

Module 5.5 : Nesterov Accelerated Gradient Descent

Question

- Can we do something to reduce these oscillations ?

Question

- Can we do something to reduce these oscillations ?
- Yes, let's look at Nesterov accelerated gradient

Intuition

- Look before you leap

Intuition

- Look before you leap
- Recall that $update_t = \gamma \cdot update_{t-1} + \eta \nabla w_t$

Intuition

- Look before you leap
- Recall that $update_t = \gamma \cdot update_{t-1} + \eta \nabla w_t$
- So we know that we are going to move by at least by $\gamma \cdot update_{t-1}$ and then a bit more by $\eta \nabla w_t$

Intuition

- Look before you leap
- Recall that $update_t = \gamma \cdot update_{t-1} + \eta \nabla w_t$
- So we know that we are going to move by at least by $\gamma \cdot update_{t-1}$ and then a bit more by $\eta \nabla w_t$
- Why not calculate the gradient (∇w_{look_ahead}) at this partially updated value of w ($w_{look_ahead} = w_t - \gamma \cdot update_{t-1}$) instead of calculating it using the current value w_t

Intuition

- Look before you leap
- Recall that $update_t = \gamma \cdot update_{t-1} + \eta \nabla w_t$
- So we know that we are going to move by at least by $\gamma \cdot update_{t-1}$ and then a bit more by $\eta \nabla w_t$
- Why not calculate the gradient (∇w_{look_ahead}) at this partially updated value of w ($w_{look_ahead} = w_t - \gamma \cdot update_{t-1}$) instead of calculating it using the current value w_t

Update rule for NAG

$$w_{look_ahead} = w_t - \gamma \cdot update_{t-1}$$

$$update_t = \gamma \cdot update_{t-1} + \eta \nabla w_{look_ahead}$$

$$w_{t+1} = w_t - update_t$$

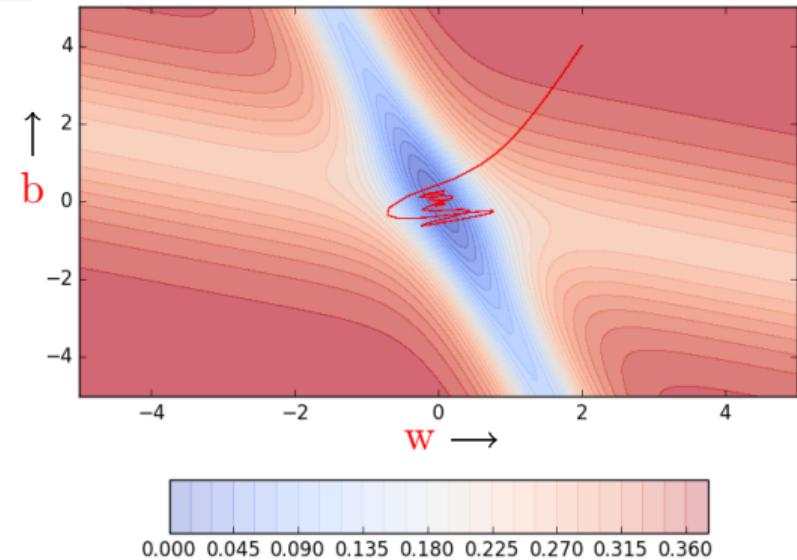
We will have similar update rule for b_t

```

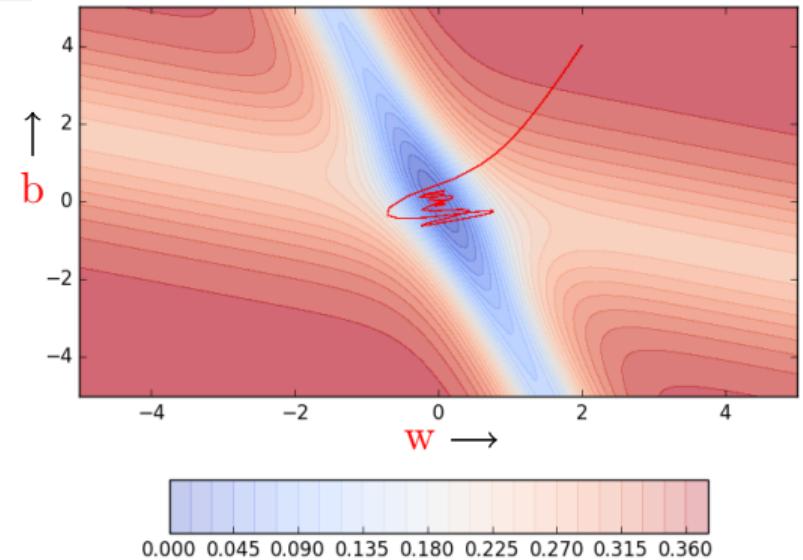
def do_nesterov_accelerated_gradient_descent() :
    w, b, eta = init_w, init_b , 1.0
    prev_v_w, prev_v_b, gamma = 0, 0, 0.9
    for i in range(max_epochs) :
        dw, db = 0, 0
        #do partial updates
        v_w = gamma * prev_v_w
        v_b = gamma * prev_v_b
        for x,y in zip(X, Y) :
            #calculate gradients after partial update
            dw += grad_w(w - v_w, b - v_b, x, y)
            db += grad_b(w - v_w, b - v_b, x, y)

        #now do the full update
        v_w = gamma * prev_v_w + eta * dw
        v_b = gamma * prev_v_b + eta * db
        w = w - v_w
        b = b - v_b
        prev_v_w = v_w
        prev_v_b = v_b

```



```
def do_nesterov_accelerated_gradient_descent() :  
    w, b, eta = init_w, init_b , 1.0  
    prev_v_w, prev_v_b, gamma = 0, 0, 0.9  
    for i in range(max_epochs) :  
        dw, db = 0, 0  
        #do partial updates  
        v_w = gamma * prev_v_w  
        v_b = gamma * prev_v_b  
        for x,y in zip(X, Y) :  
            #calculate gradients after partial update  
            dw += grad_w(w - v_w, b - v_b, x, y)  
            db += grad_b(w - v_w, b - v_b, x, y)  
  
        #now do the full update  
        v_w = gamma * prev_v_w + eta * dw  
        v_b = gamma * prev_v_b + eta * db  
        w = w - v_w  
        b = b - v_b  
        prev_v_w = v_w  
        prev_v_b = v_b
```

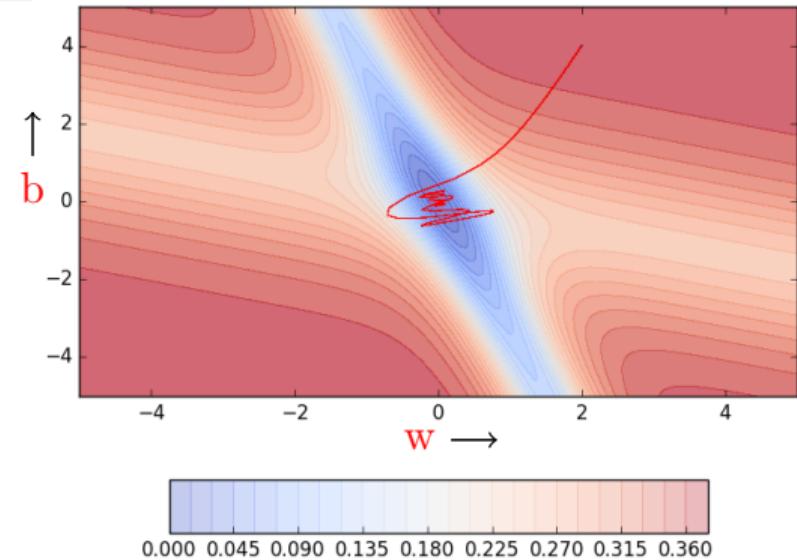


```

def do_nesterov_accelerated_gradient_descent() :
    w, b, eta = init_w, init_b , 1.0
    prev_v_w, prev_v_b, gamma = 0, 0, 0.9
    for i in range(max_epochs) :
        dw, db = 0, 0
        #do partial updates
        v_w = gamma * prev_v_w
        v_b = gamma * prev_v_b
        for x,y in zip(X, Y) :
            #calculate gradients after partial update
            dw += grad_w(w - v_w, b - v_b, x, y)
            db += grad_b(w - v_w, b - v_b, x, y)

        #now do the full update
        v_w = gamma * prev_v_w + eta * dw
        v_b = gamma * prev_v_b + eta * db
        w = w - v_w
        b = b - v_b
        prev_v_w = v_w
        prev_v_b = v_b

```

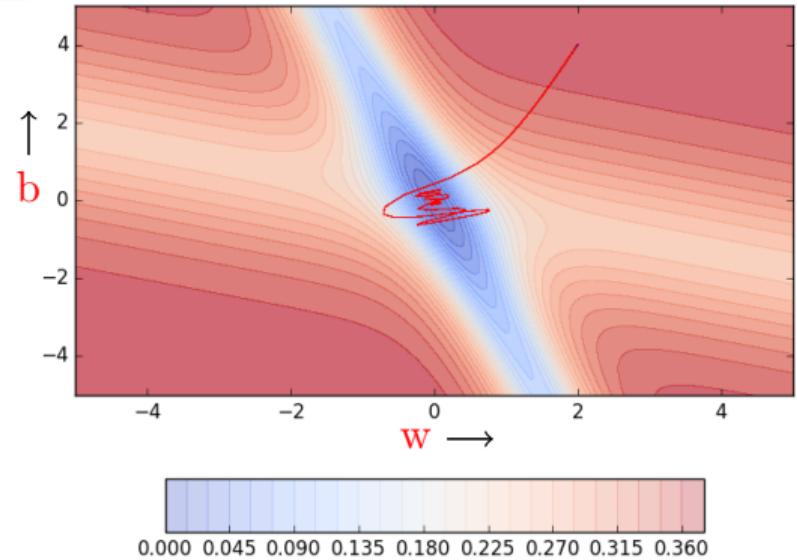


```

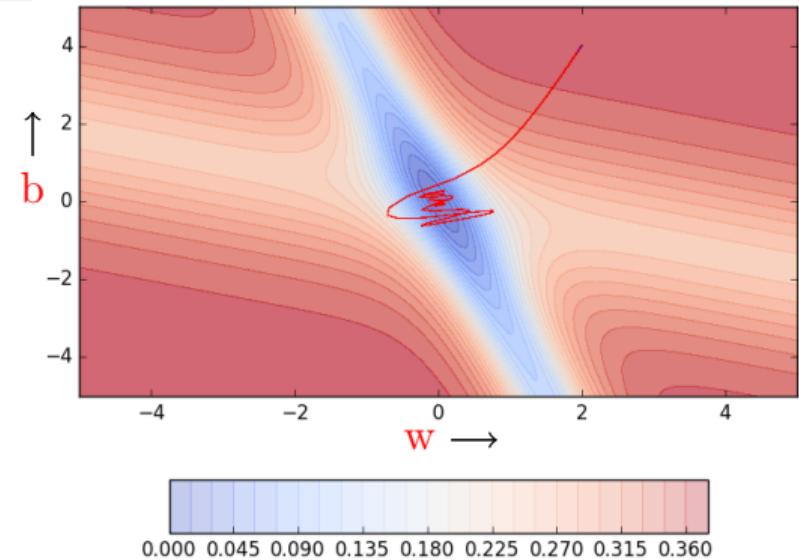
def do_nesterov_accelerated_gradient_descent() :
    w, b, eta = init_w, init_b , 1.0
    prev_v_w, prev_v_b, gamma = 0, 0, 0.9
    for i in range(max_epochs) :
        dw, db = 0, 0
        #do partial updates
        v_w = gamma * prev_v_w
        v_b = gamma * prev_v_b
        for x,y in zip(X, Y):
            #calculate gradients after partial update
            dw += grad_w(w - v_w, b - v_b, x, y)
            db += grad_b(w - v_w, b - v_b, x, y)

        #now do the full update
        v_w = gamma * prev_v_w + eta * dw
        v_b = gamma * prev_v_b + eta * db
        w = w - v_w
        b = b - v_b
        prev_v_w = v_w
        prev_v_b = v_b

