Lecture 4 – Images, colors, and design

Images are the central to the visual design of web pages, and no course on the web can be complete without a discussion of images. This lecture gives you a brief introduction into colors, images used by web pages, and web page design.

Colors

Light is electromagnetic radiation whose wavelength (and frequency) is within a certain range: from just below 400 nanometers for violet to slightly above 700 nanometers for red[1]. The colors of the rainbow (red, orange, yellow, green, blue, (indigo,) and violet) correspond to the wavelengths that make up the visible portion of the electromagnetic spectrum. Ambient light is a mix that contains photons with various wavelengths from the visible spectrum. Objects reflect some wavelengths and absorb others. White objects reflect all wavelengths. Black objects absorb all wavelengths (hence in the dark everything is black because objects have no light they can reflect). Colored objects reflect some wavelengths and absorb others. Thus a red object will reflect red light and absorb other wavelengths. The light that reaches the eye is typically an unbalanced mix of the wavelengths of the visible spectrum and the proportion of various wavelengths determines the color the observer perceives. The light that contains a single wavelength is called monochromatic light.

Digital representation of color does not faithfully record the exact mix of wavelengths, but it focuses instead on recording the impression the light mix makes on the human eye. The human eye has three types of color-sensitive cells called cone cells[2]. Each type of cell is most sensitive around a typical wavelength with its sensitivity degrading as the wavelengths are further away. Some wavelengths elicit some level of response from two or even all three of the types of color receptor cells. Computer displays emit a mix of three monochromatic wavelengths: red, green and blue. By controlling the intensity of the light emitted in these three wavelengths, the monitor can simulate the effect of any color on the human eye. 256 levels of intensity for each of the three wavelengths are enough to make smaller distinctions between colors than the human eye can perceive. Thus each color can be expressed as a vector of 3 numbers between 0 and 255 representing the intensity of the red, green and blue. This vector is often encoded as a six digit hexadecimal number with two digits for each color. Thus `ff0000` is red, `00ff00` is green, `0000ff` is blue, `000000` is black and `ffffff` is white.

The RGB color space described above is not the only one used. HSL which stands for hue, saturation, and lightness characterizes colors differently. The hue captures which

[1] This range is defined by what the human eye can see. There are animals that can sense ultraviolet light (wavelengths shorter than violet), and there are also animals that can sense infrared light (wavelengths longer than red).

[2] Many animals (e.g. dogs) see a single type of color, some (e.g. some monkeys) see 2, some (e.g. primates) see 3 like humans, some see 4 or 5 (e.g. some birds).

color dominates, the saturation captures how pure the color is (high values indicate a "pure" color, low values indicate that the color is mixed with a lot of grey), and the lightness indicates how light or dark the color is. The CMY or the CMYK color spaces are used for printing (C stands for cyan, Y for yellow, M for magenta, and K for black). The reason is that the physics of colors in printing is different from the physics for active displays such as LCDs and CRTs. Whereas computer displays start with no light (black) and add more and more light, printing starts with a white sheet that reflects all light and the role of the ink is to block certain wavelengths. Thus for paper we use a subtractive color model and for computer screens an additive color model. In principle three colors are enough but printers also use black ink because black appears frequently in documents and because combining cyan, magenta, and yellow does not give a nice black but a dark brown.

Images

Digital images have two important forms. One form is that of a sequence of bytes stored in a file. Another form is that displayed on the screen: a rectangular matrix of *pixels* (short for picture elements) each with its own RGB color information associated with it. This second form requires significantly more memory, but it is used by the rendering algorithms that generate the pixels displayed in the browser window. The different image formats in use apply various compression and encoding techniques to produce the much more compact on-disk representation of the image. The three most common image formats on the web are JPEG (targeted at encoding photographs), GIF (an early format targeted computer images developed by a company called CompuServe that ran an early alternative to the Internet), and PNG (a modern format specifically designed for the web).

For photographs, it is enough to associate an RGB colors with each pixel. But for images used in web pages which are often layered above other page elements, it is useful to also associate transparency information with each pixel. JPEG, which is targeted at representing photographs, does not support transparency at all. GIF uses a single bit per pixel to convey whether the pixel is transparent or not. PNG supports varying degrees of transparency. Each pixel has a fourth byte associated with it (the so-called alpha channel) with a value of 0 indicating that the pixel is entirely opaque and 255 indicating that it is entirely transparent.

GIF and PNG have a few features useful to web pages. Both can store images in an "interlaced" format which helps when images are downloaded over slow links. Many other image formats store the matrix of pixels row by row (or groups of rows together) with the rows in order in the image file. The browser can start rendering the image while it is being downloaded, based on the data that it has already received. For non-interlaced formats the rendering starts with a sliver at the top of the image that grows thicker and thicker until the entire image is shown when the download finishes. With interlaced GIF or PNG pixels are stored out of order so that even after receiving a small portion of the file the browser can produce a complete but blurry version of the image that that gets clearer and clearer as the image downloads. GIF also supports simple animations that work by storing multiple frames and timing information in the image file. Extensions to PNG supporting animations are not universally deployed.

Compression

Image formats combine various techniques to compress the large matrix of pixel information (three bytes of RGB information or 4 bytes if the alpha channel is also present). We discuss here some of these techniques. Most of them allow the browser to reconstruct exactly the information associated with each pixel (they are called "lossless") and they are also used by compression algorithms not specifically targeted at images. Compressing a large sequence of per pixel records is similar to compressing a long text that uses a large alphabet. Think of the information for a given pixel as one out of 2^{24} (or 2^{32}) possible letters. We will use this alternative way of looking at the compression problem in discussing the first two techniques.

