

## Section 9. Gradient Descent

1. Gradient Decent
2. Stochastic Gradient Decent
3. Newton's Method
4. More descent methods

Instructor: He Wang  
Department of Mathematics  
Northeastern University

## ➤ Taylor Expansion

- Taylor Expansion of  $f: \mathbb{R} \rightarrow \mathbb{R}$

$$f(a + s) = f(a) + sf'(a) + \frac{1}{2!} s^2 f''(a) + \frac{1}{3!} s^3 f'''(a) + \dots$$

- Taylor Expansion of  $f: \mathbb{R}^d \rightarrow \mathbb{R}$

$$f(\vec{a} + \vec{s}) = f(\vec{a}) + \vec{s}^T \nabla f(\vec{a}) + \frac{1}{2!} \vec{s}^T H(f(\vec{a})) \vec{s} + \dots$$

$$= f(\vec{a}) + \sum s_i \frac{\partial f}{\partial x_i} + \sum \frac{\partial^2 f}{\partial x_i \partial x_j} s_i s_j + \dots$$

- Taylor Expansion of  $F: \mathbb{R}^d \rightarrow \mathbb{R}^m$

$$F(\vec{a} + \vec{s}) = F(\vec{a}) + \left( \frac{\partial F(\vec{a})}{\partial \vec{x}} \right)^T \vec{s} + \frac{1}{2!} \begin{bmatrix} \vec{s}^T H(F_1(\vec{a})) \vec{s} \\ \vdots \\ \vec{s}^T H(F_m(\vec{a})) \vec{s} \end{bmatrix} + \dots$$

## ➤ Gradient Descent

**Goal:** find the local/global minimum of the cost function  $J(\vec{\theta})$ .

Examples:

$$J(\vec{\theta}) = RSS(\vec{\theta}) = \|X\vec{\theta} - \vec{y}\|^2$$
$$J^{Ridge}(\vec{\theta}) = RSS(\vec{\theta}) + \lambda \|\vec{\theta}\|^2$$
$$\bullet J^{Lasso}(\vec{\theta}) = RSS(\vec{\theta}) + \lambda \|\vec{\theta}\|_1^2$$

$\vec{\theta} = (X^T X)^{-1} X^T \vec{y}$  (SciKit-Statmodel)

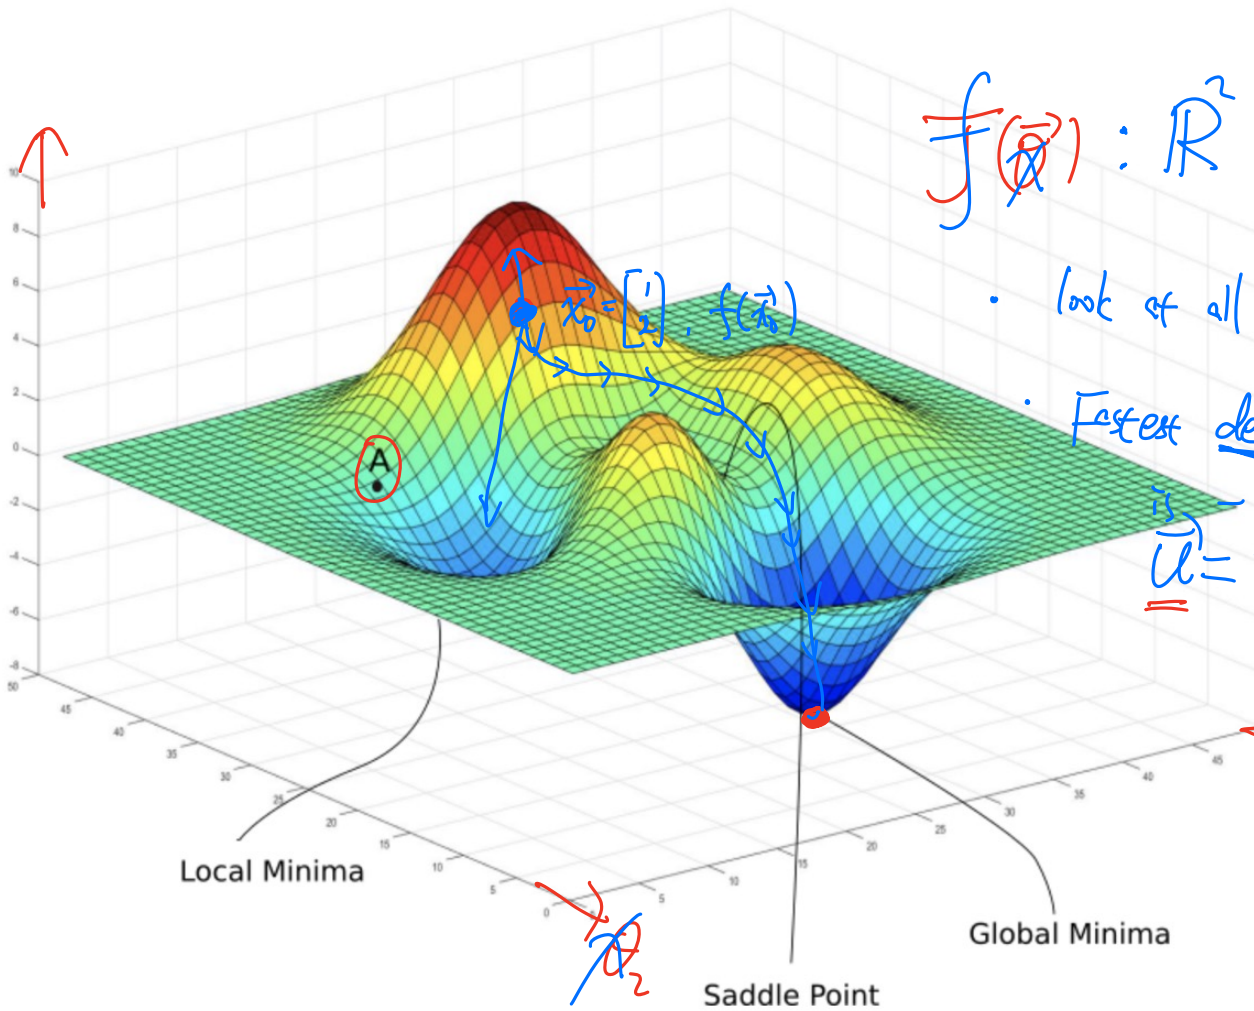
$\vec{\theta} = (X^T X + \lambda I)^{-1} X^T \vec{y}$

