

## Lab 1 review

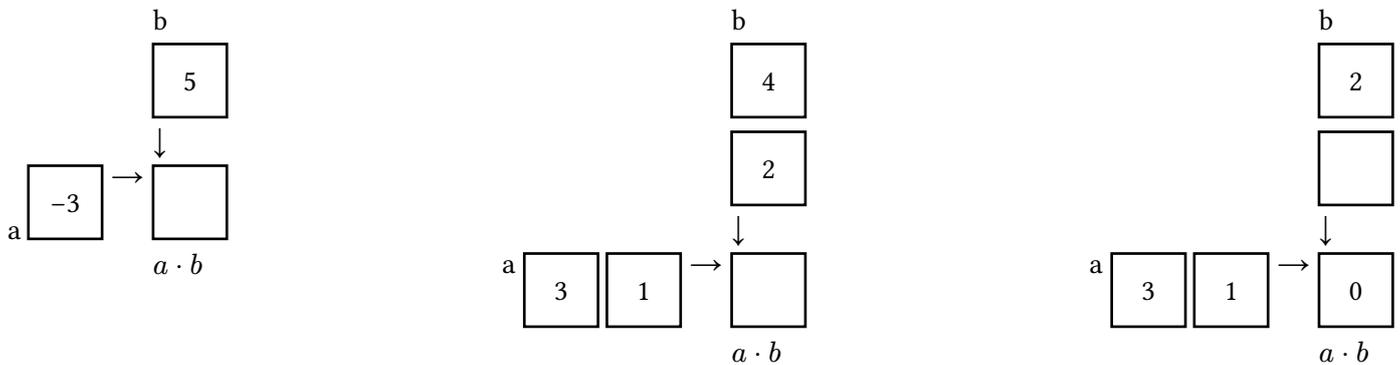
Partner up. Open up your Image Classifier notebook to answer the following questions. (You may need to run the code again.)

- How many images were in the training set? \_\_\_\_\_
- How many images were in the validation set? \_\_\_\_\_
- How many classes were we classifying? \_\_\_\_\_
- How many classes could each image belong to? \_\_\_\_\_
- What did the output of the model look like for a single image? \_\_\_\_\_
- Look at the last cell, for trying on your own image. Notice that it calls `model(input_tensor)`. In a separate cell, run `imshow(input_tensor[0].cpu())`. What do you see? \_\_\_\_\_
- In a separate cell, run `input_tensor.shape`. What do you get? Can you explain any of those numbers?  
\_\_\_\_\_

Now, draw a diagram showing the *inputs* and *outputs* of the model (treat the model itself as a black box), focusing on what are the numbers that go into the model and what are the numbers that come out.

## Dot products

To compute the dot product of two vectors *a* and *b*, we multiply corresponding entries and sum. Try it:



Now make up your own 3 examples, where the results are a positive number, a negative number, and zero.

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?