```



```
def do_nesterov_accelerated_gradient_descent() :  
    w, b, eta = init_w, init_b , 1.0  
    prev_v_w, prev_v_b, gamma = 0, 0, 0.9  
    for i in range(max_epochs) :  
        dw, db = 0, 0  
        #do partial updates  
        v_w = gamma * prev_v_w  
        v_b = gamma * prev_v_b  
        for x,y in zip(X, Y) :  
            #calculate gradients after partial update  
            dw += grad_w(w - v_w, b - v_b, x, y)  
            db += grad_b(w - v_w, b - v_b, x, y)  
  
        #now do the full update  
        v_w = gamma * prev_v_w + eta * dw  
        v_b = gamma * prev_v_b + eta * db  
        w = w - v_w  
        b = b - v_b  
        prev_v_w = v_w  
        prev_v_b = v_b
```

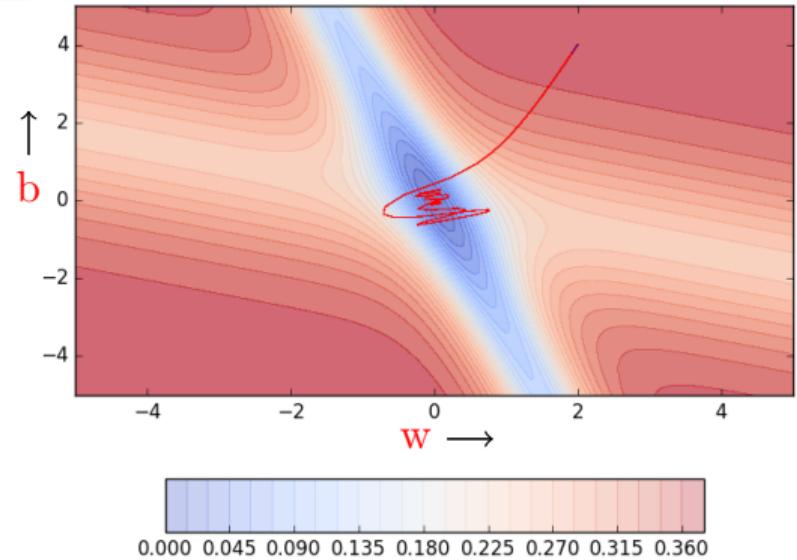


```

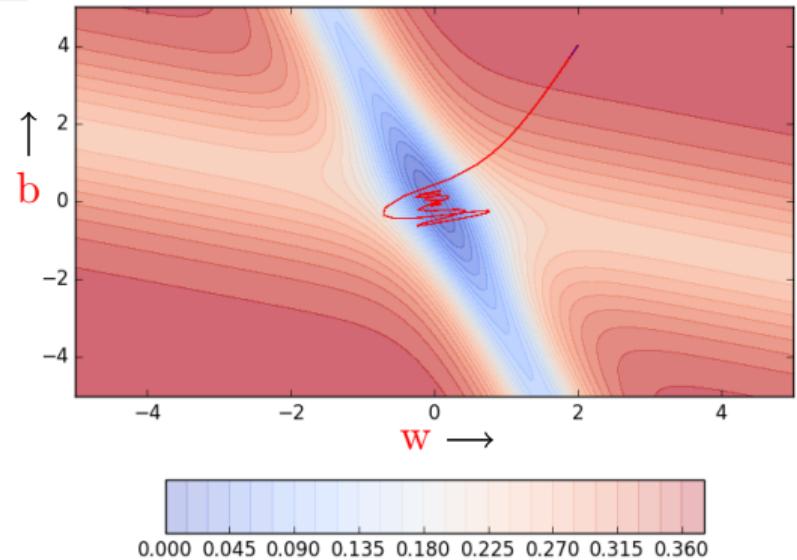
def do_nesterov_accelerated_gradient_descent() :
    w, b, eta = init_w, init_b , 1.0
    prev_v_w, prev_v_b, gamma = 0, 0, 0.9
    for i in range(max_epochs) :
        dw, db = 0, 0
        #do partial updates
        v_w = gamma * prev_v_w
        v_b = gamma * prev_v_b
        for x,y in zip(X, Y) :
            #calculate gradients after partial update
            dw += grad_w(w - v_w, b - v_b, x, y)
            db += grad_b(w - v_w, b - v_b, x, y)

        #now do the full update
        v_w = gamma * prev_v_w + eta * dw
        v_b = gamma * prev_v_b + eta * db
        w = w - v_w
        b = b - v_b
        prev_v_w = v_w
        prev_v_b = v_b

```



```
def do_nesterov_accelerated_gradient_descent() :  
    w, b, eta = init_w, init_b , 1.0  
    prev_v_w, prev_v_b, gamma = 0, 0, 0.9  
    for i in range(max_epochs) :  
        dw, db = 0, 0  
        #do partial updates  
        v_w = gamma * prev_v_w  
        v_b = gamma * prev_v_b  
        for x,y in zip(X, Y) :  
            #calculate gradients after partial update  
            dw += grad_w(w - v_w, b - v_b, x, y)  
            db += grad_b(w - v_w, b - v_b, x, y)  
  
        #now do the full update  
        v_w = gamma * prev_v_w + eta * dw  
        v_b = gamma * prev_v_b + eta * db  
        w = w - v_w  
        b = b - v_b  
        prev_v_w = v_w  
        prev_v_b = v_b
```

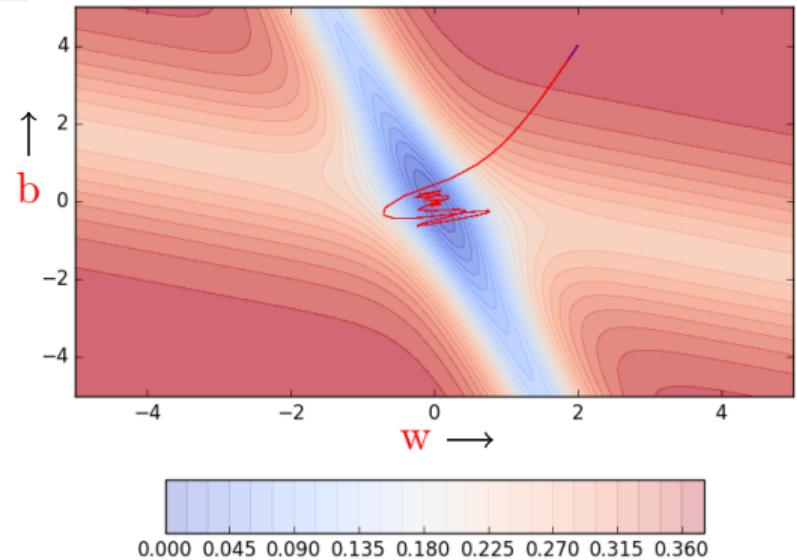


```

def do_nesterov_accelerated_gradient_descent() :
    w, b, eta = init_w, init_b , 1.0
    prev_v_w, prev_v_b, gamma = 0, 0, 0.9
    for i in range(max_epochs) :
        dw, db = 0, 0
        #do partial updates
        v_w = gamma * prev_v_w
        v_b = gamma * prev_v_b
        for x,y in zip(X, Y) :
            #calculate gradients after partial update
            dw += grad_w(w - v_w, b - v_b, x, y)
            db += grad_b(w - v_w, b - v_b, x, y)

        #now do the full update
        v_w = gamma * prev_v_w + eta * dw
        v_b = gamma * prev_v_b + eta * db
        w = w - v_w
        b = b - v_b
        prev_v_w = v_w
        prev_v_b = v_b

```

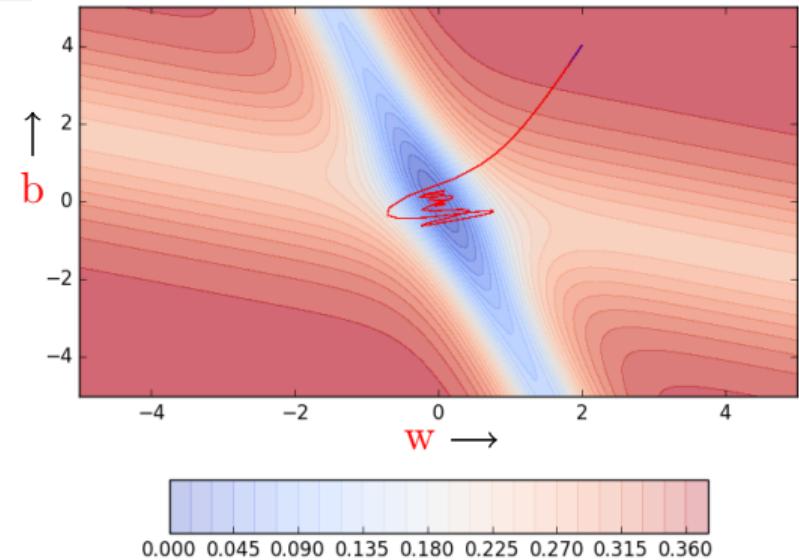


```

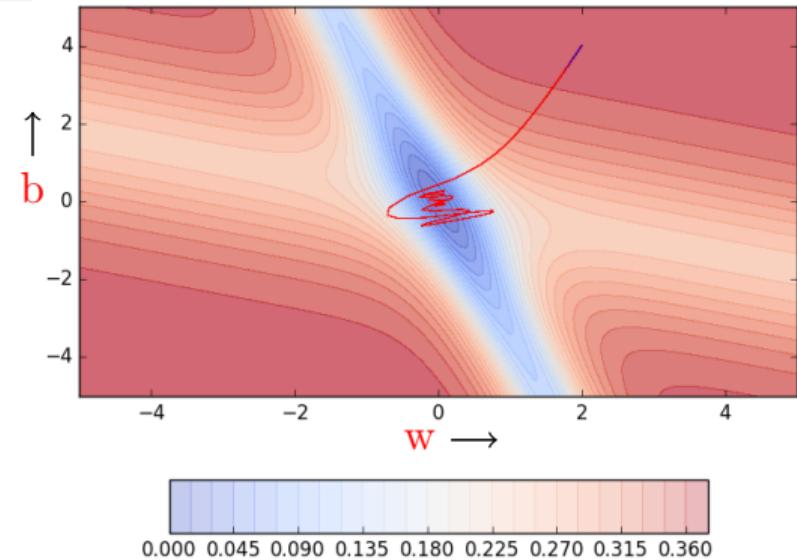
def do_nesterov_accelerated_gradient_descent() :
    w, b, eta = init_w, init_b , 1.0
    prev_v_w, prev_v_b, gamma = 0, 0, 0.9
    for i in range(max_epochs) :
        dw, db = 0, 0
        #do partial updates
        v_w = gamma * prev_v_w
        v_b = gamma * prev_v_b
        for x,y in zip(X, Y) :
            #calculate gradients after partial update
            dw += grad_w(w - v_w, b - v_b, x, y)
            db += grad_b(w - v_w, b - v_b, x, y)

        #now do the full update
        v_w = gamma * prev_v_w + eta * dw
        v_b = gamma * prev_v_b + eta * db
        w = w - v_w
        b = b - v_b
        prev_v_w = v_w
        prev_v_b = v_b