**Method:** find critical points by solving

$$\nabla J(\vec{\theta}) = 0$$

**Difficulty:**

1. No closed formula or too complicated to find a closed formula for the minimum.
2. Too complicated to compute even we have a formula, as the inverse.



$$f(\vec{x}) : \mathbb{R}^2 \rightarrow \mathbb{R}$$

- look at all directions  $\vec{u} \in \mathbb{R}^2$
- Fastest decrease direction

$$\vec{u} = -\nabla f(\vec{x}_i)$$

Local Minima

Global Minima

Saddle Point

$$\vec{J}(\vec{\theta})$$

$$\vec{u}_2$$


$$\vec{u}_1$$

Suppose  $f(\vec{x})$  is a differentiable function  $\mathbb{R}^d \rightarrow \mathbb{R}$ .

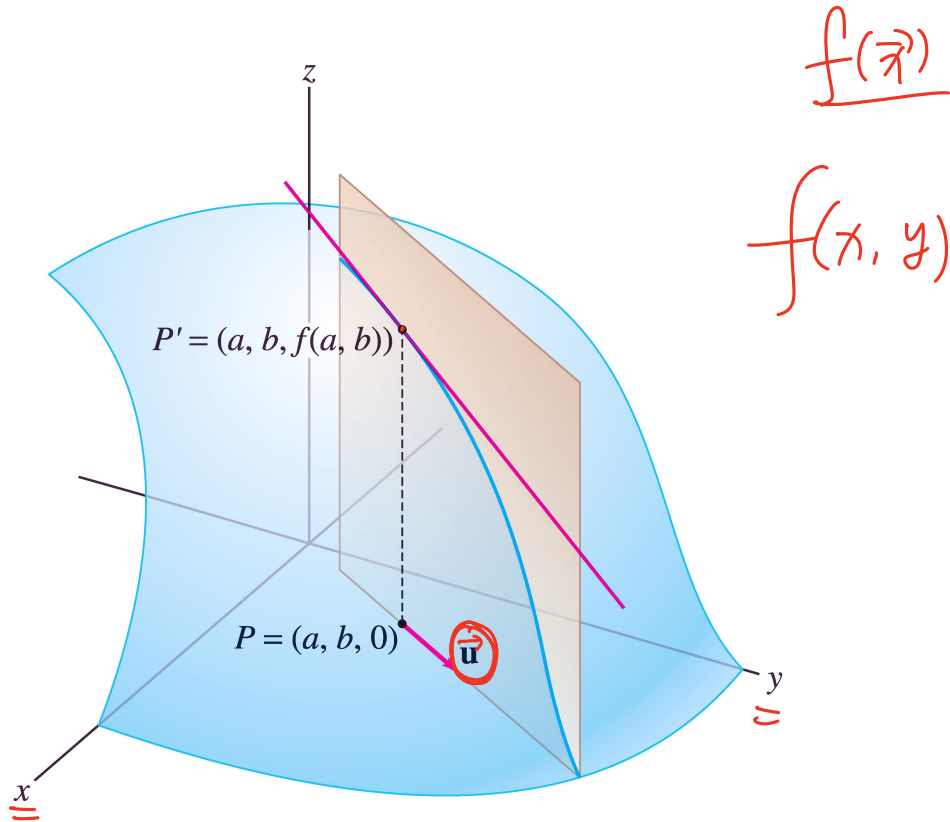
**Question:** Which **direction** has the largest rate of change?

