Adam AdamW

Adam??????

1. □□□□□
   - □□□□□□□□□□□□□ m0=0 □□□□□□□□□□□□□□□□□□□□ v0=0□
   - □□□□□□□□□ η□□□ 0.001□□□□□□□□ β1□□□ 0.9□□□□□□□□□ β2 □□□ 0.999□□□□□□□□ ϵ□□□ 1e-8□
2. □□□□□□ □□□□ t□□□□□□□□□□ θ_{t} □□□ g_{t}
3. □□□□□□□□□□ □□□□□□□□□□□□□□□□□□□□□□□□ m_{t}=β1·m_{t−1}+(1−β1)·g_{t}
4. □□□□□□□□□□□□□□□□ □□□□□□□□□□□□□□□□□□□□□□□□ v_{t}=β2·v_{t−1}+(1−β2)·g_{t}^{2}
5. □□□□□ **Bias Correction**□ □□□□□ mt □ vt □□ 0□□□□□□□ t □□□ m_{t}=m_{t}/(1−β_{1}^{t}) v^{t}=v^{t}/(1−β_{2}t)
6. □□□□ □□□□□□□□□□□□□□□□□□□□

```
class AdamOptimizer:
    def __init__(self, learning_rate=0.01, beta1=0.9, beta2=0.999, epsilon=1e-8):
        self.lr = learning_rate
        self.beta1 = beta1
        self.beta2 = beta2
        self.epsilon = epsilon
        self.m = 0
        self.v = 0
        self.t = 0
    def update(self, params, grads):
        self.t += 1
        for param, grad in zip(params, grads):
            self.m[param] = self.beta1 * self.m[param] + (1 - self.beta1) * grad
            self.v[param] = self.beta2 * self.v[param] + (1 - self.beta2) * grad**2
            m_hat = self.m[param] / (1 - self.beta1**self.t)
            v_hat = self.v[param] / (1 - self.beta2**self.t)
            param = param - self.lr * m_hat / (np.sqrt(v_hat) + self.epsilon)
        return params
```

weight decay

□□□□□□□□□□□□□□□□□□□□□□□□□□□□

```
w = w - lr * w.grad - lr * weight_decay * w
```

L2正则化，表示对权重的衰减比例和权重绝对值的平方成正比

```
final_loss = loss + wd * all_weights.pow(2).sum() / 2
```

Adam在实现 L2正则化的时候是直接修改梯度，随后梯度会参与动量的计算，导致 L2 正则化的梯度就会被抵消掉，导致效果不佳

```
grad_w = grad_w + weight_decay * w
```

AdamW

AdamW 的实现方法与传统的 L2 正则化不同，它直接对权重进行衰减，而不是将其加到损失函数上（改变 **grad_w**）

```
grad_w = grad_w
# update step
w = w - learning_rate * grad_w - learning_rate * weight_decay * w
```

Revision #4
Created 2025-03-12 08:09:34 UTC by Colin
Updated 2026-04-29 07:34:10 UTC by Colin