

CS7150 Deep Learning

Jiaji Huang

<https://jiaji-huang.github.io>

02/17/2024

Recap of Last Lecture


- Encoder: understanding, e.g., BERT
- Decoder: generation, e.g., GPT
- Encoder-Decoder: seq-to-seq tasks, e.g., translation, ASR

Recap of Gradient Descent (GD)

For step $s = 0, 1, \dots$


$$\mathbf{w}_{s+1} = \mathbf{w}_s - \eta \nabla L(\mathbf{w}_s),$$

till $\nabla L(\mathbf{w}_s) \approx 0$ (stationary point reached)



Problem with GD

- Local minima and saddle point (both are stationary points)



- Convergence may be slow

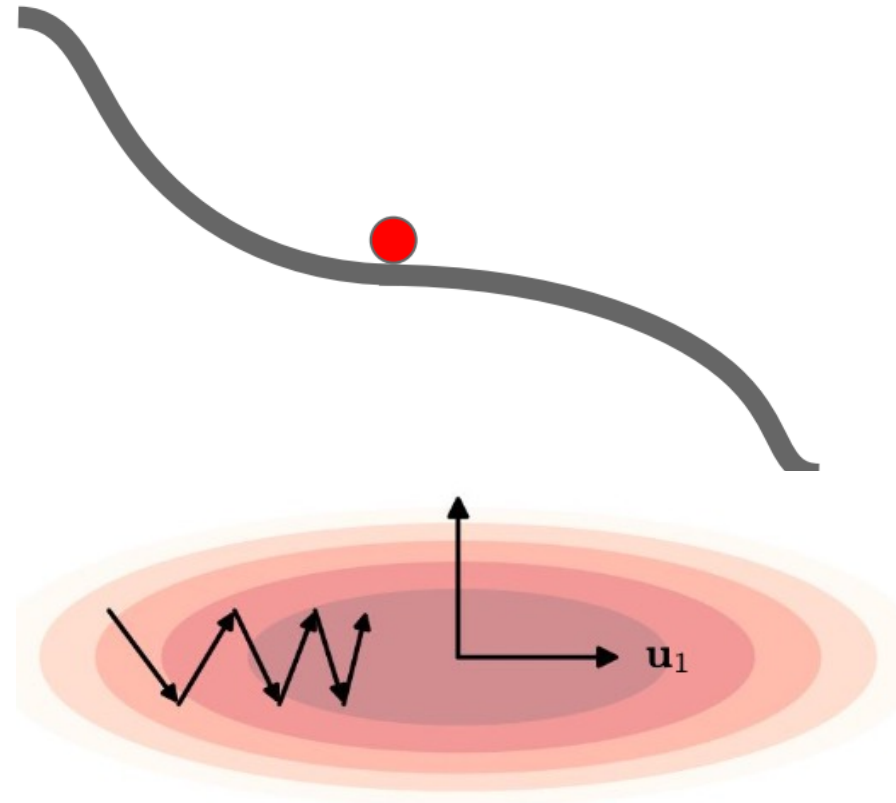



Illustration from [CS231n](#) and [Bishop book chapter 7](#)

Agenda

- Optimization
 - Initialization
 - Loss landscape and normalization
 - Momentum and higher order methods
- Generalization: old and new
- Parallelism


Initialization matters

- Different initializations land at different stationary points



A not so good idea

- Initialize all parameters as constant c ?




- $\frac{\partial \ell}{\partial v_1} = \frac{\partial \ell}{\partial v_2}, \quad \frac{\partial \ell}{\partial w_{11}} = \dots = \frac{\partial \ell}{\partial w_{32}}$

- Parameters in the same layer keep the same along gradient updates!

Break the Symmetry

- Initialize with random weights, draw from some distribution
- Control the variance



- Going deeper, activation's variance (scale) can further increase

Xavier/Glorot Initialization

- Consider MLP network with $\sigma(\cdot)$ as sigmoid or tanh
- Key idea: activation and gradient with constant variance
- Key assumption:
 - Weights and inputs centered around 0, i.i.d distributed
 - Bias initialized as all zero
- Key property:
 - $\sigma(z) \approx z$ around 0

Xavier/Glorot Initialization

- For l -th layer:

linear units $z_i^l = \sum_{j=1}^{n_{l-1}} w_{i,j} a_j^{l-1}$, activations $a_i^l = \sigma(z_i^l)$

- $Var[a_i^l] \approx Var[z_i^l]$ as $\sigma(z) \approx z$ around 0, so

$$\begin{aligned} Var[a_i^l] &= Var\left[\sum_{j=1}^{n_{l-1}} w_{i,j} a_j^{l-1}\right] = \sum_{j=1}^{n_{l-1}} Var[w_{i,j} a_j^{l-1}] \\ &= \sum_{j=1}^{n_{l-1}} Var[w_{i,j}] Var[a_j^{l-1}] = n_{l-1} Var[w_{i,j}] Var[a_j^{l-1}] \\ &\Rightarrow Var[w_{i,j}] = \frac{1}{n_{l-1}} \end{aligned}$$

