



Neural Networks

Lecture 5

Gradient-based learning and optimization

Igor Farkaš

2021

Optimization vs NN learning

- Although optimization provides a way to minimize the loss function for NN learning, the goals are fundamentally different:
- goal of optimization = **to reduce the training error**
- goal of NN learning (statistical inference) = **to reduce expected generalization error** (risk) => ML acts indirectly
- We have only access to a finite training data sample (not the whole data distribution)
- Empirical risk minimization: $E_{(x,d) \sim p(\text{data})}[Loss(f(x;w),d)]$
- ... is based on a finite training sample $\{x,d\}$, rather than known data distribution, hence is prone to overfitting.

(Goodfellow et al, 2015)

2


Surrogate loss function

- Sometimes, the loss function we actually care about (e.g. classification error) is not one that can be optimized efficiently.
- Exactly minimizing expected 0-1 loss is typically intractable (exponential in the input dimension)
- In such situations, one typically optimizes a surrogate loss function instead, which acts as a proxy, but has advantages:
- e.g. the negative log-likelihood of the correct class is used ($-\log P(y_i)$)
- test set 0-1 loss often continues to decrease for a long time after the training set 0-1 loss has reached zero, which improves the robustness of the classifier by further pushing the classes apart from each other
- This leads to extracting more information from the training data (than would have been possible by simply minimizing the average 0-1 loss on training set).

3


Problems in gradient-based NN learning

Local minima




x = weights,
 F = loss f.

Saddle points



Vanishing gradients



Ill-conditioning of Hessian matrix \mathbf{H} , i.e. rate of its change for small $\Delta \mathbf{w}$

- given by condition number (CN) = ratio of its max/min eigenvalues
- for large CN, \mathbf{H}^{-1} is particularly sensitive to error in the input



These problems slow down or hinder convergence.

4

Early modifications of gradient descent learning

Adding a **momentum**: $\Delta \mathbf{w}(t) = -\alpha \nabla E(\mathbf{w}(t)) + \gamma \Delta \mathbf{w}(t-1)$ $0 \leq |\gamma| < 1$


- helps speed up SGD and dampen oscillations



Nesterov accelerated gradient (Nesterov, 1983)

$$\Delta \mathbf{w}(t) = -\alpha \nabla E[\mathbf{w}(t) + \gamma \Delta \mathbf{w}(t-1)] + \gamma \Delta \mathbf{w}(t-1)$$

- helps adjust speed by looking into the near future



5

Role of the Hessian matrix

- \mathbf{H} plays an important role in supervised training of neural networks:

$$\mathbf{H} = \left[\frac{\partial^2 E(\mathbf{w})}{\partial w_i \partial w_j} / \mathbf{w}_0 \right]_{ij}$$

- Eigenvalues of \mathbf{H} have a profound influence on the dynamics of back-propagation learning (condition number = ratio of max/min eigenvalues)
- The inverse of \mathbf{H} provides a basis for pruning (i.e., deleting) insignificant synaptic weights from a multilayer perceptron.
- \mathbf{H} is basic to the formulation of second-order optimization methods as an alternative to BP learning.
- Typical profile of \mathbf{H} in BP learning (LeCun et al., 1998): a few small eigenvalues, many medium-sized eigenvalues, and a few large eigenvalues \Rightarrow a wide spread in the eigenvalues of the Hessian.

6

Towards second-order optimization methods

$$E(\mathbf{w}) = E(\mathbf{w}_0) + \mathbf{g}^T(\mathbf{w}_0) \Delta \mathbf{w} + 1/2 \Delta \mathbf{w}^T \mathbf{H}(\mathbf{w}_0) \Delta \mathbf{w} + O^{3+}(\Delta \mathbf{w})$$

$$\Delta \mathbf{w} = \mathbf{w} - \mathbf{w}_0$$

Taylor expansion:

Gradient vector:

$$\mathbf{g}(\mathbf{w}_0) = \nabla E(\mathbf{w}_0) = \left[\frac{\partial E}{\partial w_1} / \mathbf{w}_0, \dots, \frac{\partial E}{\partial w_{|w|}} / \mathbf{w}_0 \right]^T$$

$$1D: f(x) = \sum_{i=0}^{\infty} \frac{f^{(i)}(x_0)}{i!} (x - x_0)^i$$

- Error back-propagation is a linear approximation of E : $\Delta \mathbf{w}(t) = -\alpha \mathbf{g}(t)$
- Quadratic approx. of $E(\mathbf{w}) \rightarrow$ **Newton's method**: $\Delta \mathbf{w} = -\mathbf{H}^{-1}(t) \mathbf{g}(t)$
- Quasi-Newton method** approximates $\mathbf{H}^{-1}(t)$ with a positive definite matrix
- Conjugate-gradients methods** are intermediate between the steepest descent and the Newton's method, by achieving faster convergence (than the former) and lower computational complexity (than the latter).

7

Conjugate-gradient methods


- 2nd order optimization methods
- minimize the quadratic function $f(\mathbf{x}) = 1/2 \mathbf{x}^T \mathbf{A} \mathbf{x} + \mathbf{b}^T \mathbf{x} + c$ (*)
- \rightarrow set of linear equations: $\mathbf{A} \mathbf{x} = \mathbf{b}$ (\mathbf{A} = positive definite and symmetric)
- Solution: $\mathbf{x}^* = \mathbf{A}^{-1} \mathbf{b}$
- Given the matrix \mathbf{A} , a set of nonzero vectors $s(0), s(1), \dots$, (up to $\dim(\mathbf{A})$) is \mathbf{A} -conjugate (i.e., non-interfering with each other in the context of \mathbf{A}) if: $s(i)^T \mathbf{A} s(j) = 0$. (for $\mathbf{A} = \mathbf{Id}$, conjugacy = orthogonality).

- Example: $\mathbf{x} = [x_0, x_1]$

- Let $\mathbf{v} = \mathbf{A}^{1/2} \mathbf{x}$

- Iterative CG method:
 $\Delta \mathbf{x} = \eta(t) \cdot s(t)$

(Hestenes & Stiefel, 1952)



plot of eq. *

8

Regularization


Risk function: $R(\mathbf{w}) = E(\mathbf{w}) + \lambda C(\mathbf{w})$ [performance + complexity]

- Explicit:**

$$L_1(\mathbf{w}) = \epsilon \sum_i |w_i| \quad L_2(\mathbf{w}) = \frac{\epsilon}{2} \|\mathbf{w}\|^2$$

- Implicit:**


- weight decay $w_i^{\text{new}}(t) = \epsilon \cdot w_i^{\text{new}}(t)$, $0 \ll \epsilon < 1$,
- ...leads to L_2 -regul.
- dropout (Hinton, 2012): random turning off neurons during training
- data augmentation – increasing the size of the training set, e.g. by elastic distortions



9

Dropout

- Applied only during training
- Helps to avoid overfitting
- Free parameter = number of (randomly) dropped units



(Zhang et al, 2019)

10

AdaGrad algorithm

- Introduces the variable that accumulates gradient variance (vector)
 $\mathbf{g}(t) = \nabla E(\mathbf{w}(t)) \quad \mathbf{s}(t) = \mathbf{s}(t-1) + \mathbf{g}^2(t) \quad \Delta \mathbf{w}(t) = -\frac{\alpha}{\sqrt{\mathbf{s}(t) + \epsilon}} \cdot \mathbf{g}(t)$
- decreases the learning rate dynamically on per-coordinate basis
- uses the magnitude of the gradient as a means of adjusting how quickly progress is achieved – coordinates with large gradients are compensated with a smaller learning rate.
- First-order method (the gradient can be a useful proxy)
- On deep learning problems Adagrad can sometimes be too aggressive in reducing learning rates. Mitigating strategies exist.

