Colour Bit-Stealing to Enhance the Luminance Resolution of Digital Displays on a Single Pixel Basis

Christopher W. Tyler

Smith-Kettlewell Eye Research Institute
2232 Webster Street, San Francisco, CA 94115

Phone: (415) 561-1640
Fax: (415) 561-1610
E-mail: cwt@skivs.ski.org
http://www.ski.org/cwt.html

A look-up table algorithm is described for enhancing the luminance resolution at the expense of the colour resolution in digital displays. For colour displays with a look-up table resolution of 8 bits/gun, the algorithm provides a luminance resolution of 11-12 bits without loss of spatial or temporal resolution. This improvement can reduce the minimum luminance step in the mid luminance range from 1.5% to > 0.2% with no additional hardware.

2/9/96

Introduction

Much psychophysical testing is the measurement of contrast sensitivity for spatial patterns specified with some mathematical luminance profile and a controlled time course. Desk-top computers would appear to be ideally suited for such a task and yet most of the software for such computers requires special-purpose hardware modifications and software that requires programming to achieve the desired experimental configuration. The difficulty is to present stimuli with sufficient resolution of fine contrast variations to allow measurement of human contrast sensitivity, which can be as low as a 0.2% difference between light and dark regions of the stimulus. This resolution is not possible with conventional stimulus profiles on a standard 3 x 8-bit Macintosh display, whose smallest luminance step at mid-luminance is about 1.5% due to the accelerating phosphor nonlinearity.

A technique that allows increased luminance resolution on any digital computer with a colour monitor is the "bit-steading" procedure introduced by Tyler et al. (1992), in which the colour values within each pixel are jittered to provide a wider range of luminance combinations from which to chose. This approach amounts to relaxing the accuracy of the chromatic signal to provide greater accuracy on the luminance signal at the individual pixel. Under a wide range of conditions, sensitivity to chromatic contrast is lower than that for luminance contrast. A trade of bit-resolution between these two modalities can therefore allow increased luminance precision within individual display pixels without detectable loss of colour fidelity under many conditions. Unlike colour dither techniques, bit-stealing does not sacrifice either spatial or temporal resolution in improving the luminance specification.

Color Bit-Stealing

The bit-stealing technique was originally conceptualized for improving the grey-scale resolution of colour monitors with a particular ration of R:G:B luminances (Tyler et al., 1992). Specialized schemes designed for particular luminance ratios were described, but these are not applicable to realistic video displays in which the luminance ratio of the colour guns varies with luminance level as a result of variations in the nonlinear output function of the three colour guns. Thus, a generalized concept was developed for optimizing at any point in the colour space defined by the three colour outputs.

Normally, for every pixel of an image to be displayed, the three colour components of the image colour data are specified independently. The target luminance to be derived from the image colour data is given by an additive combination of the luminance values of the independent colour data. The bit-stealing technique, on the other hand, jitters among the available colour combinations close to the target colour to select a colour combination whose overall luminance gives an optimal match to the target luminance. The technique thus allows some degradation in the accuracy of the colour specification in order to improve the accuracy of the luminance specification.

This bit-stealing process for a specific pixel is depicted schematically for a two-colour space in Fig. 1. The two abscissa represent the levels in the look-up tables for the R and G guns at one pixel of a display. The stepped surface represents the resultant output luminance for each combination of R and G levels. The two marginal lines represent the target luminances for the cases of zero output for one of the two guns, showing how the actual luminances jitter above and below these targets along the two lines. The semivertical oblique line through 0,0 represents a particular target hue assumed for this pixel. Accurate colour representation would require the combination of R and G values to fall on this line. Bit-stealing allows greater variation away from this line, as depicted by the ellipse around a specific target luminance and chromaticity. The semihorizontal oblique line represents the locus of R/G combinations of the same luminance. The process of bit-stealing is to search for actual outputs within the range of the ellipse that provide the best approximation to the target luminance. Note that, in practice, the 0,0 position depicted would correspond to a level of about 160,160 in the look-up table (assuming that it was at the mid-luminance output level of the monitor), so that the extent of the ellipse would still represent only a small range of chromaticities.