Xavier/Glorot Initialization

- Consider backward

$$\frac{\partial L}{\partial a_i^{l-1}} = \sum_{j=1}^{n_l} \frac{\partial L}{\partial a_j^l} \cdot \frac{\partial a_j^l}{\partial z_j^l} \cdot \frac{\partial z_j^l}{\partial a_i^{l-1}} = \sum_{j=1}^{n_l} \frac{\partial L}{\partial a_j^l} \cdot w_{i,j}$$

$$\text{Var} \left[\frac{\partial L}{\partial a_i^{l-1}} \right] = n_l \text{Var} \left[\frac{\partial L}{\partial a_j^l} \right] \text{Var} [w_{i,j}]$$

$$\Rightarrow \text{Var} [w_{i,j}] = \frac{1}{n_l}$$

- Putting together, set

$$\text{Var} [w_{i,j}] = \frac{2}{n_{l-1} + n_l}$$

Xavier/Glorot Initialization

- If normal distribution

$$w_{i,j} \sim \mathcal{N} \left(0, \frac{2}{n_{l-1} + n_l} \right)$$

- If uniform distribution

$$w_{i,j} \sim \text{Uniform} \left(-\frac{\sqrt{6}}{\sqrt{n_{l-1} + n_l}}, \frac{\sqrt{6}}{\sqrt{n_{l-1} + n_l}} \right)$$

He initialization

- Consider $a = \text{ReLU}(z)$
- $\text{Var}[a_i^l] \approx \frac{1}{2} \text{Var}[z_i^l]$
- Follow the same derivation, we can get
 - $\text{Var}[w_{i,j}] = \frac{2}{n_{l-1}}$ for forward variance preservation
 - $\text{Var}[w_{i,j}] = \frac{2}{n_l}$ for backward variance preservation

Question

- If the l -th layer is convolution, what would be the n_{l-1} and n_l ?

(C_{out}, C_{in}, m, n) filter

$$n_l = C_{out}$$

$$n_{l-1} = C_{in} \times m \times n$$

Agenda

- Optimization
 - Initialization
 - Loss landscape and normalization
 - Momentum and higher order methods
- Generalization: old and new
- Parallelism

Loss landscape


- Consider Taylor's expansion near a stationary point \mathbf{w}^*
- $$L(\mathbf{w}) = L(\mathbf{w}^*) + \nabla L(\mathbf{w}^*)(\mathbf{w} - \mathbf{w}^*) + \frac{1}{2}(\mathbf{w} - \mathbf{w}^*)^T \mathbf{H}(\mathbf{w} - \mathbf{w}^*)$$
$$= L(\mathbf{w}^*) + \frac{1}{2}(\mathbf{w} - \mathbf{w}^*)^T \mathbf{H}(\mathbf{w} - \mathbf{w}^*)$$
- $\mathbf{H} = \left. \frac{\partial^2 L}{\partial \mathbf{w} \partial \mathbf{w}^T} \right|_{\mathbf{w}=\mathbf{w}^*}$ is Hessian
- Eigen-decomposition $\mathbf{H} = \mathbf{U}\mathbf{\Lambda}\mathbf{U}^T$, and rewrite
$$\frac{1}{2}(\mathbf{w} - \mathbf{w}^*)^T \mathbf{H}(\mathbf{w} - \mathbf{w}^*) = \frac{1}{2} \sum \lambda_i [\mathbf{u}_i^T (\mathbf{w} - \mathbf{w}^*)]^2$$

Loss landscape


- So near local minimum \mathbf{w}^* , loss is

$$L(\mathbf{w}) = L(\mathbf{w}^*) + \frac{1}{2} \sum \lambda_i [\mathbf{u}_i^T (\mathbf{w} - \mathbf{w}^*)]^2$$

- If locally convex, all $\lambda_i \geq 0$



The contour of $L(\mathbf{w}) - L(\mathbf{w}^*) = 1$
([Bishop et. al, chapter 7](#))



Gradient steps through the contours
([Bishop et. al, chapter 7](#))

Loss Landscape

Local approximation

$$L(\mathbf{w}) = L(\mathbf{w}^*) + \frac{1}{2} \sum \lambda_i \underbrace{[\mathbf{u}_i^T (\mathbf{w} - \mathbf{w}^*)]}_{\alpha_i}]^2$$

- Gradient

$$\nabla L(\mathbf{w}) = \sum \lambda_i \mathbf{u}_i^T (\mathbf{w} - \mathbf{w}^*) \mathbf{u}_i$$

- Take a gradient step

$$\mathbf{w}_{s+1} = \mathbf{w}_s - \eta \nabla L(\mathbf{w}_s)$$

- $\alpha_{i,s+1} = (1 - \eta \lambda_i) \alpha_{i,s}$

Loss Landscape

Inspect the excessive loss term $\frac{1}{2} \sum \lambda_i \alpha_i^2$

- After s -th gradient step, $\alpha_{i,s+1} = (1 - \eta \lambda_i) \alpha_{i,s}$

