

The Effect of Quantisation Error on the Accuracy of Colour Transforms.

Stefan Livens
Agfa-Gevaert N.V.
Septestraat 27, B-2640 Mortsel, Belgium.

Abstract

Colour data always has a limited accuracy, which is caused by two separate effects. Experimental random error is induced by the actual measurement. Quantisation error results from the quantised representation of the data. Both error types propagate through the processing, while quantisation in intermediate steps can cause additional error. Existing methods deal only with the propagation of random error, or with the effect of initial quantisation.

We study the combined effect of both error types and arrive at the result that their error variances can be added. For the case of successive quantisations, we propose a computational method for calculating the output error. We illustrate this by numerical data for various transforms.

1. Introduction

Processing of colour information usually starts from colour measurements. They have one common property: as experimental data they always contain some noise. This noise causes uncertainty about the exact values, called experimental error. Depending on its origin, the noise can have different characteristics.

When the measurements are processed, the experimental errors propagate and yield errors on the output values. These can become much larger than the input errors, especially when non-linear transforms are involved.

Determining the output error is important because it is a measure of the reliability and robustness (invariance for small input variations) of different colour transforms. The basic mathematical tools for error propagation analysis are well established, although seldom used in practice [5].

In these methods an important type of errors is not considered, namely those introduced by representing data in a quantised way. The initial quantisation is usually chosen in such a way that the additional error is small. However, intermediate processing steps act on quantised data and requantise them, which can cause much larger errors.

We will propose a method for dealing with quantisation errors and combining them with experimental errors. We will establish a computational method for evaluating the effect of quantisation in intermediate processing steps.

2. Error Propagation

2.1. Error Propagation Methods

The goal of error propagation is to determine the uncertainty or error on the output data, given the errors on the input data.

2.1.1. Analytical Error Propagation

Often, the measurement of a quantity χ is affected by an additive error e_χ (a random variable with zero mean and variance σ_e^2). The output of the measurement is also a random variable:

$$x = \mu_\chi + e_\chi$$

For an unbiased estimation, the expectation value μ_χ equals the true value χ. When x is transformed to $y = f(x)$, y is also a random variable. When the function f is sufficiently smooth, the error variance of y can be computed analytically using a Taylor series expansion [5]. This is the principle of analytical error propagation.

When more than one variable is involved, the random errors and their correlations can be described by a covariance matrix. The main problem of analytical error propagation is that its use is limited to continuous functions, which excludes quantisation operations.

2.1.2. Interval Computation

When the error is confined to a well defined interval, interval computations can be used. The error bounds on the input are processed to yield an output error interval, which should be interpreted as worst case error interval. This method is well suited for calculating the effects of input quantisation. For multiple, correlated variables, the error estimates often become unreliable.

2.1.3. Stochastic Methods

In stochastic error propagation, the distribution of the input errors is modelled. From this, a large number of possible inputs are generated and processed. The distribution of the outputs yields an estimate for the output errors.

2.2. Error Propagation in the Colour Literature

Analytical error propagation is applied to colour transforms in [1]. Based on sensor uncertainty, the uncertainty in the transformed co-ordinates is derived. Results are presented as ellipsoids of statistical confidence in CIE L*a*b* space. A practical application of this, starting from a camera signal through colorimetric transforms, is given in [2].

The best known example of how non-linear calculations can lead to large output errors is the instability of hue co-ordinates near the achromatic axis. An analysis of this effect is given in [7], where a sensor noise model is given, and an analytical derivation is given for normalised $r\ g\ b$ values and for hue.

Quantisation effects have been studied in [6]. Only errors due to the quantised input space are considered. Calculations are made that show how a uniform input quantisation results in non-uniform output errors. For the case of XYZ to CIEL*a*b*, formulae are derived for determining the maximal errors. A striking illustration of the effects of quantisation is given in [3]. Geometric patterns in output colour spaces are shown. It is demonstrated that the application of different γ values yields different quantisation patterns.

3. Quantisation Error Modeling

In order to address the effects of experimental and quantisation errors simultaneously, we combine methodology from analytical error propagation and interval computations.

3.1. Error Measures

We first need a way to quantify error that is meaningful for both experimental and quantisation errors. The maximal error (the size of the error interval) is the natural error measure for interval computations. However, for normally distributed errors (a reasonable model for experimental errors), the error interval is always infinite. The variance, basic error measure for normal errors, can also be used to describe quantisation errors. Therefore, we will use it as error measure in all following computations.