$f'(x) > 0$  left -  
 $f'(x) < 0$  right +

$d = 1$



# Directional derivative:



**Definition:** Let  $\vec{u}$  be a unit vector in  $\mathbb{R}^d$ . The directional derivative of  $f(\vec{x})$  at point  $\vec{a} \in \mathbb{R}^d$  in direction  $\vec{u}$  is

$$D_{\vec{u}}f(\vec{x}) = \lim_{t \rightarrow 0} \frac{f(\vec{a} + t\vec{u}) - f(\vec{a})}{t}$$

This is just using the Chain Rule on the composition of  $f(\vec{x})$  and the path

$$\frac{d f(\vec{a} + t\vec{u})}{dt}$$

$$\vec{x}(t) = \vec{a} + t\vec{u}$$

**Theorem:** The directional derivative of  $f(\vec{x})$  in direction  $\vec{u}$  is computed by

$$D_{\vec{u}}f(\vec{x}) = \nabla f \cdot \vec{u}$$

$$\vec{u} = -\nabla f$$

$$\begin{pmatrix} \frac{\partial f}{\partial x_1} \\ \vdots \\ \frac{\partial f}{\partial x_n} \end{pmatrix} = \nabla f$$

$\alpha = \tau$


the

$u$


**Theorem:** The **maximum** value of the directional derivative  $D_{\vec{u}}f(\vec{x})$  is  $\|\nabla f(\vec{x})\|$  and it occurs when  $\vec{u}$  has the same direction as the gradient vector  $\nabla f(\vec{x})$ .

$$D_{\vec{u}}f(\vec{x}) = \nabla f \cdot \vec{u} = \|\nabla f(\vec{x})\| \|\vec{u}\| \cos \alpha = \|\nabla f(\vec{x})\| \cos \alpha$$

$$D_{\vec{u}}f(\vec{x}) = \begin{cases} \|\nabla f(\vec{x})\| & \text{when } \alpha = 0 \\ -\|\nabla f(\vec{x})\| & \text{when } \alpha = \pi \end{cases}$$



The **absolute minimum** value of the directional derivative  $D_{\vec{u}}f(\vec{x})$  occurs when  $\vec{u}$  has the same direction  $-\nabla f(\vec{x})$ .



$$f(x) = 3x^2 + 1$$


$$f'(x) = 6x \quad f'(-10) = -60$$

Fastest decreasing is  $-(60)$

$$x_0 = -10 \quad x_1 = x_0 + \alpha(-f'(x_0)) = x_0 - \alpha f'(x_0)$$



Example:  $f(\theta) = \theta^2$

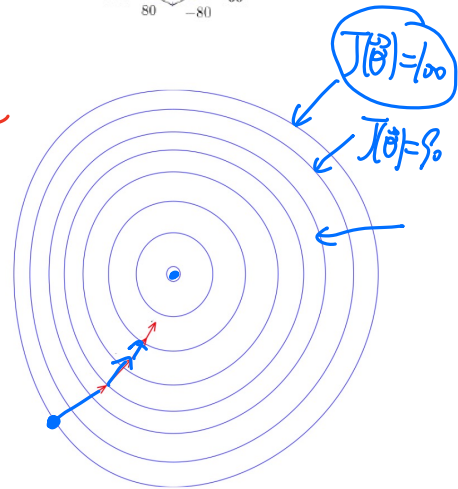




learning rate.

