Understanding the Gamma Adjustment of Images

Jakkarin Singnoo, Graham D. Finlayson, School of Computing Sciences, University of East Anglia, Norwich, UK

Abstract

Gamma adjustment is one of the simplest global tone reproduction operators. If an image is too bright or too dark the image can be made pleasing by applying a gamma greater than one (leading to a darker image) or less than one (leading to a brighter image) respectively. In recent theoretical work, the 'optimal' gamma in an information theoretic sense has been derived. The starting point of this paper is to ask the question: in adjusting gamma in images do observers make a similar choice to the information theoretic optimum?

Experimentally, we investigate the user's choice of gamma parameter by conducting double staircase psychophysical experiment on a wide range of monochrome images. Two staircases beginning with bright and dark images with respect to which gamma adjustments are made. The user progressively darkens and lightens the respective images until the staircases converge (we have the same image). The pilot experiment indicates that there is a linear relationship between the maximum entropy of image and the chosen gamma from the experiment: our experiment provides prima facie evidence that image that observers adjust images to bring out information. Moreover, the combination of entropy calculation together with our regression line we effectively provide an automatic algorithm for gamma adjustment.

Finally, we also discuss the relationship between our assumption to the chosen gamma, a modified non-linear masking operator and two versions of CIECAM, and found that all of the operators give the similar trends, but slightly poorer fits, for predicting the gamma parameter. Put another way our work indicates that existing formulae for gamma adjustment can also be related to the concept of entropy maximization.

Introduction

Gamma adjustment in the context of tone reproduction operator provides contrast adjustment. The simplest form of the operator is defined by the following power-law expression:

$$V_{out} = V_{in}^{\gamma} \tag{1}$$

where the input (V_{in}) and output (V_{out}) values are non-negative real values, typically in the normalized range of [0,1]. If γ is larger than 1, the output image will be brightened. In contrast, if γ is smaller than 1 the output will be darken. In this way, gamma adjustment can be thought as the contrast adjustment operator.

Gamma might be adjusted by a user in a package such as Photoshop or, the concern of this paper, automatically using some sort of formula: relating an image statistic to (hopefully) the appropriate gamma. So, how might we find such a gamma adjustment formula? In the standard approach we investigate gamma adjustment in a purely empirical manner. That is, run an experiment where observers adjust gamma and then fit a formula to predict observer results. Indeed, this is a perfectly valid strategy and one that was extensively employed in the development of CIECAM.

In contradistinction to this approach we adopt a purely theoretical formula (derived only from mathematical argument) for choosing gamma. Then we relate our predictions to observer choices to validate our approach. To derive our theoretical gamma, we ask the following question. Assuming that we wish to choose a gamma to best bring out image detail which gamma should we choose. We answered this question using the concept of entropy from information theory.

Shannon Entropy [11] encapsulates the idea of how expensive it is to code data. As an example a language where every character occurs with equal frequency is more expensive to encode than one where a few characters occur more frequently. In English for example we know that the letter 'e' occurs often but that 'z' is rare. We can exploit this information in coding letters with binary strings. We might code 'e' with a small number of bits and 'z' with a larger number. Remarkably, the optimal encoding (assignment of bits) to characters results in an average expected code length proportional to entropy defined as:

$$H(X) = \sum_{i=1}^{n} p(x_i)I(x_i) = -\sum_{i=1}^{n} p(x_i)log_b p(x_i) \tag{2}$$

where $I(x_i)$ denotes information is an estimate of the number of bits that would have to be used to encode a character x_i, with a given probability $p(x_i)$, b is the base of logarithm used.

In the context of an image we can think of each brightness level as a distinct character. When we raise an image to the power of gamma we are mapping one set of characters with an inherent probability distribution to another where the original and gammed image will have different distributions of brightnesses. It follows then that the original image and gammed counterpart will have different entropies.

Importantly, the idea of entropy relates to the conspicuity of detail in images. To understand why entropy relates to visible detail in images we remind the reader about histogram equalization (HE). An image with a given, say mostly bright, distribution of brightnesses is mapped by HE, to a corresponding image with a flat probability distribution by histogram equalization. If we choose two random pixels from the original image they are both likely to be bright and so visually close to one another. After equalization the same pixel pair will, on average, have quite different brightnesses. That is we will have stretched the bright image brightnesses so they occupy more of the available brightness range. And, detail will become more apparent. These ideas are illustrated in figure 1.

