CS1114 Section 3/14/2012 Exercises: Convolution

1 Convolution

Convolution is an important operation in signal and image processing. Convolution operates on two signals (in 1D) or two images (in 2D): you can think of one as the "input" signal (or image), and the other (called the kernel) as a "filter" on the input image, producing an output image (so convolution takes two images as input and produces a third as output). Convolution is an incredibly important concept in many areas of math and engineering (including computer vision, as we'll see later).

Definition. Let's start with 1D convolution (a 1D "image," is also known as a signal, and can be represented by a regular 1D vector in Matlab). Let's call our input vector f and our kernel g, and say that f has length n, and g has length m. The convolution $f * g$ of f and g is defined as:

$$(f * g)(i) = \sum_{j=1}^{m} f(j) \cdot g(i - j)$$

One way to think of this operation is that we're sliding the kernel over the input image. For each position of the kernel, we multiply the overlapping values of the kernel and image together, and add up the results. This sum of products will be the value of the output image at the point in the input image where the kernel is centered. Let's look at a simple example. Suppose our input 1D image is:

$$f = \boxed{\begin{array}{|c|c|c|c|c|c|c|} \hline 10 & 50 & 60 & 10 & 20 & 40 & 30 \\ \hline \end{array}}$$

and our kernel is:

$$g = \boxed{\begin{array}{|c|c|c|} \hline 1/3 & 1/3 & 1/3 \\ \hline \end{array}}$$

Let's call the output image h. What is the value of $h(3)$? To compute this, we slide the kernel so that it is centered around $f(3)$:

10	50	60	10	20	40	30
	1/3	1/3	1/3			

For now, we'll assume that the value of the input image and kernel is 0 everywhere outside the boundary, so we could rewrite this as:

10	50	60	10	20	30	40
0	1/3	1/3	1/3	0	0	0

We now multiply the corresponding (lined-up) values of f and g, then add up the products. Most of these products will be 0, except for at the three non-zero entries of the kernel. So the product is:

$$\frac{1}{3}50 + \frac{1}{3}60 + \frac{1}{3}10 = \frac{50}{3} + \frac{60}{3} + \frac{10}{3} = \frac{120}{3} = 40$$

Thus, $h(3) = 40$. From the above, it should be clear that what this kernel is doing is computing a windowed average of the image, i.e., replacing each entry with the average of that entry and its left and right neighbor. Using this intuition, we can compute the other values of h as well:

$$h = \boxed{\begin{array}{|c|c|c|c|c|c|c|} 20 & 40 & 40 & 30 & 20 & 30 & 23.333 \end{array}}$$

(Remember that we're assuming all entries outside the image boundaries are 0—this is important because when we slide the kernel to the edges of the image, the kernel will spill out over the image boundaries.)

We can also apply convolution in 2D—our images and kernels are now 2D functions (or matrices in Matlab; we'll stick with intensity images for now, and leave color for another time). For 2D convolution, just as before, we slide the kernel over each pixel of the image, multiply the corresponding entries of the input image and kernel, and add them up—the result is the new value of the image. Let's see the result of convolving an image with some example kernels. We'll use this image as our input: One very simple kernel is just a single
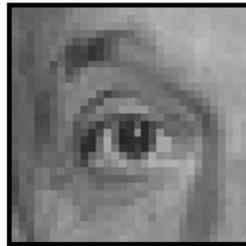
pixel with a value of 1. This is the identity kernel, and leaves the image unchanged:

