Luminance-Model-Based DCT Quantization for Color Image Compression

Albert J. Ahumada, Jr.
NASA Ames Research Center, Human Interface Research Branch
Moffett Field, California 94035
and
Heidi A. Peterson
IBM T. J. Watson Research Center
Yorktown Heights, NY

ABSTRACT

A model is developed to approximate visibility thresholds for discrete cosine transform (DCT) coefficient quantization error based on the peak-to-peak luminance of the error image. Experimentally measured visibility thresholds for R, G, and B DCT basis functions can be predicted by a simple luminance-based detection model. This model allows DCT coefficient quantization matrices to be designed for display conditions other than those of the experimental measurements: other display luminances, other veiling luminances, and other spatial frequencies (different pixel spacings, viewing distances, and aspect ratios).

1. INTRODUCTION

1.1 Discrete cosine transform-based image compression

The discrete cosine transform (DCT) has become a standard method of image compression[1,2]. Frequently the image is divided into 8×8-pixel blocks that are each transformed into 64 coefficients for reconstruction from the DCT basis functions. The DCT transform coefficients $I_{m,n}$, of an $N \times N$ block of image pixels $i_{j,k}$, are given by

$$I_{m,n} = \sum_{j=0}^{N-1} \sum_{k=0}^{N-1} i_{j,k}\, c_{j,m}\, c_{k,n}, \quad m,n = 0, \ldots, N-1, \tag{1}$$

where

$$c_{j,m} = \sqrt{\frac{1}{N}}, \quad m = 0,$$
$$= \sqrt{\frac{2}{N}} \cos\left(\frac{\pi m}{2N}[2j+1]\right), \quad m > 0.$$

The block of image pixels is reconstructed by the inverse transform, which is the same as the forward transform for this normalization,

$$i_{j,k} = \sum_{m=0}^{N-1} \sum_{n=0}^{N-1} I_{m,n}\, c_{j,m}\, c_{k,n}, \quad j,k = 1, \ldots, N-1. \tag{2}$$

Quantization of the DCT coefficients $I_{m,n}$ results in image compression. Typically, compression standards allow the user to set the level of quantization for each coefficient. A quantization matrix $Q_{m,n}$ specifies the quantization of the DCT coefficients, where the m,nth entry in the matrix is the step size of the uniform quantizer for the m,nth DCT coefficient.

1.2 Threshold-based DCT coefficient quantization

Last year at this meeting Peterson, Peng, Morgan, and Pennebaker[3] presented experimental measurements of the detection thresholds for individual DCT basis functions in each of the R, G, and B color dimensions.

Visibility thresholds were measured for purely R, G, or B basis functions on a black background, masked by a pedestal containing the same color at the luminance that resulted from the mid-range digital value 128 (plus a small veiling luminance). From these measurements, Peterson *et al.* [3] derived DCT coefficient quantization matrices for separate compression of R, G, and B images. Their quantization matrices produced reasonable compression with barely visible artifacts. However, the viewing conditions in many applications may not be similar enough to the viewing conditions of their experiments to yield good compression without visible artifacts using their specific quantization matrices. The goal is to develop a formula to generate appropriate quantization matrices for different viewing conditions. In particular, we would like to interpolate and extrapolate the experimental results to other display luminances, other veiling luminances, and other spatial frequencies (different pixel spacings, viewing distances, and aspect ratios).

2. THEORY

The theoretical basis of the model developed here is the assumption that the detectability of distortion in the decoded RGB image can be predicted based on the luminance error caused by quantization of an individual DCT coefficient. Although quantization of DCT coefficients causes both luminance and chrominance errors, we assume that for all the coefficients other than the DC coefficient, the spatial frequency is high enough that the luminance error dominates. The model only takes into account the luminance of the error; it ignores the chrominance.

The display parameters are characterized as follows:

L_0 is the background or veiling luminance on the screen,
L_i is the luminance contributed by the image,
L is the total luminance, $L_0 + L_i$,
w_x is the horizontal width of a pixel in degrees of visual angle,
w_y is the vertical height of a pixel in degrees of visual angle.