Normally, the accelerating nonlinearity of phosphor luminance as a function of input voltage from the look-up table implies that the minimum step is about 1.5% at the mid-range of screen luminances (Tyler et al., 1992). The precise improvement available by bit-stealing depends on the relative intensities of the three colour guns and the range of colour variation allowed in the implementation. For, a ± 2-step routine, we have estimated the improvement to be of the order of a factor of 10, implying a minimum step of about 0.15%, which is close to the limit of human contrast sensitivity under optimum conditions. Bit-stealing thus allows a full range of psychophysical experiments with spatiotemporally modulated stimuli to be conducted with a standard computer and monitor configuration.

Color Jitter Ellipse for Optimizing to Target Luminance

Fig. 1. Depiction of the colour jitter ellipse within which luminance output values are selected to find the best fit to the target luminance for a particular pixel within a range of similar chromaticities.

Application to Image Display

Many computer applications require the display of extended static images containing smooth gradients. Bit-stealing is equally applicable to such conditions. The problem is that only 256 steps are available for the representation of extended gradients in look-up mode. If the gradient extends over a limited range of luminances, even fewer steps are usable. This limitation can result in the appearance of stepped contours in graded images. Based on the lateral inhibitory filtering of the visual system, such contours are much more visible than would be expected from the 1-2% limit implied by the physical steps of light. Bit-stealing can be applied to such images to provide finer gradations in luminance and eliminate the perceptible contours. Because there is typically a maximum of 256 steps available to output, this procedure should be done in conjunction with a range optimizing technique, to distribute the available steps over the luminance range so as to maximize the number of steps in regions of graded luminance changes and minimize them where rapid or step-wise changes are present.

The kind of improvement that is obtainable by bit-stealing is shown for a 35-pixel sequence in Fig. 2. The solid stepped line shows the poor resolution with which a luminance ramp could be displayed if the colours had to be present at equal look-up table levels. The bar graph depicts the smooth ramp that can be obtained by bit-stealing, when the colours are allowed to jitter around equality; combinations can be found that generate a much more accurate representation of the desired luminance ramp.

Ramp by Color Bit-Stealing

Fig. 2. Display output for a luminance ramp showing the crude representation with accurate colour specification (solid stepped line) and the smooth ramp obtained with application of the bit-stealing

algorithm (bar graph). Note that the bit-stealing operates at each individual pixel, although it is shown here over a 35-pixel ramp.

Residual Color Changes

Although it is generally the case that human sensitivity for chromatic differences is poorer than that for luminance differences, the opposite is true for some types of stimuli. For variations of moderate spatial and temporal frequency, the lowest contrast threshold has been reported to be 3% (van der Horst, 1967). For large fields, on the other hand, chromatic contrast thresholds may run as low as 0.4% (Stromeyer et al., 1995). Bit-stealing is not designed for application in these situations.

However, one of the stimuli of choice in modern psychophysics is the local Gabor patch, the product of a sinusoidal test modulation and a Gaussian window function. Bit-stealing works well in the strongly modulated region of the patch because the rate of chromatic variation is spatially rapid in relation to the rate of luminance modulation. The chromatic signal is thus of relatively high spatial frequency, for which the visual system has poor contrast sensitivity. In the low contrast flanking regions, however, there may be large areas in which the stimulus waveform is represented by a single value in the look-up table. Similar problems are present for presentation of any low-frequency function such as a Gaussian blob. In the relatively uniform regions, the original bit-stealing algorithm will specify a single chromaticity that is different from that of the background, allowing their detection by means of the colour variations alone. This artifact limits the utility of the original bit-stealing algorithm to functions without large uniform regions.

Quick Fixes for the Problem of Visible Color Differences in Broad Regions

Since the dominant paradigm for measuring psychophysical sensitivity is the two-alternative forced-choice (2AFC) method, the key aspect to control is not so much the absolute properties of the image but the differences between the test and null stimuli of the 2AFC presentation. If it can be arranged, therefore that the chromatic signal is the same between the test and null intervals, it cannot act as a discriminative variable in the task. In the specific case of the Gabor function, this result could be achieved by identifying (from the profile map) the region in the outer skirts of the Gabor function where the displayed chromaticity is constant over a number of pixels. The signal in this region is then presented in the null interval as well as the test interval so that it does not provide a cue to discrimination.

