ELSEVIER

Signal Processing: *Image Communication* 12 (1998) 33–47

SIGNAL PROCESSING: IMAGE COMMUNICATION

Improving the performance of edge localization techniques through error compensation

Federico Pedersini, Augusto Sarti*, Stefano Tubaro

Dip. Elettronica e Informazione, Politecnico di Milano, Piazza Leonardo da Vinci 32, 20133 Milano, Italy

Received 10 January 1996

Abstract

The performance of subpixel edge localization (EL) techniques can often be improved through the compensation of the systematic portion of the localization error. In order to prove this fact, we propose and analyze a method for estimating the EL characteristic of a given subpixel localizator through statistical analysis of appropriate test images. Such estimate can be used to characterize the compensator independently of the subpixel EL technique that is being used.

In order to evaluate the impact of the proposed compensation technique on the performance of a subpixel edge localizator, we have embedded it into a camera calibration procedure and compared the accuracy of the calibration with and without compensation. The improvement in the calibration precision has been proven to be significant (44%), which can be of crucial importance especially in applications of low-cost photogrammetry and 3D reconstruction from multiple views.

Keywords: Feature extraction; Edge localization; Subpixel detection; Camera resolution enhancement

1. Introduction

In a variety of image applications such as 3D reconstruction, photogrammetry, remote sensing and automatic inspection, a precise localization of image edges can be of critical importance [25]. The importance of edges, i.e. abrupt changes in the luminance function, lies in the fact that they usually carry significant information about the imaged scene [17, 5].

The accuracy of typical edge detection algorithms [5, 7] is limited by the CCD camera resolution which can be increased only at very high cost. Alternatively, one could look into the possibility of using subpixel feature localization algorithms in order to reach super-resolution performance with low-cost CCD cameras [25, 24].

It is important to notice that, in order to solve the edge localization problem with subpixel accuracy, it is necessary to have some a priori information about the nature of both image edges and acquisition system. In fact, as the CCD camera performs image sampling, its model is not invertible for all signals. In particular, signals having high-frequency components such as abrupt luminance transitions (edges) cannot be retrieved beyond pixel resolution. In order to overcome the limitation represented by Shannon's sampling theorem, we may look at the subpixel edge localization problem as that of inverting the camera model for

* Corresponding author. Tel.: 39 2 2399 3647; fax: 39 2 2399 3413.

0923-5965/98/$19.00
PII S0923-5965(97)00034-9

a very specific class of signals. This corresponds to determining the parameters of a specified edge transition model that, when cascaded with the camera model, produces the available digital image.

In principle, lack of proper filtering before sampling causes aliasing. This phenomenon could be reduced or eliminated through optical blurring in order to perform an accurate edge localization [21, 23]. Intentional blurring, however, reduces the capability of resolving details that are close to each other, therefore it is seldom applicable. In normal operative conditions, we are thus forced to operate in the presence of aliasing.

In order to understand how to approach the edge localization problem in the presence of aliasing, let us consider a straight edge belonging to the family of step-like functions [25, 5], and let us assume that a good model for the camera lens is the pin-hole. If the image plane was partitioned into a collection of regularly spaced photosensitive elements, then it would be possible, in condition of low noise, to exactly estimate the edge position at subpixel precision. In fact, the luminance associated to the pixels that the edge cuts through would exactly tell us the position and the orientation of the sharp transition corresponding to the edge itself.

The reliability of the adopted models can critically affect the performance of a subpixel edge localization technique. In real CCD cameras, for example, the photosensitive area is actually smaller than the pixel area itself. As a consequence, as we will see later on, believing that the CCD plane is partitioned into photosensitive regions [24] could result in a systematic edge localization error (ELE). Other physical sources of non-ideal behavior that need to be taken into account when dealing with super-resolution edge localization are the aberrations and the finite aperture of real optical lenses, or the non-ideality of the transfer and readout processes associated to the CCD sensor. The accuracy of their models may be a critical factor in the overall performance of subpixel localization techniques.

In this article we propose and evaluate a method for improving the performance of sub-pixel edge localization techniques, which is based on the correction of the ELE associated to nearly horizontal or nearly vertical edges in low-noise images. The method is particularly useful in those applications where the accuracy of edge localization is more important than noise suppression [18]. In Section 2 we describe an accurate model of the acquisition system and, in particular, that of a CCD camera, while a characterization of the ELE is given in Section 3. In Section 4, we propose and describe a method for estimating the EL function and show how to derive an ELE compensation map from it. Such a method is based on a statistical analysis of appropriate test images, therefore, we do not need any a priori information either on the camera system or on the adopted subpixel EL technique.

Tests have been performed in order to evaluate the impact of the proposed technique on concrete situations. In particular, we have embedded the ELE compensator into a complete camera calibration procedure. The task of estimating intrinsic and extrinsic camera parameters from the analysis of known image targets [22, 20] is comparatively performed with subpixel detectors with and without ELE-compensation. The results of such experiments, reported in Section 5, show that the performance of the calibration procedure improves significantly when ELE compensation is being employed in the localization algorithm.