- To make sure excessive loss decays, we want

$$|1 - \eta \lambda_i| \leq 1 \text{ for all } i$$
$$\Rightarrow \eta \leq \frac{2}{\lambda_{max}}$$

- The loss term that decays slowest has $1 - \eta \lambda_{min} \geq 1 - 2 \cdot \frac{\lambda_{min}}{\lambda_{max}}$

Inverse Condition number

The effect of unnormalized feature

- Consider linear regression

$$\min_{\mathbf{w}} \|\mathbf{X}\mathbf{w} - \mathbf{y}\|^2$$

- Hessian is $\mathbf{X}^T \mathbf{X}$
- e.g., 2D case, x_1 : blood platelet count, $\sim 10^5$
 x_2 : body height in meters, $\sim 10^0$
- $\mathbf{X}^T \mathbf{X}$ would $\sim \begin{bmatrix} 10^{10} & 10^5 \\ 10^5 & 1 \end{bmatrix}$, $\lambda_{max} \sim 10^{10}$, $\lambda_{min} \sim 0$
- Very skewed landscape, and $\frac{\lambda_{min}}{\lambda_{max}} \sim 0$, extremely slow convergence

Normalization

- Mean and standard deviation

$$\mu = \frac{1}{N} \sum_{n=1}^N \mathbf{x}_n,$$

$$\sigma_i = \sqrt{\frac{1}{N} \sum_{n=1}^N (x_{n,i} - \mu_i)^2}$$

- Normalization:

$$\overline{x_{n,i}} = \frac{x_{n,i} - \mu_i}{\sigma_i}$$

Illustration of the effect of input data normalization. The red circles show the original data points for a data set with two variables. The blue crosses show the data set after normalization such that each variable now has zero mean and unit variance across the data set.


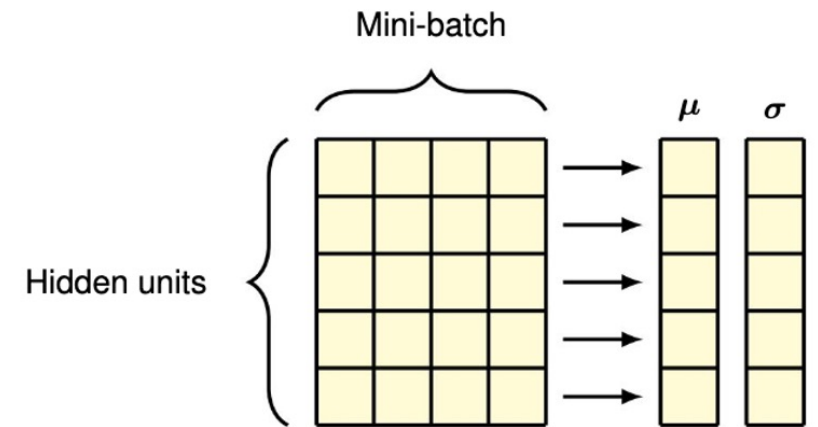


Illustration from [Bishop book Chapter 7](#)

Batch Normalization

- For the activations \mathbf{x} at a layer
- Training phase:
 - Estimate μ and σ_i 's for a minibatch, update moving average estimate
 - Normalize \mathbf{x} to $\bar{\mathbf{x}}$
 - Update to $\tilde{\mathbf{x}}_i = \gamma_i \bar{\mathbf{x}}_i + \beta_i$
- Inference phase
 - Use the moving average estimated μ and σ_i 's



[Ioffe et. al, 2015](#)

Illustration: [Bishop book Chapter 7](#)

Discussion

- Why we set the learnable multipliers γ_i and β_i ?
- Drawback of batch normalization?
 - Batch size = 1?
 - Determinism

Layer Norm

- Do this for each sample \mathbf{x} in the minibatch
- Compute μ by averaging all dimensions of \mathbf{x}
- Compute σ as the norm of $\mathbf{x} - \mu$
- Normalize by $\tilde{\mathbf{x}} = (\mathbf{x} - \mu) / \sigma$
- Compute $\tilde{\mathbf{x}}_i = \gamma_i \bar{\mathbf{x}}_i + \beta_i$
- No difference between training and testing phase




Illustration: [Bishop book Chapter 7](#)

Agenda

- Optimization
 - Initialization
 - Loss landscape and normalization
 - Momentum and higher order methods
- Generalization: old and new
- Parallelism

Recap of Stochastic Gradient Descent (SGD)

- $\min_{\mathbf{w}} \{L = \frac{1}{N} \sum_{n=1}^N \ell(\mathbf{w}, \mathbf{x}_n)\}$
- If N is big, Compute $\nabla \ell(\mathbf{w}, \mathbf{x}_n)$ for all n 's can be costly
- So we **sample a mini-batch \mathcal{B}** of n 's each step


For step $s = 0, 1, \dots$

$$\mathbf{w}_{s+1} = \mathbf{w}_s - \eta \frac{1}{|\mathcal{B}|} \sum_{n \in \mathcal{B}} \nabla \ell(\mathbf{w}_s, \mathbf{x}_n)$$

till stopping criteria met

Problems with SGD

- Same as GD
 - big condition number hamper convergence
 - Local minima and saddle point