Assuming ideal interpolation of a spatially sampled signal displayed on a linear monitor, the image of the error $E_{m,n}(x,y)$, due to quantization of the m,nth DCT basis function coefficient $I_{m,n}$ can be written

$$E_{m,n}(x,y) = L_e \cos\left(\frac{\pi m}{2N\mathrm{w}_x}[2x+1]\right)\cos\left(\frac{\pi n}{2N\mathrm{w}_y}[2y+1]\right), \quad 0 \le x < N\mathrm{w}_x,\ 0 \le y < N\mathrm{w}_y, \tag{3}$$

$$= 0, \quad \text{elsewhere,}$$

where L_e is the luminance amplitude of the error image. We assume that the luminance error due to quantization of a particular DCT coefficient is not visible if L_e is less than a threshold $T_{m,n}$, which is a function only of the coefficient indexes, m,n, and the above display parameters.

2.1 Single-Orientation Case

For either m or n equal to zero (but not both), and ignoring windowing (that is, the limited spatial extent of the test pattern in the experiment), the problem becomes one of predicting $T_{m,n}$ for a vertical grating with spatial frequency

$$f_{m,0} = \frac{m}{2N\mathrm{w}_x} \tag{4a}$$

or a horizontal grating with spatial frequency

$$f_{0,n} = \frac{n}{2N\mathrm{w}_y}\ . \tag{4b}$$

In our model we approximate the log of $T_{m,n}$ by a parabola in log spatial frequency,

$$\log T_{m,n} = \log(T_{\min}) + K(\log f_{m,n} - \log f_{\min})^2, \quad m = 0 \text{ or } n = 0, \tag{5}$$

where the parameters $T_{\min}$, K, and $f_{\min}$ are functions of the total luminance L. $T_{\min}$ is the luminance threshold at $f_{\min}$, the frequency where the threshold is smallest; and K determines the steepness of the parabola.

Figure 1 illustrates that contrast sensitivity (luminance divided by threshold luminance) for luminance-only sine wave gratings can be fit by parabolic functions of log spatial frequency over a range of luminances, when only the fit to high (> 1 cycle per degree) spatial frequencies is considered. The data are from the work of van Nes and Bouman[4], as reported by Olzak and Thomas[5]. The parabolas are least-squared error fits to the individual sets of luminance data. The parabola model is not valid for low spatial frequencies, since the decrease in sensitivity at low frequencies is slow for imagery that is not stabilized on the retina[6]. This model is not applicable for $T_{0,0}$. The results presented in Peterson *et. al.*[3] indicate a conservative estimate for $T_{0,0}$ is the smaller of $T_{1,0}$ and $T_{0,1}$.

Figure 1. Contrast sensitivity data from van Nes and Bouman[4]. The data were measured at luminances spaced one decade apart from $2.9\times10^{-4}\, cd/m^2$ to $2.9\times10^{-2}\, cd/m^2$. Parabolas were fit by least squares to the log data with spatial frequencies of 1 cycle/degree or greater.

2.2 Two-Orientation Case

When neither m nor n are equal to zero, the trigonometric identity

$$2\cos(a)\cos(b) = \cos(a+b)+\cos(a-b) \tag{6}$$

shows that the DCT basis functions are the sum of two frequency components with the same spatial frequency

$$f_{m,n} = \sqrt{f_{m,0}{}^2+f_{0,n}{}^2}, \tag{7}$$

but different orientations. The angle θ between the two components is

$$\theta = \arcsin\frac{2 f_{m,0} f_{0,n}}{f_{m,n}{}^2}. \tag{8}$$

There are two extreme cases. If either m or n is zero, the result is the single orientation case discussed above. If m equals n, the components are orthogonal, and detection can be predicted either by probability summation, or more simply, by summing the fourth powers of the amplitudes and taking the fourth root. The latter leads to the threshold being increased by the factor $2^{0.75} = 1.68$. This prediction for the effect of the summation of the two Fourier components follows the work of Phillips and Wilson[7] and that of Watson[8]. The effect of intermediate positions of the two Fourier components can be approximated as a multiplicative threshold factor

$$\frac{1}{r+(1-r)\cos^b\theta}, \tag{9}$$

where b represents the orientation tuning of the visual system in grating detection. We assume $b = 2$. In addition, when f_m and f_n are equal, the two Fourier components are in the oblique direction, which usually results in a further increase in threshold. The oblique effect can be included by decreasing the value of r. Incorporating the factor of Equation (9) in the model of Equation (5), the complete model for predicting the log visibility thresholds becomes

$$\log T_{m,n} = \log\left(\frac{T_{\min}}{r+(1-r)\cos^2\theta}\right) + K(\log f_{m,n} - \log f_{\min})^2, \quad m, n = 0, \ldots, N-1. \tag{10}$$