2. A model for the acquisition system

The digital images that we are interested in are those acquired either with a digital camera or with a standard TV-resolution CCD camera cascaded with a frame-grabber. In order to model this system we may refer to Fig. 1 (see [6]), where all basic camera elements that contribute to the global system transfer function are displayed.

2.1. Basic elements of a CCD camera

In order to model the acquisition system of Fig. 1 we need to characterize all of its building blocks.

2.1.1. Optical lens

The simplest model of optical lens that is normally employed in the literature is the 'pin-hole', which corresponds to a simple perspective projection $I_g(x, y)$ of the scene onto the image plane (CCD sensor).

Such a model, however, is not close enough to the reality for the applications that interest us. Real lenses, in fact, always have a low-pass effect on the image, even when the subject is perfectly focused (the depth of field [8, 3] includes the whole imaged object). As

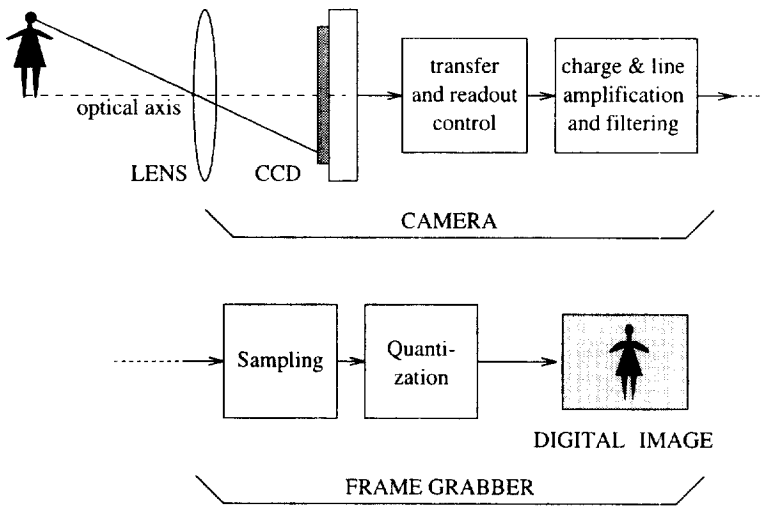

Fig. 1. Block diagram of a standard TV-resolution CCD acquisition system.

a matter of fact, there always is a bandwidth limitation on the optical lens due to its finite aperture, so that the image $I(x, y)$ that actually forms on the CCD surface is a low-pass version of the ideal image $I_g(x, y)$:

$$I(x, y) = h(x, y) \bigstar I_g(x, y),$$

$h(x, y)$ being the impulse intensity response of the lens,[1] whose Fourier transform is known as modulation transfer function (MTF).

The MTF is intimately related to the shape of the lens aperture; in fact it can be obtained by computing, through an appropriate change of variables, the auto-correlation of the *pupil function*.[2] The MTF [8] is thus a circularly symmetric low-pass response, which can be modeled in a parametric form as

$$H(\rho) = \begin{cases} \dfrac{2}{\pi}\left(\arccos\left(\dfrac{\rho}{2\rho_0}\right) - \dfrac{\rho}{2\rho_0}\sqrt{1-\left(\dfrac{\rho}{2\rho_0}\right)^2}\right) & \text{for } \rho \leqslant 2\rho_0, \\ 0 & \text{for } \rho > 2\rho_0, \end{cases}$$

ρ being the frequency associated to polar spatial co-ordinates, and

$$\rho_0 = \frac{l}{2\lambda d} = \frac{F}{\lambda}$$

being the corresponding cut-off frequency. The parameter d is the distance between the lens aperture and the image plane, λ is the average wavelength of the light, l is the lens aperture diameter (iris diaphragm), and F is the so-called *stop number*, i.e. the *relative*[3] *aperture* of the lens.

Besides lens aperture, a variety of other physical phenomena concur to prevent the lens from behaving like a simple pin-hole. The types of aberrations that are commonly encountered when dealing with TV-resolution CCD camera lenses can be roughly classified in those that cause a local shift of image points (distortion), and those that cause blurring (such as curvature of field, astigmatism, coma, etc.).

Lens distortion causes image points to be shifted from the positions predicted through paraxial approximation. This positional shift mainly occurs along the radial direction from the optical center,[4] in which case it is called *radial distortion* [22]. Radial distortion can be quite accurately described in a compact parametric

[1] The lens is assumed not to be affected by chromatic aberration, as it normally occurs with good quality lenses [15].

[2] The pupil function is equal to one in a circular region that corresponds to the iris diaphragm, and is zero outside this region.

[3] Relative to the focal length.

[4] The optical center is defined as the intersection between the optical axis and the image plane.

form by truncating the series expansion that expresses the undistorted radial coordinate ζ_u as a function of the distorted one ζ_d:

$$\zeta_u = \zeta_d(1 + k_3\zeta_d^2 + k_5\zeta_d^4 + \cdots).$$

A truncation of the above series to the third term (fifth-order) is usually sufficient for an accurate description of the positional shift of the image points.