Example:  $f(\vec{\theta}) = \theta_1^2 + \theta_2^2$

$$\nabla f = \begin{bmatrix} \frac{\partial f}{\partial \theta_1} \\ \frac{\partial f}{\partial \theta_2} \end{bmatrix} = \begin{bmatrix} 2\theta_1 \\ 2\theta_2 \end{bmatrix}$$

$$\vec{\theta}_1 = \vec{\theta}_0 - \alpha \nabla f(\vec{\theta}_0) = \begin{bmatrix} 1 \\ 1 \end{bmatrix} - \alpha \begin{bmatrix} 2 \\ 2 \end{bmatrix}$$




➤ Gradient Descent:

**Goal:** find the local/global minimum of the cost function  $J(\vec{\theta})$ .

Gradient Descent Algorithm:




- Start with  $\vec{\theta}^{\text{ini}}$  = some initial value. =  $\vec{\theta}_0$
- Repeat  $\vec{\theta}^{\text{next}} = \vec{\theta}^{\text{current}} - \alpha \nabla J(\vec{\theta}^{\text{current}})$  until converge.  
    *learning rate.*

$$\begin{bmatrix} \theta_0 \\ \theta_1 \\ \vdots \\ \theta_d \end{bmatrix}^{\text{next}} = \begin{bmatrix} \theta_0 \\ \theta_1 \\ \vdots \\ \theta_d \end{bmatrix} - \alpha \begin{bmatrix} \frac{\partial J(\vec{\theta})}{\partial \theta_0} \\ \vdots \\ \frac{\partial J(\vec{\theta})}{\partial \theta_d} \end{bmatrix}$$




Key points:

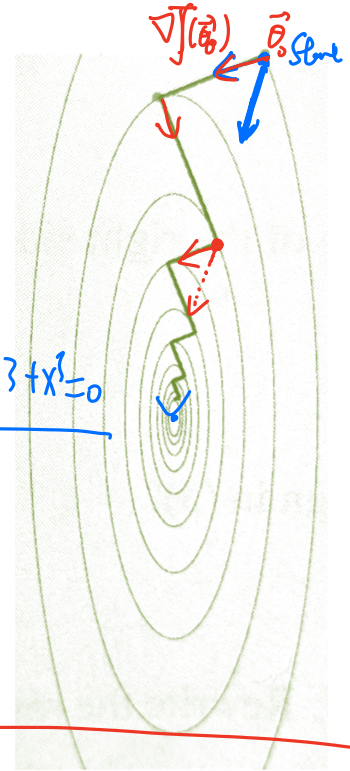
- Compute  $\nabla J(\vec{\theta})$
- Set initial value  $\vec{\theta} = \vec{\theta}_0$
- Set a good learning rate  $\alpha$ 
  - Set different  $\alpha$  and recording the cost
  - Start from large  $\alpha_0$ , then smaller  $\alpha$ .
  - Set  $\alpha_k = \frac{1}{\sqrt{k}} \alpha_0$  or  $\alpha_k = \frac{1}{k} \alpha_0$
  - ...



$J(\theta)$



$$z^T + e^{x^2} + \sin x + \log x + 3 + x^T = 0$$



• Data  $(\vec{x}^{(i)}, y^{(i)}) \sim (X, \vec{y})$

• Model  $h_{\theta}(\vec{x})$

• Cost function  $\begin{cases} \rightarrow \frac{\|h_{\theta}(X) - \vec{y}\|}{\text{cross-entropy}} \end{cases} = J(\vec{\theta})$

$$\cdot \boxed{\text{Solve } \nabla J(\vec{\theta}) = \vec{0}}$$

• Minimize  $J(\theta)$

Minimiere  $J(\vec{\theta})$ .  $\Leftrightarrow$  Find.  $\vec{\theta} = \arg \min J(\vec{\theta})$   $\leftarrow$  Gradient descent.