- In addition: noisy steps






Figure from [CS231n](#)

Momentum

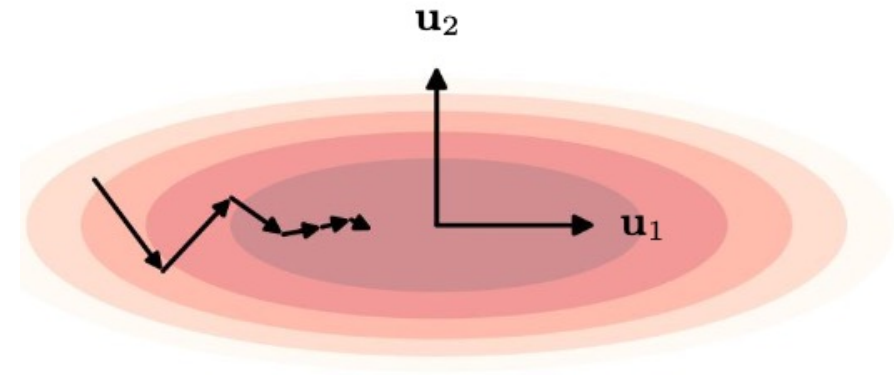
- Smooth the steps by moving average of historical gradients

$$V_s = \alpha V_{s-1} - \eta \nabla L(\mathbf{w}_s)$$
$$\mathbf{w}_s = \mathbf{w}_{s-1} + V_s$$

- α : momentum parameter



Gradient steps without momentum



Gradient steps with momentum

Skip local minima

- The inertia helps




Illustration from [paperwithcode](#)

In pyTorch

```
>>> optimizer = torch.optim.SGD(model.parameters(), lr=0.1, momentum=0.9)
>>> optimizer.zero_grad()
>>> loss_fn(model(input), target).backward()
>>> optimizer.step()
```


- Usually use momentum > 0.9

Agenda

- Optimization
 - Initialization
 - Loss landscape and normalization
 - Momentum
 - Higher order methods
- Generalization: old and new
- Parallelism

2nd order Optimization

- Approximate with quadratic function, and move to its minimum



Newton's method

- $L(\mathbf{w}) \approx L(\mathbf{w}_s) + (\mathbf{w} - \mathbf{w}_s)^T \nabla L(\mathbf{w}_s) + (\mathbf{w} - \mathbf{w}_s)^T \mathbf{H}(\mathbf{w} - \mathbf{w}_s)$

- RHS has minimum at

$$\mathbf{w}_{s+1} = \mathbf{w}_s - \mathbf{H}^{-1} \nabla L(\mathbf{w}_s)$$

- No learning rate required, faster convergence than GD


- Damped Newton's method:

$$\mathbf{w}_{s+1} = \mathbf{w}_s - \eta \cdot \mathbf{H}^{-1} \nabla L(\mathbf{w}_s)$$

Where $0 < \eta < 1$

- Inverse Hessian is costly, $O(M^3)$, where M is number of model parameters

Image source: notes [here](#)



Black: Gradient Descent
Blue: Newton's method

Stochastic Version

- Approximate the $\mathbf{H}^{-1} \nabla L(\mathbf{w}_s)$, where $L(\mathbf{w}) = \frac{1}{N} \sum_{i=1}^N \ell(\mathbf{w}, \mathbf{x}_i)$
- Sample a batch \mathcal{B}

$$\nabla \widehat{L}(\mathbf{w}_s) = \frac{1}{|\mathcal{B}|} \sum_{n \in \mathcal{B}} \nabla \ell(\mathbf{w}_s, \mathbf{x}_n)$$

$$\widehat{\mathbf{H}} = \frac{1}{|\mathcal{B}|} \sum_{n \in \mathcal{B}} \nabla^2 \ell(\mathbf{w}_s, \mathbf{x}_n)$$

- Still not practical for deep learning, because of $\widehat{\mathbf{H}}^{-1}$

Scale the Gradient: Adagrad

- Normalize each dimension of the gradient by accumulated magnitude
 $G = 0$

For $s=1,2,\dots$

sample batch \mathcal{B}

compute stochastic gradient $G = \frac{1}{|\mathcal{B}|} \sum_{n \in \mathcal{B}} \nabla \ell(\mathbf{w}_s, \mathbf{x}_n)$

update $G += g * g$

$$\mathbf{w}_{s+1} = \mathbf{w}_s - \eta \cdot \frac{g}{\sqrt{G}}$$

Scale the Gradient: RMSProp

$$G = 0$$

For $s=1,2,\dots$

sample batch \mathcal{B}

compute stochastic gradient $g = \frac{1}{|\mathcal{B}|} \sum_{n \in \mathcal{B}} \nabla \ell(\mathbf{w}_s, \mathbf{x}_n)$

