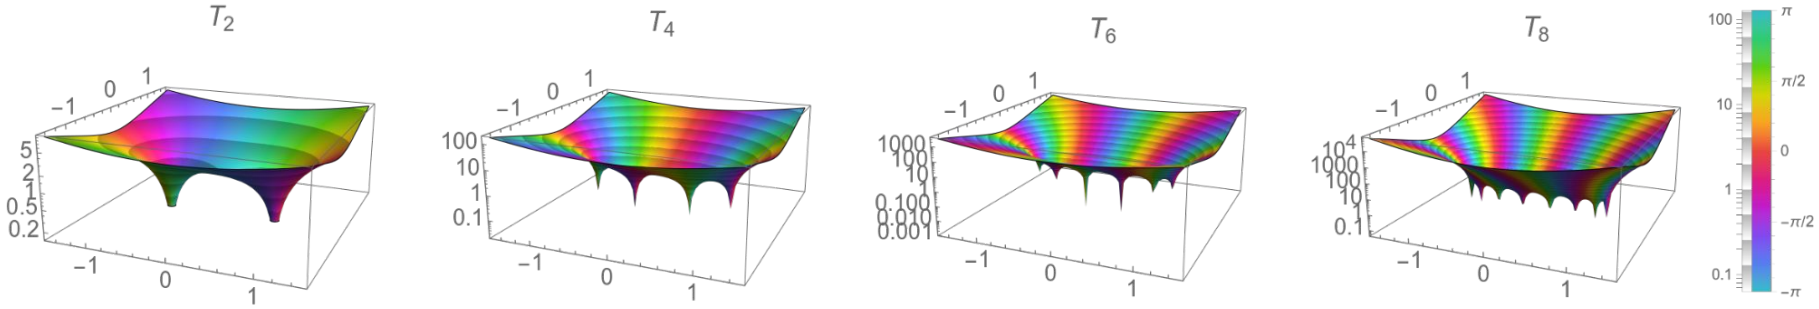

Chebyshev polynomials grow exponentially in $\mathbb{C}\backslash[-1, 1]$

Google Research

Link function

Pre-image no longer interval.

union of two intervals

Robust region: If

$[\mu, L] \subseteq \sigma^{-1}([-1, 1])$

Then $\|x_t - x_\star\| = \mathcal{O}(\sqrt{m}^t)$

Constant step-size link function

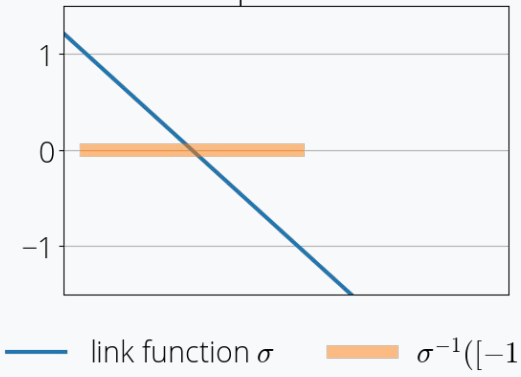

Cyclical step-size link function

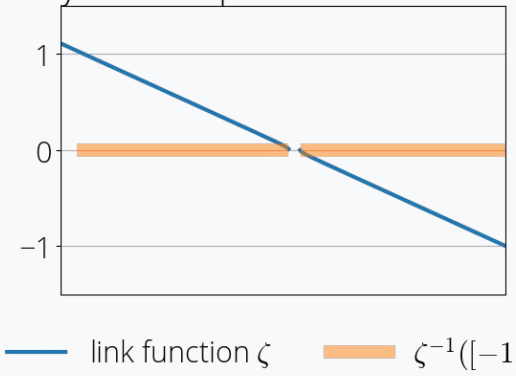

Google Research

A finer model for the Hessian eigenvalues

Consider eigenvalues in union of two disjoint intervals

$$\Lambda = [\mu_1, L_1] \cup [\mu_2, L_2], \underbrace{L_1 - \mu_1 = L_2 - \mu_2}_{\text{same size}}.$$

The ratio *R* will play an important role:

$$R \triangleq \frac{\mu_2 - L_1}{L_2 - \mu_1}$$

R = 0, one interval
R=1, all eigenvalues are at extremes.

