


Convolved network


9. (8 points) We will explore how convolutional neural networks operate by designing one. Our objective is to be able to locate the pattern



in an image. Throughout this problem, treat dark squares as having value +1 and light squares as having value -1.

- (a) Consider the image that would result from convolving the image below with a filter that is the same as the pattern above. (Use our definition of convolution, in which we slide the filter over the image and compute the dot product.) Assume that the edges are padded with -1 and that use a stride of 1.

Indicate which pixel in the resulting image will have the maximum value by writing the resulting pixel value in the appropriate cell of the image on the right below.



3		1		1
5			9	
3		1		1
1				
3	3	1		

- (b) In order to detect this pattern, we would create a network that has
- a convolutional layer with a single filter, corresponding to the desired pattern,
 - a max-pooling layer with input size equal to the image size, and finally
 - a single ReLU unit.

Provide a value for the offset W_o on the input to the ReLU that, for any image, would guarantee the output of the ReLU is positive if and only if there is a perfect instance of this pattern in the image.

Solution: -8

A perfect score is 9. The next best match would be 8 correct and 1 wrong, which would total to 7. Any value between 7 and 9 would be correct here.


Name: _____

- (c) Kanye Volution thinks that instead of having this single convolution layer with a single filter matching the whole desired pattern, it would be better to start with a convolutional layer with four smaller filters, shown below:



The following images are the result of convolving the original image with these 4 simple filters and running through a ReLU. Black squares have value +1, grey squares have value +0.5, and the rest have value 0.

It is slightly unusual to have 2 x 2 filters (usually they have odd dimension). When we apply them, we place the upper-left pixel of the filter on top of the image pixel whose value we are computing.



The next layer of Kanye's network now takes an input of depth 4 and applies a single 2 x 2 x 4 filter. Specify a filter on the output of the simple filters that will generate an image with a high value at the pixel located at the upper left corner of the pattern and lower values elsewhere. Fill weight values (either +1 or -1) into the squares below.