```



```
def do_nesterov_accelerated_gradient_descent() :  
    w, b, eta = init_w, init_b , 1.0  
    prev_v_w, prev_v_b, gamma = 0, 0, 0.9  
    for i in range(max_epochs) :  
        dw, db = 0, 0  
        #do partial updates  
        v_w = gamma * prev_v_w  
        v_b = gamma * prev_v_b  
        for x,y in zip(X, Y) :  
            #calculate gradients after partial update  
            dw += grad_w(w - v_w, b - v_b, x, y)  
            db += grad_b(w - v_w, b - v_b, x, y)  
  
        #now do the full update  
        v_w = gamma * prev_v_w + eta * dw  
        v_b = gamma * prev_v_b + eta * db  
        w = w - v_w  
        b = b - v_b  
        prev_v_w = v_w  
        prev_v_b = v_b
```

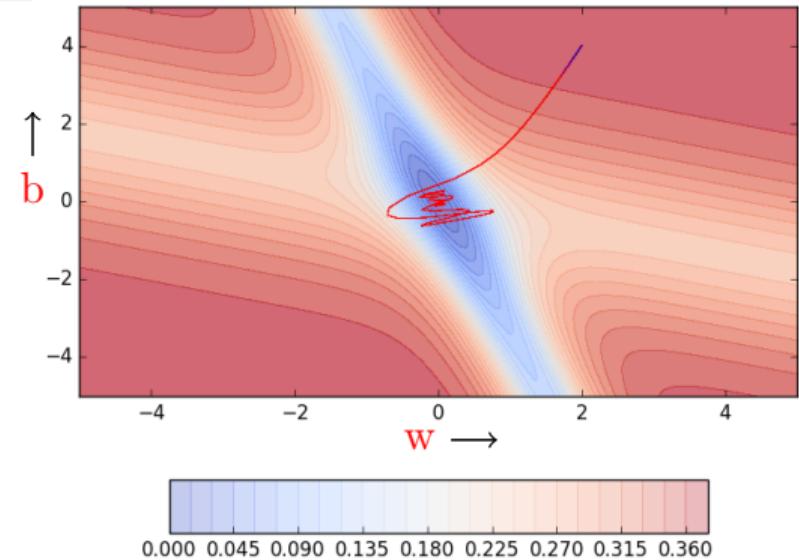


```

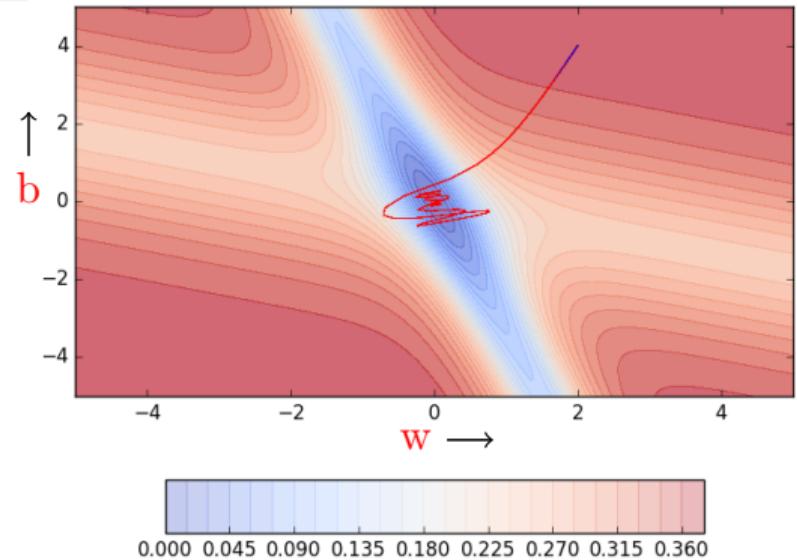
def do_nesterov_accelerated_gradient_descent() :
    w, b, eta = init_w, init_b , 1.0
    prev_v_w, prev_v_b, gamma = 0, 0, 0.9
    for i in range(max_epochs) :
        dw, db = 0, 0
        #do partial updates
        v_w = gamma * prev_v_w
        v_b = gamma * prev_v_b
        for x,y in zip(X, Y) :
            #calculate gradients after partial update
            dw += grad_w(w - v_w, b - v_b, x, y)
            db += grad_b(w - v_w, b - v_b, x, y)

        #now do the full update
        v_w = gamma * prev_v_w + eta * dw
        v_b = gamma * prev_v_b + eta * db
        w = w - v_w
        b = b - v_b
        prev_v_w = v_w
        prev_v_b = v_b

```



```
def do_nesterov_accelerated_gradient_descent() :  
    w, b, eta = init_w, init_b , 1.0  
    prev_v_w, prev_v_b, gamma = 0, 0, 0.9  
    for i in range(max_epochs) :  
        dw, db = 0, 0  
        #do partial updates  
        v_w = gamma * prev_v_w  
        v_b = gamma * prev_v_b  
        for x,y in zip(X, Y) :  
            #calculate gradients after partial update  
            dw += grad_w(w - v_w, b - v_b, x, y)  
            db += grad_b(w - v_w, b - v_b, x, y)  
  
        #now do the full update  
        v_w = gamma * prev_v_w + eta * dw  
        v_b = gamma * prev_v_b + eta * db  
        w = w - v_w  
        b = b - v_b  
        prev_v_w = v_w  
        prev_v_b = v_b
```

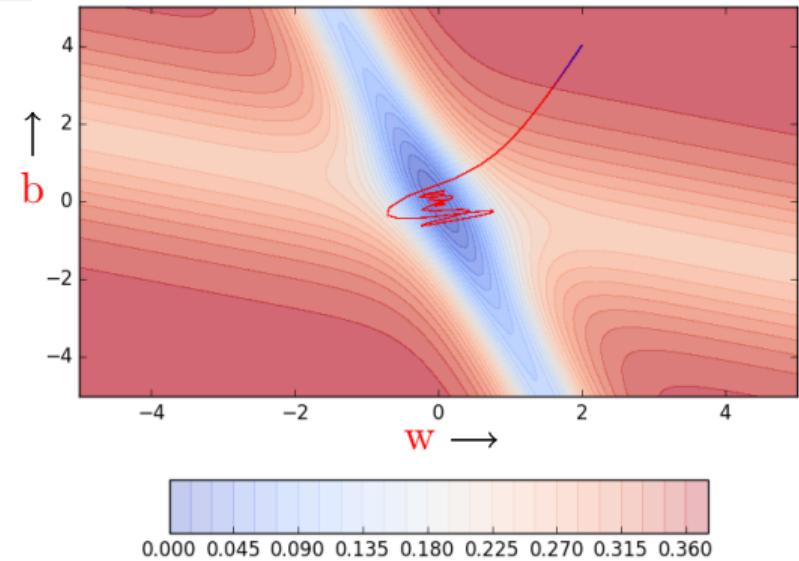


```

def do_nesterov_accelerated_gradient_descent() :
    w, b, eta = init_w, init_b , 1.0
    prev_v_w, prev_v_b, gamma = 0, 0, 0.9
    for i in range(max_epochs) :
        dw, db = 0, 0
        #do partial updates
        v_w = gamma * prev_v_w
        v_b = gamma * prev_v_b
        for x,y in zip(X, Y) :
            #calculate gradients after partial update
            dw += grad_w(w - v_w, b - v_b, x, y)
            db += grad_b(w - v_w, b - v_b, x, y)

        #now do the full update
        v_w = gamma * prev_v_w + eta * dw
        v_b = gamma * prev_v_b + eta * db
        w = w - v_w
        b = b - v_b
        prev_v_w = v_w
        prev_v_b = v_b

```

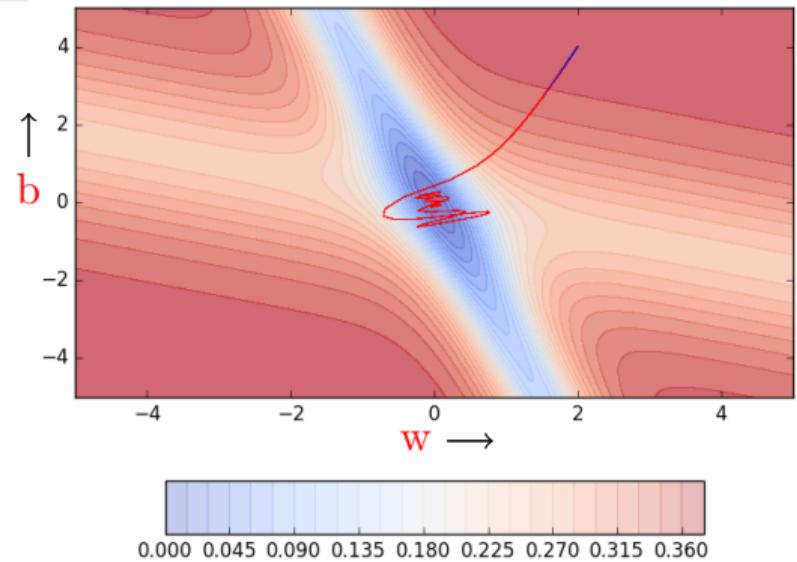


```

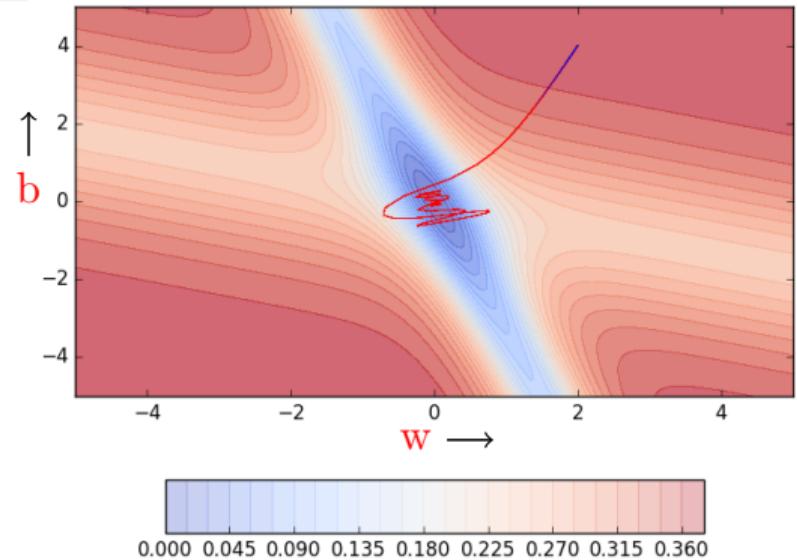
def do_nesterov_accelerated_gradient_descent() :
    w, b, eta = init_w, init_b , 1.0
    prev_v_w, prev_v_b, gamma = 0, 0, 0.9
    for i in range(max_epochs) :
        dw, db = 0, 0
        #do partial updates
        v_w = gamma * prev_v_w
        v_b = gamma * prev_v_b
        for x,y in zip(X, Y) :
            #calculate gradients after partial update
            dw += grad_w(w - v_w, b - v_b, x, y)
            db += grad_b(w - v_w, b - v_b, x, y)

        #now do the full update
        v_w = gamma * prev_v_w + eta * dw
        v_b = gamma * prev_v_b + eta * db
        w = w - v_w
        b = b - v_b
        prev_v_w = v_w
        prev_v_b = v_b