3. PARAMETER ESTIMATION AND GOODNESS OF FIT

A luminance-only model for DCT basis function visibility thresholds requires the model parameters $T_{\min}$, $f_{\min}$, and K each be expressed as a function of luminance. The van Nes and Bouman[4] data were collected over a wide range of luminances. The experiments of Peterson *et al.*[3] had a narrow luminance range, but were shown to lead to good quantization matrices. Therefore, our strategy for estimating the functions is first to use the van Nes and Bouman data to determine the form of these functions, and then to fit these form templates to the Peterson *et al.* data. Each function is assumed to be a power law (linear in log-log coordinates), clipped at some maximal value. The exponent of the power law is found by a least squares fit to the van Nes and Bouman parabola parameters in the log domain. The maximal value for each function is also based on that data. The resulting templates are the dashed lines in Figure 2. The van Nes and Bouman peak contrast sensitivities have been shifted downwards by approximately 0.3 log units relative to those in Figure 1. This shift attempts to account for the difference in test pattern size for the two experiments, and corresponds to the log of the fourth root of the ratio of the two areas, the same summation rule discussed above [7,8]. (Peterson *et al.* used 50×50 pixel test patterns that occupied 1.47 degrees horizontally and 1.27 degrees vertically, while van Nes and Bouman used a 8.25 by 4.5 degree signal.)

Figure 2. $T_{\min}$ (top), $f_{\min}$ (middle), and K are shown as a function of L. Open circles, $\times$'s, and squares represent parameter values of the predictions for the Peterson *et al.* data in Figure 2 for R, G, and B, respectively. Solid circles represent parameter values for the curves fit to the van Nes and Bouman data in Figure 1. The solid lines show Equations (12), (13), and (14), developed to extrapolate the measured DCT basis function thresholds.

These templates are then used to extrapolate the data of Peterson *et. al.* by finding the best vertical shift for fitting the templates to the Peterson *et. al.* data. The quality of the model fit to the data is evaluated by the chi-square measure

$$\chi^2 = \sum \frac{N\,(\log T_{obs} - \log T_{pred})^2}{\sigma^2}, \tag{11}$$

where N is the number of subjects (15), T_{pred} is the threshold predicted by the model (Equation (10)), and T_{obs} is the experimental threshold. The standard deviation was estimated from the between-subjects standard deviations, excluding thresholds whose value was limited by the range of available stimulus levels. The standard deviations of the log thresholds do not appear to vary with basis function, but a different standard deviation was used for the B ($\sigma = 0.31$) thresholds than for the R and G ($\sigma = 0.21$). The sum is over the 58 non-DC basis functions (18 R, 22 G, 18 B) included in the experiment.

3.1 Template fitting procedure

Figure 2 shows model parameters estimated directly from the Peterson *et al.* data along with the templates fit to that data. To assess quantitatively the fit of the templates a sequence of nested hypotheses was formed that successively added constraints to the model parameters. Then $T_{\min}$, $f_{\min}$, K, and r were estimated by finding the values minimizing the chi-square measure of fit, subject to their constraints. Initially, all four model parameters were allowed to vary arbitrarily with luminance (*i.e.* color). The results appear in the first column of Table 1. Next all parameters were re-estimated with r assumed constant. These results appear in the second column of Table 1 and are also the points in Figure 2 representing the estimates unconstrained by the templates. Then all parameters were re-estimated with r constant and K assumed to follow its template (with a vertical shift allowed). In this case the model has eight parameters: r, a vertical shift for the K template, three $f_{\min}$'s (one for each luminance), and three $T_{\min}$'s (one for each luminance). In turn, $f_{\min}$ and $S_{\max} = L/T_{\min}$ were additionally constrained, so that the rightmost column in Table 1 gives estimates of the model parameters consistent with all the templates. Table 1 also gives the minimum χ^2 fit of the model of Equation (10) for this series of nested constraints on the parameters. Differences between successive fits would be distributed approximately as χ^2 if the constraints were valid (and other assumptions held).

