

The Accuracy of Sub-Pixel Localisation in the Canny Edge Detector

Peter Rockett

Department of Electronic & Electrical Engineering

University of Sheffield

Sheffield S1 3JD, UK

p.rockett@shef.ac.uk

Abstract

We examine the accuracy of the sub-pixel location of edges in the Canny detector using a discrete step edge model and Monte-Carlo simulation. Comparison is made with previously published analysis in the continuous domain. In this paper we identify potential systematic errors due to a combination of the width of the smoothing kernel and the quadratic interpolation scheme which we show can be reduced to less than one-fiftieth of a pixel with look-up table. We present data which permits prediction of the likely error in the sub-pixel estimate as a function of grey level step height in the presence of noise.

Keywords: Edge detection, Performance benchmarking, Canny algorithm, Sub-pixel interpolation, Monte-Carlo simulation.

1 Introduction

The interpolation of edge locations to sub-pixel accuracy is important in a number of applications of computer vision such as: the estimation of disparities in edge-based stereo, the calculation of optic flow as well as gauging in industrial inspection systems. Probably the most widespread edge detection algorithm used for obtaining sub-pixel estimates is due to Canny [1] since the image gradient estimates for this algorithm are held to be sufficiently stable to make interpolating to sub-pixel locations worthwhile. Limited work, however, has been published on the accuracy of these estimates and since edge detection is frequently used as the first in a series processing stages, a knowledge of the errors due to the Canny algorithm is needed to allow errors to be propagated through the whole processing sequence. The objective of the present paper has been to produce error estimates which will be useful in practice to assess quantitatively the performance of vision algorithms for which the Canny detector forms a first stage.

Although Canny's original paper has been extremely influential, his original implementation of convolving the image with each of a large number of orientated filter kernels in order to find the maximum response is very inefficient and this has led to many alternatives. The implementation studied here is, we believe, highly typical of the "Canny" algorithms employed in most computer vision laboratories and follows the implementation described by Jain *et. al.* [2] and Faugeras [3]:

BMVC99

- (a) Convolution of the image with a Gaussian smoothing mask.
- (b) Estimation of x - and y -gradients using simple gradient operators $G_x = [-1 \ 0 \ -1]$ and $G_y = [-1 \ 0 \ -1]$ followed by non-maximal suppression to locate the local maxima.
- (c) Hysteresis thresholding on the gradient magnitude values.
- (d) Finally, sub-pixel interpolation by least-squares fitting of a quadratic polynomial to the gradient (approximately) normal to the detected edge.

For brevity, we refer throughout this paper to this sequence of processing stages as "the Canny algorithm" although we acknowledge other groups may employ variants on this. For the purpose of what follows, we note that in the case of an ideal step edge, processing stages (a) and (b) produce a filter output normal to the edge which is Gaussian with a peak corresponding to the edge location.

The errors on the sub-pixel location obtained have been studied previously by a number of authors [4,5,6] all of whom have carried-out purely mathematical analyses. Importantly, all of these authors have worked: (i) solely in the continuous (as opposed to the discrete image) domain, and (ii) have defined the estimate of the edge location as the stationary (maximum) point of the linear filter's output. This latter point is reasonable in the continuous domain but how this relates to estimates interpolated by (d) above in the discrete domain is not considered. Kakarala & Hero [4] comment that their error bound derived in the continuous domain will be degraded in passing to the discrete domain.

In this paper we consider the errors in estimating the location of a step edge defined in the discrete domain by Monte-Carlo simulation under typical experimental conditions. This allows us to make direct estimates of the sub-pixel location error applicable to digital images. We find that for a commonly-used size of smoothing kernel the quadratic interpolation scheme exhibits systematic errors which are amenable to correction with a look-up table. We present results as a function of step height which permit estimation of the sub-pixel location error in practical situations. Further, we also observe that for certain pathological edge orientations and sub-pixel positions, the Canny detector - indeed any detector which locates locally maximal gradient peaks - breaks down due to the image sampling.

The experimental methodology is described in Section 2 and the results for various conditions in Section 3; in Section 4 we discuss some of the ramifications of the present work.

2 Experimental Method

Although it is possible to produce closed-form results for the distribution of sub-pixel errors in the continuous domain, translating these to the discrete image domain is not straightforward. We have therefore chosen to generate an image patch directly in the discrete domain and process this with the Canny algorithm. Since we know the 'ground truth' location and orientation of the edge we can readily estimate errors.

