

- We're classifying images as one of 5 types of flowers: [daisy, dandelion, roses, sunflowers, tulips]. A feature extractor gives us 1280 numbers for each image, now we need to make the final predictions using a linear classifier parameterized by weights W and biases b . For a batch of 16 images:
 - X .shape = _____
 - W .shape = _____
 - b .shape = _____
 - output = _____, output.shape = _____
 - For one image, if our model outputs scores [1.0, 1.1, -1.0, 0.5, 0.0] for the 5 flower types, what probabilities would these scores correspond to? (you'll need a calculator/computer, use base e).
 - If the true label for this image is "dandelion", what's the model's **accuracy** and **loss** for this image?
 - Write two different prediction vectors that would have the same accuracy as the one in part 2, but a lower loss: first, one vector that changes **only** the score for "dandelion", and second, one vector that does *not* change the score for "dandelion".
-

- We're classifying images as one of 5 types of flowers: [daisy, dandelion, roses, sunflowers, tulips]. A feature extractor gives us 1280 numbers for each image, now we need to make the final predictions using a linear classifier parameterized by weights W and biases b . For a batch of 16 images:
 - X .shape = _____
 - W .shape = _____
 - b .shape = _____
 - output = _____, output.shape = _____
- For one image, if our model outputs scores [1.0, 1.1, -1.0, 0.5, 0.0] for the 5 flower types, what probabilities would these scores correspond to? (you'll need a calculator/computer, use base e).
- If the true label for this image is "dandelion", what's the model's **accuracy** and **loss** for this image?
- Write two different prediction vectors that would have the same accuracy as the one in part 2, but a lower loss: first, one vector that changes **only** the score for "dandelion", and second, one vector that does *not* change the score for "dandelion".

ReLU warmup

$\text{relu}(x) = \max(0, x)$ clips negative values to 0. Positive values are unchanged.

1. $\text{relu}([-1, 0, 2]) =$ _____
2. Find two possible vectors x such that $\text{relu}(x) = [1, 0, 0]$. Make them interestingly different.
 - a. $x =$ _____
 - b. $x =$ _____

Check-In

Before you leave, pick a couple of these questions to react to:

- What was the most important concept from today for you?
 - What was the muddiest concept today?
 - How does what we did today connect with what you've learned before?
 - What would you like to review or clarify next time we meet?
 - What are you curious, hopeful, or excited about?
-

ReLU warmup

$\text{relu}(x) = \max(0, x)$ clips negative values to 0. Positive values are unchanged.

1. $\text{relu}([-1, 0, 2]) =$ _____
2. Find two possible vectors x such that $\text{relu}(x) = [1, 0, 0]$. Make them interestingly different.
 - a. $x =$ _____
 - b. $x =$ _____

Check-In

Before you leave, pick a couple of these questions to react to:

- What was the most important concept from today for you?
- What was the muddiest concept today?
- How does what we did today connect with what you've learned before?
- What would you like to review or clarify next time we meet?
- What are you curious, hopeful, or excited about?