Original   Identical image
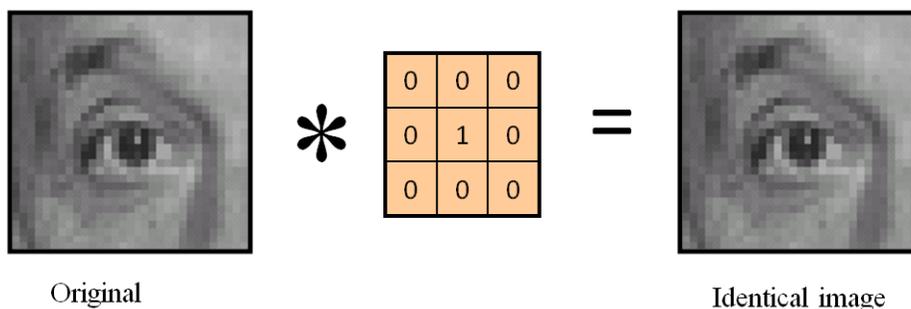
Another useful 2D kernel is an averaging or mean filter. Here's what convolving the image with a 3×3 mean filter looks like:

Original   Blur (with a mean filter)
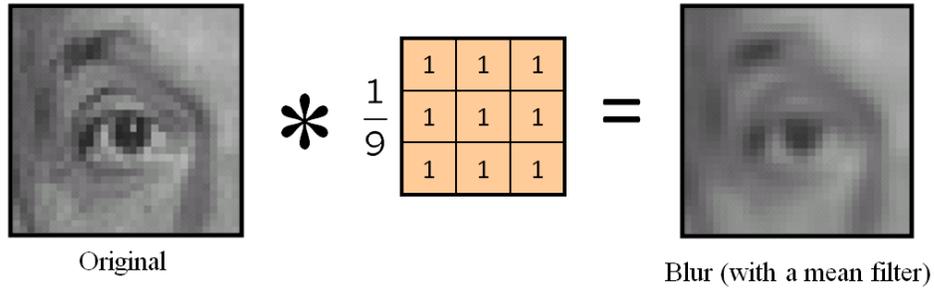
In this case, the mean kernel is just a 3×3 image where every entry is $1/9$ (why is it a good thing for the entries to sum to one?). This kernel has the effect of *blurring* the images—an image where the intensities are smoothed out.

What happens if we convolve the image with an identity kernel but where the one is shifted by one pixel? Then then output image is also shifted by one pixel:

Original   Shifted left
By 1 pixel
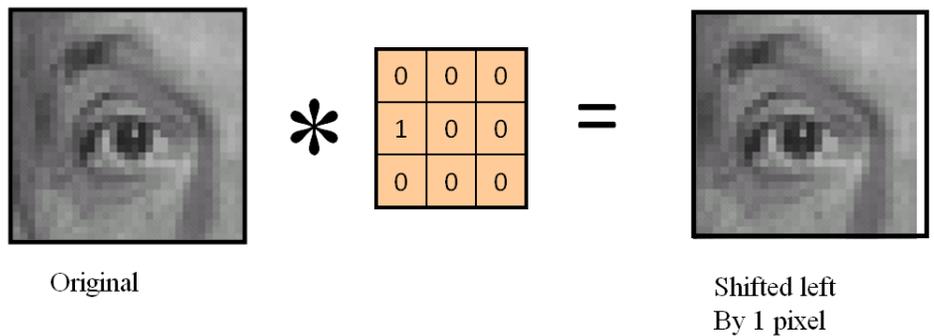
We can create filters that do many other things, as well. An identity kernel that has a 2 in the center instead of a 1 multiplies all the image intensities by 2. Something interesting happens if we combine such a kernel with one that *subtracts* the average of the intensities in a 3×3 window: (Note that some of the entries in the resulting kernel will be negative.) This