➤ Example: (linear regression)  $h(\vec{x}) = \vec{\theta}^T \vec{x} = \theta_0 + \theta_1 x_1 + \dots + \theta_d x_d$

$$J(\vec{\theta}) = \left( \frac{1}{n} \sum_{i=1}^n (h(x^{(i)}) - y^{(i)})^2 \right) = \frac{1}{n} \text{RSS}(\vec{\theta}) = \frac{1}{n} \|h(X) - \vec{y}\|^2$$

For each  $j = 0, 1, \dots, d$


$$= \frac{1}{n} (X\vec{\theta} - \vec{y})^T (X\vec{\theta} - \vec{y})$$

~~$$\begin{aligned} \frac{\partial}{\partial \theta_j} J(\vec{\theta}) &= \frac{\partial}{\partial \theta_j} \left( \frac{1}{n} \sum_{i=1}^n (h(x^{(i)}) - y^{(i)})^2 \right) \\ &= \frac{1}{n} \left( \sum_{i=1}^n \frac{\partial}{\partial \theta_j} (h(x^{(i)}) - y^{(i)})^2 \right) \\ &= \frac{1}{n} \cdot \sum_{i=1}^n \left( 2(h(x^{(i)}) - y^{(i)}) \cdot \frac{\partial}{\partial \theta_j} (h(x^{(i)}) - y^{(i)}) \right) \\ &= \frac{2}{n} \sum_{i=1}^n \left( (h(x^{(i)}) - y^{(i)}) \cdot \frac{\partial}{\partial \theta_j} (\theta_0 + \theta_1 x_1^{(i)} + \theta_2 x_2^{(i)} + \dots + \theta_d x_d^{(i)} - y^{(i)}) \right) \\ &= \frac{2}{n} \sum_{i=1}^n (h(x^{(i)}) - y^{(i)}) \cdot x_j^{(i)} \end{aligned}$$~~

Repeat until converge

$$\theta_j := \theta_j - \alpha \cdot \left( \frac{2}{n} \sum_{i=1}^n (h(x^{(i)}) - y^{(i)}) \cdot x_j^{(i)} \right)$$

★ ➤ Example: (linear regression, vector notation)



$$h(\vec{x}) = \vec{\theta}^T \vec{x} = \theta_0 + \theta_1 x_1 + \dots + \theta_d x_d$$

$$J(\vec{\theta}) = \frac{1}{n} \text{RSS}(\vec{\theta}) := \frac{1}{n} \|\mathbf{X}\vec{\theta} - \vec{y}\|^2 = \frac{1}{n} (\vec{\theta}^T \mathbf{X}^T \mathbf{X} \vec{\theta} - 2\vec{y}^T \mathbf{X} \vec{\theta} + \vec{y}^T \vec{y})$$

★  $\nabla_{\vec{\theta}} J = \frac{2}{n} (\mathbf{X}^T \mathbf{X} \vec{\theta} - \mathbf{X}^T \vec{y})$  *der.*


Gradient descent formula:  $\vec{\theta}^{\text{next}} = \vec{\theta} - \alpha \frac{2}{n} \mathbf{X}^T (\mathbf{X} \vec{\theta} - \vec{y})$

$\vec{\theta} = (\mathbf{X}^T \mathbf{X})^{-1} \mathbf{X}^T \vec{y}$

Python (broadcast):  $\vec{\theta}^{\text{next}} = \vec{\theta} - \alpha \frac{2}{n} \text{sum}[(\mathbf{X} \vec{\theta} - \vec{y}) * \mathbf{X}]$  ----->

Golden Rule: If you can use vector, never use a for loop.


We ran the update rule for all the training examples  $(X, \vec{y})$  at once, which is called **batch gradient descent**.



$$f(x_1) = \theta_0 + \theta_1 x_1$$

Find a good learning rate:

For different learning rate  
 Use a small data set  
 Repeat 100 times



$$\vec{\theta}^{\text{next}} = \vec{\theta} - \alpha \nabla J$$

Stochastic Gradient Descent (SGD):

For each step, we use only one data point  $(\vec{x}^{(i)}, y^{(i)})$  to find descent direction.

$$\vec{\theta}^{\text{next}} = \vec{\theta} - \alpha \nabla J(\vec{\theta}; \vec{x}^{(i)}, y^{(i)})$$

For example, in linear regression,


$$\vec{\theta}^{\text{next}} = \vec{\theta} - \alpha \vec{x}^{(i)} (\vec{x}^{(i)T} \vec{\theta} - y^{(i)})$$

Remark:


1. Randomly with replacement, or use a random order on the data.
2. It is fast.
3. It may achieve global minimum.
4. We call an epoch for repeating a data set

$$\nabla J$$

$$\nabla J = \frac{2}{n} (\underline{X}^T X - X^T \underline{y})$$



# iterations






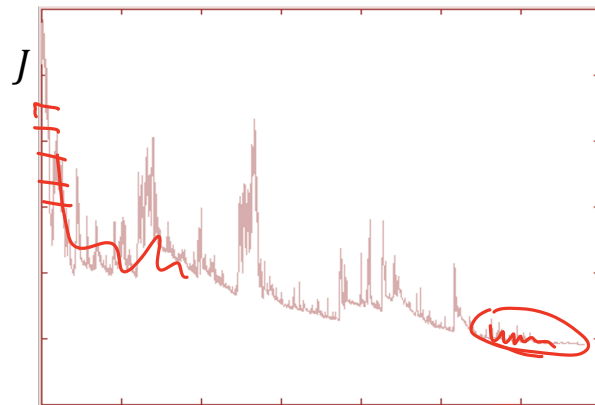
➤ Mini-batch Gradient Descent:

For each step, we use only a subset of data points  $D_j \subset D$  to find descent direction  $\nabla J(\vec{\theta}; D_j)$ .

- $\vec{\theta}^{\text{next}} = \vec{\theta} - \alpha \nabla J(\vec{\theta}; D_j)$