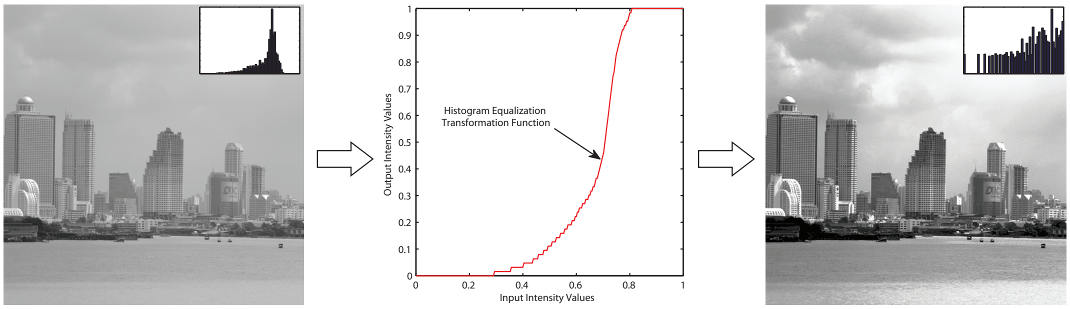

Figure 1. *Illustration of histogram equalization. Left, shows the input and it corresponding histogram. Middle, shows the histogram equalization transformation function. Right column shows the output image along with the corresponding histogram. As can be seen, histogram equalization makes image histogram flatter; these indicate the state that the stimulus has a higher entropy and hence more information.*

The Optimal Gamma

In companion work to this paper [5] we solved for the gamma that resulted in the image with the maximum entropy. Remarkably, assuming a continuous probability distribution of brightness in the interval [0,1], the optimal gamma can be calculated as:

$$\gamma_{optimal} = -\frac{1}{mean(log(x))} \tag{3}$$

Simple algebraic manipulation will convince the reader that once the optimal gamma is applied there is no benefit in finding the optimal gamma for the result (the second gamma will always be 1 the 'identity' gamma adjustment).

Our gamma adjustment to maximize entropy works similarly to HE though the amount the image changes will be less. This we argue is to our advantage. After all, a gamma adjustment in log space is just a scalar multiplication (i.e. contrast adjustment). The mapping relating input and output log values (pre and post gamma adjustment) is just a straight line (whose slope equals gamma) In contrast, HE might apply a very wiggly 'high frequency' tone curve and this can result in the well known problems of there being too much contrast and countering (edges appearing that were not visible in the original).

In the context of this work we are interested in whether the gamma adjustments made by observers match those that are optimal in an information theoretic sense.

Experiments

So, if we adjust an image using the information theoretic optimal choice of gamma, does it correspond with the adjustments made by observers? To test our assumption, we first need to design a psychophysical experiment to acquire observer gamma choices and then, in a second stage, relate these adjustments to the optimal gamma. Here we use a double (interleaved) staircase psychophysical experiment [4, 6, 12]. Informally, the idea is that a user will adjust a pair of reproductions of an image that are respectively too bright and too dark. They progressively and respectively darken and lighten the pair (this is the double staircase) until the image appearance convergence. A single staircase is not used here because often an observer will 'overshoot' the preferred gamma setting and make an image too bright (and they will only realize this when a much darker counterpart is shown). This darker counterpart is an informal invocation of the second staircase.

Experiment Methodology

Staircase or up-and-down method is one of the widely used methods in modern psychophysics. In the experiment, the experimenter decides what the next stimulus will be depending upon the observer response. That is, the response on the last trial (trial n) determines the stimulus on the current trial (trial n+1).

In addition, three things to think about before we can conduct this experiment are; where to start the stimulus (how bright/dark the image is), how big the step size is (what is the minimum step size of gamma), and when to stop the trials (how do we decide that the two staircases have converged). If we use a single staircase the observers can easily guess what is going on, since the next stimulus is depended on the observer response; this can lead to the unreliable results. Double staircases can solve this problem: two staircases (bright becoming darker and dark becoming brighter image staircases) are run simultaneously where we randomly switch from one to the other when asking observers to adjust gamma. In this way we remove the obvious sequential dependency of the trials. Of course, within each staircase this dependency remains; but by randomly interleaving the two staircases the dependency is concealed from the observer.