The step edge model employed is very similar to that used by Lyvers & Mitchell [7] and is depicted in Figure 1. The model is parameterised in terms of: the pixel displacement from the origin, ρ ; the edge orientation, θ ; the high and low grey-level values, I_H




Figure 1: Step edge model

and I_L . The patch size was chosen relative to the dimensions of the Gaussian smoothing kernel to avoid edge effects. Each pixel is assumed to exactly abut its neighbour and to have unit dimension. Individual pixel intensities were calculated as the sum of the high and low intensity contributions weighted by their respective projected areas on each pixel site.

The Canny code employed was specifically written for this work and used IEEE-754:1985-compliant single-precision floating point numbers at all stages of processing to guard against artefacts caused by truncation errors in integer arithmetic. The pixel intensities generated in floating point format by the step edge model above were rounded to one of 256 levels to emulate the analogue-to-digital conversion process of a framegrabber. The width of the Gaussian smoothing filter was fixed at 1.0 since this appears a sensible and popular choice for suppressing image noise, as opposed to carrying-out genuine multi-scale processing; the Gaussian filter was truncated at a value of 0.01.

The sub-pixel interpolation scheme was identical to that described by Jain *et. al.* [2] and quantises the pixel orientation estimate into one of the eight principal directions of the compass (*i.e.* N-S, SW-NE, E-W, *etc.*) A quadratic polynomial is then fitted by minimising a conventional least-squares measure to the appropriate three neighbouring gradient values and the location of the edge determined by identifying the stationary point of this polynomial. It is the distribution of this edge location interpolated to sub-pixel acuity that concerns us here.

For studying the effects of noise, the image patch was corrupted with independent zero-mean Gaussian-distributed noise with a variance of 4.0, this value being typical of a range of commercial CCD cameras examined in this laboratory. One hundred thousand trials were used throughout for the Monte-Carlo simulations.

3 Results

Due to symmetries, the response of the detector is periodic; results are therefore only presented for θ in the range $[0..45^\circ]$ since this is sufficient.

BMVC99




Figure 2: Systemetic errors in the sub-pixel location estimate vs. true edge orientation




Figure 3: Illustration of detector breakdown. (a) shows the gradient map and (b) the non-maximally suppressed image. Note the mislabelling of the central pixel in (b).

BMVC99




Figure 4: Region in which detector breakdown occurs

3.1 Systematic Effects

Initially we examined the errors in the computed estimates, $\hat{\rho}$ and $\hat{\theta}$ in the absence of noise to determine whether the Canny algorithm displays any systematic biases. The results are shown in Figure 2 where it can be seen that while there is no bias when $\rho = 0$, increasing values of ρ result in greater peak values of bias which reach maxima at $\theta = 45^\circ$. There is also a discontinuous increase in error caused by the switch at $\theta = 20.5^\circ$ of the direction of sub-pixel interpolation from E-W to SW-NE¹. For the width of the smoothing kernel employed here ($\sigma = 1$) the assumed quadratic function is an acceptable fit close to the peak in the E-W direction since the data point spacing is unity - the pitch of the pixel lattice. For diagonal NE-SW interpolation, the data point spacing is $\sqrt{2}$ - too far down towards the tails of the underlying Gaussian function to yield an acceptable fit. When ρ is very small the quadratic still predicts the peak location well although the predicted (although unimportant) value of peak gradient is seriously in error. For larger values of ρ the peak location estimate is significantly in error reaching a maximum of about 0.11 pixels when $\theta = 45^\circ$. We have verified that using larger kernel widths reduces this peak error.

Since the underlying function which describes the spatial variation of gradient values is a Gaussian, it would be thought that fitting a Gaussian rather than a quadratic would yield a bias-free estimate of ρ regardless of value. Indeed, Kakarala & Hero [4] demonstrated theoretically that such a procedure would achieve the Cramer-Rao bound although did not demonstrate this on real data. Attempts to fit a Gaussian function to the present data proved unsuccessful since the normal equations generated by the least-squares problem proved to be too ill-conditioned even when using the Levenburg-Marquardt algorithm - a notably robust technique. Clearly more than three gradient values are required to fit a Gaussian in practice.