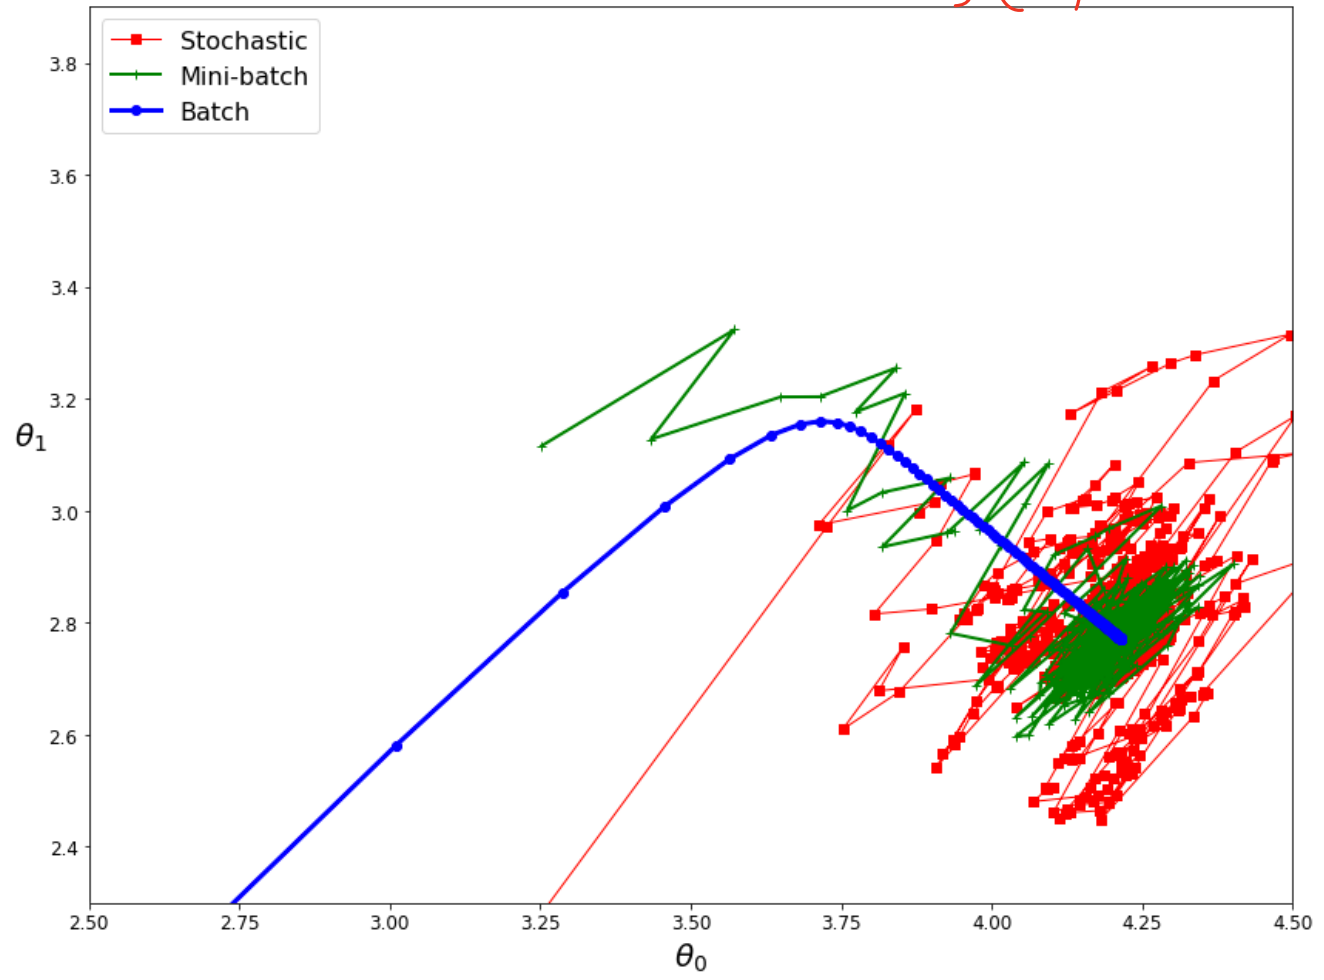


If each minibatch  $D_j$  contains one point, it is Stochastic Gradient Descent.  
If each minibatch  $D_j$  contains all points, it is batch Gradient Descent.



# iterations  $\Rightarrow 1000$


$$J(\vec{\theta}) = c$$



Remarks:

1. Normal equation
2. Stochastic gradient descent
3. Batch gradient descent
4. Mini batch gradient descent

Scale the features first: normalization or standardization



• Our problem: Solve  $\nabla J(\vec{\theta}) = 0$ .

## ➤ Newton's method

Find **root** of a function  $f: \mathbb{R} \rightarrow \mathbb{R}$ .

Solve  $f(x) = 0$

### Newton's method Algorithm

1. Make a guess  $x_0$
2. Repeat



$$x_{k+1} = x_k - \frac{f(x_k)}{f'(x_k)}$$

Reason:

$$x_2 = x_1 - \frac{f(x_1)}{f'(x_1)}$$

$$f(x_1 + s) \approx f(x_1) + sf'(x_1) = 0$$

$$s = -\frac{f(x_1)}{f'(x_1)}$$



High dimension Newton's method for  $F: \mathbb{R}^m \rightarrow \mathbb{R}^m$

Repeat  $\vec{x}_{k+1} = \vec{x}_k - B^{-1}F(\vec{x}_k)$

where,  $B = \left( \frac{\partial F(\vec{x}_k)}{\partial \vec{x}} \right)^T = \begin{bmatrix} \frac{\partial f_1}{\partial x_1} & \dots & \frac{\partial f_1}{\partial x_m} \\ \vdots & \ddots & \vdots \\ \frac{\partial f_m}{\partial x_1} & \dots & \frac{\partial f_m}{\partial x_m} \end{bmatrix}$

$\nabla J = 0$   
 $F = \begin{bmatrix} f_1 \\ \vdots \\ f_m \end{bmatrix}$

Application of Newton's method to

**Goal:** find the local/global minimum of the cost function  $J(\vec{\theta})$ .

Find  $\nabla J(\vec{\theta}) = 0$

Let  $F(\vec{\theta}) = \nabla J(\vec{\theta}) = \begin{bmatrix} \frac{\partial J(\vec{\theta})}{\partial \theta_0} \\ \vdots \\ \frac{\partial J(\vec{\theta})}{\partial \theta_d} \end{bmatrix}$  and apply Newton's method.

$$\vec{\theta}^{\text{next}} = \vec{\theta} - \underline{B}^{-1} F$$

$$\vec{\theta}_{k+1} = \vec{\theta}_k - H^{-1} \nabla J(\vec{\theta}_k)$$

→ arg min  $J(\vec{\theta})$

Here  $H$  is the Hessian matrix  $H = \begin{bmatrix} \frac{\partial^2 J}{\partial \theta_1^2} & \dots & \frac{\partial^2 J}{\partial \theta_1 \partial \theta_d} \\ \vdots & \ddots & \vdots \\ \frac{\partial^2 J}{\partial \theta_d \partial \theta_1} & \dots & \frac{\partial^2 J}{\partial \theta_d^2} \end{bmatrix}$

G.D.  $\vec{\theta}_{k+1} = \vec{\theta}_k - \alpha \nabla J(\vec{\theta}_k)$

Example. Linear Regression.

$$J(\vec{\theta}) = (X\vec{\theta} - \vec{y})^T (X\vec{\theta} - \vec{y})$$

$$\nabla J = 2X^T X \vec{\theta} - 2X^T \vec{y}$$

$$H J(\vec{\theta}) = 2X^T X$$

$X_{n \times d}$   
 $\text{rank}(X) = d$

$$\begin{aligned} \vec{\theta}^{\text{next}} &= \vec{\theta}^{\text{current}} - H^{-1} \nabla J = \vec{\theta} - (2X^T X)^{-1} (2X^T X \vec{\theta} - 2X^T \vec{y}) \\ &= \vec{\theta} - \vec{\theta} + (X^T X)^{-1} X^T \vec{y} \\ &= (X^T X)^{-1} X^T \vec{y} \end{aligned}$$

Remark: Newton's method is faster, since it depends on the second derivative. However, sometimes it is hard to calculate or it is not invertible.