		next parameter modelled				
		none	r	K	$f_{\min}$	$S_{\max}$
fit χ^2		107.80	108.14	112.02	133.58	233.11
test χ^2		107.80	0.34	3.88	21.56	99.53
df		46	2	2	2	2
signif. level		<0.001	0.85	0.15	<0.001	<0.001
	B	0.65	0.70	0.69	0.68	0.68
r	*R*	0.65	0.70	0.69	0.68	0.68
	G	0.70	0.70	0.69	0.68	0.68
	B	2.10	2.12	2.28	2.56	2.39
K	*R*	2.62	2.67	2.44	2.64	2.56
	G	2.41	2.41	2.64	2.97	2.77
	B	2.82	2.82	3.02	3.72	3.40
$f_{\min}$	*R*	4.36	4.36	4.17	4.43	4.04
	G	4.68	4.68	4.90	5.44	4.96
	B	62.6	62.6	61.1	52.0	74.5
$S_{\max}$	*R*	91.3	91.3	89.2	93.4	94.7
	G	101.3	101.3	103.7	108.6	94.7

Table 1. Parameter estimates and goodness of fit.

These test χ^2 values indicate that parameters r and K are well fit by their templates, but $f_{\min}$ and $T_{\min}$ are not well fit statistically. The luminances for B, R and G are 6.8, 17.4, and 53.9 cd/m^2, respectively.

3.2 $\mathbf{T_{min}}$, $\mathbf{f_{min}}$, and K luminance functions

The relation between $T_{\min}$ and luminance with all model parameters constrained to their templates is

$$T_{\min} = \frac{L^{a_T}\, {L_T}^{1-a_T}}{S_0}, \quad L \le L_T, \qquad (12)$$
$$= \frac{L}{S_0}, \qquad L > L_T,$$

where $L_T = 13.45\, cd/m^2$, $S_0 = 94.7$, and $a_T = 0.649$. Figure 2 shows this relation in contrast sensitivity $S_{\max} = L/T_{\min}$. Despite the adjustment for area, there is still a difference of about 0.25 log units between the two data sets.

The two data sets agree very closely on the relation of $f_{\min}$ to L. Figure 2 shows that both sets of data are fit well by the power function

$$f_{\min} = f_0\,(L/L_f)^{a_f}, \quad L \le L_f, \qquad (13)$$
$$= f_0, \qquad L > L_f,$$

where $f_0 = 6.78\, cycles/deg$, $a_f = 0.182$, and $L_f = 300\, cd/m^2$. This function is clipped at high luminances because van Nes and Bouman[4] found no difference between the contrast sensitivity functions for luminances of 290 and 1880 cd/m^2.

Although χ^2 tests could not reject the hypothesis that the same value of K could be used for R, G, and B, the van Nes and Bouman data suggest that K is also approximately a power function of L, as shown in Figure 2. Accordingly, we approximate K by the power function

$$K = K_0\,(L/L_K)^{a_K}, \quad L \le L_K, \qquad (14)$$
$$= K_0, \qquad L > L_K,$$

where $K_0 = 3.125$, $a_K = 0.0706$, and $L_K = 300\, cd/m^2$. This function is clipped at high luminances for the same reason as the $f_{\min}$ function.

Equations (12), (13), and (14) give $T_{\min}$, $f_{\min}$, and K as functions of L. When they are substituted into Equation (10), the result is a luminance-only model for DCT basis function visibility thresholds that extrapolates the Peterson *et al.* data.

3.3 Orientation correction

Figure 3 shows threshold predictions using the templates described above in the model of Equation(10) (dashed curves from the rightmost column in Table 1). These curves appear to fit the experimental thresholds nearly as well as those estimated directly from the data (solid curves from the second column in Table 1). The m or n equal zero experimental data of Peterson *et al.* are indicated by small symbols in Figure 3. The two-orientation data, indicated by large symbols in Figure 3, are shown divided by the correction factor of Equation (9). The need for this correction factor is demonstrated in Figure 4. Figure 4 plots the ratio of the experimentally measured thresholds for two-orientation basis functions ($m, n \neq 0$) to the uncorrected model predictions (Equation (5)) as a function of θ, the angle between the orientation components. These ratios show that the experimentally measured two-orientation thresholds are consistently higher than those predicted from just the spatial frequency of the basis function (that is, the model of Equation (5)). The solid line in Figure 4 shows Equation (9) with $r = 0.70$. This value was used to correct the two-orientation data points in Figure 3.