¹The change in the direction of interpolation takes place at 20.5° as opposed to the expected value of 22.5° because the direction is determined from the detector's *estimate* of orientation, not the true orientation value. Although not shown here, the orientation error for a true θ value of 22.5° and is around 2° leading to the observed switchover.

The estimates of systematic orientation error obtained here agree closely with the results of Lyvers & Mitchell [7].

During the process of mapping the behaviour of the Canny detector we observed that for certain pathological combinations of ρ and θ the detector exhibited a breakdown originating in the non-maximal suppression stage. An example of this phenomenon is shown in Figure 3 for $\rho = 0.47$ and $\theta = 20^\circ$. Figure 3(a) shows the gradient magnitudes of the smoothed image and (b) shows the non-maximally suppressed image. It can be seen that the central pixel of Figure 3(b) is not labelled as a local maximum - rather the pixel to the right is erroneously labelled. In fact, the difference between these two gradient values is only $\sim 0.5\%$ but this is sufficient to bring about a labelling breakdown. Additional work is required to establish whether this effect is significant in the presence of noise which may well mask the phenomenon. Figure 4 maps-out the regions (ρ, θ) which exhibit this breakdown behaviour.

3.2 Effects of Noise




Figure 5: Influence of noise on accuracy for a step edge ($\rho = 0.4$)

Since edge labelling is a form of symbol detection, clearly the performance of any algorithm will be influenced by the signal-to-noise ratio; for edge detection this can be expressed as the quotient of the step magnitude and the standard deviation of the pixel noise. Figure 5 shows the effects of noise for a high-contrast step with a height of 150 intensity units equivalent to a signal-to-noise ratio of around 38dB for $\rho = 0.4$. Clearly the systematic deficiencies are the largest source of error for $\theta \geq 20^\circ$ which suggests that using a look-up table to correct for the systematic effects would be worthwhile.

3.3 Correction of Systematic Errors

We have sought to reduce the systematic bias in the detector in the by employing a look-up table indexed by both the estimated orientation $\hat{\theta}$, and the estimated location $\hat{\rho}$, obtained

BMVC99




Figure 6: Systemetic location error after correction ($\theta = 45^\circ$)

from the quadratic interpolation process described above. In this way we have reduced the magnitude of the maximum bias in the sub-pixel location to less than one-fiftieth of a pixel; Figure 6 shows the bias after correction as a function of true pixel location for $\theta = 45^\circ$. Other combinations show smaller deviations than this.

3.4 Edge Location Variability under Noise




Figure 7: Standard deviation of sub-pixel location estimate as a function of step height

Although correction of the systematic error is independent of the step height, in the presence of noise the edge location accuracy will be subject to some variability which will be

affected by step height. Assuming a constant noise variance, we can define signal-to-noise ratio as the quotient of the step height and the standard deviation of the noise. Since one of the objectives of the present work has been to obtain estimates of the error of sub-pixel interpolation, we have to consider step height as a parameter. Figure 7 shows the standard deviation of the sub-pixel estimate, $\hat{\rho}$ averaged across all edge orientations together with error bars on this measure indicating the degree of variability for a given step height. It can be seen that for high contrast edges sub-pixel estimates of edge location are accurate to around 0.01 pixels while the error for weaker edges increases rapidly for reducing step height reaching around a tenth of a pixel for step heights of 10 intensity units. This latter point represents a signal-to-noise ratio of $\sim 14\text{dB}$.

4 Discussion and Future Work

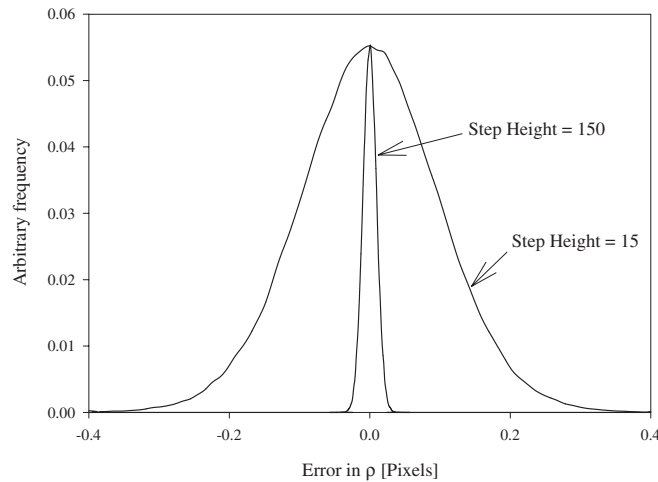


Figure 8: Distribution of sub-pixel location error for step height of 150 and 15 intensity units. Note: the normalisations are arbitrary solely to permit ready visual comparison.