Regarding the step-size (the delta between different gammas), if it too large, then the observer's response will oscillate between the two stimuli resulting in no real final preference. In contrast, if the step-size is too small, then the experiment becomes inefficient because the observer might worry about their choice of the same judgment for long consecutive trials. One way to find out the appropriate step-size is to run pilot experiments.

Also, when dealing with large different starting positions (bright and dark images), we might vary the step size over the experiment. In our experiment, we also use this strategy in order to speed up the process and also reduce the problems from the observer's tiredness.

Finally, we measure the observer's final preference by looking at the reversals in gamma choices. Considering the bright image staircase, the observer will, initially, always make the image darker. However, as the image converges towards a final preferred gamma, the observer will tend to oscillate: making the image brighter then darker. This kind of adjustment is called a reversal. The experiment is terminated when a fixed number of reversals is observed, then the preferred gamma is calculated by averaging these reversals. This means that all the observers will have the same number of reversals but different number of trials. Because some observers have more variable, so they need more trials to reach their criterion.

Experiment Design

The stimulus in the experiment is the gamma adjustment (value of γ). As described in the introduction, if we increase the value of gamma, this will darken the image. In contrast, if we decrease the gamma, we will brighten the image. With this in mind, in the experiment, the participants were asked to choose whether an image shown on the screen appear either "too dark" or "too bright". If the response is "too dark", the next displayed stimulus will be two steps brighter to the current one ($\gamma_{(n+1)} = \gamma_n - 2 * accuracy$). In contrast, if the response is "too bright", the next stimulus will be two steps darker ($\gamma_{(n+1)} = \gamma_n + 2 * accuracy$). Once the two staircases converge, then the step-size will be halved ($\gamma_{(n+1)} = \gamma_n \pm accuracy$). The accuracy of the experiment is the γ of 0.1. The number of reversals to be collected in the experiment is six.

We use two image sets in our experiments: ours (27) and the (15) standard images from Kodak [10] (42 images in total). Our images were chosen to have different average intensities ranges from very dark to very bright images were incorporated in the experiment. The Kodak images are often used as references for photographic reproduction and have a much more limited brightness range. For our experiments we use only the Luminance channel of the Kodak images. The starting points for the bright and dark staircase are images 'gammed' to be overly bright and dark. Specifically we raise each image to the gamma which makes the average log value equal to -0.5 (overly bright image) and -5 (overly dark). All images are displayed to observers in random order.

All the original images in our set are 'gamma corrected' for display (to deal with the display non linearity). We assume the images are coded as sRGB and invert this gamma as a first step. That is, all our images are assumed to be linear. Thus when an observer chooses a gamma of, say, 1.5 the image we display is equal to the linear image raised to the power of 1.5/2.2. Here the 1.5 changes the image contrast and 1/2.2 applies the display gamma.

The total number of participants in the pilot experiment are 12 (eight males and four females) with normal color vision, naïve (the participants have not seen the original before doing the experiment) for the goal of the experiment under the control environment conditions. The whole procedure per participant took approximately 60 minutes (two 25 minute sessions with a 10 minute break in the middle).

Results

To derive the gamma predicting model based on the image's entropy, we first pre-calculate the information theoretic gamma ($\gamma_{optimal} = -1/mean(log(x))$). From mathematical argument alone, this gamma will maximize the image entropy (make details theoretically most visible). We then plot them against the

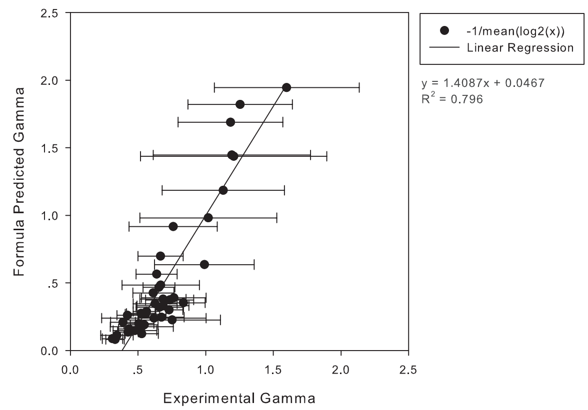

Figure 2. *Average preferred gamma from 12 observers plotted against the theoretical optimal gamma (-1/mean(log2(x))) for 42 testing images.*

chosen gamma received from the experiment. Figure 2 shows this relation. The error bar in the plot is the standard deviation.