update $G = \gamma G + (1 - \gamma) g * g$

$$\mathbf{w}_{s+1} = \mathbf{w}_s - \eta \cdot \frac{g}{\sqrt{G}}$$

Scale the Gradient: Adam

$$G_1 = 0, G_2 = 0$$

For $s=1,2,\dots$

sample batch \mathcal{B}

compute stochastic gradient $g = \frac{1}{|\mathcal{B}|} \sum_{n \in \mathcal{B}} \nabla \ell(\mathbf{w}_s, \mathbf{x}_n)$

$$\text{update } G_1 = \beta_1 G_1 + (1 - \beta_1)g$$

$$\text{update } G_2 = \beta_2 G_2 + (1 - \beta_2)g * g$$

bias correction: $G_1 = G_1 / (1 - \beta_1^t)$, $G_2 = G_2 / (1 - \beta_2^t)$

$$\mathbf{w}_{s+1} = \mathbf{w}_s - \eta \cdot \frac{G_1}{\sqrt{G_2}}$$

In pytorch

ADAGRAD

```
CLASS torch.optim.Adagrad(params, lr=0.01, lr_decay=0, weight_decay=0,  
initial_accumulator_value=0, eps=1e-10, foreach=None, *, maximize=False,  
differentiable=False) \[SOURCE\]
```

Implements Adagrad algorithm.

RMSPROP

```
CLASS torch.optim.RMSprop(params, lr=0.01, alpha=0.99, eps=1e-08, weight_decay=0, momentum=0,  
centered=False, foreach=None, maximize=False, differentiable=False) \[SOURCE\]
```

Implements RMSprop algorithm.



ADAM

```
CLASS torch.optim.Adam(params, lr=0.001, betas=(0.9, 0.999), eps=1e-08, weight_decay=0,  
amsgrad=False, *, foreach=None, maximize=False, capturable=False, differentiable=False,  
fused=None) \[SOURCE\]
```

Implements Adam algorithm.

Learning Rate Scheduler

- Infeasible to compute ideal learning rate
- Often decay the learning rate along training
- *Warmup* is sometimes used to skip local minima



cosine learning rate scheduler. Left: no warmup; right: with warmup

Check pytorch [page](#), figures from this [post](#)

Discussion


- How to choose batch size in practice?
- How do we scale learning rate w.r.t. batch size?

Agenda

- Optimization
 - Initialization
 - Loss landscape and normalization
 - Momentum and higher order methods
- Generalization: old and new
- Parallelism

Recap: Generalization

- The ability to predict well on unseen samples
- Model family $\mathcal{F} = \{f_{\theta}(\mathbf{x})\}$
- Generalization error: expected error/loss on a test input
- \approx modeling error + estimation error



Reduce the Variance

- Inject **inductive bias** (encodes our prior knowledge)
- e.g., small variation doesn't change object category

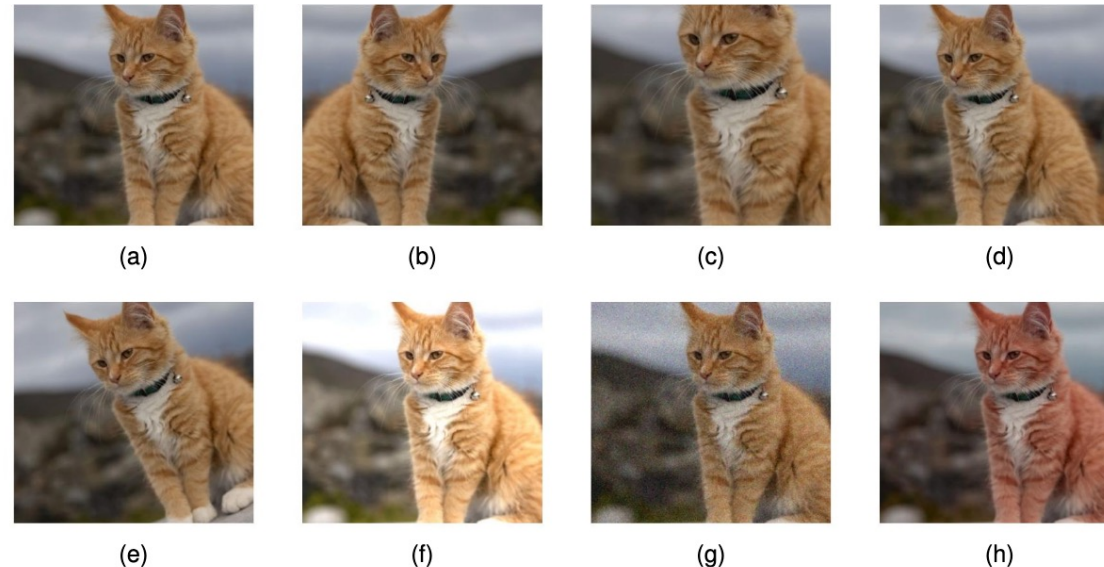



Figure 9.1 Illustration of data set augmentation, showing (a) the original image, (b) horizontal inversion, (c) scaling, (d) translation, (e) rotation, (f) brightness and contrast change, (g) additive noise, and (h) colour shift.