Modern high-quality lenses are designed in such a way that blurring due to aberration [5] results as being negligible with respect to that due to its limited bandwidth. [6] Unfortunately, commercial low-cost CCD cameras seldom belong to this category. On the other hand, aberrations such as curvature of field, astigmatism, coma, etc., are combined with the limited lens and CCD-sensor aperture in such a way to behave globally as a low-pass, mostly space-invariant [8, 3], transfer function.

2.1.2. The CCD sensor

The light coming from the lens focuses on the image plane, i.e. the CCD sensor surface, and forms an analog image, which is sampled over time and space.

The sampling process is not ideal as the photosensitive regions have a non-negligible size and the time-integration period is not zero. The total light intensity falling on each photosensitive area, for example, is integrated over a *shutter* period, which is less than or equal to a field period (20 ms for CCIR-Standard Interline Transfer CCDs). As far as spatial sampling is concerned, the light is integrated over the photosensitive area of each pixel. This operation can be modeled by a low-pass filter followed by an ideal 2-D spatial sampling over a grid defined by the geometry of the CCD sensor. The spatial impulse response of the cell, which is normally called *CCD aperture function*, corresponds to the light sensitivity map associated to a pixel cell [6], and is normally assumed to be $W_{\text{cell}}(x, y) = 1$ inside the photosensitive region and zero outside (see Fig. 2). The luminance sample $I_{h,k}$ associated to a pixel is thus proportional to the amount of energy absorbed in one shutter period by its photosensitive area $\mathscr{P}_{h,k}$, i.e.

$$I_{h,k} = \iint_{\mathscr{P}_{h,k}} I(x, y)\,\mathrm{d}x\,\mathrm{d}y$$

$$= \iint I(x, y)\, W_{\text{cell}}(x - h\Delta_x, y - k\Delta_y)\,\mathrm{d}x\,\mathrm{d}y,$$

where pixels are assumed to be organized on a rectangular grid whose horizontal and vertical sampling intervals correspond to their relative inter-cell distances Δ_x and Δ_y, as shown in Fig. 2.

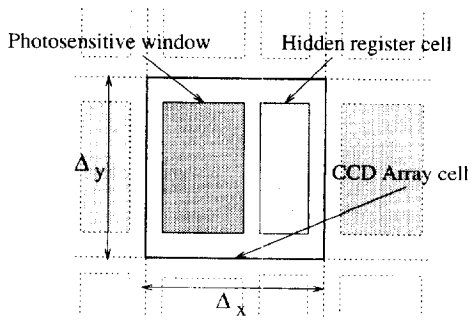

Fig. 2. Typical cell layout of Interline Transfer CCDs. The photosensitive region is only part of the total area of the pixel cell.

Notice that, in order to compute the transfer function of the image sensor, we need to know the actual size of the photosensitive area, which is normally quite different from the size of the pixel cell. In most Interline Transfer CCDs, for example, the photosensitive area is approximately one-third of the whole cell area [13].

2.1.3. CCD charge transfer and readout

In standard Interline Transfer (ILT) CCDs, while one field is being collected by the CCD sensors, the other one must be temporarily stored and transferred, pixel by pixel, to the output. This operation is performed by transferring each sample from the photosensitive region into a hidden region of the cell, as shown in Fig. 2. A column of such hidden cells constitutes a shift-register. All column shift-registers feed a faster shift register whose aim is to "push" out all image lines, one by one.

The transferal of charges through a shift-register is not an ideal process. In fact, the charge contained in one cell is only partially transferred to the next one [4] while a small portion of it remains in the cell.

[5] Under paraxial approximation [3].

[6] Lenses whose dominant deviation from ideality is the finite aperture are said to be *diffraction-limited* [15].

Transfer inefficiency is due to a variety of physical phenomena that tend to affect the diffusion process of charge carriers (barriers of potential energy between wells, carrier traps at interfaces, etc.).

Transfer inefficiency can be described in terms of an *inefficiency factor* α, that represents the fractional loss occurring when a charge is being moved from one cell to the next one. In other words, if Q is the charge in one cell, only the portion $(1-\alpha)Q$ will transfer to the next one, while the rest αQ will be left behind. As the charge scattering along the shift register depends on the number of shifts,[7] it is quite clear that transfer inefficiency is a space-varying phenomenon.

In most of the recently designed CCD image sensors the inefficiency factor α is well below 10^{-4} [10]; therefore the impact of transfer inefficiency is negligible when compared with other undesired effects.

2.1.4. Analog signal filtering and resampling

After the readout process, the analog stream of luminance samples of the CCD array is processed first by a charge amplifier and then by a line amplifier. Image digitization is then performed by a frame grabber through resampling and quantization. The cascade of the *charge* amplifier and the *line* amplifier can be modeled as a low-pass filter whose cut-off frequency is half the pixel clock frequency [11].