MNIST

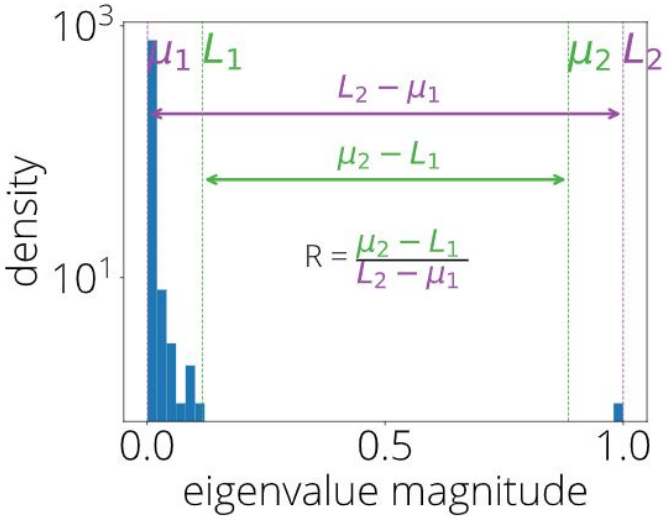

Google Research

Eigengaps Everywhere

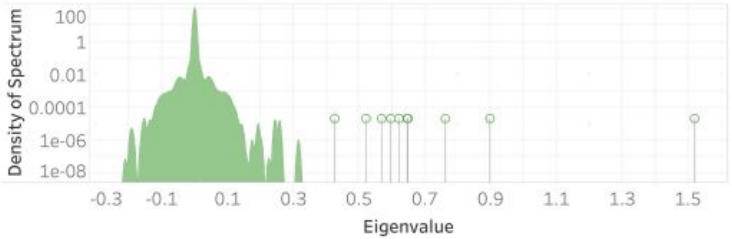

(a) MNIST, train

(b) MNIST, test

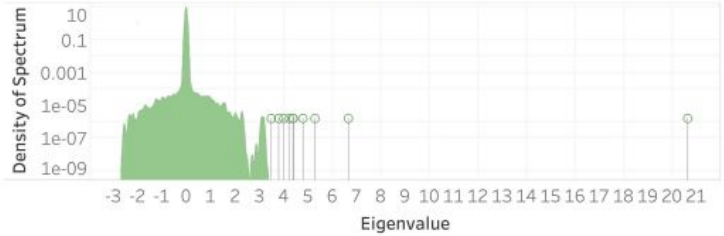

(e) CIFAR10, train

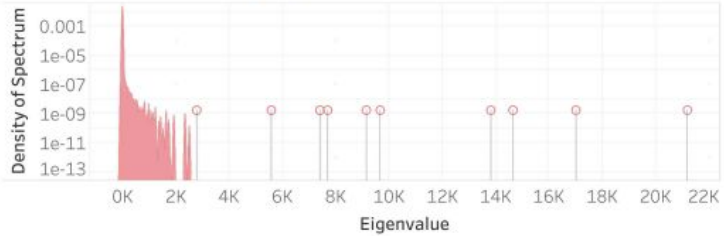

(f) CIFAR10, test

[(Papyan 2020)]

Google Research

Optimal parameters

Minimize ***m*** s.t. $[\mu,L] \subseteq \sigma^{-1}([-1,1])$

robust region

Solution

$$m = \left(\frac{\sqrt{\rho^2 - R^2} - \sqrt{\rho^2 - 1}}{\sqrt{1 - R^2}} \right)^2 \quad \text{with} \quad \rho \overset{\text{def}}{=} \frac{L_2 + \mu_1}{L_2 - \mu_1}$$

- R=0 we recover Polyak HeavyBall
- Decreasing in R

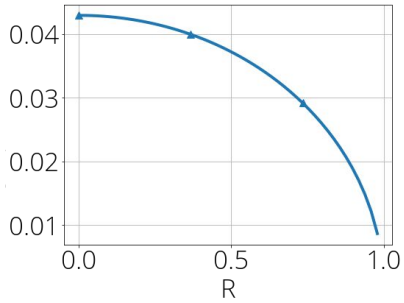

Google Research

Optimal step-sizes

$$h_t = \frac{1+m}{L_1} \text{ if } t \text{ is odd and } h_t = \frac{1+m}{\mu_2} \text{ otherwise}$$