3.2. Quantisation Error

We consider quantisation into levels on a regular grid, with distance between levels δ. This maps continuous data χ (without any random error) onto discrete levels x, hereby effectively binning the data.

$$x = Q(\chi) = \chi + q(\chi)$$

Because x is always mapped to the nearest discrete level for the bias $q(\chi)$ and its variance it holds that:

$$-\delta/2 < q(\chi) < \delta/2$$

$$0 < \sigma_q^2 < (\delta/2)^2$$

Since the alignment of the quantisation grid is independent of the measured value, we can assume all values of $q(\chi)$ to be equally probable [4]. By averaging over all values that fall into the same bin, we obtain the average error:

$$< \sigma_q^2 >=< \sigma_i^2 >= \frac{1}{\delta}\int_{x_i-\frac{\delta}{2}}^{x_i+\frac{\delta}{2}} (x_i-\chi)^2 d\chi = \frac{1}{3}(\frac{\delta}{2})^2$$

3.3. Combining Random and Quantisation Error

Whenever the measurement of the true value is corrupted with noise, the situation becomes more complicated:

$$x = Q(\chi + e_\chi) = \chi + q(\chi, e_\chi)$$

For different realisations of e_χ, the measured value can fall into different bins. Therefore we calculate the probabilities P_i that the measured value ends up in every bin. To simplify the notation, we introduce the notation $\alpha = \chi - x_i$, where x_i is the nearest level to χ.

$$P_i(\alpha) = \Phi((i+1/2)\delta - \alpha) - \Phi((i-1/2)\delta - \alpha)$$

with Φ the cumulative normal distribution function with mean μ_n and variance σ_e^2:

$$\Phi(x) = \frac{1}{\sigma_n\sqrt{2\pi}}\int_0^x e^{\frac{(t-\mu_n)^2}{2\sigma_e^2}} dt$$

The error for one position is:

$$\sigma_i^2(\alpha) = (x_i - \alpha)^2$$

The total error becomes:

$$\sigma_{tot}^2(\alpha) = \sum_{i=0}^{n} P_i(\alpha)\sigma_i^2(\alpha) \approx \sum_{i=-\infty}^{\infty} P_i(\alpha)\sigma_i^2(\alpha)$$

For the second equality, it is assumed that the total probability is confined within the measurable range.

We average over α to obtain the average error, :

$$< \sigma_{tot}^2 >= \int_{-\frac{\delta}{2}}^{\frac{\delta}{2}} \sum_{i=-\infty}^{\infty} P_i(\alpha)\sigma_i^2(\alpha)$$

It can be shown that:

$$< \sigma_{tot}^2 >= \sigma_e^2 + < \sigma_q^2 >$$

This means that quantisation error and random error can be combined by adding their variances.

Besides average error values, we can compute extreme values. The minimum occurs for $\alpha = 0$ and the maximum for $\alpha = \delta/2$. The evolution of these error bounds, together with the average error, is shown in Fig. 1.

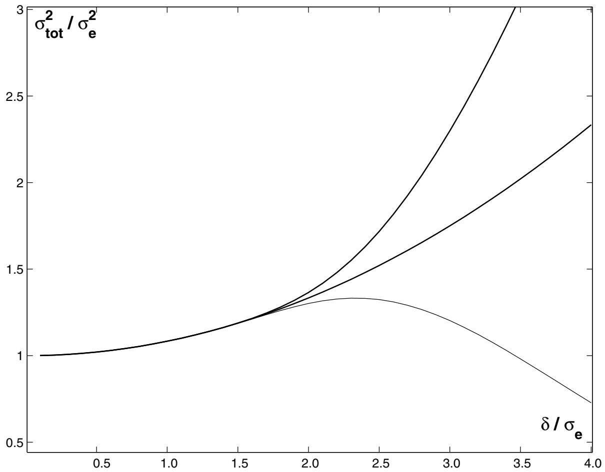

Figure 1: Relative increase in error variance $\sigma^2_{tot}/\sigma^2_n$ due to quantisation for varying ratio δ/σ^2. The average, minimal and maximal curves are shown.

3.4. Successive Quantisations

The previous averaging calculations were based on the random and equally probable occurrence of measured values with respect to the alignment of the quantisation grid. When several quantisations are performed successively, this assumption is no longer valid. The combined quantisation effect is not the sum of the individual quantisation variances, has to be calculated in a different way.

