

## » Stochastic Gradient Descent (SGD)

- \* Recall our general iterative minimisation algorithm:

```
x=x0
for k in range(num_iters):
    step = calcStep(fn,x)
    x = x - step
```

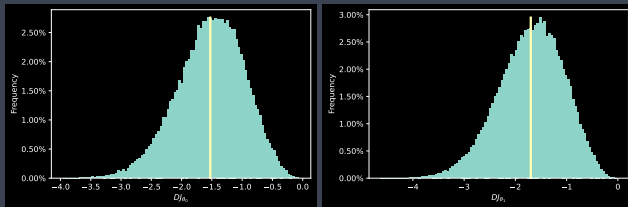
and one way to choose the step, namely:

$$step = \alpha \left[ \frac{\partial f}{\partial x_1}(x), \frac{\partial f}{\partial x_2}(x), \dots, \frac{\partial f}{\partial x_n}(x) \right]$$

- \* In SGD we use approximate derivatives  $DJ_{x_1}(\theta)$  etc instead of exact derivatives  $\frac{\partial f}{\partial x_1}(x)$  etc

## » SGD “Noise”

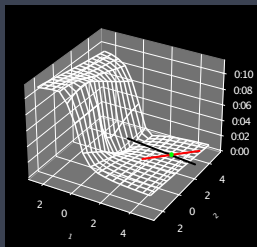
- \* Recall  $DJ_{\theta_i} = \frac{\partial J}{\partial \theta_i} + \text{noise}$
- \* E.g. Quadratic loss and  $y = \theta^T x + \text{noise}$ ,  $\theta = [3, 4]$ , repeatedly draw random samples, calc  $DJ_{\theta_1}(\theta)$  and plot the values:



$$\frac{\partial J}{\partial \theta_0} = -1.526, \quad \frac{\partial J}{\partial \theta_1} = -1.699$$

## » Impact of SGD “Noise”

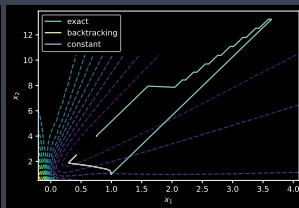
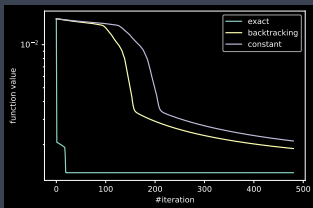
- \* How does the noise on the approx derivatives it affect line search?
- \* Toy neural net:



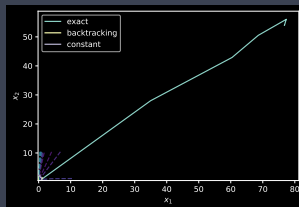
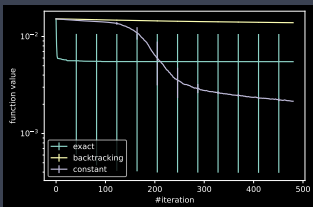
- \* Suppose black line indicates derivative and red line the approximate derivative  $\rightarrow$  since the surface is quite flat, small errors in approx derivative can cause linesearch to select very different points

## » Impact of SGD “Noise”: Line Search

- \* Toy neural net, starting point  $x = [1, 1]$ , constant  $\alpha_0 = 0.75$
- \* Gradient descent:



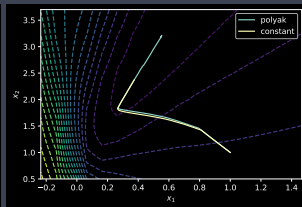
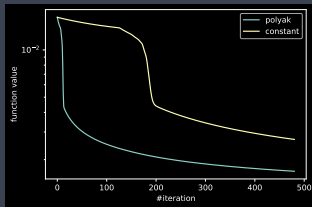
- \* Mini-batch SGD, batch size 5