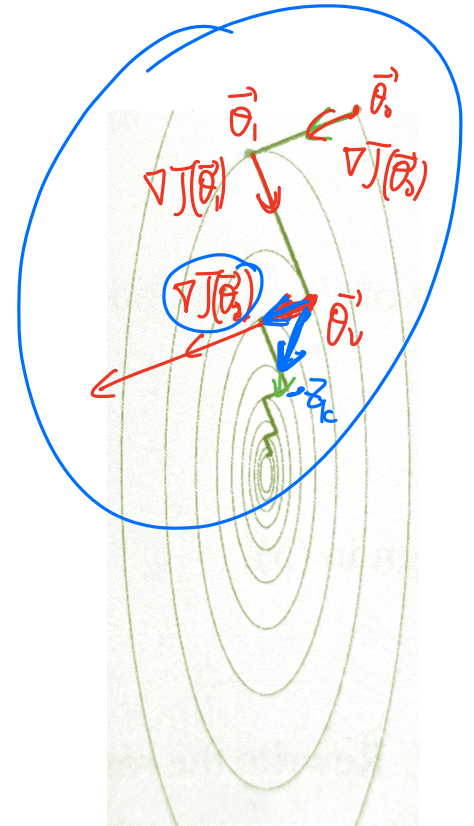
More gradient methods:

$$\text{Recall GD: } \vec{\theta}^{\text{next}} = \vec{\theta} - \alpha \nabla J(\vec{\theta})$$

1. Descent with momentum(memory)

$$\vec{\theta}_{k+1} = \vec{\theta}_k - \alpha \mathbf{Z}_k$$

$$\text{Here } \mathbf{Z}_k = \nabla J(\vec{\theta}_k) + \beta \mathbf{Z}_{k-1}$$





## 2. Adaptive Stochastic Gradient Descent

Recall SGD:  $\vec{\theta}^{\text{next}} = \vec{\theta} - \alpha \nabla J(\vec{\theta}; \vec{x}^{(i)}, y^{(i)})$

Adaptive:  $\vec{\theta}_{k+1} = \vec{\theta}_k - \alpha_k D_k$

Here  $\alpha_k = \alpha(\nabla J_k, \nabla J_{k-1}, \dots, \nabla J_0)$

$D_k = D(\nabla J_k, \nabla J_{k-1}, \dots, \nabla J_0)$

For example, **ADAGRAD (2011)**

$$\alpha_k = \frac{\alpha}{\sqrt{k}} \left( \frac{1}{k} \text{diag} \sum_{i=1}^k \|\nabla J_i\|^2 \right)^{\frac{1}{2}} \quad \text{and} \quad D_k = \nabla J(\vec{\theta}_k)$$

## ADAM (2015)

Recursive formula:

$$\begin{aligned} D_k &= \delta D_{k-1} + (1 - \delta) \nabla J(\vec{\theta}_k) \\ \alpha_k^2 &= \beta \alpha_{k-1}^2 + (1 - \beta) \|\nabla J(\vec{\theta}_i)\|^2 \end{aligned}$$

More explicitly,

$$\begin{aligned} D_k &= (1 - \delta) \sum_{i=1}^k \delta^{k-i} \nabla J(\vec{\theta}_k) \\ \alpha_k &= \frac{\alpha}{\sqrt{k}} \left( (1 - \beta) \text{diag} \sum_{i=1}^k \beta^{k-i} \|\nabla J(\vec{\theta}_i)\|^2 \right)^{\frac{1}{2}} \end{aligned}$$

Diederik P. Kingma and Jimmy Lei Ba. Adam: a Method for Stochastic Optimization. International Conference on Learning Representations, pages 1–13, 2015.

**An overview of gradient descent optimization algorithms**

<https://arxiv.org/abs/1609.04747>