From the pilot results, we found the linear relation between image's entropy and the chosen gamma that we have got from the experiment. As can be seen, the chosen gamma increases as the image's entropy increase. We believe this is a highly significant result as it provides evidence that the observer is behaving in as an 'optimal information processor'. There are several psychophysical papers which propose that aspects of vision can be explained by appealing to the notion of information theory and optimal information processing. For example, Buchsbaum and Gottschalk [3] proposed that colour opponency could be explained by efficient information coding. Here we are proposing gamma adjustment is proportional to the gamma that maximizes entropy.

The reader might be interested to consider why the relation has a slope and intercept (as opposed to just a slope). The equation relating observer gamma to optimal gamma can be written as $\gamma_{display} = 0.62\gamma_{optimal} + 0.38$ (we invert the line equation relating x to y). The intercept 0.38 effectively imposes a minimum gamma (a gamma below this number is not possible). We believe this bound makes sense when we think about image content. First, our dark images clearly looked like night time scenes so an observer would not make these scenes so bright as to resemble daylight conditions. Also, there were typically two illumination fields in each scene e.g. outdoor dark and indoor artificial lights. Thus the gamma has to compromise brightening the shadows and not destroying the highlight detail.

In addition, the reader might doubt why there are few images that have the entropy value above the value of 0.5, this is because actually our entropy calculation is calculated in log-2 space. This results in the cluster at the origin of the coordinate. Furthermore, results show that image with narrow histogram tends to have wider range of chosen gamma (larger error bars) than the ones with broaden histogram and because most of the images with high value of log-mean normally have narrow histogram that condenses in the high value of intensities. In the next section, we are going to provide more evidences that support this assumption by investigating the three related works.

Discussion

The Non-Linear Masking operator by Moroney [9] performs local gamma correction. The operator uses power function where the exponent is computed pixel-by-pixel from a mask that derived from a negative low-pass filtered monochrome version of the input image. This can be written as the following equation:

$$Output = Input^{2^{\frac{0.5-Mask}{0.5}}} \quad (4)$$

To make the operator acts as a global gamma predictor, we modify the equation 4 by simply replace the mask with a mean intensity of image μ, and replace the base of exponent of 2.0 with 2.5, since it gives a wider range of possible exponent. The modified equation can be written as:

$$Output = Input^{2.5^{\frac{\mu-0.5}{0.5}}} \quad (5)$$

If μ is greater than 0.5, the exponent will be more than 1, in contrast, if μ are less than 0.5, the exponent will be less than 1. The range of possible exponent is between 0.4 and 2.5.

In the same sense, the lightness calculation found in the two versions of the well known Color Appearance Model; CIECAM97s and CIECAM02 [8], are very similar to Moroney's operator, since the equations found in both versions are power functions. The Lightness J is calculated from the following equation:

$$J = 100\left(\frac{Y_b}{Y_W}\right) \quad (6)$$

where A/A_w is the ratio of the achromatic response of the sample to the response of the adopted white point, c is the predefined surround factor and z is the base exponential nonlinearity. The different in lightness calculation between the two versions of CIECAM is the calculation of z:

$$z_{CIECAM97s} = 1 + F_{LL}n^{\frac{1}{2}} \quad (7)$$

$$z_{CIECAM02} = 1.48 + n^{\frac{1}{2}} \quad (8)$$

$$n = \left(\frac{Y_b}{Y_W}\right) \quad (9)$$

where F_{LL} found in CIECAM97s is a lightness contrast factor, n is the background induction factor and equal to the ratio of a luminance of the source background Y_b to a luminance of the white Y_w, and ranges from 0 for a background luminance factor of zero to 1 for a background luminance factor equal to the luminance factor of the adopted white point [2].

