

1. Let's use an MLP with one hidden layer to **classify** handwritten digits as 0, 1, ..., 9. We'll give the model a 28x28-pixel image of a digit, flattened into a 784-dimensional vector. Fill in the blanks below with reasonable values.

```

W1 = torch.randn(size=(_____))
b1 = torch.randn(size=(_____))
W2 = torch.randn(size=(_____))
b2 = torch.randn(size=(_____))

for x_batch, y_batch in training_data:
    # forward pass, starting with linear layer
    linear_out_1 = _____
    activations_1 = _____
    logits = _____
    probs = _____
    loss = cross_entropy_loss(probs, y_batch)

    # backward pass
    loss.backward() # grads now stored in .grad
    for param in [W1, b1, W2, b2]:
        param += _____

# Fill in the shapes below:
x_batch.shape = (N, 784); y_batch.shape = (N,)
linear_out_1.shape = _____
activations_1.shape = _____
logits.shape = _____
probs.shape = _____
loss.shape = _____

W1.grad.shape = _____
b1.grad.shape = _____
W2.grad.shape = _____
b2.grad.shape = _____

```

---

1. Let's use an MLP with one hidden layer to **classify** handwritten digits as 0, 1, ..., 9. We'll give the model a 28x28-pixel image of a digit, flattened into a 784-dimensional vector. Fill in the blanks below with reasonable values.

```

W1 = torch.randn(size=(_____))
b1 = torch.randn(size=(_____))
W2 = torch.randn(size=(_____))
b2 = torch.randn(size=(_____))

for x_batch, y_batch in training_data:
    # forward pass, starting with linear layer
    linear_out_1 = _____
    activations_1 = _____
    logits = _____
    probs = _____
    loss = cross_entropy_loss(probs, y_batch)

    # backward pass
    loss.backward() # grads now stored in .grad
    for param in [W1, b1, W2, b2]:
        param += _____

# Fill in the shapes below:
x_batch.shape = (N, 784); y_batch.shape = (N,)
linear_out_1.shape = _____
activations_1.shape = _____
logits.shape = _____
probs.shape = _____
loss.shape = _____

W1.grad.shape = _____
b1.grad.shape = _____
W2.grad.shape = _____
b2.grad.shape = _____

```

## For next time

Open figure 2.2 of <https://udlbook.github.io/udlfigures/> (linear regression with least-squares loss). Sketch the following plots by hand (accuracy isn't critical, but try to capture the shape including curvature). Use the initial values (intercept = 1.20, slope = 0.2) for each plot while you sweep the other parameter. *If you have time, repeat this with  $w_2$  from the ReLU interactive notebook*

- a. Plot **intercept** on the x-axis and **loss** on the y-axis.      b. Plot **slope** on the x-axis and **loss** on the y-axis.

2. On both plots above, mark the points corresponding to intercept = 1.20 and slope = 0.2. Sketch the tangent lines at those points and use them to compute estimates of  $\frac{\delta \text{ loss}}{\delta \text{ intercept}}$  and  $\frac{\delta \text{ loss}}{\delta \text{ slope}}$ . Write down your estimates below. Be prepared to discuss next class what this tells you about how you might adjust the parameters to reduce the loss.

## For next time

Open figure 2.2 of <https://udlbook.github.io/udlfigures/> (linear regression with least-squares loss). Sketch the following plots by hand (accuracy isn't critical, but try to capture the shape including curvature). Use the initial values (intercept = 1.20, slope = 0.2) for each plot while you sweep the other parameter. *If you have time, repeat this with  $w_2$  from the ReLU interactive notebook*

- a. Plot **intercept** on the x-axis and **loss** on the y-axis.      b. Plot **slope** on the x-axis and **loss** on the y-axis.

2. On both plots above, mark the points corresponding to intercept = 1.20 and slope = 0.2. Sketch the tangent lines at those points and use them to compute estimates of  $\frac{\delta \text{ loss}}{\delta \text{ intercept}}$  and  $\frac{\delta \text{ loss}}{\delta \text{ slope}}$ . Write down your estimates below. Be prepared to discuss next class what this tells you about how you might adjust the parameters to reduce the loss.