(Duchi et al, 2011)

11

RMSprop

- decouples rate scheduling from coordinate-adaptive learning rates
 $\mathbf{s}(t) = \gamma \mathbf{s}(t-1) + (1-\gamma) \mathbf{g}^2(t) \quad \Delta \mathbf{w}(t) = -\frac{\alpha}{\sqrt{\mathbf{s}(t) + \epsilon}} \cdot \mathbf{g}(t) \quad \epsilon = 10^{-6}$
- coefficient γ determines how long the history is when adjusting the per-coordinate scale.
- RMSprop shares with momentum the leaky averaging. However, RMSprop uses the technique to adjust the coefficient-wise preconditioner (for reducing the condition number).

(Tieleman & Hinton, 2012)

12

AdaDelta

- Yet another variant of AdaGrad: it decreases the amount by which the learning rate is adaptive to coordinates
- It does not literally have a learning rate since it uses the amount of change itself as calibration for future change:

$$s(t) = \rho s(t-1) + (1-\rho) g^2(t) \quad g'(t) = \sqrt{\frac{\Delta w(t-1) + \epsilon}{s(t) + \epsilon}} \cdot g(t)$$

$$w(t) = w(t-1) - g'(t)$$

$$\Delta w(t) = \rho \Delta w(t-1) + (1-\rho) w^2(t)$$

(Zeiler, 2012)

13

Adam algorithm

- Combines 3 preceeding techniques into one efficient algorithm
- uses leaky averaging to obtain an estimate of both the momentum and also the second moment of the gradient

$$\begin{aligned} v(t) &= \beta_1 v(t-1) + (1 - \beta_1) g(t) & v'(t) &= v(t) / (1 - \beta_1^t) & \beta_1 &= 0.9 \\ s(t) &= \beta_2 s(t-1) + (1 - \beta_2) g^2(t) & s'(t) &= s(t) / (1 - \beta_2^t) & \beta_2 &= 0.999 \end{aligned}$$

$$\Delta w(t) = -\frac{\alpha}{\sqrt{s'(t) + \epsilon}} \cdot v'(t) \quad (\text{Kingma \& Ba, 2014})$$

- Still, gradients with significant variance may hinder convergence ($s(t)$ can blow up)

- Yogi algorithm addresses this: $s(t) = s(t-1) + (1 - \beta_2) (g^2(t) - s(t-1))$

(Zaheer et al, 2018)

$$s(t) = s(t-1) + (1 - \beta_2) g^2(t) \cdot \text{sgn}(g^2(t) - s(t-1))$$

14

Natural gradient learning

- use **Fisher information**: a positive semidefinite matrix ($|w| \times |w|$), defines a Riemannian metric (-> information geometry) (Amari, 1998)
- look at p.d.f. via $KL(f(x;w) || f(x;w+\Delta w)) = \dots \approx \frac{1}{2} (\Delta w)^T \mathbf{F} \Delta w$
- matrix \mathbf{F} is the negative expected Hessian of $\log f(x;w)$
- $\Delta w^* = \arg \min_{\Delta w} \{L(w+\Delta w) + \lambda \cdot KL(f(x;w) || f(x;w+\Delta w)) - c\}$
- $\Delta w(t) = -\alpha \mathbf{F}^{-1}(w(t)) g(t)$, i.e. **natural gradient** $g_{nat}(t) = \mathbf{F}^{-1}(w) g(t)$
- can be interpreted as curvature of the log likelihood function f
- in NG descent, we control movement in **prediction space** (rather than parameter space)
- Approximations of \mathbf{F}^{-1} possible (Amari et al, 2019)

15

Summary

- NN learning and classical optimization have different objectives
- NN goal = minimize generalization error (sometimes using surrogate loss functions)
- Various known problems hinder first-order gradient methods
- Second-order methods provide more informaton but are much more costly
- Earlier methods focused on approximating the Hessian
- Recent methods foces only on gradients and its adaptive versions
- Natural gradient learning uses Riemannian metric
- Further improvements possible (found useful in deep learning), to be mentioned later

16