The first quantisation bins the data into m levels. The second quantisation acts on the quantised levels x_i and requantises them into n levels $Q_n(x_i)$. The average error can still be calculated as an average over an original bin:

$$< \sigma_i^2 >= \frac{1}{\delta_m}\int_{x_i-\frac{\delta_m}{2}}^{x_i+\frac{\delta_m}{2}} (Q_n(x_i)-\chi)^2 d\chi$$

Using $\alpha = \chi - x_i$, this can be rewritten to:

$$< \sigma_i^2 >= \frac{1}{\delta_m}\int_{-\frac{\delta_m}{2}}^{\frac{\delta_m}{2}} (Q_n(x_i)-x_i+\alpha)^2 d\alpha$$

$$< \sigma_i^2 >= \frac{1}{3}(\frac{\delta_m}{2})^2 + \frac{1}{\delta_m}(Q_n(x_i)-x_i)^2$$

The first term is the error caused by the initial quantisation. The additional term describes the increase in error caused by the *shifting* of levels x_i away from their position in the middle of the bin. This term is only zero when all levels remain unchanged, which only occurs when n is a multiple of m. This includes the trivial case of two identical quantisations.

The effect can be broken down into numbers that are relatively prime. E.g., a conversion from 256 to 200 levels can be treated as a conversion from 32 to 25 levels, repeated 8 times.

In the more general case a function f is applied between the first and second quantisation. The error becomes:

$$< \sigma_i^2 >= \frac{1}{\delta_m}\int_{x_i-\frac{\delta_m}{2}}^{x_i+\frac{\delta_m}{2}} (Q_n(f(x_i)-f(\chi))^2 d\chi$$

Because f can be an arbitrary function, this error can only be computed numerically. However, when f is sufficiently smooth, meaning that it can be approximated linearly over the size of a single bin, we can expand $f(\chi)$ in a Taylor series around x_i:

$$< \sigma_i^2 >= \frac{1}{\delta_m}\int_{\frac{\delta_m}{2}}^{\frac{\delta_m}{2}} (Q_n(f(x_i)-(f(x_i)+f'(x_i)\alpha+...))^2 d\alpha$$

$$= \frac{1}{3}(\frac{\delta_m}{2})^2 f'(x_i)^2 + \frac{1}{\delta_m}(Q_n(f(x_i))-f(x_i))^2 + ...$$

Again, the first term is the initial quantisation error per bin, now rescaled with $f'(x_i)^2$, the squared derivative of f at x_i. The second term expresses the additional error due to shifting caused by the second quantisation.

3.5. Multiple Variables

The previous results were all obtained for a single variable. For a multiple variable space the Euclidean distance is the obvious choice as a measure of the total error. It allows to calculate the error variances of the different variables separately and combine them afterwards.

When multiple variables and their errors are uncorrelated, they obviously can be treated exactly like single variables. It can be shown that this also holds for correlated variables. The variances due to experimental and quantisation errors can be calculated separately and combined by adding variances.

To compute the error associated with a certain output variable that is a mixture of input variables we proceed as follows. We use conventional covariance error propagation to compute the error stemming from the input experimental error. We compute the quantisation error of the input-output system, and finally add both errors. This gives all necessary tools needed to study realistic colour transforms.

From a practical point of view the calculations become very computationally expensive if we use a 3 dimensional space with more than 8 bits per co-ordinate. This is the case for colour transforms in practice. E.g. in a conversion from CIE XYZ to CIE L*a*b*, in order to get input quantisation errors of the order of $1\Delta E$ one has to start

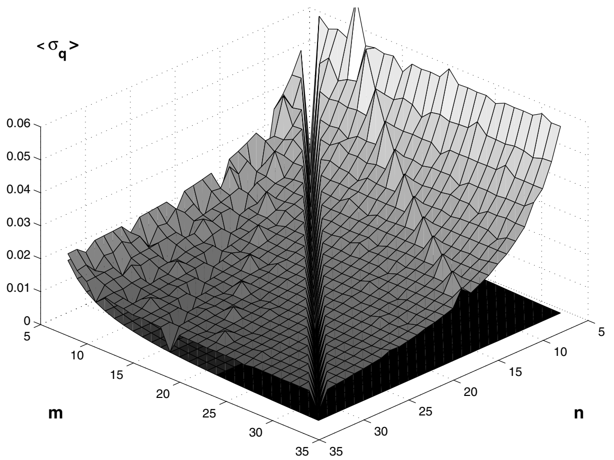

Figure 2: $< \sigma_{q(m,n)} >$ *for various* (m, n) *combinations in a* $m \to n$ *requantisation. Underneath is* $< \sigma_{q(m)} >$*, the error of the initial quantisation.*