In CIECAM the lightness ration A/A_W is somewhat laborious to calculate and not easy to relate to the original image (which is a key concern for us in our experiments). We solve this by relating CIECAM lightness to Colorimetric Luminance. Specifically, we generated an achromatic ramp of intensities and then calculated the ratios of Y/Y_W and A/A_W for the ramp. We then solved for the gamma which when applied to Luminance, resulted in values closest to CIECAM Lightness. We found that this relation is roughly equal to the square root of the ratio of input intensity:

$$\frac{A}{A_W} \approx \left(\frac{Y}{Y_W}\right)^{\left(\frac{1}{2}\right)} \quad (10)$$

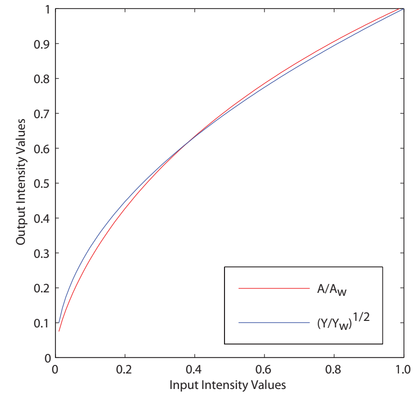

***Figure 3.** The approximation of A/A_W by the square root of the ratio of input intensities Y/Y_W.*

Figure 3 illustrates this approximation.

This implies that if we apply a square root to the lightness equation, the exponent of this new equation is now comparable with the other predictors. The new lightness equation can be written as:

$$J \approx \left(\frac{Y}{Y_W}\right)^{\left(\frac{cz}{2}\right)} \quad (11)$$

For the average surround ($c = 0.69$, $F_{LL} = 1.0$), this result in the exponent of between 0.35 and 0.69 for CIECAM97s and between 0.51 and 0.86 for CIECAM02.

Figure 4 shows the plot between the observer's preferred gamma and the predicted gamma of the four predictors; ours, modified Moroney, CIECAM97s and CIECAM02. As can be seen, the four predictors have linear fittings to the perceived gamma from the experiment and positive slopes. In addition, although there are fluctuations among all of the four predictors, compared with the three predictors, our predictor has highest correlation to the linear fitting. Furthermore, the modified version of Moroney and our predictor cover a broad range of useful gamma, indicates that both of them are appropriate to predict the chosen gamma, whereas the two CIECAM fail for predicting the gamma, since a little different in prediction can make a large different to the observer's preferred gamma. Table 1 summarizes the four gamma predicting functions derived from the plot in Figure 4.

	Predicting equation	(R^2)
Ours	y = 1.62x - 0.62	0.86
Modified Moroney	y = 1.41x + 0.05	0.80
CIECAM97s	y = 0.17x + 0.61	0.69
CIECAM02	y = 0.17x + 0.45	0.69

Gamma predicting equations along with the correlation coefficient (R^2) of the four models.

From the experiment to an automatic gamma adjustment operator

There are numerous possible applications of the model. One application is to use the derived equations from the previous section as an automatic global tone reproduction operator. We will

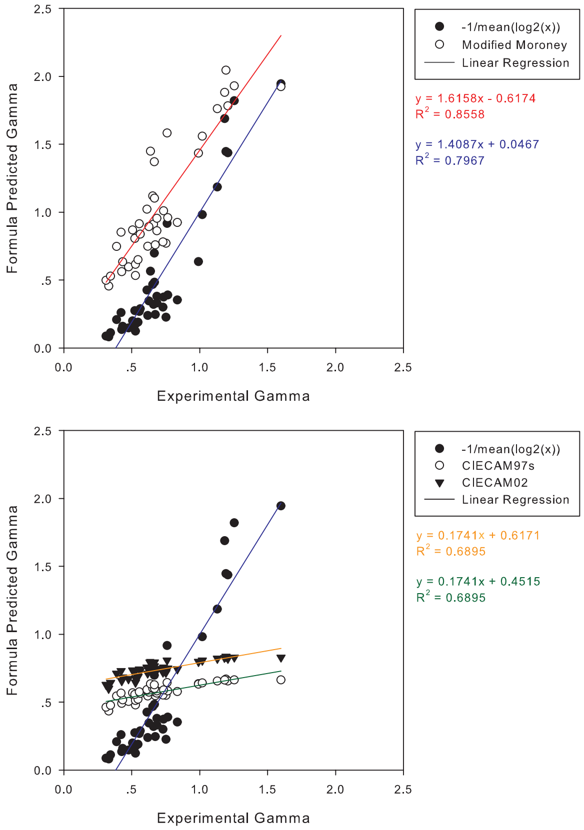

Figure 4. *Different gamma predictors plotted against the average chosen gamma from the experiment, the corresponding lines to each dataset are the linear fitting to each of the predictor. Top, the modified version of Moroney and our predictor ($-1/mean(log2(x))$) are plotted. Bottom, the two versions of CIECAM (97s and 2002) are plotted along with our predictor.*