```



```
def do_nesterov_accelerated_gradient_descent() :  
    w, b, eta = init_w, init_b , 1.0  
    prev_v_w, prev_v_b, gamma = 0, 0, 0.9  
    for i in range(max_epochs) :  
        dw, db = 0, 0  
        #do partial updates  
        v_w = gamma * prev_v_w  
        v_b = gamma * prev_v_b  
        for x,y in zip(X, Y) :  
            #calculate gradients after partial update  
            dw += grad_w(w - v_w, b - v_b, x, y)  
            db += grad_b(w - v_w, b - v_b, x, y)  
  
        #now do the full update  
        v_w = gamma * prev_v_w + eta * dw  
        v_b = gamma * prev_v_b + eta * db  
        w = w - v_w  
        b = b - v_b  
        prev_v_w = v_w  
        prev_v_b = v_b
```

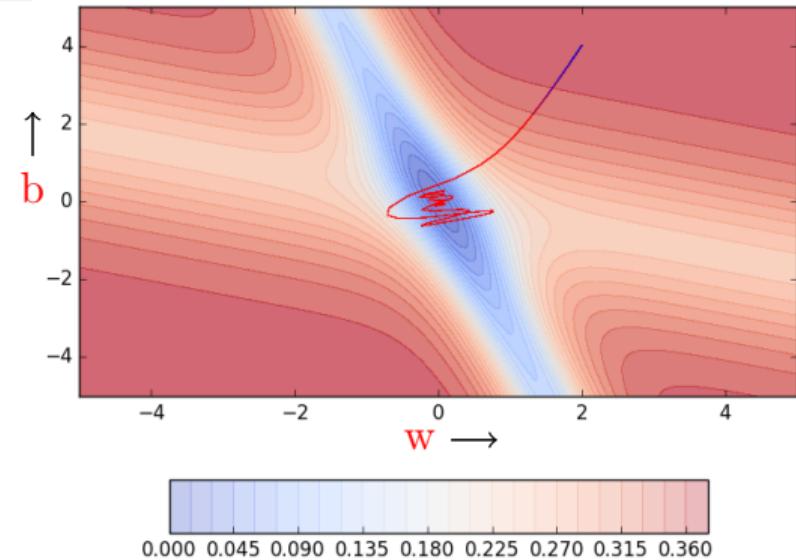


```

def do_nesterov_accelerated_gradient_descent() :
    w, b, eta = init_w, init_b , 1.0
    prev_v_w, prev_v_b, gamma = 0, 0, 0.9
    for i in range(max_epochs) :
        dw, db = 0, 0
        #do partial updates
        v_w = gamma * prev_v_w
        v_b = gamma * prev_v_b
        for x,y in zip(X, Y) :
            #calculate gradients after partial update
            dw += grad_w(w - v_w, b - v_b, x, y)
            db += grad_b(w - v_w, b - v_b, x, y)

        #now do the full update
        v_w = gamma * prev_v_w + eta * dw
        v_b = gamma * prev_v_b + eta * db
        w = w - v_w
        b = b - v_b
        prev_v_w = v_w
        prev_v_b = v_b

```

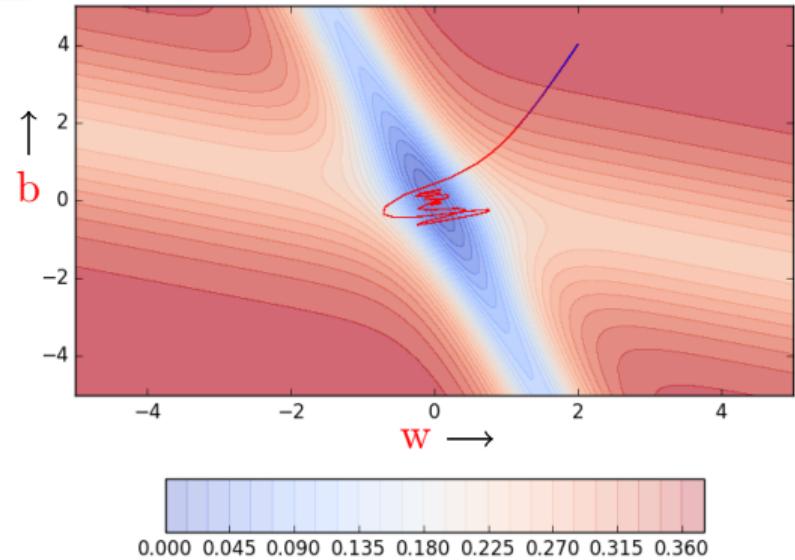


```

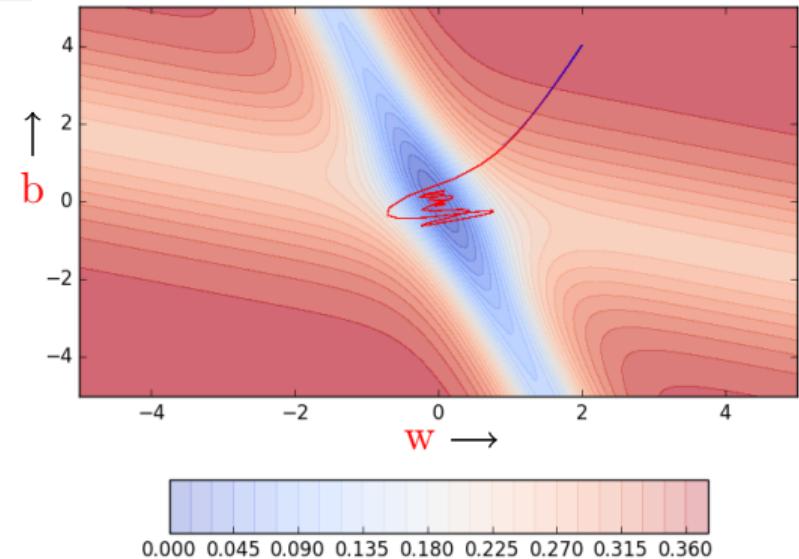
def do_nesterov_accelerated_gradient_descent() :
    w, b, eta = init_w, init_b , 1.0
    prev_v_w, prev_v_b, gamma = 0, 0, 0.9
    for i in range(max_epochs) :
        dw, db = 0, 0
        #do partial updates
        v_w = gamma * prev_v_w
        v_b = gamma * prev_v_b
        for x,y in zip(X, Y) :
            #calculate gradients after partial update
            dw += grad_w(w - v_w, b - v_b, x, y)
            db += grad_b(w - v_w, b - v_b, x, y)

        #now do the full update
        v_w = gamma * prev_v_w + eta * dw
        v_b = gamma * prev_v_b + eta * db
        w = w - v_w
        b = b - v_b
        prev_v_w = v_w
        prev_v_b = v_b

```



```
def do_nesterov_accelerated_gradient_descent() :  
    w, b, eta = init_w, init_b , 1.0  
    prev_v_w, prev_v_b, gamma = 0, 0, 0.9  
    for i in range(max_epochs) :  
        dw, db = 0, 0  
        #do partial updates  
        v_w = gamma * prev_v_w  
        v_b = gamma * prev_v_b  
        for x,y in zip(X, Y) :  
            #calculate gradients after partial update  
            dw += grad_w(w - v_w, b - v_b, x, y)  
            db += grad_b(w - v_w, b - v_b, x, y)  
  
        #now do the full update  
        v_w = gamma * prev_v_w + eta * dw  
        v_b = gamma * prev_v_b + eta * db  
        w = w - v_w  
        b = b - v_b  
        prev_v_w = v_w  
        prev_v_b = v_b
```

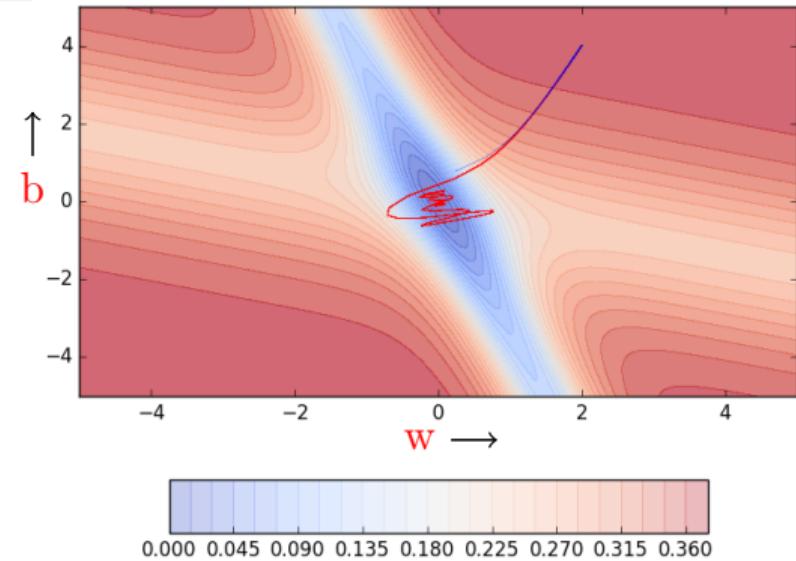


```

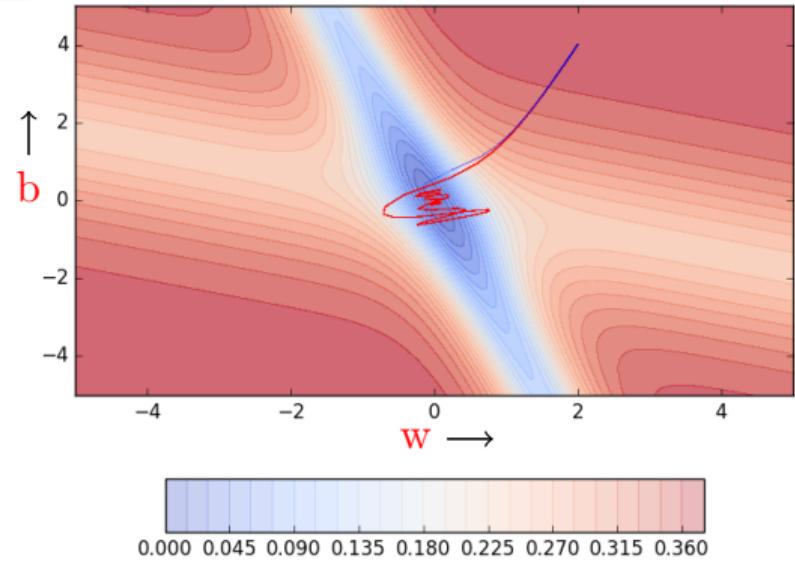
def do_nesterov_accelerated_gradient_descent() :
    w, b, eta = init_w, init_b , 1.0
    prev_v_w, prev_v_b, gamma = 0, 0, 0.9
    for i in range(max_epochs) :
        dw, db = 0, 0
        #do partial updates
        v_w = gamma * prev_v_w
        v_b = gamma * prev_v_b
        for x,y in zip(X, Y) :
            #calculate gradients after partial update
            dw += grad_w(w - v_w, b - v_b, x, y)
            db += grad_b(w - v_w, b - v_b, x, y)

        #now do the full update
        v_w = gamma * prev_v_w + eta * dw
        v_b = gamma * prev_v_b + eta * db
        w = w - v_w
        b = b - v_b
        prev_v_w = v_w
        prev_v_b = v_b

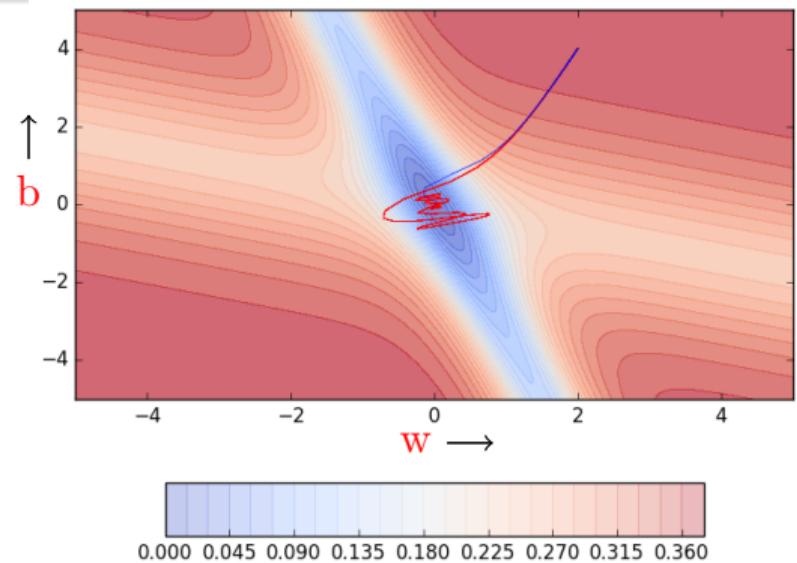
```