Original   Sharpening filter
(accentuates edges)
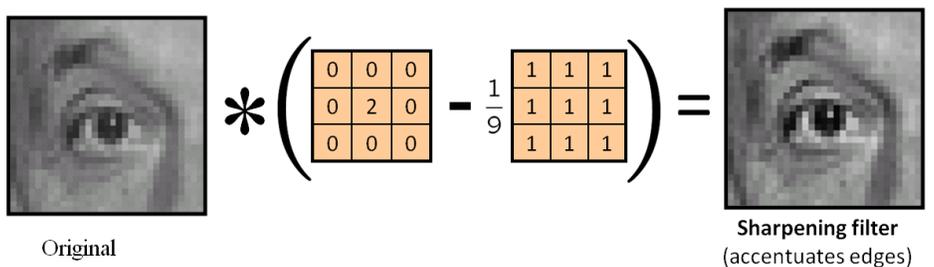
results in an image where edges are accentuated; this is known as a *sharpening* operation, which is kind of like the opposite of blurring. This and many other kernels are built into image editing software such as Photoshop.

3

2   Exercises

Now it's your turn to play with image convolution. In Matlab, 2D convolution can be done with the `conv2` function. Copy over the files in the `/courses/cs1114/sections/conv/` directory, and open up Matlab. There are a couple of images that you should have copied over, (`wires.png`) and (`otter.png`). Open up one of the images (and convert to a matrix of doubles, as the Matlab convolution routines assume doubles). For instance:

```
>> f = im2double(imread('wires.png'));
```

Now create a kernel as a small matrix. For instance, you might create a 3×3 mean filter as:

```
>> g1 = [ 1 1 1 ; 1 1 1 ; 1 1 1 ] / 9;
```

Try creating `f` and `g1`, and convolving them to form an image `h`:

```
>> h = conv2(f, g1);
>> imshow(h);
```

How does the output image `h` look compared to `f`?
   ⟹Now try to create some new kernels. Try creating `g2`, a filter that multiplies the image intensities by 2, and `g3`, a filter that sharpens the image (you should define `g3` in terms of `g1` and `g2`). Try applying both of these to an image, and take a look at the results.
   ⟹Try out this kernel:

```
>> g4 = [ -1 -1 0 ; -1 3 0 ; 0 0 0 ];
```

Try applying this to an image (the otter image works better in this case). When you use imshow to view the result, the image will look mostly black—this is because the entries of this kernel sum to 0, instead of 1. To fix this, add 0.5 to the resulting image, e.g.

```
>> imshow(conv2(f, g4) + 0.5);
```

This is known as an *emboss* kernel.
   ⟹Come up with your own kernel, and see what happens when you apply it to an image.

Let's consider the blur kernel `g1` again. What if we want to blur more? We could blur the image with `g1` once, then blur the result again with `g1`. This will give a "twice-blurred" image. This is equivalent to the operation: $((f * g_1) * g_1)$. Try this out, and see what the result looks like.

Alternatively, we could first filter `g1` with *itself*, then blur `f` with the result. This works because convolution is associative, so $((f * g_1) * g_1) = (f * (g_1 * g_1))$. Create a kernel `g11` that is `g1` filtered with itself, then convolve `f` with this new filter. We could blur more and more by convolving `g1` with itself multiple times. You can visualize the resulting kernels using `imagesc` (e.g., `imagesc(g11)`). Try this out.

3   Object recognition

Notice that there are two other images in the directory, `scene.png` and `butterfly.png`. Read in these two images (remembering to convert to double using `im2double`). You'll notice that the butterfly image appears in the scene image in two places. Suppose we don't know where the butterfly is in the image, and we want to find it—this is a simple version of the *object recognition* problem. It turns out that we can use convolution to solve this. Actually, we need a very similar operation called *normalized cross correlation* (the `normxcorr2` function in Matlab). The differences between this and convolution are subtle— you can think of it as a modified convolution, though the details are not too important right now.

Try running `normxcorr2` on the butterfly and scene images (the butterfly image will be the first argument to `normxcorr2` in this case), and look at the results using `imagesc`. You should see two bright spots corresponding to the two large butterflies in the scene image. It turns out that convolving one image with another, "real" subimage, will fire strongly where the two images are similar, and less strongly where they are different. So to detect a particular object in a big image, one simple way is to convolve the big image with the object image.