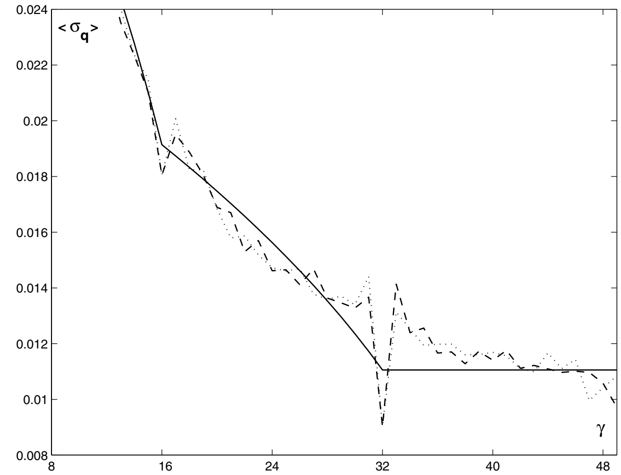

Figure 3: $\sigma_{q(m,n,p)}$ *a transform* $m \to n \to 32$ *with* $8 < n < 50$*. The solid line is for m = 64, the others for* $m = 63$ *and* $m = 65$*.*

for a 13 bit quantisation in the XYZ space (cfr. [6]). Our future work will focus on developing efficient computational methods which will enable more practical evaluation of colour conversions.

4. Simulation Experiments

We now apply the formulae of the previous section in a series of simulation experiments. We compute the increase in average quantisation error due to requantisation. For this we evaluate the expressions for $< \sigma_i^2 >$ numerically. All results are for uniform input densities and for transforms acting on the unity interval.

4.1. Transform from m to n levels

The simplest case is a transform with $F(x) : x \to x$, where only the number of quantisation levels changes. Results are shown in Fig. 2. The smallest errors lie on the diagonal (identical quantisers). The error also increases when the number of levels is increased ($n > m$) and the order of the quantisations makes a difference: $m \to n$ yields a different result than $n \to m$.

4.2. Consecutive Transforms ($m \to n \to p$)

More than two quantisations can be applied in succession. The result does not only depend on the number of input and output levels, but also on the intermediate quantisation. We present results for a transform series for which the number of initial input and final output levels are both fixed. A single intermediate quantisation is considered, for which the number of levels is varied.

Results are shown in Fig. 3. When m is a multiple of p, the error decreases smoothly with increasing n, untill the value of the direct transform ($m \to p$ is achieved at $n = p$. For all $n > p$ the error remains constant. When m and p are relatively prime, the evolution shows the same tendency but the curve is much more erratic, with a sharp dip at $n = p$.

4.3. Transform with γ

Non-linear functions in colour transforms very often take the form of an exponential, with γ as exponent: $x \to x^\gamma$. The result of this is that part of the range is expanded, and the other part compressed. When the output is requantised, the new distribution makes that the errors are no longer equal for all bins. For the expanded part, the error increases slightly because of the repositioning. For the compressed part, bins are pooled and the increase is more drastic. The result is shown in Fig. 4. The error has a sharp minimum for $\gamma = 1$ and oscillates with a frequency proportional to n. Interestingly, a small deviation of γ from 1 can cause an error increase as large as that of a large deviation.

4.4. Transform with varying slope

Often the transformation function is not a mathematical function, but an interpolation between some measured values. As an example of this, we consider a piecewise linear mapping. We use a transform with 5 equidistant points. The first and last points are fixed to 0 and 1, the three middle points can vary according to a normal distribution with standard deviation σ_p.