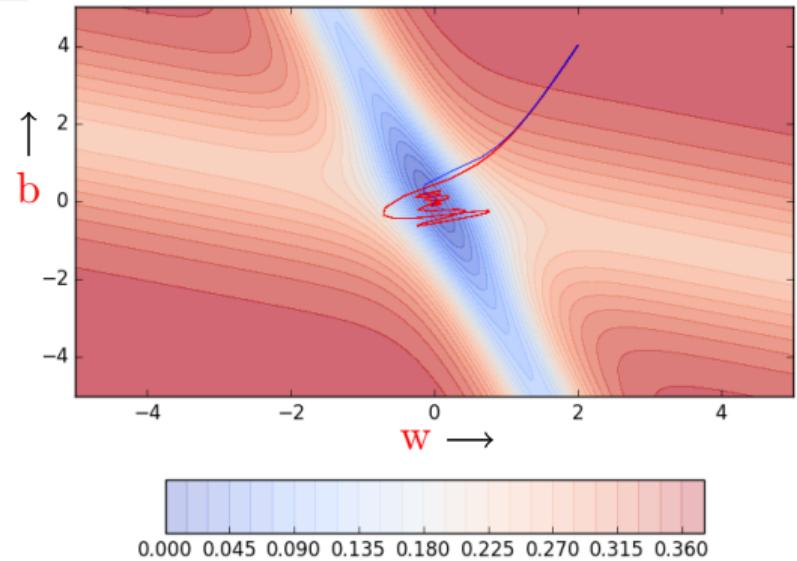
```
def do_nesterov_accelerated_gradient_descent() :  
    w, b, eta = init_w, init_b , 1.0  
    prev_v_w, prev_v_b, gamma = 0, 0, 0.9  
    for i in range(max_epochs) :  
        dw, db = 0, 0  
        #do partial updates  
        v_w = gamma * prev_v_w  
        v_b = gamma * prev_v_b  
        for x,y in zip(X, Y) :  
            #calculate gradients after partial update  
            dw += grad_w(w - v_w, b - v_b, x, y)  
            db += grad_b(w - v_w, b - v_b, x, y)  
  
        #now do the full update  
        v_w = gamma * prev_v_w + eta * dw  
        v_b = gamma * prev_v_b + eta * db  
        w = w - v_w  
        b = b - v_b  
        prev_v_w = v_w  
        prev_v_b = v_b
```



```
def do_nesterov_accelerated_gradient_descent() :  
    w, b, eta = init_w, init_b , 1.0  
    prev_v_w, prev_v_b, gamma = 0, 0, 0.9  
    for i in range(max_epochs) :  
        dw, db = 0, 0  
        #do partial updates  
        v_w = gamma * prev_v_w  
        v_b = gamma * prev_v_b  
        for x,y in zip(X, Y) :  
            #calculate gradients after partial update  
            dw += grad_w(w - v_w, b - v_b, x, y)  
            db += grad_b(w - v_w, b - v_b, x, y)  
  
        #now do the full update  
        v_w = gamma * prev_v_w + eta * dw  
        v_b = gamma * prev_v_b + eta * db  
        w = w - v_w  
        b = b - v_b  
        prev_v_w = v_w  
        prev_v_b = v_b
```



```
def do_nesterov_accelerated_gradient_descent() :  
    w, b, eta = init_w, init_b , 1.0  
    prev_v_w, prev_v_b, gamma = 0, 0, 0.9  
    for i in range(max_epochs) :  
        dw, db = 0, 0  
        #do partial updates  
        v_w = gamma * prev_v_w  
        v_b = gamma * prev_v_b  
        for x,y in zip(X, Y) :  
            #calculate gradients after partial update  
            dw += grad_w(w - v_w, b - v_b, x, y)  
            db += grad_b(w - v_w, b - v_b, x, y)  
  
        #now do the full update  
        v_w = gamma * prev_v_w + eta * dw  
        v_b = gamma * prev_v_b + eta * db  
        w = w - v_w  
        b = b - v_b  
        prev_v_w = v_w  
        prev_v_b = v_b
```

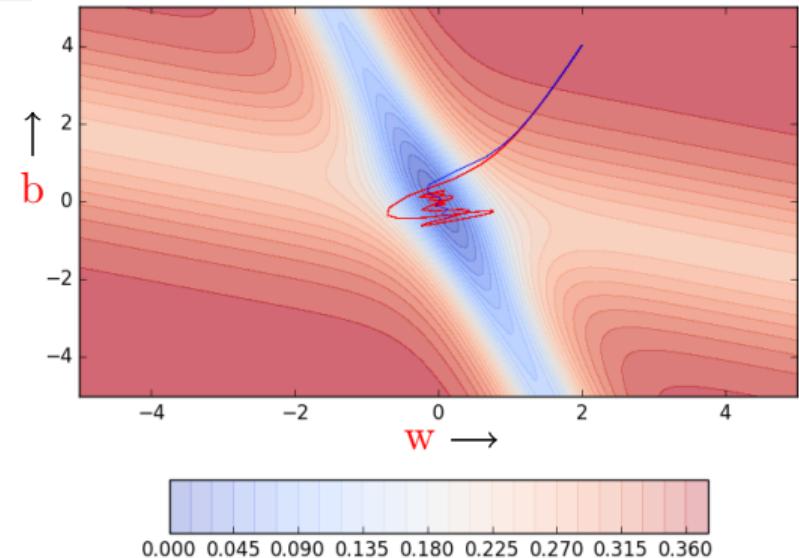


```

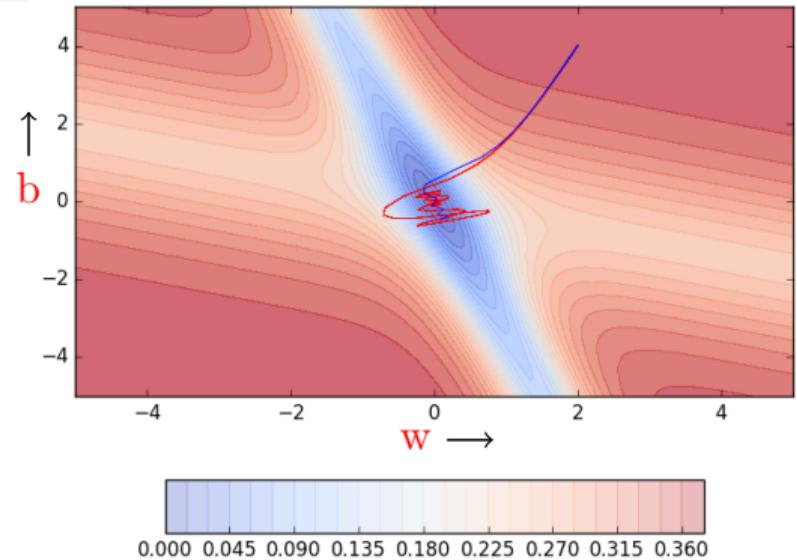
def do_nesterov_accelerated_gradient_descent() :
    w, b, eta = init_w, init_b , 1.0
    prev_v_w, prev_v_b, gamma = 0, 0, 0.9
    for i in range(max_epochs) :
        dw, db = 0, 0
        #do partial updates
        v_w = gamma * prev_v_w
        v_b = gamma * prev_v_b
        for x,y in zip(X, Y) :
            #calculate gradients after partial update
            dw += grad_w(w - v_w, b - v_b, x, y)
            db += grad_b(w - v_w, b - v_b, x, y)

        #now do the full update
        v_w = gamma * prev_v_w + eta * dw
        v_b = gamma * prev_v_b + eta * db
        w = w - v_w
        b = b - v_b
        prev_v_w = v_w
        prev_v_b = v_b

```



```
def do_nesterov_accelerated_gradient_descent() :  
    w, b, eta = init_w, init_b , 1.0  
    prev_v_w, prev_v_b, gamma = 0, 0, 0.9  
    for i in range(max_epochs) :  
        dw, db = 0, 0  
        #do partial updates  
        v_w = gamma * prev_v_w  
        v_b = gamma * prev_v_b  
        for x,y in zip(X, Y) :  
            #calculate gradients after partial update  
            dw += grad_w(w - v_w, b - v_b, x, y)  
            db += grad_b(w - v_w, b - v_b, x, y)  
  
        #now do the full update  
        v_w = gamma * prev_v_w + eta * dw  
        v_b = gamma * prev_v_b + eta * db  
        w = w - v_w  
        b = b - v_b  
        prev_v_w = v_w  
        prev_v_b = v_b
```

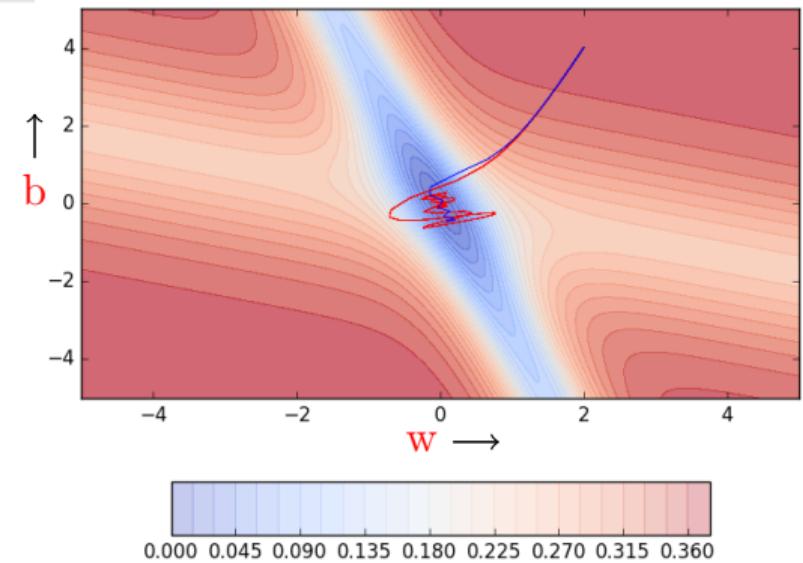


```

def do_nesterov_accelerated_gradient_descent() :
    w, b, eta = init_w, init_b , 1.0
    prev_v_w, prev_v_b, gamma = 0, 0, 0.9
    for i in range(max_epochs) :
        dw, db = 0, 0
        #do partial updates
        v_w = gamma * prev_v_w
        v_b = gamma * prev_v_b
        for x,y in zip(X, Y):
            #calculate gradients after partial update
            dw += grad_w(w - v_w, b - v_b, x, y)
            db += grad_b(w - v_w, b - v_b, x, y)

        #now do the full update
        v_w = gamma * prev_v_w + eta * dw
        v_b = gamma * prev_v_b + eta * db
        w = w - v_w
        b = b - v_b
        prev_v_w = v_w
        prev_v_b = v_b