A second quick fix is to design the stimuli employed in a study to avoid regions controlled by a single value in the look-up table. For some stimuli, such as true Gaussian blobs at low contrast, this may be impossible. Other stimuli, such as Gabor patches, are amenable to a simple manipulation. If the Gaussian windowing function of the Gabor patch is replaced by a half-cosine bell, for example (Fig. 3), the contrast gradient becomes asymptotic to a constant slope near the outer edge of the function, so there are no extended uniform regions to trigger the colour sensitivity. This small change should not affect the threshold measurements of such stimuli since they are dominated by the higher contrast regions of the luminance modulation, so that the precise waveform in the lower contrast regions probably has little influence on detectability.

Fig. 3. Comparison of the waveform of a classic Gabor function (upper waveform; the product of a sinusoid and the Gaussian window shown above it) and a steep-flanked function (lower waveform; the product of a sinusoid and the half-cosine bell window shown above it).

The equation of the 2D half-cosine bell waveform, which avoids extended shallow gradients around the flanks, is:

$$\begin{array}{ll} \cos(2\pi r), & r < \pi / 2 \\ 0, & r > \pi / 2 \end{array}$$

where $r = \sqrt{x^2 + y^2}$.

<u>Chromatic Noise Jittering</u>

A useful technique to avoid chromatic artifacts introduced by bit-stealing that may act as a discriminative cue is to add a low level of chromatic noise to the look-up table. The concept of colour noise jittering is to add random colour jitter to divert the bit-stealing routine from choosing the same colour over large truly uniform areas of the field. Two components are necessary. One is to eliminate large uniform sub-regions by using waveforms without extended shallow slopes, such as the outer skirts of a Gaussian. The other is to add random colour jitter to the test image. Within the range of randomization, each subregions in the image will have a slightly different colour, preventing the tendency to colour regions of similar target intensity to be selected to have the same colour when the bit-stealing process is operating near the bottom of its range. Moreover, the null stimulus of a 2AFC pair will have the identical randomized structure, so that any visible colour contours will not provide a discriminative cue because they are similarly randomized between the two intervals. The randomization will be different in each interval, but the colour changes will be random and therefore will provide no information about the presence of the stimulus in one or the other interval.

The presence of colour differences across the image does not necessarily imply that they will be visible as colour fringes. The colour jitter may have a sufficiently small range that they are invisible, or at least not noticeable, in the narrow contour strips of the image in which they are present. By "not noticeable", I mean that they may be classed as one of many small deviations from uniformity in the stimulus as it is perceived by the observer (i.e., including optical defects in the eye and neural noise) but are not noticed as a discriminative stimulus in the task.

Look-up Table Implementation

Specifically, the problem of detectable colour in bit-stealing arises for the lowest contrast images, where the optimal luminance match is likely to result in the same choice of colour for many adjacent stimulus regions, producing large regions of the same colour (and luminance. The chromatic noise solution is implemented by adding random colour jitter to the target colours for the look-up table. Within the range of randomization, each look-up table level will then have a different target colour, so that the search for the best match to the target luminance will be unlikely to result in the same choice of colour of many adjacent image fields. In fact, adjacent repeats can be prohibited in the algorithm so as to ensure a colour change between each pair of levels.

Quantitatively, the ideal is to ensure that there are as many of the 256 entries in the look-up table represented in the stimulus as possible. (If the stimulus is a image such as a mirror-symmetric Gabor patch of 100 x 100 pixels, for example, then there can be no more than 50 different luminance contour levels within it, assuming that no contour can be less than one pixel. Even on the diagonal, the symmetry implies that there are no more than 100/2 contours. Ensuring that no two adjacent contours are the same colour requires jittering the values on the three colour guns across a maximum of 5-6 levels, since $3^5 = 243$, close to the total of 256 available. Fewer jitter levels would be required for patterns with a smaller number of discrete levels. This situation would results in a time saving for the search process in computation of the look-up tables, which may be significant if many look-up tables are precomputed so as to allow rapid colour cycling among them.

Thus, for each look-up table value, its target chromaticity is jittered within the prespecified range to set a new, random target chromaticity offset slightly from the original (corresponding to the random colour noise value for that level). The bit-stealing search for the best match to the original target luminance is then conducted by generating a sequence of nearby colour combinations are determining whether their calibrated luminance sum is within the prespecified criterion of the target luminance. If so, the algorithm moves to the next look-up table level until the requisite number has been computed.