Clearly the present work assumes an ideal step edge. How close this comes to reality depends on effects such as diffraction/aberration in the imaging optics compared to the physical dimensions of the pixel sites in the CCD sensor. The degree of coherence of the illumination also affects the intensity profile of a projected 'knife' edge. The methodology described here is amenable to treating such situations since the input step edge patch could be convolved with an appropriate point spread function to simulate, for example, diffraction or defocus effects [8]. Additionally, the present methodology can be extended to quantify errors in the detection of edges; this area of work will be addressed in the future.

It is apparent from Figure 7, the error associated with sub-pixel estimates is a function of step height which corresponds to edge strength. The propagation of errors could potentially be automated using, say, interval arithmetic if Figure 7 were replotted in terms of edge strength giving an 'online' estimate of error.

The breakdown in the detection algorithm was unexpected. In fact, similar breakdown will occur in any detector which locates local peaks by non-maximal suppression. From Figure 4 it is apparent that the problem is most severe for edges which are almost vertical - just the sort of edges from which stereo disparity estimates may be calculated. It is clear that this defect will contribute a certain number of outliers to any subsequent processing stage although the severity of the problem needs to be established.

Finally, Koplowitz & Greco [6] concluded that for poor signal-to-noise ratios (*i.e.* small step heights in the present work) the error distribution has such heavy tails that its variance is undefined. In Figure 8 we show the (arbitrarily scaled) distributions of the sub-pixel estimates for step heights of 150 and 15 intensity units corresponding to signal-to-noise ratios of $\sim 38\text{dB}$ and $\sim 17\text{dB}$, respectively. We observe that in both cases the distributions are Gaussian. In fact, the example distribution with heavy tails shown by Koplowitz & Greco is for a signal-to-noise ratio of $\sim 6\text{dB}$. It is questionable whether such small changes in grey level can be judged statistically significant and we conclude that for step edges which are detectable in practice, the error distribution can be taken as normal.

5 Conclusions

In this paper we have examined the errors in the sub-pixel location of edges in the Canny detector by Monte-Carlo simulation using a discrete step edge model; we suggest that the use of a discrete model gives more direct results than carrying-out the analysis in the continuous domain. We identify potential systematic errors due to a combination of the smoothing kernel width and the quadratic interpolation which can be reduced to less than 0.02 pixels using a correcting look-up table. We present data which predict the likely error in the sub-pixel estimate as a function of step height. Finally, we conclude that the distribution of the sub-pixel estimates is Gaussian for all signal-to-noise ratios of practical interest.

References

- [1] J.F.Canny, "A computational approach to edge detection", IEEE Trans. Pattern Analysis & Machine Intelligence, 8(6), 679 (1986)
- [2] R.Jain, R.Kasturi & B.G.Schunk, "Machine Vision", McGraw-Hill, 1995.
- [3] O.D.Faugeras, "Three-dimensional computer vision: a geometric viewpoint", MIT Press, 1993.
- [4] R.Kakarala & A.O.Hero, "On the achievable accuracy in edge location", IEEE Trans. Pattern Analysis & Machine Intelligence, 14(7), 777 (1992)
- [5] H.D.Tagare & R.J.P.deFigueiredo, "On the localization performance measure and optimal edge detection", IEEE Trans. Pattern Analysis & Machine Intelligence, 12(12), 1186 (1990)

BMVC99

- [6] J.Koplowitz & V.Greco, "On the edge location error for local maximum and zero crossing edge detectors", IEEE Trans. Pattern Analysis & Machine Intelligence, 16(12), 1207 (1994)
- [7] E.P.Lyvers & O.R.Mitchell, "Precision edge contrast and orientation estimation", IEEE Trans. Pattern Analysis & Machine Intelligence, 10(6), 927 (1988)
- [8] M.Born & E.Wolf, "Principles of optics: Electromagnetic theory of propagation, interference and diffraction of light", 5th ed., Pergamon , 1975