Gap is small

Gap is large

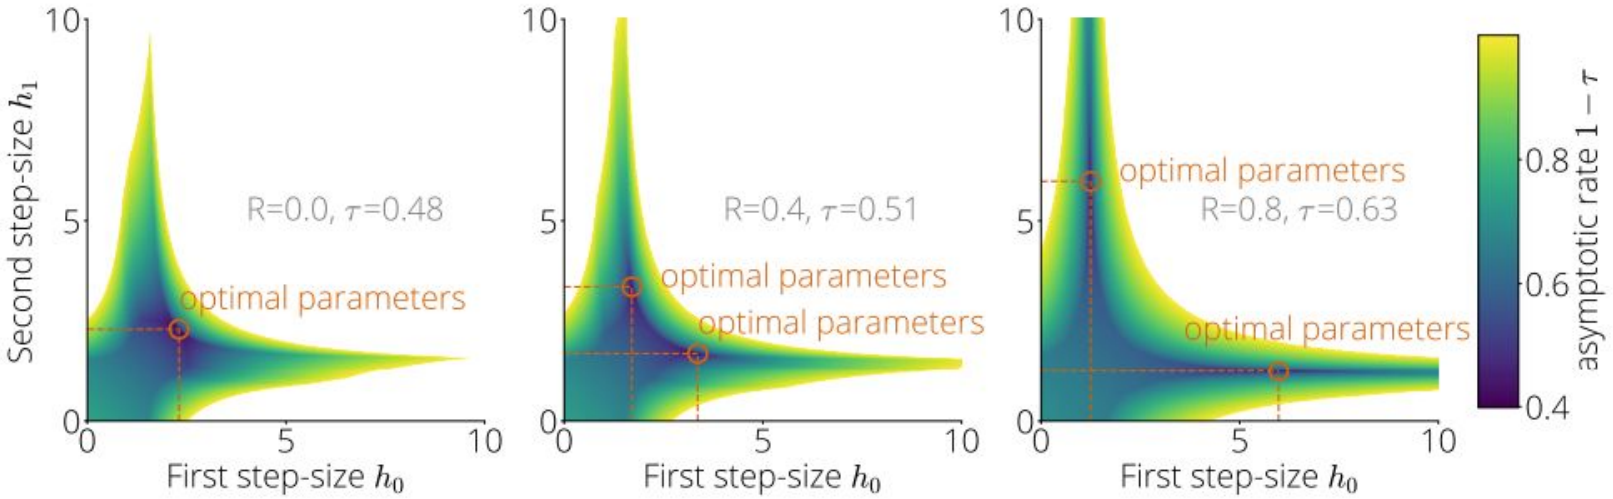

Google Research

Convergence Rates

Asymptotic rate = $\sqrt{m} = \dfrac{\sqrt{\rho^2 - R^2} - \sqrt{\rho^2 - 1}}{\sqrt{1 - R^2}}$