It is common to use a different way to represent the letters from the alphabet. Two popular variants of this technique are using a smaller alphabet, and Huffman coding. If only a few of the 2^{24} possible letters occur in the actual image, say less than 256 of them, it is enough to store one byte per pixel and store somewhere which actual color each byte corresponds to. This is the approach used by GIF and it is the reason why GIF images can never have more than 256 colors, but they can have any 256 colors (each image has a color palette of size 256). Huffman coding associates variable length bit strings called codes with letters from the alphabet. Assume that the text we want to compress is the word "lossless". Since there are 4 distinct letters (e,l,o, and s), we can use an alphabet of size 4 (2 bits per letter) and encode the string in 8*2=16 bits. But we can achieve savings by associating shorter codes to more frequently occurring letters. As long as none of the codes is a prefix of another code, there are no ambiguities at decoding. If we associate "0" with s, "10" with l, "110" with e, and "111" with o we can encode the word as "10111001011000" (broken into codes for individual letters as "10 111 0 0 10 110 0 0") using 1*4+2*2+3*1+3*1=14 bits. In the general case the length of the code of a letter is close to the base ½ logarithm of its frequency, and savings are often more significant than for this toy example.

Another common technique is run length encoding. If there are many consecutive positions with the same letter, we can store the letter just once and add a repetition count. For example we can encode AAABBCDDDD as 3A2B1C4D. This technique works particularly well if there are large areas in the picture using the exact same color (very common for faxes where there are only two colors to choose from – black and white).

Lossy compression techniques do not guarantee that the individual pixels of the rendered image will match those in the original, but they ensure that the overall image will be similar. The advantage of such algorithms is that they can achieve better compression than lossless algorithms. JPEG uses lossy compression techniques. With most image editors you can explicitly control the tradeoff between quality and size: you can get a very small file that will render as a blurrier version image, or a larger file that will be crisper. With both lossy and lossless compression techniques, the size of the image file depends on the actual image, not just on the size of the pixel matrix: images with a lot of detail are harder to compress and result in larger files.

Using images in a web page

The `<img />` tag instructs the browser to insert an image into the text of the document. The `src` attribute has the URL of the image and the `alt` tag contains an alternative text that summarizes the image for browsers incapable of graphics. By default the bottom of the image is aligned with the baseline of the text and the amount of spacing between the line the image is on and the one above it is increased to accommodate the image. Various CSS attributes give the author more control about how the image is aligned. Putting an image inside a table cell or a separate div, gives the author even more options for controlling the placement of the image.

The `<img />` tag also has the `width` and `height` attributes that tell the browser how large to render the image. The image file contains the actual width and height of the image, but these HTML attributes can override them causing the browser to stretch or shrink the image. If only one of the attributes is specified the browser maintains the aspect ratio of the image. The `width` and `height` can specified as percentages. If you can avoid it, do not shrink images using these attributes. Having the browser download a large version of the image and then shrink it can increase the time it takes for the page to render; it is better to use an image editor and create a smaller version of the image instead that downloads faster.

Another powerful use for images is as background for various elements of the page (or the entire page) by using the `background-image` CSS attribute. The background image can be rendered once, repeated along the horizontal axes, along the vertical axis, or along both depending on the value of the `background-repeat` CSS attribute which can be set to `no-repeat`, `repeat`, `repeat-x` or `repeat-y`. With `background-position-x` and `background-position-y` you can control where the background image goes. For example the declaration `background-image:url(bar.png)` tells the browser to use bar.png as background and by default it will be repeated both horizontally and vertically. The most common use of background images is not to achieve wallpaper-like effects but to add small visual details such as rounded corners or colors blending into each other in the background. For example this second effect is used by the Yahoo home page for various tabs and headings of various content blocks. The effect relies on the use of a one-pixel-wide background image that is repeated throughout the width of the heading. Since the background image is very small and it can be shared across many headings and tabs in the page, this is a very efficient way of achieving the desired visual effect.

On web site design

When designing a web site one has to make decisions about layout, color schemes, textures, font choices, and various visual elements typically implemented with custom images. The best web designers have not just an understanding of web technologies (HTML, CSS, the problems of cross-browser compatibility) but also extensive domain knowledge about design parameters listed above together with a refined aesthetic sense and an awareness of the history and the recent trends of design. We discuss here just a few basic, simple and relatively uncontroversial design rules, you are encouraged to read one of the many existing books on web design if you want to learn more.

Most good web site designs use a limited number of colors (a small color palette), fonts, and visual elements. Like most design rules, this is not a strict rule either. For example it is quite common to not just use the 2 to 6 main colors from a palette, but also have more or less saturated versions of the main colors and images or backgrounds where two colors from the palette gradually blend into each other. Also photographs naturally have a large number of actual colors (but only some of them are dominant). Re-using fonts and visual elements in different places in the page and the site (e.g. using the same font and colors for tabs in various places on the page, or using rounded corners for all content areas) gives your design unity and makes it easier for the reader to recognize the roles of the various elements on your page.

You rarely design a web page on its own; typically you design an entire site that consists of multiple pages that work together. For the pages to work together as a site it is important for them to have the same “look-and-feel”. Thus the layout, the colors used, the font selection, logos or other identifying visual elements for all pages should be mostly the same. Of course one often uses different layouts for different pages on the same site, but keeping some elements the same (e.g. header or the navigation bar) is in general a good idea.