Notice that, as the 1D analog video signal is a time-sequential scanning of the 2D image, filtering the analog 1D signal corresponds to filtering the 2D signal only in the horizontal direction. As a consequence, the bandwidth of CCD images is normally wider along columns than along rows.

The output of low-cost CCD cameras usually consists of an analog signal and, sometimes, of a clock signal which is synchronized with the pixel rate. As a consequence, even though the CCD array is characterized by its own sampling grid, the resampling process introduced by the frame grabber may virtually modify the horizontal size of the pixels. This happens, in particular, when the sampling rate f_s of the frame-grabber is different from the camera pixel-clock f_{pc}, in which case the horizontal size d_x^{image} of one pixel of the final digital image does not correspond to the horizontal size d_x^{CCD} of the CCD cells. Such a difference

[7] The pixel (i,j) of an ILT CCD will have to undergo a total of $N=i+j$ transfers.

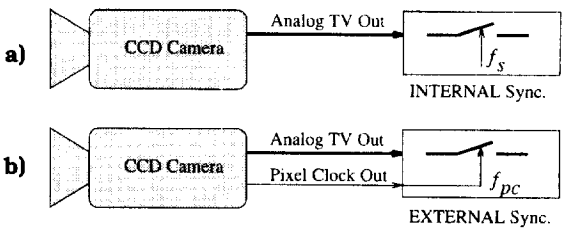

Fig. 3. Digital image acquisition system (a) without and (b) with pixel-clock. When the frame grabber is synchronized with the camera pixel-clock, an image pixel actually corresponds to a CCD cell.

is normally taken into account through an appropriate scale factor [22, 14]

$$s_x = \frac{d_x^{\text{image}}}{d_x^{\text{CCD}}} = \frac{f_{\text{pc}}}{f_s}\,.$$

When the camera is equipped with the clock output and when the frame-grabber is capable of using such a signal for synchronization, then s_x becomes a known parameter. In the scheme of Fig. 3(b), for example, the frame-grabber sampling rate is the same as that used to scan the CCD grid. As a consequence, image pixels and CCD cells end up having exactly the same size. If, conversely, the frame-grabber synchronization is made from the synchronism pulses contained in the video signal itself, then s_x results as being affected by an error (line jitter). For this reason the scheme of Fig. 3(b) is often preferred for videometric applications [2]. In fact, pixel-clock synchronization makes the scheme of Fig. 1 equivalent to that of a fully digital camera with a reasonably good approximation.

2.1.5. Sources of noise or distortion

The CCD sensor is characterized mainly by three types of noise [1, 12]: *dark current noise*, caused by the hole–electron pairs that are generated by thermal vibration of the silicon lattice, *'fat-zero' input noise*, caused by the random fluctuations that the CCD cell bias charge is affected by,[8] and *reset noise*, caused by the charge amplifier. The first type of noise can be modeled as a Poisson random process and depends on

[8] A modest bias charge is injected into the potential well of each CCD cell in order for the charge transferal to occur in conditions of linearity and prevent transfer inefficiency from being nonlinearly dependent on the charge magnitude.

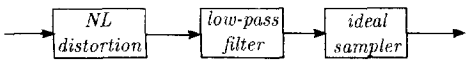

Fig. 4. Overall model of the acquisition system.

the sensor's temperature. Moreover, the spectral density of the third type of noise is a decreasing function of the frequency in the low-pass range, and tends to flatten to a constant as the frequency increases [12]. For these reasons, we are justified in modeling such noises altogether like an additive Gaussian white noise.

The non-homogeneous sensitivity of the photosensitive cells of the CCD sensor is an important source of image distortion. Some applications of photogrammetry take the cell sensitivity into account by accurately measuring their gain, and compensating for their non-homogeneity afterwards. Another cause of image distortion is the *blooming* effect, which takes place especially at the border between very light and very dark areas of the image. This effect, due to problems with charge diffusion in the CCD grid, manifests itself as an *erosion* of dark areas on the behalf of light areas, and its impact can be minimized by reducing the luminance dynamics.

Finally, we should not forget that the digitization process introduces a certain *quantization error* in the analog-to-digital conversion. It is quite clear that such a noise can be neglected if an adequate number of quantization levels is available.

2.2. *An overall view*

Considering the characterization of the basic elements of the acquisition system, given in Section 2.1, we can summarize the structure of the acquisition system model as shown in Fig. 4. As we can see, the nonlinear (radial/tangential) distortion is all included in the first block, while all sources of blurring (lens aperture, aberrations, CCD sensor aperture, etc.) are all included in the second block.

As already said in Section 2.1, lens distortion can be thought of as a nonlinear stretching of the image plane, which can be accurately described by a very limited number of parameters. It is not difficult to realize that step-like luminance transitions undergoing lens distortion remain step-like. As a consequence, if our goal is that of recovering the location of the transition with subpixel accuracy, we can compensate for the distortion *after* edge localization. This corresponds to 'incorporating' lens distortions into the imaged scene and localizing the edges of a 'distorted universe', while leaving the task of unwarping the universe to afterwards. This operation is possible if we have a reliable estimate of the distortion coefficients, which can be obtained through camera calibration [22, 20].