Regularization

- Constrain the model family \mathcal{F}
- e.g., weight decay $\min_{\mathbf{w}} L(\mathbf{w}) + \frac{\lambda}{2} \|\mathbf{w}\|^2$
- e.g., sparsity $\min_{\mathbf{w}} L(\mathbf{w}) + \frac{\lambda}{2} \|\mathbf{w}\|_1$

Question: what are the gradients for these two cases?


Weight Decay: Loss Landscape View




- Suppressing the magnitude of weights
- Especially those to which loss is less sensitive

Illustration from [Bishop book, chapter 9](#)

Earlier Stopping



[Bishop book Chapter 9](#)



From the perspective of weight decay

Drop out


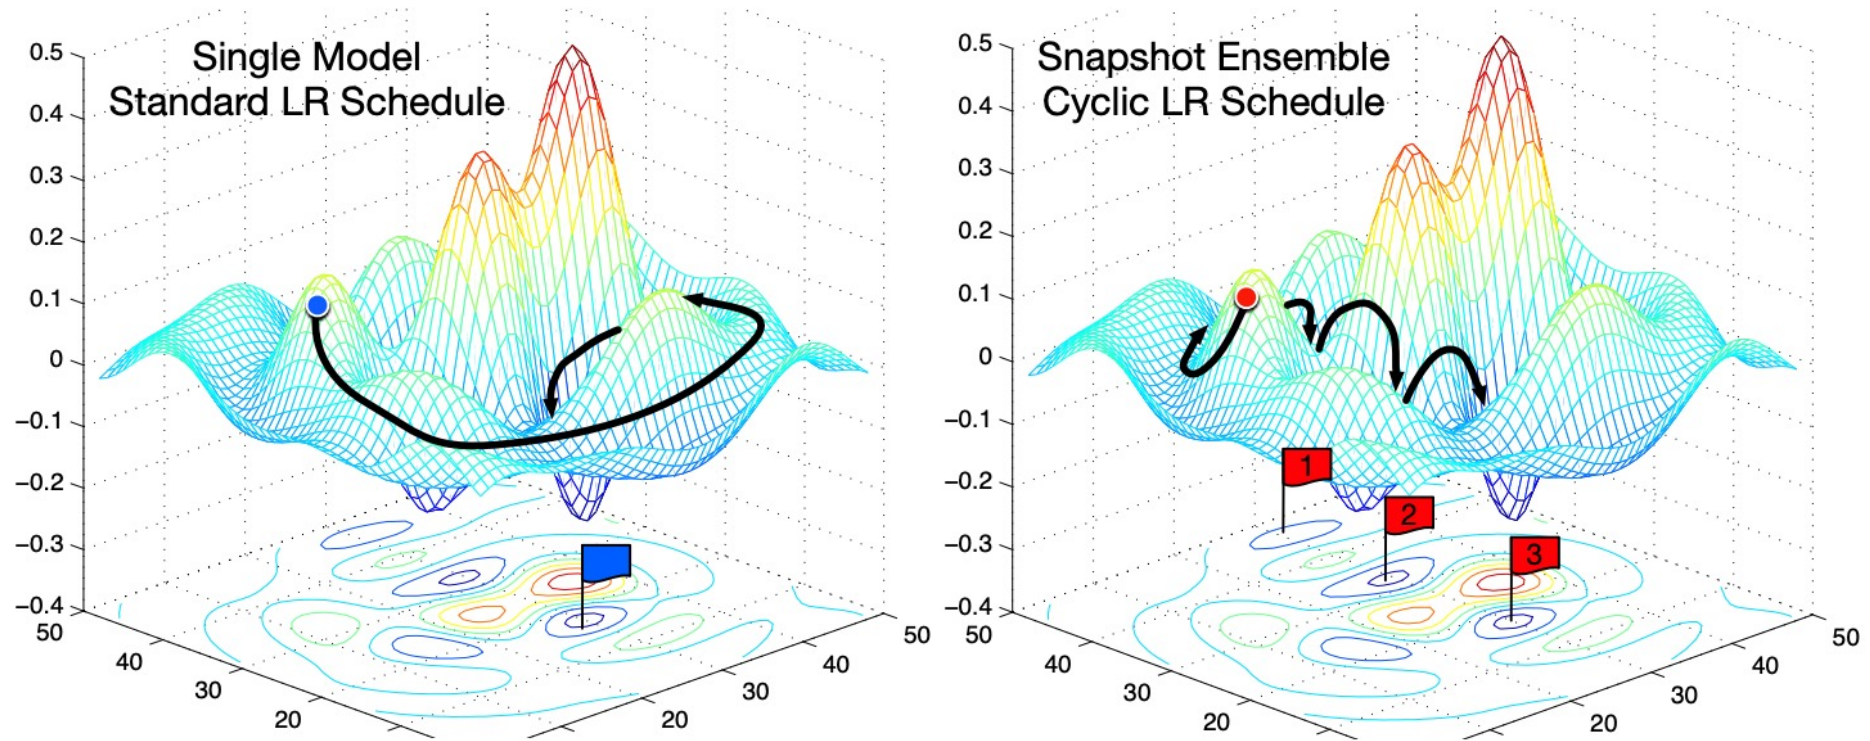


Figure 9.17 A neural network on the left along with two examples of pruned networks in which a random subset of nodes have been omitted.