Figure 3. Visibility thresholds T for DCT basis functions measured by Peterson *et al.*[3] are represented as contrast sensitivities $S = L/T$, as a function of their spatial frequency f. Circles, $\times$'s, and squares correspond to R, G, and B experimental thresholds, respectively. Small symbols indicate single-orientation basis function thresholds. Large symbols indicate two-orientation basis function thresholds, after correction by Equation (9). Solid lines show the model predictions using the parameters $T_{\min}$, $f_{\min}$, and K estimated directly from the data by minimum χ^2 (the second column of Table 1). Dashed curves show model predictions using Equations (12), (13), and (14). The top, middle, and bottom (at the right edge of the graph) curves are fit to the G, R, and B, respectively.

Figure 4. Ratios of uncorrected basis function visibility thresholds to their single-orientation model predictions are shown as a function of θ, the angle between the orientation components. The solid line shows the the graph of Equation (9) with $r = 0.70$. Points are labeled as in Figure 3.

4. GENERATION OF QUANTIZATION MATRICES

To convert visibility thresholds measured in luminance to quantization units, $T_{m,n}$ values are first multiplied by 2. This accounts for the maximum quantization error being one half of the quantizer step size. Two more factors are needed to complete the conversion. First we need ΔL_g, the change in luminance caused by a change of one quantization unit for gun g, at luminance L. For a linear display this is the difference between the minimum luminance $L_{g,\mathrm{min}}$ and maximum luminance $L_{g,\mathrm{max}}$ from gun g, divided by the total number of quantization units M,

$$\Delta L_g = \frac{(L_{g,\mathrm{max}} - L_{g,\mathrm{min}})}{M}. \tag{15}$$

The second factor accounts for the normalization constants of the DCT,

$$\begin{aligned} \alpha_m &= \sqrt{\frac{1}{N}}, \quad m = 0, \\ &= \sqrt{\frac{2}{N}}, \quad m > 0. \end{aligned} \tag{16}$$

The quantization matrix elements for that gun are then

$$Q_{m,n} = \frac{2\, T_{m,n}}{\alpha_m\, \alpha_n\, \Delta L_g}. \tag{17}$$

5. DISCUSSION

There remains the question of why the shape and the peak of the contrast sensitivity functions of this data are so different from the van Nes and Bouman[4] data. Two factors that would force the results in this direction are the spatial frequency transfer functions of the displays and of the observers. No attempt was made by Peterson *et al.* to check the optical correction of the observers. In Figure 3 the small symbols corresponding to horizontal and vertical gratings appear in pairs, with the horizontal grating frequency just to right of the vertical grating frequency. At high spatial frequencies the sensitivity for vertical gratings is much less than that for horizontal gratings, a result consistent in direction with cone spacing[9], but also with some low pass filtering by the display electronics. Other possible factors include the experience and selection of the observers. The above factors seem generally in the direction of making the results more practical, but if one wants to insure that artifacts are not visible, one could use the templates as fit to the van Nes and Bouman data, at the cost of less compression. One major difference between the two studies was that the small area signals of Peterson *et al.* were on a dark surround instead of an adapting field of the same mean luminance. It is well known that adaptation near the masking level is optimal for brightness discrimination[10]. This raises the interesting question of how the luminance term of the model should be computed from the image luminance distribution on the display to optimize compression.

A luminance-based model appears to adequately account for the variations between the R, G, and B sensitivities found by Peterson *et al.*[3]. There seems to be no need to introduce chromatic mechanisms to account for their data. An advantage of separate compression for R, G, and B images is that only the luminance of the output guns is needed to generate the quantization matrices. Transformation to luminance and two iso-luminant color coordinates will give better compression, but then it may be critical that the reconstruction be done so that the iso-luminant color planes are accurately reconstructed at the display. Whatever color transformation is used, at least one dimension will be luminance dominated at high spatial frequencies. Our model provides an experimentally based method for generating quantization matrices for such dimensions for a range of display parameters and viewing conditions.