3. Edge localization error

Quite a variety of subpixel edge localization (EL) techniques is currently available in the literature. Each of them is characterized by a different level of accuracy and noise-rejection. Depending on the strategy adopted by the subpixel method, the edge localization error (ELE) associated to it may exhibit a certain *systematic* character [16]. It is quite evident that, if we can completely characterize and predict the ELE, then we can also compensate for it.

It is not difficult to realize that dealing with image edges that are either nearly horizontal or nearly vertical greatly simplifies the analysis and the characterization of the ELE. We will see later on, however, that limiting our analysis to the case of nearly horizontal or nearly vertical edges does not represent a serious loss of generality, as the edges that predominantly suffer from a systematic ELE are the horizontal and the vertical ones.

The 1-D ELE corresponding to an abrupt luminance transition is the distance between the sharp transition that would form on the image plane when using an *ideal* optical lens [9] and the edge that has been actually detected. It is quite clear that, besides depending on the acquisition system, the ELE critically depends on the subpixel edge localization (EL) technique under examination.

In order to clarify the concept, let us consider the simplified situation shown in Fig. 5, where a sharp luminance transition occurring at the coordinate p on the image axis (horizontal or vertical coordinate of the image plane) is being localized at subpixel precision by using of a simple linear interpolation technique.

[9] With *ideal* lens we mean a lens with no aberrations and with unlimited bandwidth, which allows us to adopt a model based on geometrical optics.

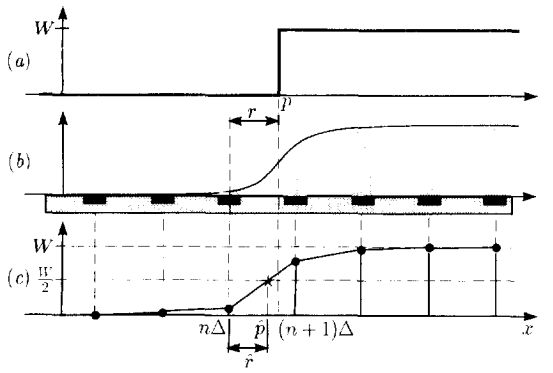

Fig. 5. Subpixel edge detection based on linear interpolation. (a) Ideal luminance profile, (b) luminance profile incident on the image plane, (c) linear interpolation of the image samples.

The luminance profile of Fig. 5(a) is a section of what would be imaged on the image plane if an *ideal* lens were used instead of the real one. Due to the limited aperture of the lens, the actual luminance profile is a filtered version of the ideal one, as shown in Fig. 5(b). When the light reaches the array of photosensors of the CCD camera, the image is spatially sampled. The luminance samples that are actually collected from the CCD array, however, are not the result of an ideal sampling of the luminance profile of Fig. 5(b), as they depend on the light that falls on the whole photosensitive area of the pixel. In fact, assuming that the photosensitive area is equally sensitive to the light, each luminance sample is given by the area of the shaded regions in Fig. 5(b), as shown in Fig. 5(c).

A simple way of estimating the subpixel location p of the ideal edge from the samples collected from the CCD array consists in linearly interpolating (see Fig. 5(c)) the collected samples, and determining the intersection between the resulting piecewise linear profile and an appropriate threshold. The threshold is set equal to half the amplitude W of the luminance discontinuity, and the resulting intersection can be taken as an estimate of the edge location. Such an example of subpixel EL method is simple enough to visualize the ELE associated to it, in fact the estimated edge location $\hat{p}$ differs from the ideal location p of a quantity called edge localization error e. It is not difficult to realize that the ELE is a periodic function of the edge location, provided that some conditions of regularity in the acquisition system are satisfied (e.g. homogeneous CCD array and small radiometric distortions). In this case it is convenient to refer ideal and estimated edge locations to the center of the pixel immediately before the ideal edge location and limit the description to one period of it.

In what follows, the function that maps the ideal relative edge location r into the estimated one $\hat{r}$ is called edge localization function (ELF), $\hat{r} = F_{\mathrm{EL}}(r)$, and the ELE can be written in terms of the ELF as follows:

$$e = \hat{r} - r = E_{\mathrm{ELE}}(r) = F_{\mathrm{EL}}(r) - r. \tag{1}$$

We will see in the next section how it is possible to estimate the ELF from the analysis of appropriate test images, and how to use this estimate for compensating the ELE of a generic subpixel EL technique. In order to be able to do so, we will have to make two basic assumptions:

- the ELE is approximately *space-invariant*, therefore we only need to measure *one* ELF for the whole image,
- the ELF is *invertible*, and its inverse can be used as a compensation function.

4. Error compensation

As already mentioned in Section 3, if the ELF $\hat{r} = F_{\mathrm{EL}}(r)$ is an invertible function of the local coordinate r of the 1-D (horizontal or vertical) image axis, then we can compensate for the ELE provided that a reliable ELF is available.