```

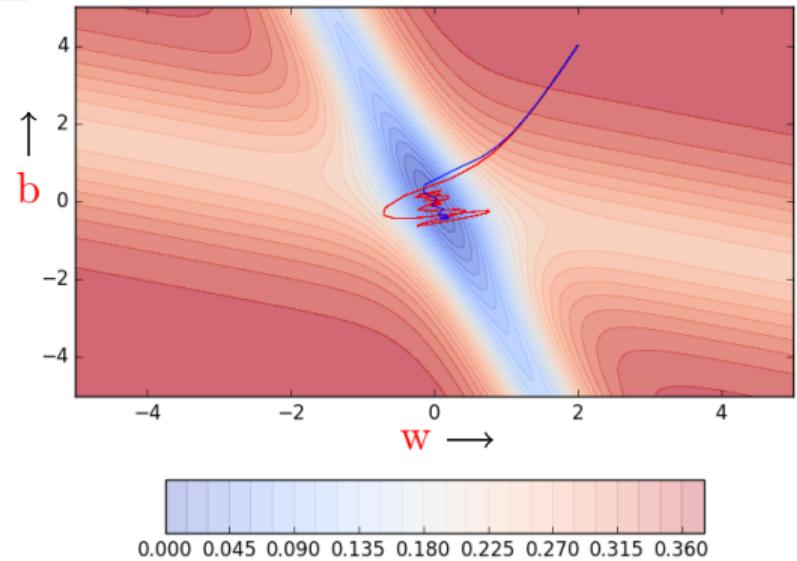


```

def do_nesterov_accelerated_gradient_descent() :
    w, b, eta = init_w, init_b , 1.0
    prev_v_w, prev_v_b, gamma = 0, 0, 0.9
    for i in range(max_epochs) :
        dw, db = 0, 0
        #do partial updates
        v_w = gamma * prev_v_w
        v_b = gamma * prev_v_b
        for x,y in zip(X, Y) :
            #calculate gradients after partial update
            dw += grad_w(w - v_w, b - v_b, x, y)
            db += grad_b(w - v_w, b - v_b, x, y)

        #now do the full update
        v_w = gamma * prev_v_w + eta * dw
        v_b = gamma * prev_v_b + eta * db
        w = w - v_w
        b = b - v_b
        prev_v_w = v_w
        prev_v_b = v_b

```

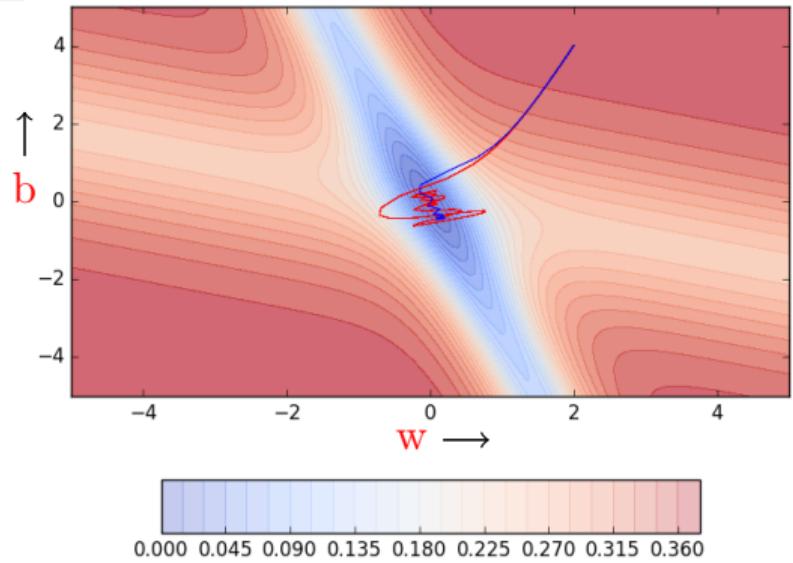


```

def do_nesterov_accelerated_gradient_descent() :
    w, b, eta = init_w, init_b , 1.0
    prev_v_w, prev_v_b, gamma = 0, 0, 0.9
    for i in range(max_epochs) :
        dw, db = 0, 0
        #do partial updates
        v_w = gamma * prev_v_w
        v_b = gamma * prev_v_b
        for x,y in zip(X, Y) :
            #calculate gradients after partial update
            dw += grad_w(w - v_w, b - v_b, x, y)
            db += grad_b(w - v_w, b - v_b, x, y)

        #now do the full update
        v_w = gamma * prev_v_w + eta * dw
        v_b = gamma * prev_v_b + eta * db
        w = w - v_w
        b = b - v_b
        prev_v_w = v_w
        prev_v_b = v_b

```

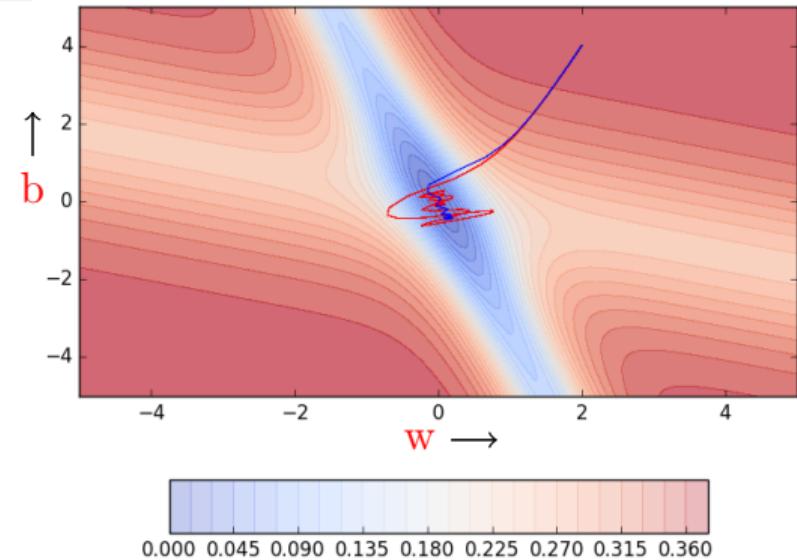


```

def do_nesterov_accelerated_gradient_descent() :
    w, b, eta = init_w, init_b , 1.0
    prev_v_w, prev_v_b, gamma = 0, 0, 0.9
    for i in range(max_epochs) :
        dw, db = 0, 0
        #do partial updates
        v_w = gamma * prev_v_w
        v_b = gamma * prev_v_b
        for x,y in zip(X, Y) :
            #calculate gradients after partial update
            dw += grad_w(w - v_w, b - v_b, x, y)
            db += grad_b(w - v_w, b - v_b, x, y)

        #now do the full update
        v_w = gamma * prev_v_w + eta * dw
        v_b = gamma * prev_v_b + eta * db
        w = w - v_w
        b = b - v_b
        prev_v_w = v_w
        prev_v_b = v_b

```

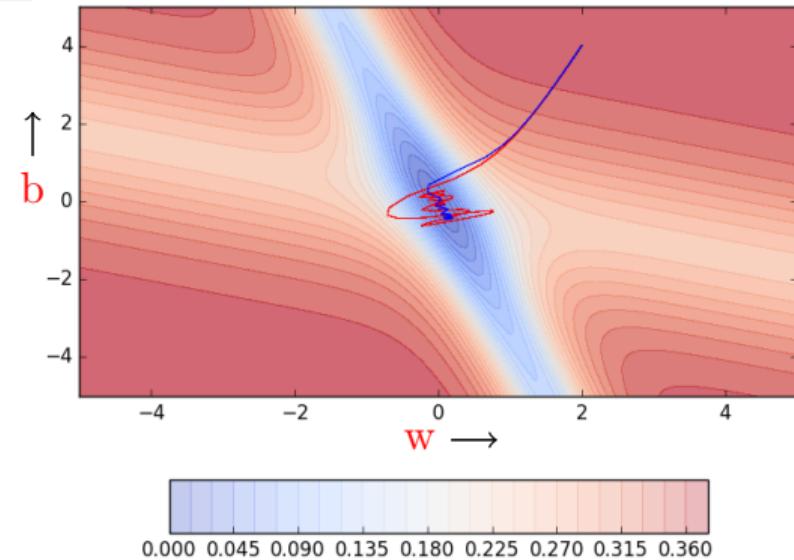


```

def do_nesterov_accelerated_gradient_descent() :
    w, b, eta = init_w, init_b , 1.0
    prev_v_w, prev_v_b, gamma = 0, 0, 0.9
    for i in range(max_epochs) :
        dw, db = 0, 0
        #do partial updates
        v_w = gamma * prev_v_w
        v_b = gamma * prev_v_b
        for x,y in zip(X, Y) :
            #calculate gradients after partial update
            dw += grad_w(w - v_w, b - v_b, x, y)
            db += grad_b(w - v_w, b - v_b, x, y)

        #now do the full update
        v_w = gamma * prev_v_w + eta * dw
        v_b = gamma * prev_v_b + eta * db
        w = w - v_w
        b = b - v_b
        prev_v_w = v_w
        prev_v_b = v_b

```

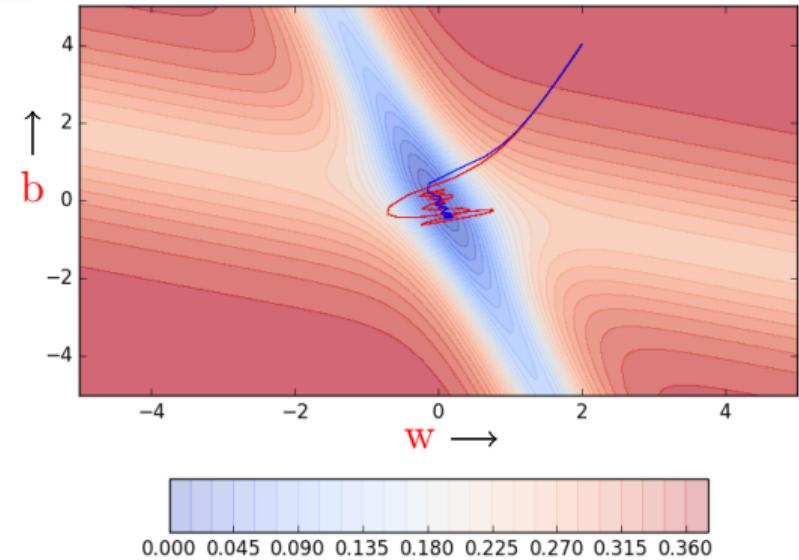


```

def do_nesterov_accelerated_gradient_descent() :
    w, b, eta = init_w, init_b , 1.0
    prev_v_w, prev_v_b, gamma = 0, 0, 0.9
    for i in range(max_epochs) :
        dw, db = 0, 0
        #do partial updates
        v_w = gamma * prev_v_w
        v_b = gamma * prev_v_b
        for x,y in zip(X, Y) :
            #calculate gradients after partial update
            dw += grad_w(w - v_w, b - v_b, x, y)
            db += grad_b(w - v_w, b - v_b, x, y)

        #now do the full update
        v_w = gamma * prev_v_w + eta * dw
        v_b = gamma * prev_v_b + eta * db
        w = w - v_w
        b = b - v_b
        prev_v_w = v_w
        prev_v_b = v_b