evaluate experimentally how well this idea works in future research. As a taster, Figure 5 shows five of our images and the automatically chosen display gamma.

Conclusion

A simple and common way to adjust an image, making it brighter and darker, is to raise the image to the power of gamma: to carry out gamma adjustment. Recent theoretical work has provided a formula that, for a given image, calculates the gamma that maximizes entropy. That is, it returns the gamma which, theoretically, is the optimal in an information theoretic sense. In this paper, we provide evidence that there is a simple linear relationship between the information theoretically inspired optimal gamma the gamma adjustment made by observers. Further, a similar (though, lesser linear prediction is delivered by Moroney's non linear tone masking formula and via CIECAM-type formulae. Plausibly, these functions are also increasing the information content of image. Finally, we propose that the linear prediction formula, relating the theoretical optimum result to our experimental data, can be used to adjust gamma in images.

References

[1] Chang, C.-I.; Du, Y.; Wang, J.; Guo, S.-M.; Thouin, P.D. (2006) Survey and comparative analysis of entropy and relative entropy thresholding techniques. In Vision, Image and Signal Processing, Volume 153, Issue 6, 837 - 850, doi:10.1049/ip-vis:20050032.

[2] Cie Tc, The CIE 1997 Interim Colour Appearance Model (Simple Version), CIECAM97s, October 05 1998.

[3] G. Buchsbaum, A. Gottschalk, Trichromacy, Opponent Colours Coding and Optimum Colour Information Transmission in the Retina, Proceedings of the Royal Society (London), B220:89-113 (1983).

[4] G.B. Wetherill and H. Levitt, Sequential estimation of points on a psychometric function, Br. J. Math. Stat. Psychol. 18 (1965), pp. 1-10.

[5] G.D. Finlayson, R. Xu and J. Singnoo, Optimal Gamma Adjustment, School of Computing Sciences Technical Report, 2010.

[6] H. Levitt, Transformed up-down methods in psychoacoustics, J. Acoust. Soc. Am. 49 (Suppl. 2) (1971) 467-477.

[7] M. D. Fairchild, Color Appearance Models, Wiley, February 2005.

[8] Nathan Moroney and Mark D. Fairchild and Robert W. G. Hunt and Changjun Li and M. Ronnier Luo and Todd Newman, The CIECAM02 Color Appearance Model, Color Imaging Conference, pp. 23-27, IS&T - The Society for Imaging Science and Technology, 2002.

[9] Nathan Moroney, Local Color Correction Using Non-Linear Masking,Proceedings of the 8th Color Imaging Conference on Color Science and Engeneering Systems, Technologies and Applications (CIC-00), pp. 108-111, IS&T, November 7-10 2000.

[10] Rich Franzen, "Kodak Lossless True Color Image Suite". February 17, 2010. April 10, 2010. http://www.r0k.us/graphics/kodak/index.html.

[11] Shannon, Claude E.: Prediction and entropy of printed English, The Bell System Technical Journal, 30:50-64, January 1951.

[12] T.N. Cornsweet, The staircase method in psychophysics, Am. J. Psychol. 75 (1962), pp. 485-491.

Author Biography

Jakkarin Singnoo received his BSc in Imaging and Printing Technology from Chulalongkorn University, Thailand (2004) and his MSc in Advanced Computer Graphics and Applications from University of East Anglia, UK (2006). He is now a PhD student under the supervision of Prof. Graham D. Finlayson in Colour Group, School of Computing Sciences at University of East Anglia. His work has focused on HDR imaging including Tone Mapping Operators and Psychophysics.

Figure 5. *Examples of our automatic gamma adjustment. Left column shows original images. Right column shows result images.*