6. REFERENCES

1. G. Wallace, "The JPEG still picture compression standard", *Communications of the ACM*, vol. 34, no. 4, pp. 30-44, 1991.
2. D. LeGall, "MPEG: A video compression standard for multimedia applications", *Communications of the ACM*, vol. 34, no. 4, pp. 46-58, 1991.
3. H. A. Peterson, H. Peng, J. H. Morgan, W. B. Pennebaker, "Quantization of color image components in the DCT domain", in B. E. Rogowitz, M. H. Brill, J. P. Allebach, eds., *Human Vision, Visual Processing, and Digital Display II*, Proc. SPIE, vol. 1453, pp. 210-222, 1991.
4. F. L. van Nes, M. A. Bouman, "Spatial modulation transfer in the human eye", *Journal of the Optical Society of America*, vol. 57, pp. 401-406, 1967.
5. L. A. Olzak, J. P. Thomas, "Seeing spatial patterns", in K. R. Boff, L. Kaufman, J. P. Thomas, eds., *Handbook of Perception and Human Performance*, pp. 7.1-7.56, Wiley, New York, 1986.
6. D. H. Kelly, "Visual processing of moving stimuli", *Journal of the Optical Society of America A*, vol. 2, pp. 216-225, 1985.
7. G. C. Phillips, H. R. Wilson, "Orientation bandwidths of spatial mechanisms measured by masking", *Journal of the Optical Society of America A*, vol. 1, pp. 226-232, 1984.
8. A. B. Watson, "Detection and recognition of simple spatial forms", in O. J. Braddick, A. C. Sleigh, eds., *Physical and Biological Processing of Images*, Springer-Verlag, Berlin, 1983.
9. D. R. Williams, "Topography of the foveal cone mosaic in the living human eye". *Vision Research*, vol. 28, pp. 433-454, 1988.
10. K. J. W. Craik, "The effect of adaptation on differential brightness discrimination", *Journal of Physiology*, vol. 92, pp. 406-421, 1938.

Figure 1

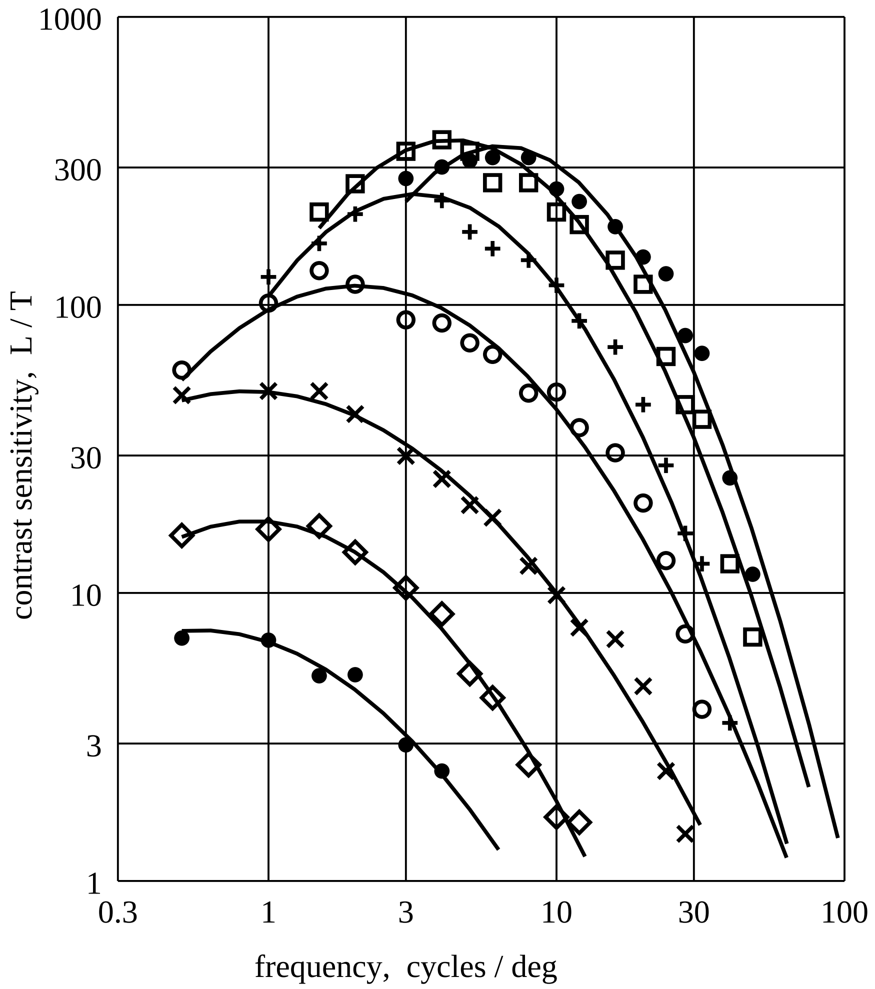
1000
300
100
30
10
3
1
contrast sensitivity, L / T
0.3
1
3
10
30
100
frequency, cycles / deg