```

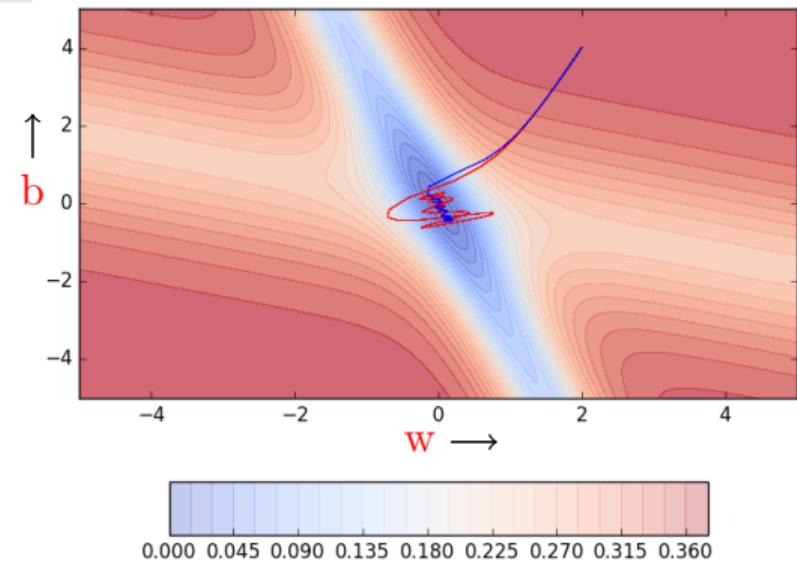


```

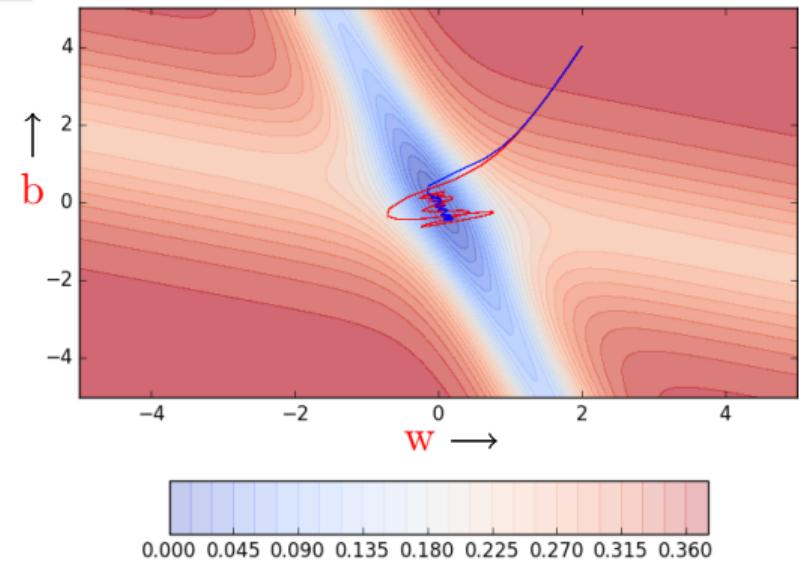
def do_nesterov_accelerated_gradient_descent() :
    w, b, eta = init_w, init_b , 1.0
    prev_v_w, prev_v_b, gamma = 0, 0, 0.9
    for i in range(max_epochs) :
        dw, db = 0, 0
        #do partial updates
        v_w = gamma * prev_v_w
        v_b = gamma * prev_v_b
        for x,y in zip(X, Y) :
            #calculate gradients after partial update
            dw += grad_w(w - v_w, b - v_b, x, y)
            db += grad_b(w - v_w, b - v_b, x, y)

        #now do the full update
        v_w = gamma * prev_v_w + eta * dw
        v_b = gamma * prev_v_b + eta * db
        w = w - v_w
        b = b - v_b
        prev_v_w = v_w
        prev_v_b = v_b

```

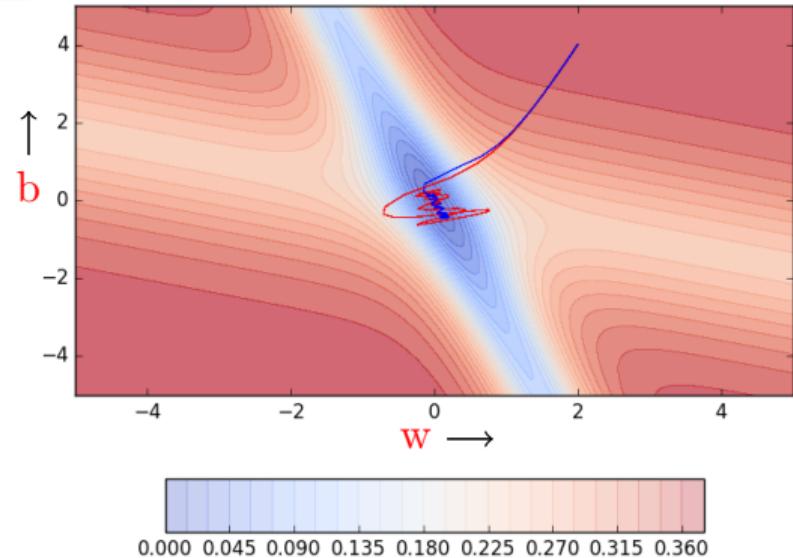


```
def do_nesterov_accelerated_gradient_descent() :  
    w, b, eta = init_w, init_b , 1.0  
    prev_v_w, prev_v_b, gamma = 0, 0, 0.9  
    for i in range(max_epochs) :  
        dw, db = 0, 0  
        #do partial updates  
        v_w = gamma * prev_v_w  
        v_b = gamma * prev_v_b  
        for x,y in zip(X, Y) :  
            #calculate gradients after partial update  
            dw += grad_w(w - v_w, b - v_b, x, y)  
            db += grad_b(w - v_w, b - v_b, x, y)  
  
        #now do the full update  
        v_w = gamma * prev_v_w + eta * dw  
        v_b = gamma * prev_v_b + eta * db  
        w = w - v_w  
        b = b - v_b  
        prev_v_w = v_w  
        prev_v_b = v_b
```



```
def do_nesterov_accelerated_gradient_descent() :
    w, b, eta = init_w, init_b , 1.0
    prev_v_w, prev_v_b, gamma = 0, 0, 0.9
    for i in range(max_epochs) :
        dw, db = 0, 0
        #do partial updates
        v_w = gamma * prev_v_w
        v_b = gamma * prev_v_b
        for x,y in zip(X, Y) :
            #calculate gradients after partial update
            dw += grad_w(w - v_w, b - v_b, x, y)
            db += grad_b(w - v_w, b - v_b, x, y)

        #now do the full update
        v_w = gamma * prev_v_w + eta * dw
        v_b = gamma * prev_v_b + eta * db
        w = w - v_w
        b = b - v_b
        prev_v_w = v_w
        prev_v_b = v_b
```

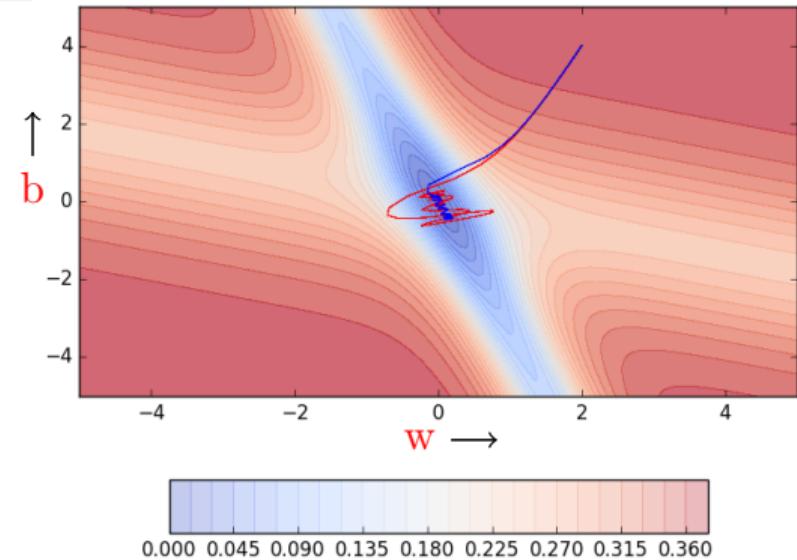


```

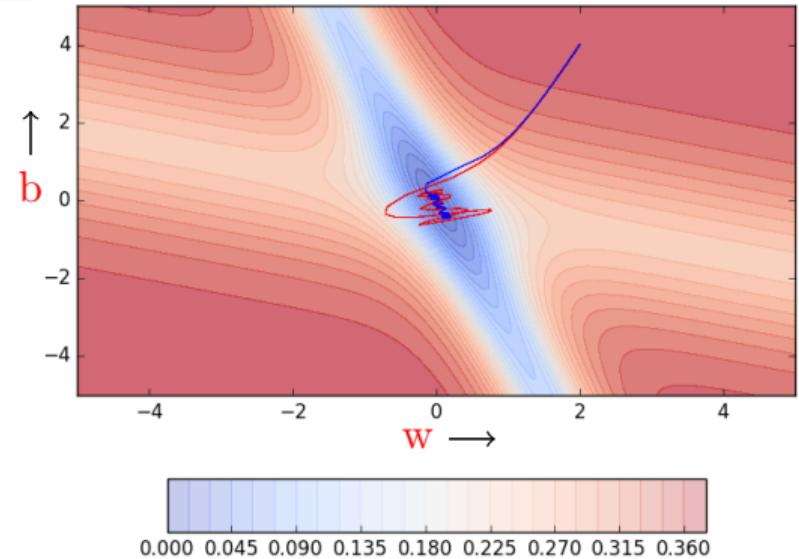
def do_nesterov_accelerated_gradient_descent() :
    w, b, eta = init_w, init_b , 1.0
    prev_v_w, prev_v_b, gamma = 0, 0, 0.9
    for i in range(max_epochs) :
        dw, db = 0, 0
        #do partial updates
        v_w = gamma * prev_v_w
        v_b = gamma * prev_v_b
        for x,y in zip(X, Y) :
            #calculate gradients after partial update
            dw += grad_w(w - v_w, b - v_b, x, y)
            db += grad_b(w - v_w, b - v_b, x, y)

        #now do the full update
        v_w = gamma * prev_v_w + eta * dw
        v_b = gamma * prev_v_b + eta * db
        w = w - v_w
        b = b - v_b
        prev_v_w = v_w
        prev_v_b = v_b

```



```
def do_nesterov_accelerated_gradient_descent() :  
    w, b, eta = init_w, init_b , 1.0  
    prev_v_w, prev_v_b, gamma = 0, 0, 0.9  
    for i in range(max_epochs) :  
        dw, db = 0, 0  
        #do partial updates  
        v_w = gamma * prev_v_w  
        v_b = gamma * prev_v_b  
        for x,y in zip(X, Y) :  
            #calculate gradients after partial update  
            dw += grad_w(w - v_w, b - v_b, x, y)  
            db += grad_b(w - v_w, b - v_b, x, y)  
  
        #now do the full update  
        v_w = gamma * prev_v_w + eta * dw  
        v_b = gamma * prev_v_b + eta * db  
        w = w - v_w  
        b = b - v_b  
        prev_v_w = v_w  
        prev_v_b = v_b
```