In this section we show how to estimate the ELF and how to derive the relative compensation function through statistical analysis of some test images.

4.1. Estimation of the error characteristic

With reference to Fig. 5, we have seen in Section 3 that the estimated (affected by ELE) relative edge location $\hat{r}$ can be seen a function $F_{\mathrm{EL}}(r)$ (EL *function*) of the actual (ideal) relative edge location r . Since the response of the CCD camera can be considered space-invariant, the ELE function $e = \hat{r} - r = E_{\mathrm{EL}}(r)$, must be periodic of period 1 pixel, therefore we can limit our analysis, for example, to any interval like $r_0 \leqslant r \leqslant 1 + r_0$. The periodicity of the ELE and Eq. (1) results in

$$1 + \hat{r} = F_{\mathrm{EL}}(1 + r).$$

If $F_{EL}(r)$ is monotonic, then it is also bijective, in which case its inverse function, $r = F_{EL}^{-1}(\hat{r})$, is bijective as well. The output range corresponding to $r_0 \leqslant r \leqslant 1 + r_0$ results as $\hat{r}_0 \leqslant \hat{r} \leqslant 1 + \hat{r}_0$, where $\hat{r}_0 = F_{EL}(r_0)$. The inverse $F_{EL}^{-1}(\cdot)$ of the ELF can thus be used as an *error compensation function*.

As the ELF is a map from the ideal edge locations onto the detected edge locations, we can derive information on it from the joint statistics of both its input and its output. The estimation of the error compensation function, in fact, can be done through statistical analysis of an appropriate test image. The statistical distribution of the estimated edge locations can be quite easily extracted from the test image, while the statistics of the ideal edge location can be inferred from the pattern characteristics in particular cases. From a practical viewpoint it is convenient to choose test images whose ideal edge points (referred to the center of the pixel area that they fall on) are uniformly distributed over pixel areas. Images that satisfy such a requirement are very common. An example of such a test image is shown in Fig. 6.

In order to avoid confusion in the notation, in what follows we will denote with capital letters (R and $\hat{R}$) the random variables that represent ideal and estimated fractional edge locations, while lowercase letters (r and $\hat{r}$) will be used for denoting the corresponding instances.

If the probability density function (p.d.f.) of the ideal edge point position R is uniform,

$$f_R(r) = \begin{cases} 1 & \text{for } r_0 \leqslant r < 1 + r_0, \\ 0 & \text{elsewhere,} \end{cases} \tag{2}$$

then the p.d.f. of $\hat{R} = F_{EL}(R)$ can be expressed as

$$f_{\hat{R}}(\hat{r}) = \frac{f_R(r)}{F'_{EL}(r)}, \quad r = F_{EL}^{-1}(\hat{r}), \tag{3}$$

where $\hat{r}_0 \leqslant \hat{r} < 1 + \hat{r}_0$, $F'_{EL}(r)$ is the first derivative of $F_{EL}(r)$, and the absolute value can be omitted in the denominator of Eq. (3) as $F_{EL}(r)$ is assumed to be a monotonically increasing function. By replacing Eq. (2) into (3) and by applying a property of derivatives, we obtain

$$f_{\hat{R}}(\hat{r}) = \left. \frac{1}{F'_{EL}(r)} \right|_{r = F_{EL}^{-1}(\hat{r})} = \frac{d}{d\hat{r}} F_{EL}^{-1}(\hat{r}), \tag{4}$$

where $r_0 \leqslant r < 1 + r_0$ and $\hat{r}_0 \leqslant \hat{r} < 1 + \hat{r}_0$.

By integrating the p.d.f. of the subpixel edge locations $\hat{R}$ detected from the image we obtain the compensation function

$$r = C(\hat{r}) = F_{EL}^{-1}(\hat{r}) = F_{EL}^{-1}(\hat{r}_0) + \int_{\hat{r}_0}^{\hat{r}} f_{\hat{R}}(a)\, da. \tag{5}$$

Notice that the value of $\hat{r}_0$ is not a known parameter, therefore all that we can obtain from the analysis of the image is the statistical distribution of $\hat{R}$, computed over an arbitrary pixel-wide interval like $(A, 1 + A)$, generally not entirely contained in the interval $(\hat{r}_0, 1 + \hat{r}_0)$. Eq. (5) could thus be expressed as follows:

$$\begin{aligned} r &= \int_{\hat{r}_0}^{A} f_{\hat{R}}(a)\, da + \int_{A}^{\hat{r}} f_{\hat{R}}(a)\, da + r_0 \\ &= \int_{A}^{\hat{r}} f_{\hat{R}}(a)\, da + K_A. \end{aligned} \tag{6}$$

It is quite clear from Eq. (6) that different choices of the interval of definition of $\hat{r}$ result in different vertical offsets K_A for the compensation function.