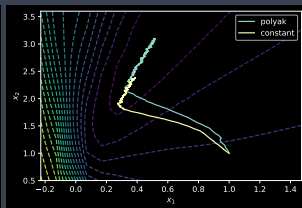
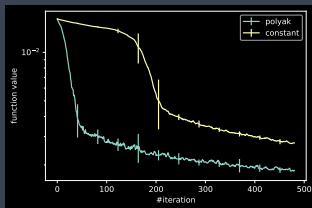
- \* Exact/backtracking line search can take v large steps, tend to greatly amplify the “noise” and so fail to converge

## » Impact of SGD “Noise”: Polyak

- \* Polyak choice of step size?
- \* Toy neural net, starting point  $x = [1, 1]$ , constant  $\alpha_0 = 0.75$
- \* Gradient descent:



- \* Mini-batch SGD, batch size 5



- \* Seems ok, but has not been much studied for ML applications<sup>1</sup>.

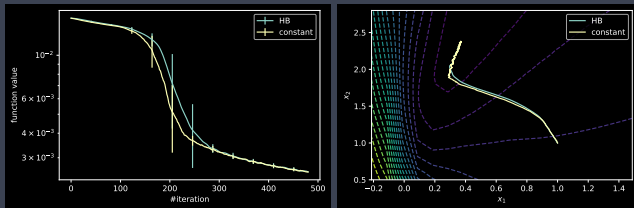
<sup>1</sup>Although see <http://proceedings.mlr.press/v130/loizou21a.html>

## » SGD and Momentum

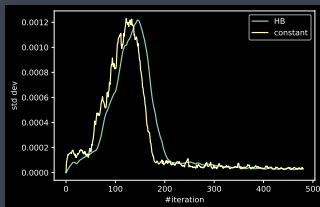
- \* Intuition: SGD uses noisy approximate gradients. Momentum averages past values of gradients, and so should average out the noise in the approx gradients → reduce impact of approx gradient noise
- \* But as usual its not quite so straightforward. Let's look at an example ....

## » SGD and Momentum

- \* Use Heavy Ball since its a simple momentum approach
- \* Toy neural net, starting point  $x = [1, 1]$ , constant  $\alpha_0 = 0.75$ , HB  $\alpha = 0.075/\beta = 0.9$
- \* *Step sizes of HB and constant strategy are matched, its just momentum that differs. Mini-batch SGD, batch size 5*

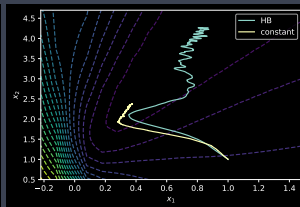
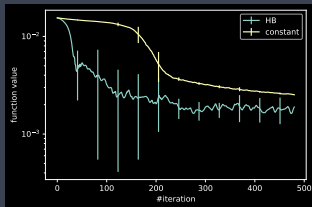


- \* Error bars (standard deviation) much the same for both algo's:



## » SGD and Momentum

- \* Momentum allows is to increase the step size and get faster convergence ... increasing step size amplifies SGD “noise”
- \* Toy neural net, starting point  $x = [1, 1]$ , constant  $\alpha_0 = 0.75$ , HB  $\alpha = 0.75/\beta = 0.9$  (HB  $\alpha$  now  $10\times$  larger):

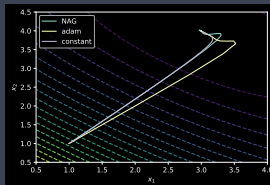
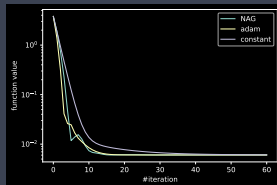


- \* Error bars (standard deviation) are now larger for HB than with constant step size
- \* Let's look at NAG and Adam next. Recall NAG schedules  $\beta$  and uses lookahead, Adam combines RMSprop+HB

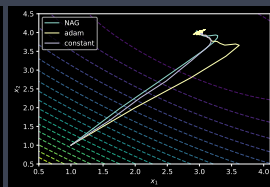
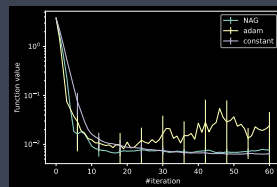


