# CS 4620 Homework 1: Memory and Gamma Correction

**Prof. Doug James**

out: Monday 31 August 2009
**due: Monday 7 September 2009**

**Problem 1: Remembering It All** (An exercise in computing image sizes)

Have you ever wished you could go back and see something again? Your first puppy? That perfect sunset? Or how about reliving your first kiss? Given the rapid increase in digital video and photography, it may be only a matter of time before you will be able to record your entire visual and auditory experience. While this may seem excessive, your memory storage costs are really not that significant, at least by 2069 standards. But how much memory is it, really? Let's assume that you record for 60 years (after reaching the age of consent), and that you sleep for 8 hours/day (on average).

Follow tradition by giving data sizes in kilobytes, megabytes, or gigabytes, and data rates in kilobits, megabits, or gigabits per second. Don't forget that these units refer to multiples of 1024, not multiples of 1000. Round your answers to three significant digits—for example, 3.52 MB.

1. **Entry-level setup:** If you just want to get your feet wet, you might consider our *Memory Glasses*$^{TM}$ rig. It uses a single glasses-mounted 720-by-480 camera with 8 bits per color channel (8bpc, for 24bpp) at 30 Hz rates—which will look a little dated in 2069. You only need to purchase your own external storage. *How much memory will you need per year? for 60 years?*

2. **Dream-level setup (not for the squeamish):** Of course, given how common retinal transplants are becoming, I strongly encourage full digital retinal transplants so that *you can see what you saw*$^{TM}$ by recording high-resolution stereo video that closely matches your day-to-day experience. The *DigiEyes*$^{TM}$ setup can produce two 2048-by-1280 stereo images with 10 bits per color channel (10bpc, for 30bpp) at 60 Hz rates. It's almost as good as the real thing. You only need to purchase your own external storage, and sign a health waiver. *How much memory will you need per year? for 60 years?*

**Problem 2: Gamma Correction Blues** (An exercise in pixel degradation)

This question is about gamma correction used for image encoding and display. Assume that an application stores each RGB pixel value in three double-precision floating point numbers, using the convention that pixel values are linearly related to display intensity, with 1.0 representing the maximum value.

Consider an image that is first converted to 8 bit integer format (8bpc, for 24bpp) with nonlinear quantization at gamma 1.8, with viewing flare ignored.

1. What must the application do to convert its pixels into framebuffer pixels for this 8bpc/24bpp framebuffer format? Give a mathematical expression that you could type into the code that

does this processing for each channel. Assume the application does not worry about viewing flare, preferring to believe that pixel value zero is exactly black. Assume the standard functions *round*, *pow*, *min*, and *max* are available.

2. Next consider converting the 24bpp, $\gamma = 1.8$ format to a 24bpp, $\gamma = 2.2$ format, thereby leading to discretization artifacts over direct conversion to the $\gamma = 2.2$ format. For each of the discrete pixel values (0, 1, 2, ... , 255), what are the relative differences in observed intensity of the $\rightarrow 1.8 \rightarrow 2.2$ image compared to the direct $\rightarrow 2.2$ image?

3. If we assume a just noticeable difference of 2%, are any of these "double discretization" artifacts noticeable?

**Hint:** You will probably want to write a short program and plot your results to illustrate the introduced discretization artifacts.