Notice that Eq. (5) can be used to compute $F_{EL}^{-1}(\hat{r})$ only for $A \leqslant \hat{r} \leqslant 1 + A$. Since $E_{EL}(r)$ is a periodic function of period 1 pixel, we extend the range of $F_{EL}(\cdot)$ and $F_{EL}^{-1}(\cdot)$ by using the relationships

$$\begin{aligned} F_{EL}(r + k) &= k + F_{EL}(r) \\ F_{EL}^{-1}(\hat{r} + k) &= k + F_{EL}^{-1}(\hat{r}) \end{aligned} \quad k = 0, \pm 1, \pm 2, \ldots; \tag{7}$$

therefore, if $A \leqslant \hat{r} \leqslant A + 1$, then we can use Eq. (5), otherwise we can always find an integer k such that $\hat{r} = \hat{r}' + k$, $A \leqslant \hat{r}' \leqslant A + 1$, and use Eq. (7).

It is worth it to emphasize that the fact that the compensation function is derived from a p.d.f. through integration gives us no information on the offset K_A, which means that we can linearize the ELF (i.e. eliminate its ripple) but we still need to determine its offset. The extra unknown can be determined by using further a priori information on the test image, or through camera calibration.

4.2. The estimation procedure

The test image shown in Fig. 6(a) is designed for the localization of edge points in the vertical direction, i.e. for image edges that are locally nearly horizontal.

Notice that the stripes are slightly tilted in order to guarantee the detected sub-pixel position of its edges to have uniform statistical distribution, as required.

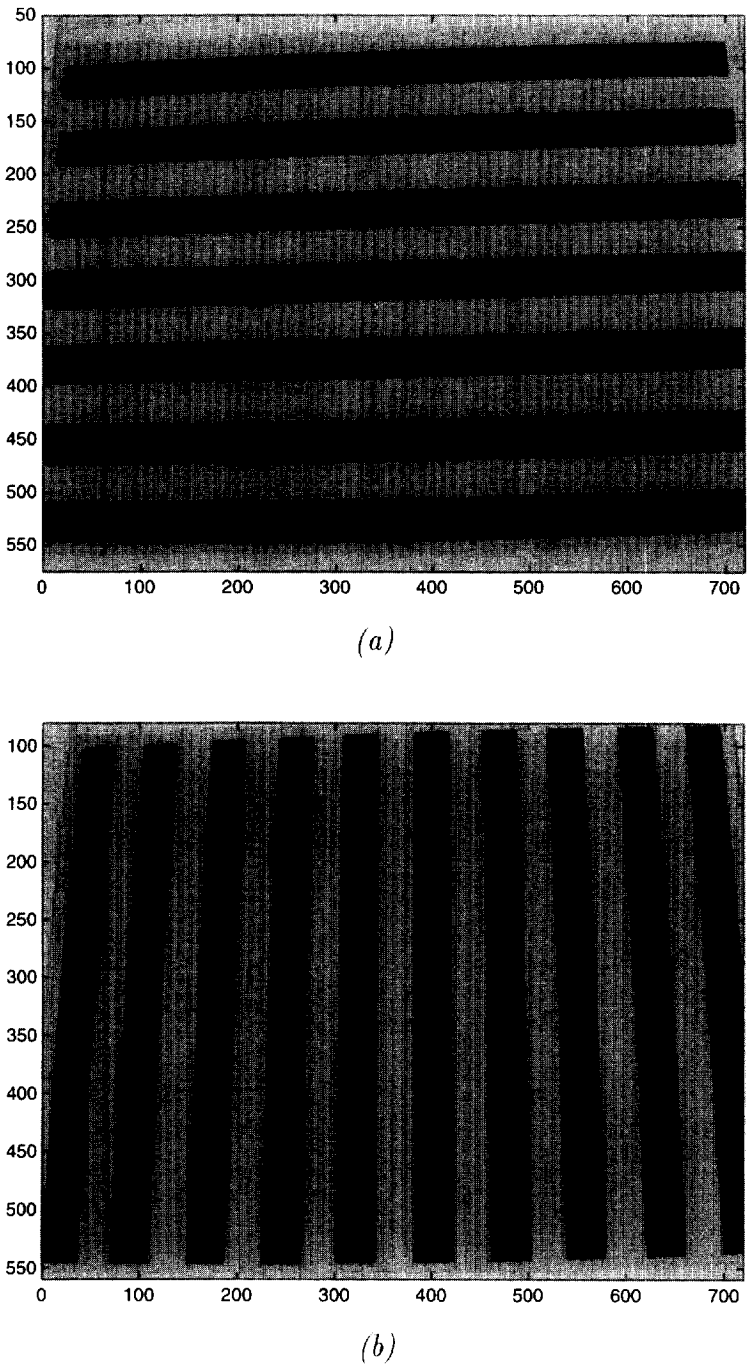

Fig. 6. Typical test images for the determination of (a) the horizontal and (b) the vertical error compensation function.