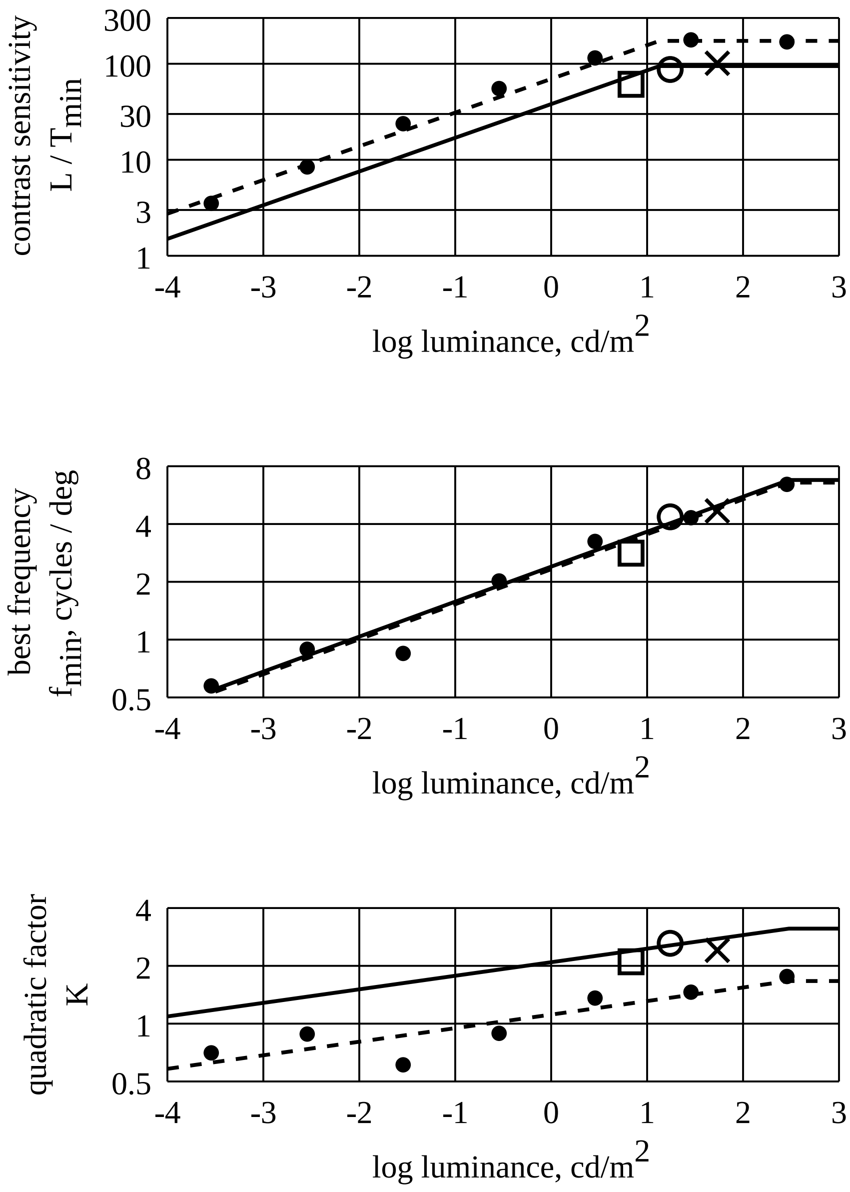
contrast sensitivity
L / T_{min}
300
100
30
10
3
1
-4
-3
-2
-1
0
1
2
3
log luminance, cd/m^2
best frequency
f_{min}, cycles / deg
8
4
2
1
0.5
log luminance, cd/m^2
quadratic factor
K
4
2
1
0.5
log luminance, cd/m^2

