

In the previous chapter we established that a digital image is simply an image that is made up of numbers. How exactly is that possible?

Here is one way: take every image that has ever existed in the world, put them all in a big stack, and assign each image an integer corresponding to its location in the stack. So the first picture in the pile would be #1, the second #2, and so on. There! Every image has a number and we've made a world full of digital images.

There are obvious problems with this approach. There is no way anyone could possibly collect every image that has ever existed, and even if they could, the complete set could not contain every conceivable image. In other words, the finite stack we assemble to "digitize" (turn into numbers) can't capture the infinite variety of all possible images.

There is another problem. Let's say I am in San Francisco and I want to transfer an image to you using this new digital representation. Perhaps it's image number 663. If I tell you to use image 663, you are required to have the entire stack of all the world's images just like I have! In other words, the number we have assigned to the image has nothing to do with the contents of the image. Ideally, our digital representation would somehow capture the visual aspects of the image in the numbers we choose.

This thought experiment establishes two very clear criteria for our system. First, we want to be able to represent any image (even ones we have never seen), and second, we want the numbers to somehow encode the visual properties of the image.

(footnote: Recall the other definition of "digital," namely, representing something with discrete units.)

Have you ever seen a painting by impressionist pointillist painter Georges Seurat? Seurat is known in particular for building up complex naturalistic images using small dots (points) of paint. Standing right up against a Seurat painting you only see the individual dots of color, standing further away you start to see people, landscapes, patterns in clothes, leaves, flowers. Somehow those small dots fuse in our eye until we cannot see them any more.

(insert Seurat pointillism image here)

(insert mosaic tile image here)

(insert image from sporting event where they stick colored cards under every seat)

All of these examples work because discrete units of color – when chosen carefully and viewed from a great enough distance – appear as a constant enough field of color that we see the whole instead of the component parts.

Lean this book against a wall and walk far away; what does the cover look like?



So let's take the following approach, inspired by pointillism, by ceramic tiles, and by football games. Let's represent an image as a rectangular array of numbers where the numbers themselves encode brightness. For the time being, let's work only with black and white images and let's assign the number 0 to "black" and the number 100 to "white." We can then use all values between 0 and 100 to represent the infinite shades of grey between black and white.

With this definition, here is an example digital image:

```
100  0
   0 100
```

Now I know it doesn't look like much, but this is a big step. We have just created a digital version of an image of a checkerboard. If we want to make it a full-size, regulation checkerboard with eight rows and eight columns, our image would look like this:

```
100  0  100  0  100  0  100  0
0   100 0  100 0  100 0  100
100  0  100  0  100  0  100  0
0   100 0  100 0  100 0  100
100  0  100  0  100  0  100  0
0   100 0  100 0  100 0  100
100  0  100  0  100  0  100  0
0   100 0  100 0  100 0  100
```

All we've really done is encoded brightness with a number between 0 and 100 and used a regularly-spaced grid to pinpoint location.

We can obviously use this simple system to create rather complicated images. Here, for instance, is a digital gradient image, smoothly varying from black to white:

```
0    20   40   60   80   100
0    20   40   60   80   100
...
```

The Pixel

The more numbers we use to represent our digital images the more visual variety we can capture. Each of these numbers is essentially a small, very simple picture, namely, one with only one color and a square shape. This is one of the fundamental assumptions behind digital imaging, namely, that a large enough number of these square, single-valued picture elements (pixels) can successfully recreate any image we want.

Pixels are the fundamental units of digital images. Clearly, an image made up of only one pixel can only have one brightness value. Our checkerboard example used an eight-by-eight grid of numbers resulting in 64 total pixels. There is no conceptual limit on the number of pixels we can use to make digital images, and it seems clear that more pixels allows for greater visual variety. But don't mistake: having a large number of pixels doesn't guarantee visual variety! After all, what's stopping you from making them all the same value?

Spatial Resolution

The total number of pixels used to make up a digital image is known as **the spatial resolution** of that image. People seem happy to refer to spatial resolutions as width by height pairs as well as the total number in an image. In other words, our checkerboard image has a spatial resolution of 8 by 8 (8 x 8) pixels, but it also has a spatial resolution of 64 pixels. Both are right answers, but I prefer the first way of expressing spatial resolutions. It more clearly describes the shape of the digital image. For instance, a 64 pixel image could be 2 x 32, 4 x 16, or even 1 x 64 pixels in spatial extent.

Width is generally the first number listed, followed by height. So a 200 x 100 pixel image is 200 pixels wide by 100 pixels high, for a total of 20,000 pixels. An image's width divided by its height is known as the **aspect ratio** of an image. Images with aspect ratios of 1 are square. Images with aspect ratios greater than 1 are wider than they are high. For instance, American movies are generally shot with aspect ratios of 1.85 or 2.4 (commonly known as 2.35). The higher the aspect ratio, the wider the image. Television in the United States has an aspect ratio of 1.33 or 4:3. The new HDTV standard promises 16:9 aspect ratios, which if you do the math are quite close to the 1.85 we're used to from Hollywood.