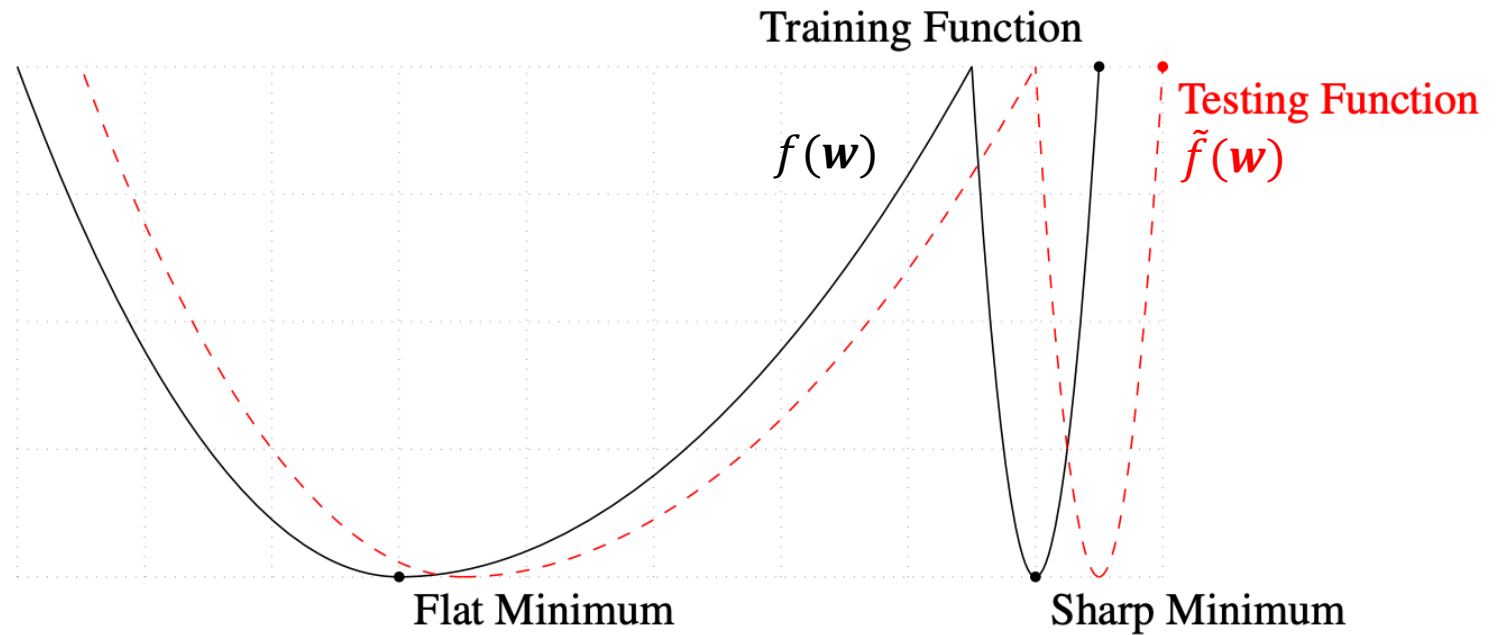
Model Averaging



[Huang et. al, 2017](#)

Sharpness

- Sharp local minima generalizes poorly
- Question: How to measure sharpness?



Agenda

- Optimization
 - Initialization
 - Loss landscape and normalization
 - Momentum and higher order methods
- Generalization: old and new
- Parallelism

