

WEEKLY PARTICIPATION 10

We saw that the output of (unpadded) convolutional layers is smaller than their input. We can force the output of convolutional layers to have the same size as the input by padding the input image with enough zeros in the x and y directions, or by repeating the bordering pixel values. The former is done in Keras by using the option `padding="same"` in `Conv2D` layers.

If the input image is $d_1 \times d_2$ and the kernel size is $k_1 \times k_2$, how wide are the borders p_x, p_y that Keras adds to ensure that the output size is the same as in the input size?