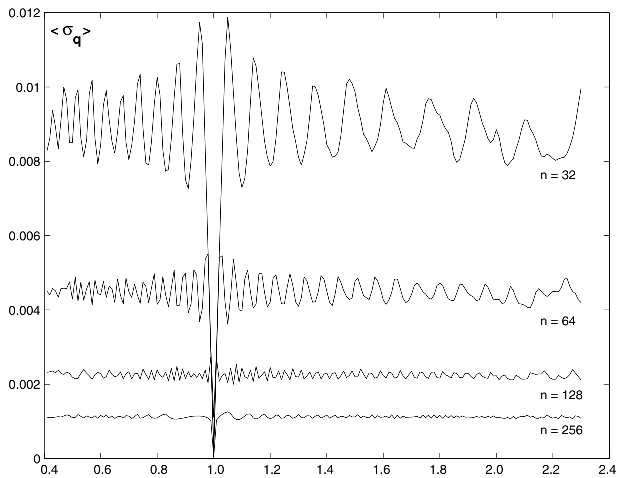

Figure 4: *σ_q for a requantisation $n \rightarrow n$ with a function $x \rightarrow x^\gamma$ ($0.4 < \gamma < 2.2$) in between.*

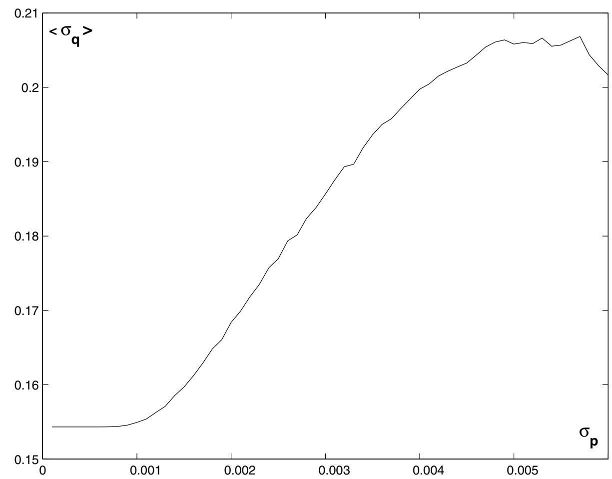

Figure 5: *σ_q for a piecewise linear transform, for which the slope varies with a normal perturbation term σ_p ($0 < \sigma_p < 0.006$). The input and output quantisations are fixed at 256 levels.*

Because of the random position of the intermediate points, it was needed to resort to stochastic error propagation method. We repeated the experiment 10000 times until a statistically large enough sample was obtained.

It can be seen from Fig. 5 that below a certain limit, no additional error is introduced as all the quantised input levels remain mapped to the same output levels. For larger deviations, σ_q increases nearly linearly with σ_p. Note that the erratic behaviour for $\sigma_p > 0.0045$ is caused by a limitation of the experiment and not a quantisation effect.

5. Conclusions

We investigated the effect of quantised representations on the accuracy of colour specifications. We derived the necessary theorical results for combining quantisation errors with experimental random errors. The quantisation error can be expressed as a variance, and can be combined with other error sources by adding variances.

Secondly we proposed a method for dealing with successive quantisations. We presented a formula which can be used for practical numerical investigation of the effect. For every initial quantised level the extra error due to requantisation can be computed. This method was illustrated with various examples. Both methods can be used together to incorporate the effects of quantisation into a general error analysis.

To analyse a complete system, the procedure should be the following:
1. The experimental errors should be estimated and propagated using conventional error propagation.
2. The quantisation errors should be computed numerically for the total system.
3. The total error is obtained by adding variances.

References

[1] P. D. Burns and R. S. Berns. Error propagation analysis in color measurement and imaging. *Color Research and Applic.*, 22:280–289, 1997.

[2] P. D. Burns and R. S. Berns. Image noise and colorimetric precision in multispectral image capture. In *Proc. Sixth Color Imaging Conf.*, Scottsdale, Arizona, November 1998.

[3] K. Godfrey. Digital quantization patterns in C I E spaces and artifacts from misaligned color planes. In *Proc. Sixth Color Imaging Conf.*, Scottsdale, Arizona, November 1998.

[4] Haralick and Shapiro. *Computer and Robot Vision*. Addision Wesley, Reading, MA, 1993.

[5] R. M. Haralick. Covariance propagation in computer vision. In *Proc. Workshop on Performance Characteristics of Vision Algorithms, ECCV-96*, Cambridge, England, April 1996.

[6] H. Ikeda, W. Dai, and Y. Higaki. A study on colorimetric errors caused by quantizing color information. *IEEE trans. on Instrumentation and Measurement*, 41(6):845–849, December 1992.

[7] H. Stokman and T. Gevers. Error propagation due to color space transforms. In *Proc. Sixth Color Imaging Conf.*, Scottsdale, Arizona, November 1998.