Figure 3

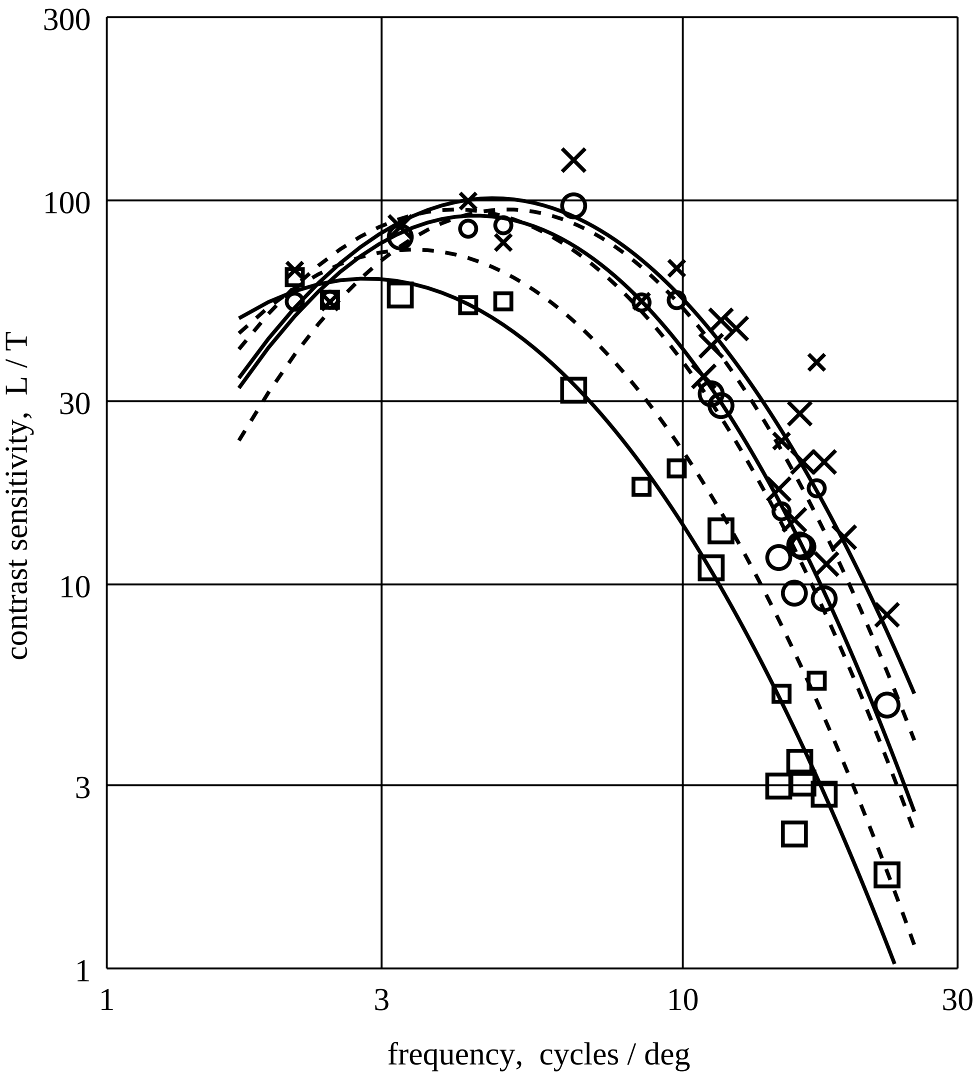

300
100
30
10
3
1
1
3
10
30
contrast sensitivity, L / T
frequency, cycles / deg

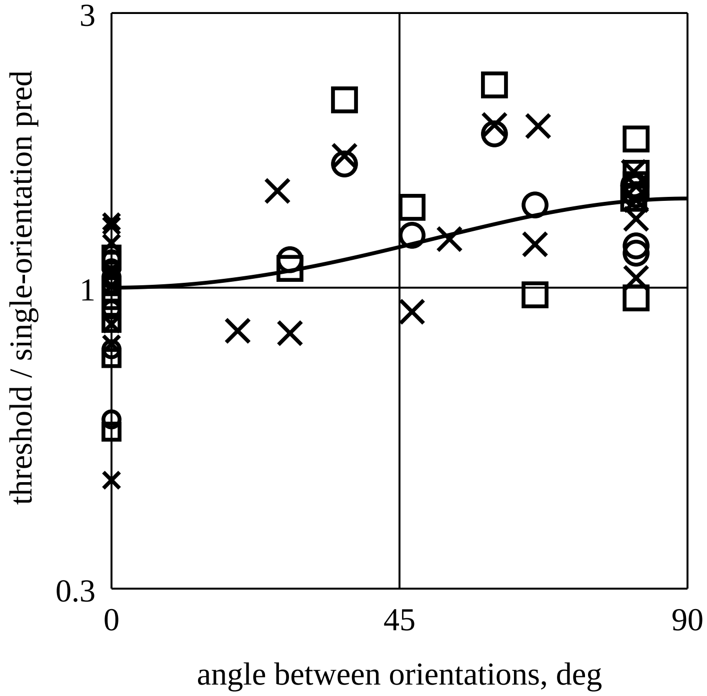
3
1
0.3
threshold / single-orientation pred
0
45
90
angle between orientations, deg

Figure 4