CSE 126/228F Lecture 6: April 17

JPEG (Part 3)

JPEG Stages

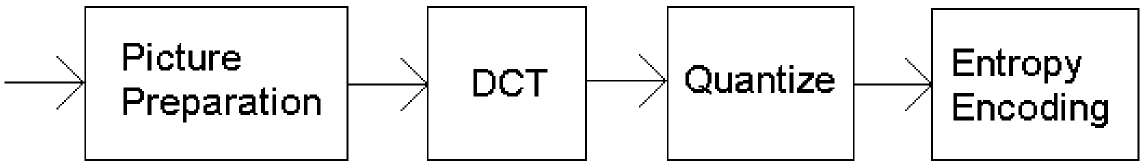

Quantization:

JPEG utilizes a *quantization table* in order to quantize the results of the DCT. Remember that JPEG is a lossy compression scheme. Quantization allows us to define which elements should receive fine quantization and which elements should receive coarse quantization. Those elements that receive a finer quantization will be reconstructed close or exactly the same as the original image, while those elements that receive coarse quantization will not be reconstructed as accurately, if at all. A quantization table is an 8x8 matrix of integers that correspond to the results of the DCT. Each entry in this table is an 8-bit integer. To quantize the data, one merely divides the result of the DCT by the quantization value and keeps the integer portion of the result. Therefore, the higher the integer in the quantization table, the coarser and more compressed the result becomes. This quantization table is defined by the compression application, not by the JPEG standard. A great deal of work goes into creating "good" quantization tables that achieve both good image quality and good compression. Also, because it is not defined by the standard, this quantization table must be stored with the compressed image in order to decompress it.

Entropy Encoding:

Last time we saw that DCT converts the spatial intensity values into the frequency domain and orders the results in a zig-zag sequence. Also, we saw that the there tends to be little or no high-frequency change in an 8x8 block. Here is where we really use this to our advantage. The first step in entropy encoding is to Run-Length Encode the zero values. Because there will be many zeroes in the high-frequency results of the DCT, we can produce long run-length encoded strings of zeroes with the zig-zag sequence. As an example, consider an image that only has non-zero values for the DC and the first two AC coefficients. If we encode row-by-row, we get DC, AC1, {0,6}; AC2, {0,7};{0,8};{0,8};… If we use the zig-zag sequence, however, we get DC,AC1,AC2,{0,61}, a tremendous space savings! Another important point to remember here is that the DC coefficients are encoded as differences from the previous DC coefficient (except for the DC coefficient of the first block, which obviously cannot do this). Because there tends to be little change from block-to-block, this gives us additional space savings. The final step in entropy encoding is to Huffman code the results. The Huffman algorithm can also have a table specified by the JPEG application, or it can create its own by examining the data. In either case, this must also be written with the compressed results in order to decompress.

JPEG Decompression

Decompressing a JPEG image is basically applying the compression process in reverse. First, the Huffman table is read and the image is entropy decoded. Then, the quantization table is read and the result is de-quantized. The process of de-quantization will naturally produce a blocky result. In order to fix this, most applications will interpolate the pixel values to bring them closer to a smooth curve. Next, the IDCT (Inverse-DCT) will be performed on this result, and from this the image will have been reconstructed.

Lossy JPEG

Loss is introduced into the JPEG compression scheme in 3 ways:

1. Quantization
2. Subsampling the U and V components
3. Cosine (used in DCT) is a transcendental function. Mathematically speaking, this means $\cos^{-1}(\cos(x)) \cong x$

Lossless JPEG

Lossless JPEG eliminates the 3 sources of loss shown above. It encodes pixels in a 4x4 block, storing values for 3 of the pixels and encoding one as the difference from another pixel. This means that this algorithm has a maximum compression of around 25%.

Multi-Resolution JPEG

In order to deal with limited bandwidth, there are multiple resolution JPEG formats.

- **Progressive JPEG:** Progressive JPEG subsamples the original image by 4x4, and encodes the result, then subsamples the original image by 2x2 and encodes the pixel differences from the previous image, and finally encodes pixel differences between the original image and the result of the 2x2 subsample.
- **Alternative 1:** It is possible to pass the image through the JPEG stages using only the four most significant bits of the pixel values first, then pass the four least significant bits second.
- **Alternative 2:** It is also possible to pass the DC and 31 AC coefficients through the stages first, then pass the remaining 32 AC coefficients through afterward.

These multi-resolution formats are commonly used on the WWW. These allow a rough representation of the image to be drawn first, then the image is refined, and then the image is refined again into the final version.

Results of JPEG Compression

In general, the following results are seen with JPEG compression:

- Good Quality uses 0.25-0.5 bits per pixel
- 0.5-0.75 bpp gives better quality
- 0.75-1.5 bpp is excellent quality
- 1.5-2 bpp gives a result almost indistinguishable from the original.

Remembering that the original image used 24 bpp, there is a lot of savings here!

MPEG

Background of MPEG

MPEG stands for the Motion, Movie, or Moving Picture Experts Group (all 3 can be used). This is not affiliated with Hollywood; rather it is a group of computer scientists trying to make a standard for digital representation of video.

There are several flavors of MPEG:

- **MPEG-1:** MPEG-1 was the original MPEG standard, designed exclusively for computer use. It allows a 320x240, 30 frame per second video.
- **MPEG-2:** MPEG-2 is a higher-resolution version of MPEG-1 designed for digital television broadcast.
- **MPEG-3:** MPEG-3 was designed for HDTV. However, HDTV is just normal TV with a higher resolution and frame rate, so this standard was folded into MPEG-2 and is no longer used.
- **MPEG-4:** MPEG-4 is a new standard for digital video scheduled for release in November 1998. It is designed for use over low-bit-rate wireless and mobile communication systems. DCT cannot provide the required compression to operate over this type of line, so MPEG-4 does not enforce an encoding method. Instead, it leaves the choice of encoding method up to the designer.
- **MPEG-7:** MPEG-7 is yet another standard for digital video that has barely begun. The focus of MPEG-7 is supposed to be designing a representation for digital video that allows it to be stored and queried by content in a video database.

Why MPEG instead of JPEG for video?

JPEG is an algorithm designed exclusively for digital images. While MPEG does utilize JPEG to some extent, motion video has some additional properties that JPEG does not consider:

- Use and synchronization of multiple media streams, such as Video, Audio, and Closed-Captioning
- Time relationship between frames

Because video is displayed at 30 frames per second, even JPEG cannot give us the compression necessary to make digital video feasible. However, if we can exploit the relationship between successive frames (there will likely be little or no change from frame to frame), we can compress even more. How does MPEG do this? Tune in next time for the answer.