Data Parallelism

- Want to use large batch size, but single GPU doesn't have enough memory

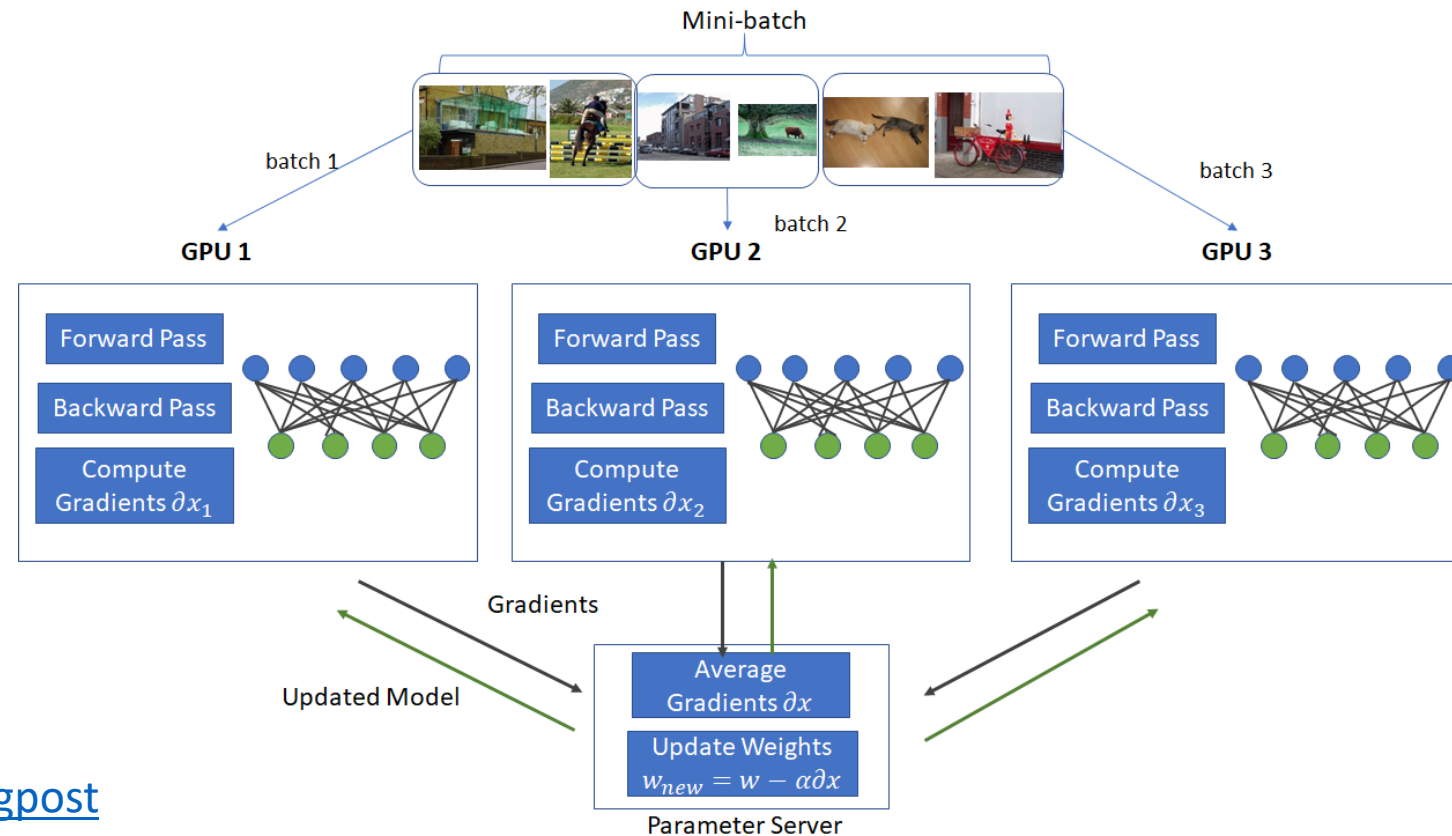
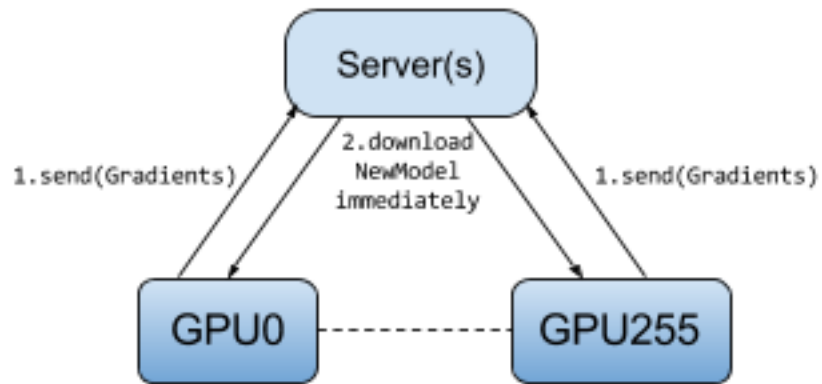


Illustration from this [blogpost](#)

SGD with multiple workers

Asynchronous SGD



Synchronous SGD





Illustration from this [blogpost](#)

Model Parallelism

- Pipeline parallel



- Tensor parallel

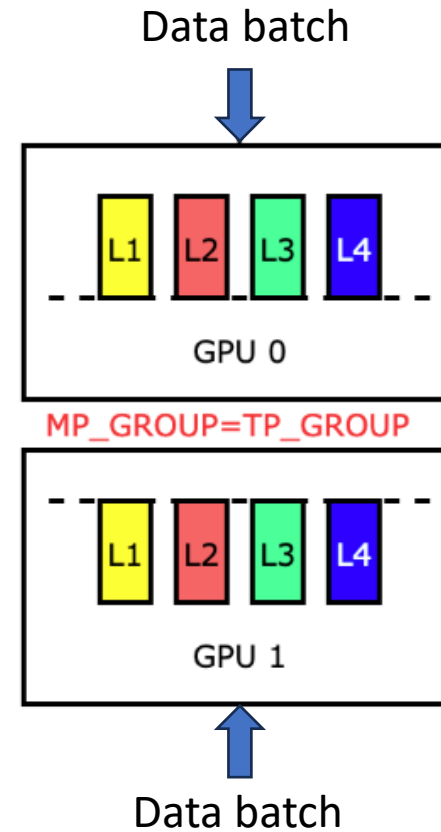


Illustration adapted from [aws sagemaker](https://aws.amazon.com/sagemaker/)