Intraobserver Variation in Spectral Efficiency

Implementation of the bit-stealing algorithm requires a precise knowledge of the luminance value for every level of the look-up tables for all three colour guns. This implies not only an accurate calibration in terms of radiant energy but also accurate information about the effect of that radiant energy on the human eye. This information is available for the average observer in the form of the V_λ curve of spectral efficiency, which is built in to the definition of luminance for photometric calibrators. There may, however, be large variations from these values in the spectral efficiency either for individual observers or across retinal locations, which could reduce the accuracy of the bit-stealing process. This is a particular problem in the case of protanopic colour blindness, which has a reduced efficiency for the long-wavelength end of the spectrum.

A simple solution to this problem is to determine the spectral efficiency for each colour gun for the particular observer and spatiotemporal stimulus conditions under test. Even though the form of the spectral efficiency curves for each cone type may vary widely across observers, their contribution to luminance for a given individual is defined by the integral of the product of the spectral efficiency and the spectral emission of each colour gun. Consequently, only two measurement values are required to determine the relative sensitivities to the three guns, those of the ratios of relative sensitivities for the R and B guns to that for the G gun. The usual approach to this problem is to perform heterochromatic flicker photometry between pairs of colours, but this approach may need to be viewed with caution because the luminance match varies with temporal frequency (Hamer and Tyler, 1992). Instead, it may be preferable to perform a minimum distinct border match between two spatially adjacent fields of the different colours (Kaiser, 1971) or some other assay. The resulting values may be entered as the relative effective luminance of the three colour guns from which to determine the luminance of their combination.

System Requirements

We have developed psychophysical software for Macintosh computers based on the bit-stealing principle (see Tyler, 1996, in this volume). The bit-stealing display software requires a colour monitor and 24 bit graphics card for its operation. Any size of monitor is supported. (Since it relies on look-up table animation, the program will not work with the "thousands" or "millions" modes selected for the graphics card, although the bit-stealing principle as such is not inherently limited in this way). The luminance scale is fully linearizable to 99% output contrast (with monitor calibration).

An Automated Calibrator for Macintosh Computers

A bit-stealing algorithm requires calibration of the output luminance levels to high accuracy for each of the 256 steps of the look-up table for each colour gun. To achieve this level of performance with minimum effort, we have developed an automated calibrator for Macintosh computers in conjunction with Dr. Anthony M. Norcia, Steven Chung (hardware engineer) and Kirk Swenson (software engineer). It allows calibration of any point on a monitor screen in about

15 minutes. This calibrator provides an exact specification of every look-up table value without curve-fitting or approximation formulas.

A photocell with a CIE photopic filter is attached to the monitor screen and sends a signal to the A/D converter. Because the digital signal is input through the Macintosh ADB (keyboard and mouse) bus of any of Macintosh II and PowerMac family of computers, the calibrator is termed a LightMouse. The look-up table levels for each colour gun are stepped through with sufficient time for the monitor output to settle before each reading is taken. The luminance of each gun is stepped first up and then down, to average over any warm-up or drift effects. The output, calibrated in candelas per meter^2, is placed in an ASCII calibration table for use by the Morphonome software or any other application.

Conflict-of-Interest Statement. The **LightMouse™** calibration package is distributed at cost on a non-profit basis through the Smith-Kettlewell Eye Research Institute for the benefit of the vision community.

Acknowledgments

We thank Lenny Kontsevich for raising the issue of individual differences in spectral sensitivity. Supported by NIH grants EY 7890 and RR 5566.

References

Kaiser P.K. (1971) The minimally distinct border as a preferred psychophysical criterion in visual heterochromatic photometry. *J. Opt. Soc. Am.* 71, 966-971.

Stromeyer C.F., Eskew R.T., Kronauer R.E. (1995) What the eye sees best in color. *Science*.

Tyler C.W., Chan H., Liu L., McBride B., Kontsevich L.L. (1992) Bit-stealing: How to get 1786 or more grey levels from an 8-bit color monitor. *Proc. SPIE* 1666. 351-364.

van der Horst G.J.C., de Weert C.M.M., Bouman M.A. (1967) Transfer of spatial chromaticity-contrast at threshold in the human eye. *J. Opt. Soc. Am.* 57, 1260-1266.

Hamer R.D., Tyler C.W. (1992) Analysis of visual modulation sensitivity. V. Faster response for G- and R-cone pathway? *J. Opt. Soc. Am.* A9, 1889-1904.