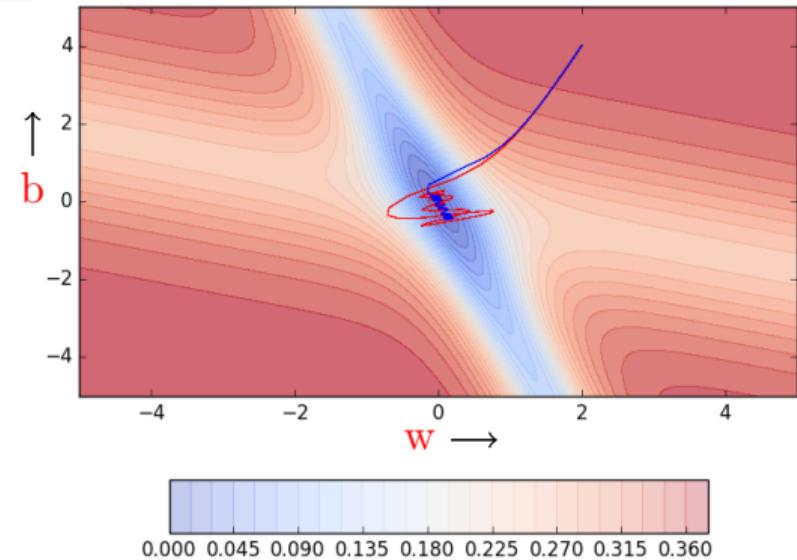


```

def do_nesterov_accelerated_gradient_descent() :
    w, b, eta = init_w, init_b , 1.0
    prev_v_w, prev_v_b, gamma = 0, 0, 0.9
    for i in range(max_epochs) :
        dw, db = 0, 0
        #do partial updates
        v_w = gamma * prev_v_w
        v_b = gamma * prev_v_b
        for x,y in zip(X, Y) :
            #calculate gradients after partial update
            dw += grad_w(w - v_w, b - v_b, x, y)
            db += grad_b(w - v_w, b - v_b, x, y)

        #now do the full update
        v_w = gamma * prev_v_w + eta * dw
        v_b = gamma * prev_v_b + eta * db
        w = w - v_w
        b = b - v_b
        prev_v_w = v_w
        prev_v_b = v_b

```

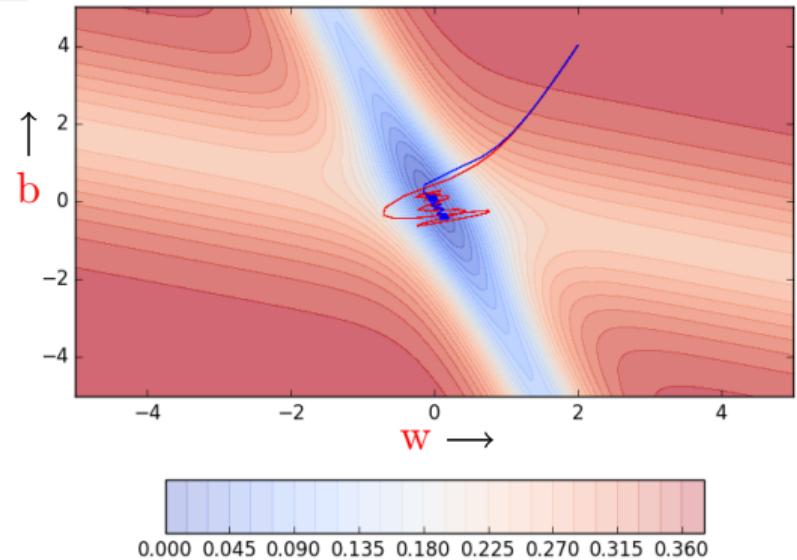


```

def do_nesterov_accelerated_gradient_descent() :
    w, b, eta = init_w, init_b , 1.0
    prev_v_w, prev_v_b, gamma = 0, 0, 0.9
    for i in range(max_epochs) :
        dw, db = 0, 0
        #do partial updates
        v_w = gamma * prev_v_w
        v_b = gamma * prev_v_b
        for x,y in zip(X, Y) :
            #calculate gradients after partial update
            dw += grad_w(w - v_w, b - v_b, x, y)
            db += grad_b(w - v_w, b - v_b, x, y)

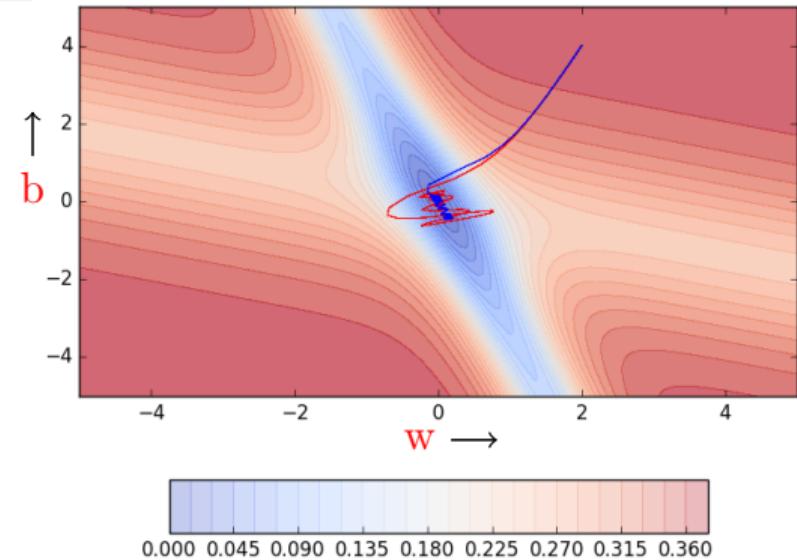
        #now do the full update
        v_w = gamma * prev_v_w + eta * dw
        v_b = gamma * prev_v_b + eta * db
        w = w - v_w
        b = b - v_b
        prev_v_w = v_w
        prev_v_b = v_b

```



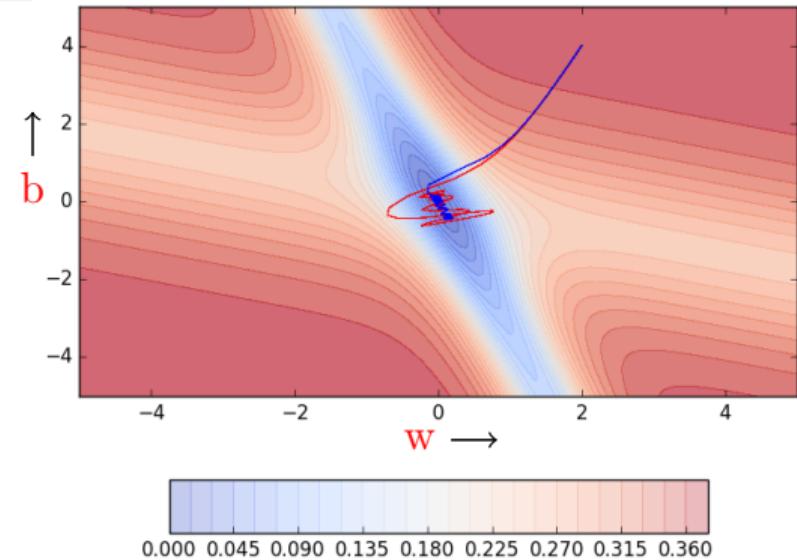
```
def do_nesterov_accelerated_gradient_descent() :
    w, b, eta = init_w, init_b, 1.0
    prev_v_w, prev_v_b, gamma = 0, 0, 0.9
    for i in range(max_epochs) :
        dw, db = 0, 0
        #do partial updates
        v_w = gamma * prev_v_w
        v_b = gamma * prev_v_b
        for x,y in zip(X, Y) :
            #calculate gradients after partial update
            dw += grad_w(w - v_w, b - v_b, x, y)
            db += grad_b(w - v_w, b - v_b, x, y)

        #now do the full update
        v_w = gamma * prev_v_w + eta * dw
        v_b = gamma * prev_v_b + eta * db
        w = w - v_w
        b = b - v_b
        prev_v_w = v_w
        prev_v_b = v_b
```



```
def do_nesterov_accelerated_gradient_descent() :
    w, b, eta = init_w, init_b, 1.0
    prev_v_w, prev_v_b, gamma = 0, 0, 0.9
    for i in range(max_epochs) :
        dw, db = 0, 0
        #do partial updates
        v_w = gamma * prev_v_w
        v_b = gamma * prev_v_b
        for x,y in zip(X, Y) :
            #calculate gradients after partial update
            dw += grad_w(w - v_w, b - v_b, x, y)
            db += grad_b(w - v_w, b - v_b, x, y)

        #now do the full update
        v_w = gamma * prev_v_w + eta * dw
        v_b = gamma * prev_v_b + eta * db
        w = w - v_w
        b = b - v_b
        prev_v_w = v_w
        prev_v_b = v_b
```

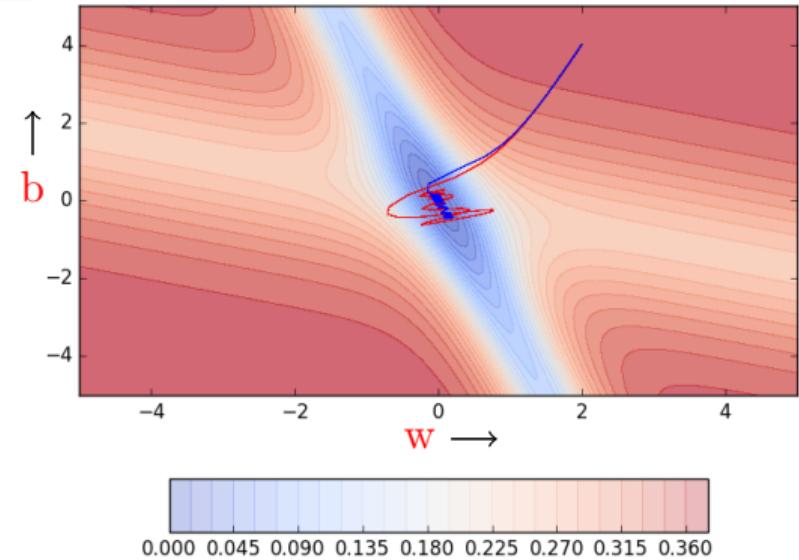


```

def do_nesterov_accelerated_gradient_descent() :
    w, b, eta = init_w, init_b , 1.0
    prev_v_w, prev_v_b, gamma = 0, 0, 0.9
    for i in range(max_epochs) :
        dw, db = 0, 0
        #do partial updates
        v_w = gamma * prev_v_w
        v_b = gamma * prev_v_b
        for x,y in zip(X, Y) :
            #calculate gradients after partial update
            dw += grad_w(w - v_w, b - v_b, x, y)
            db += grad_b(w - v_w, b - v_b, x, y)

        #now do the full update
        v_w = gamma * prev_v_w + eta * dw
        v_b = gamma * prev_v_b + eta * db
        w = w - v_w
        b = b - v_b
        prev_v_w = v_w
        prev_v_b = v_b

```



Observations about NAG

- Looking ahead helps NAG in correcting its course quicker than momentum based gradient descent

Observations about NAG

- Looking ahead helps NAG in correcting its course quicker than momentum based gradient descent
- Hence the oscillations are smaller and the chances of escaping the minima valley also smaller