## » Impact of SGD “Noise”: NAG and Adam

- \* Quadratic loss  $y = \theta^T x + \text{noise}$  with  $\theta = [3, 4]$ ,  $y$  noise std deviation 1, starting point  $x = [1, 1]$ , constant  $\alpha_0 = 0.75$ , NAG  $\alpha_0 = 0.5$ , Adam  $\alpha_0 = 1/\beta_1 = 0.25$
- \* Gradient descent:



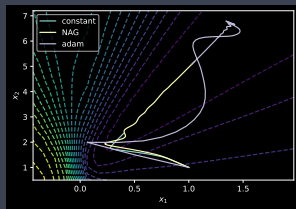
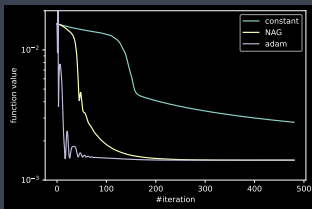
- \* Mini-batch SGD, batch size 5



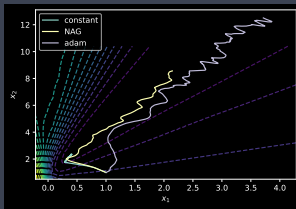
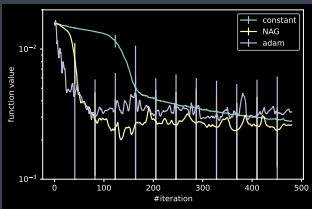
- \* Adam creates oscillations rather than removing them! NAG v similar to constant step size strategy

## » Impact of SGD “Noise”: NAG and Adam

- \* Toy neural net  $m = 100$  training data points,  $\theta = [1, 5]$ , output  $y$  noise with std dev 0.05, starting point  $x = [1, 1]$ , constant  $\alpha_0 = 0.75$ , NAG  $\alpha_0 = 0.5$ , Adam  $\alpha_0 = 0.5/\beta_1 = 0.25$



- \* Mini-batch SGD, batch size 5



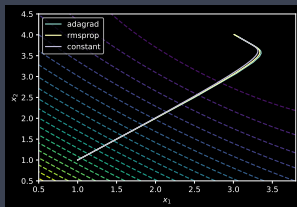
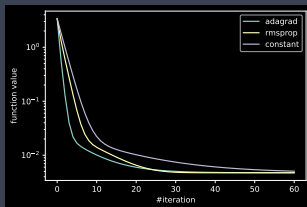
- \* SGD noise greatly amplified by NAG and Adam. Could reduce  $\alpha$  to reduce fluctuations but then convergence slows

## » SGD and Diminishing Step Size

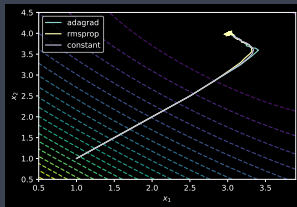
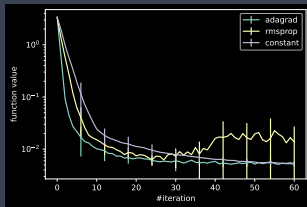
- \* Intuition:
  - \* Initially use large step size to move quickly to vicinity of minimum - noise in approx gradients is not too important.
  - \* Once near minimum noise+large step size means we'll bounce around it → so gradually reduce step size so that bounce around less and eventually converge to minimum
- \* Can implement this manually using a step size “schedule”:
  - \*  $\alpha = \alpha_0 \gamma^t$ . Reduce step size  $\alpha$  by fixed factor  $\gamma$  at end of each epoch (or after a fixed number of epochs) i.e. step sizes are  $\alpha, \gamma\alpha, \gamma^2\alpha$  etc. E.g. for  $\gamma = 0.1$  and  $\alpha = 0.1$  then the step sizes are 0.1, 0.01, 0.001, 0.0001 etc
  - \*  $\alpha = \alpha_0 / t^\gamma$ , e.g.  $\alpha = \alpha_0 / \sqrt{t}$  when  $\gamma = 0.5$
  - \* Cosine decay, see <https://arxiv.org/pdf/1608.03983.pdf>
- \* Use Adagrad/RMSprop (or Adam)