For ill-conditioned problems ($\mu \ll L$),

$$\sqrt{m} \approx \sqrt{1 - R^2}\, r^{\text{Polyak}}$$

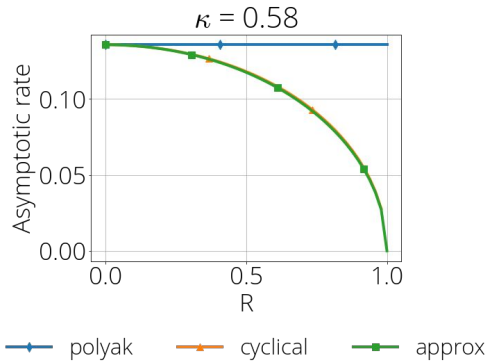

Google Research

Benchmarks

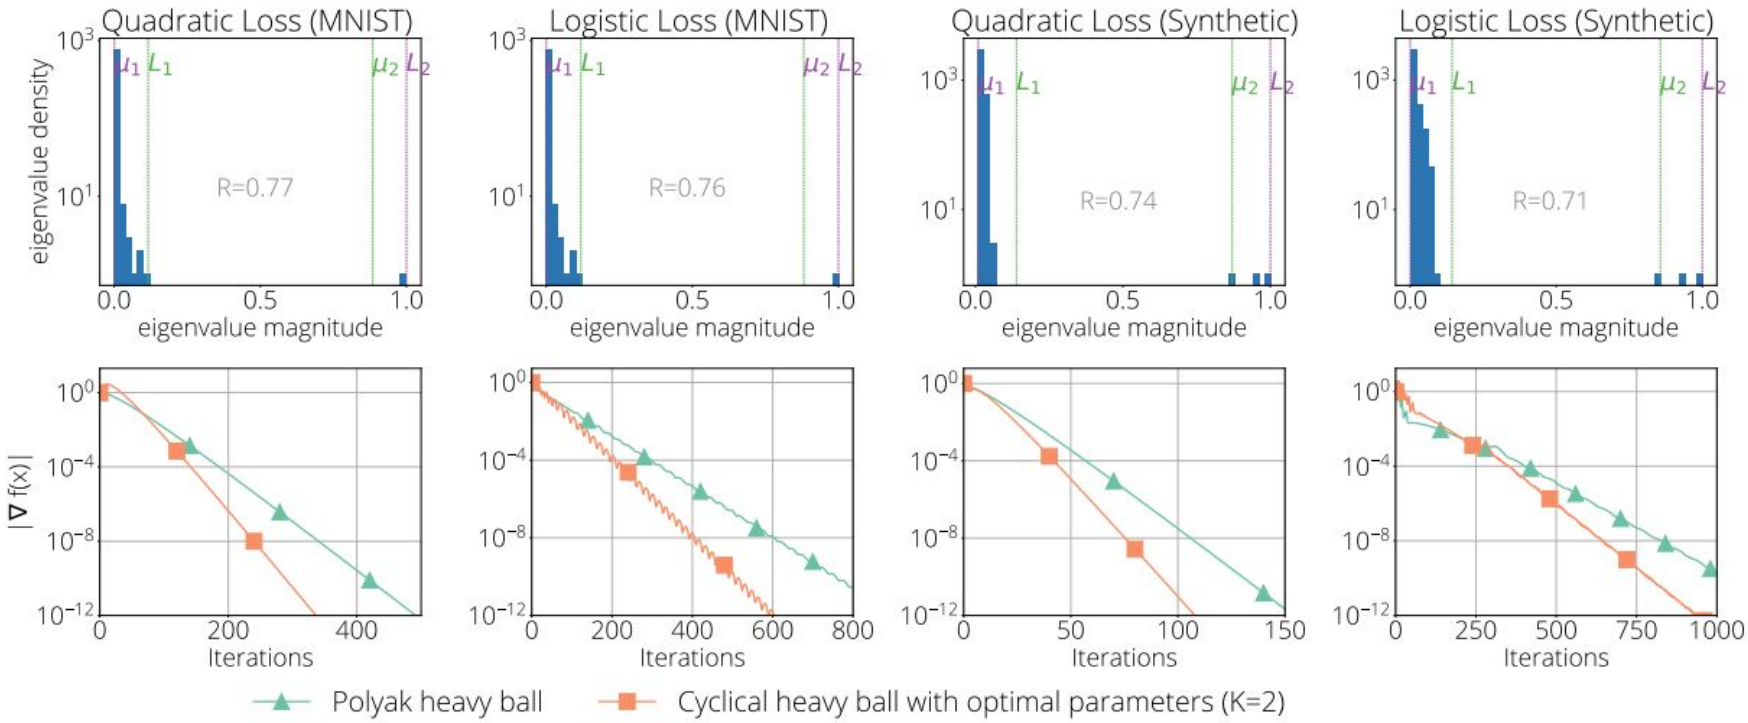

Google Research

Beyond cycles of length 2

Link functionsare optimal if 2ζ-1 hit ±1 at edges and ∉[-1, 1] outside

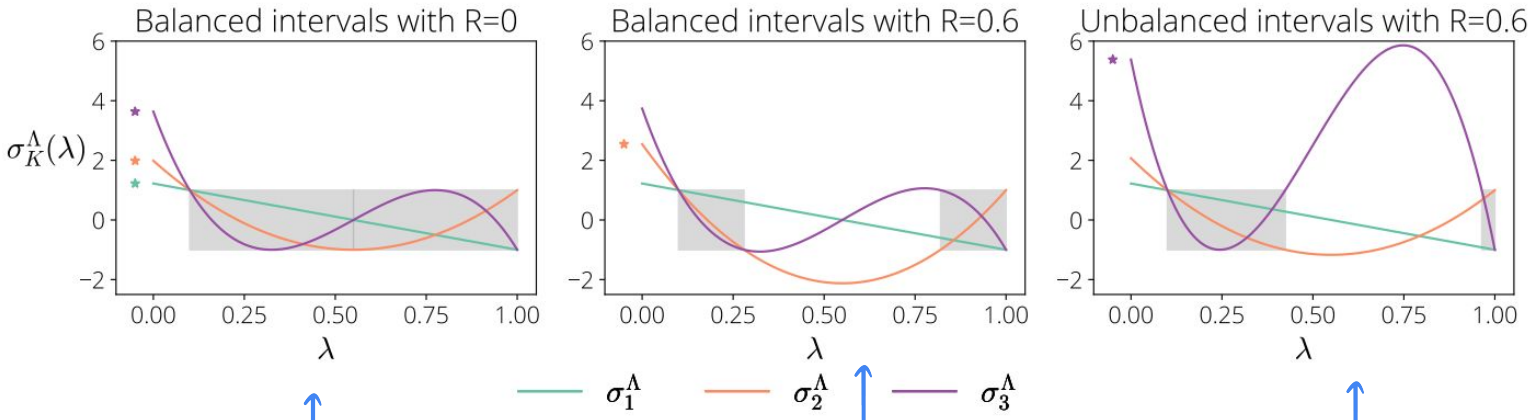

cycle 1 is optimal, cycle 2 and 3 are optimal, cycle 3 is optimal

Conclusions

Cyclical Heavy Ball converges faster in the presence of spectral gap.

Assuming knowledge of this gap, converges at a rate $\approx \sqrt{1 - R^2}\, r^{\mathrm{Polyak}}$

Speedup observed also on non-quadratic objectives.

Google Research

Open Problems

More complex Hessian support: closed form for larger cycles.

Interpolating step-sizes

How to estimate the eigen-gap?

Stochastic algorithm? Non-quadratic objectives?