## » Impact of SGD “Noise”: Adagrad/RMSprop

- \* Quadratic loss  $y = \theta^T x + \text{noise}$  with  $\theta = [3, 4]$ ,  $y$  noise std deviation 1, starting point  $x = [1, 1]$ , constant  $\alpha_0 = 0.5$ , Adagrad  $\alpha_0 = 1.5$ , RMSprop  $\alpha_0 = 0.3$
- \* Gradient descent:



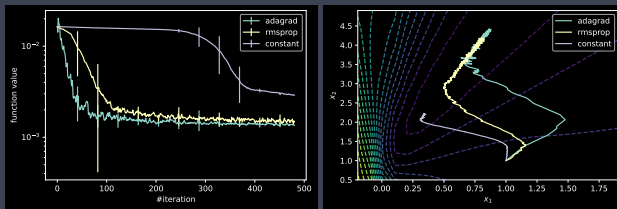
- \* Mini-batch SGD, batch size 5



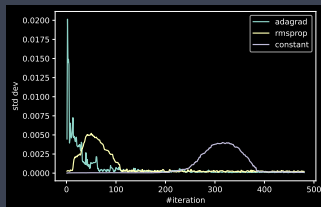
- \* RMSprop can increase step size when close to the minimum since gradient is small ( $\alpha = \alpha_0 / \text{sum}$ ,  $\text{sum} = \beta \text{sum} + (1 - \beta) \frac{df}{dx}(x_t)^2$ )

## » Impact of SGD “Noise”: Adagrad/RMSprop

- \* Toy neural net, starting point  $x = [1, 1]$ , constant  $\alpha_0 = 0.75$ , Adagrad  $\alpha_0 = 0.05$ , RMSprop  $\alpha_0 = 0.75$
- \* Mini-batch SGD, batch size 5

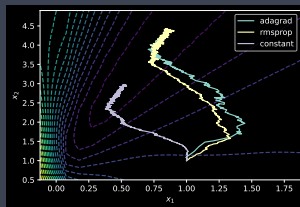
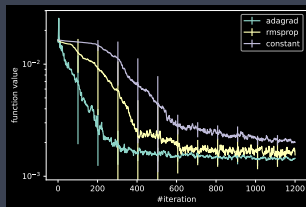


- \* Error bars (standard deviation) are similar for all algos, although constant step size has a jump around iterations 300-400:

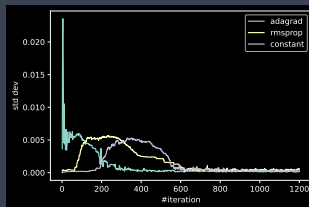


## » Impact of SGD “Noise”: Adagrad/RMSprop

- \* What if reduce mini-batch size (and so increase SD noise)?
- \* Toy neural net, starting point  $x = [1, 1]$ , constant  $\alpha_0 = 0.75$ , Adagrad  $\alpha_0 = 0.05$ , RMSprop  $\alpha_0 = 0.75$
- \* Mini-batch SGD, batch size 2



- \* Error bars (standard deviation) are again similar for all algos:



## » Summary So Far

- \* Interaction between momentum and SGD noise is complicated:
  - \* Momentum  $\beta$  can increase oscillations as well as decrease them
  - \* Increasing gain  $\alpha$  tends to increase oscillations but to accelerate convergence
  - \* → on balance, in our examples momentum tends to amplify SGD noise

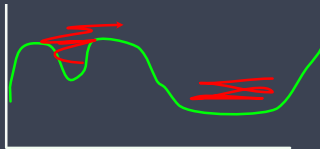
## » Summary So Far

- \* More straightforward for diminishing step size → decreasing the step size enough must eventually reduce fluctuations in  $x$  caused by SGD noise
  - \* Use of a pre-defined step size schedule is popular e.g. scale step size by 0.1 after every 10 epochs → *but choice of schedule is a bit of a black art, a current research topic*
  - \* With adaptive schedule (Adagrad, Adam) need to strike a balance between (i) large initial step size for fast convergence, (ii) not decreasing step size so quickly that “stall” before reaching minimum and (iii) reducing step size enough to reduce SGD noise
  - \* *In our examples, when  $\alpha$  is tuned for fast convergence the impact of SGD noise with Adagrad and Adam is much the same as with a constant step size*
- \* Why do we care so much about SGD noise anyway?



## » Regularisation By SGD Noise

- \* Intuition (for *nonconvex* cost functions):
  - \* SGD noise will tend to allow algorithm to escape from narrow minima ...
  - \* ... and encourage algorithm to settle on minima located in a broad bowl



- \* Broad minima tend to correspond to models that generalize well

## » Regularisation By SGD Noise

- \* Some experimental evidence for regularising role of SGD noise when training neural nets but its an active area of research → *not well understood!*
  - \* Generalisation performance observed to be poorer with larger mini-batches<sup>2</sup>
  - \* Increasing step size with larger mini-batch size may compensate<sup>3</sup>
  - \* Rather than decreasing step size over time increase the mini-batch size<sup>4</sup>
  - \* Adaptive methods (Adagrad, Adam, RMSprop) may converge to solutions that don't generalise well<sup>5</sup>
- \* *Some more intuition: “annealing”.*
  - \* Start with plenty of SGD noise - so can escape local minima and explore cost surface.
  - \* Gradually decrease noise (but not too quickly) - will tend to settle in region around a broad minimum and then converge onto minimum itself
  - \* → might be better way to think about step-size schedule, and joint mini-batch size/step-size scheduling

<sup>2</sup> On large-batch training for deep learning: Generalization gap and sharp minima, 2017 <https://arxiv.org/pdf/1609.04836.pdf>

<sup>3</sup> Accurate, Large Minibatch SGD: Training ImageNet in 1 Hour, 2018. <https://arxiv.org/pdf/1706.02677.pdf>

<sup>4</sup> Don't Decay The Learning Rate, Increase The Batch Size 2018. <https://arxiv.org/pdf/1711.00489.pdf>

<sup>5</sup> The Marginal Value of Adaptive Gradient Methods in Machine Learning, 2018. <https://arxiv.org/pdf/1705.08292.pdf>

## » How to choose step-size and mini-batch size?

- \* Common practice seems to be:
  - \* Select mini-batch size using cross-validation. Typical choices are in range 32-256.
  - \* Use a constant step size initially, then decrease when cost function starts to plateau (as measured on held-out test data)<sup>6</sup>  
→ piece-wise constant schedule, typically manually tuned
  - \* To find initial step size, guess a value → if cost rises or fluctuates a lot then reduce step size, if cost decreases slowly increase step size
- \* *But this is a fast moving topic ... this advice may well change.*

---

<sup>6</sup>E.g. see

## » Which optimisation algorithm to use?

- \* There is no one “best” optimisation algorithm, it depends on the problem being tackled:
  - \* Look in the literature for other work that has looked at similar problems/models → what algorithms do they use, and why
- \* If cost function is *convex* then optimisation is much easier:
  - \* Pretty much any standard optimisation algo will be fine, guaranteed to converge to global minimum
  - \* If *strongly convex/quadratic-like* (e.g. linear regression) → very fast convergence
  - \* If convex but not strongly convex (e.g. logistic regression) or non-smooth/has kinks (e.g. with  $L_1$  regulariser) → convergence might be slower
- \* However, neural net cost functions are *non-convex*:
  - \* *Mini-batch SGD with constant step-size or step-size schedule is a good baseline* → often performs well, hard to beat
  - \* *Adam* seems popular currently, but tune  $\alpha$ ,  $\beta$  parameters
  - \* We haven't touched on: *initialisation of NN weights, batch normalisation.*
  - \* *Be alert to potentially complex nature of cost function (flat areas, steep valleys, multiple minima etc), interactions between mini-batch size, step-size and generalisation*