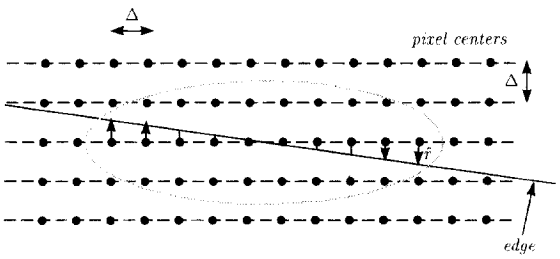

Fig. 7. Construction of the histogram of the local edge coordinates.

Notice also that the presence of barrel distortion does not affect the accuracy of the estimation of the ELE characteristic. In fact, the edge coordinates can be assumed as uniformly distributed over pixel areas when the edges can be considered as *locally* straight in the absence of ELE, which is true also in the presence of barrel distortion.

All edge points of the test image are localized with sub-pixel accuracy by using any edge localization algorithm, for example, the one based on cubic interpolation (with the edge location given by the flex point), or even the one based simply on linear interpolation. From each edge coordinate x, we compute the local edge coordinate

$$\hat{r} = \hat{x} - n\Delta,$$

where the pixel center $n\Delta$ is the nearest one to $\hat{x}$. Assuming that the above lengths are measured in pixels, we have $-\frac{1}{2} \leqslant e \leqslant \frac{1}{2}$ and $A = -\frac{1}{2}$.

The p.d.f. $f_{\hat{R}}(\hat{r})$ of the detected subpixel relative locations is estimated by building a histogram for $\hat{r}$, as shown in Fig. 7. This operation corresponds to building a piecewise constant approximation of the desired p.d.f., and then normalizing its amplitude. The number of histogram intervals depends on the number of available samples of $\hat{R}$ (e.g. 100 intervals for the test image of Fig. 6).

Finally, we integrate $f_{\hat{R}}(\hat{r})$ in order to compute the first term of Eq. (6). As far as the offset K_A is concerned, as we will see in Section 5, its determination depends on the specific application.

Fig. 8(a) shows an example of the p.d.f. $f_{\hat{r}}(b)$ that is estimated from a test image (edges localized

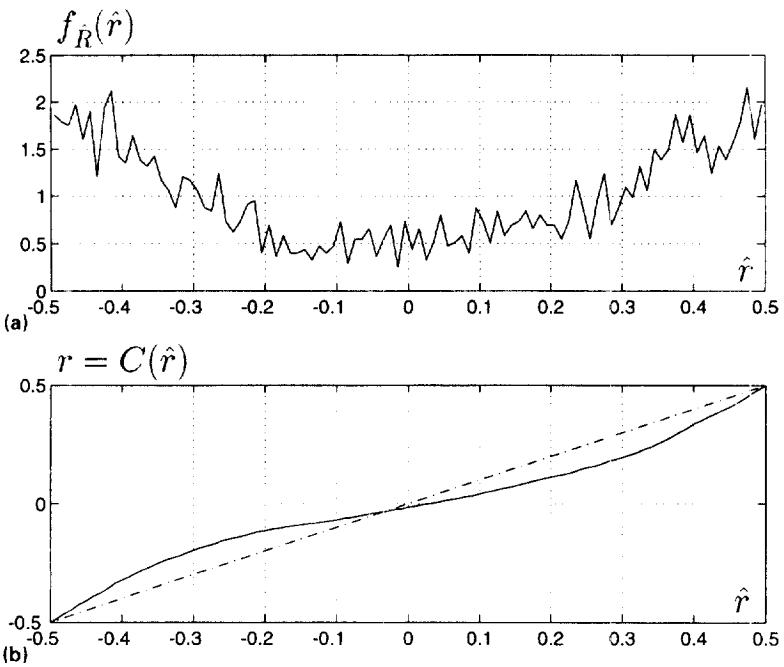

Fig. 8. Estimated p.d.f. of (a) the detected subpixel residuals and (b) the relative compensation function.

along columns). Compensation is performed by using Eqs. (6) and the resulting *compensation function*

$$C(\hat{r}) = F_{\mathrm{EL}}^{-1}(\hat{r}),$$

is shown in Fig. 8(b). The compensated edge position r is then obtained by simply applying the compensation $C(\hat{r})$ to the detected position $\hat{r}$,

$$r = F_{\mathrm{EL}}^{-1}(\hat{r}) = C(\hat{r}). \tag{8}$$

Notice that the compensation method described above does not depend on the choice of subpixel edge localization technique that is being used. However, it is reasonable to expect cubic interpolation to outperform linear interpolation in the edge placement because of a different noise rejection. In fact, a cubic interpolator averages over a larger number of samples, thus reducing the noise.

5. An example of application

The ELE compensation method of Section 4 can be used for improving the performance of an EL technique as follows:

1. Perform subpixel edge localization on the *test* image.
2. Estimate the compensation curve from the edge points of the test image.
3. Perform subpixel edge localization on the *scene* image.
4. Correct the edge coordinates through the estimated compensation function.

Notice that, if the edge points in the scene image satisfy condition (2), i.e. if the ELE is uniformly distributed, then the scene image can be used as a test image, and step 3 can be skipped.

In order to evaluate the impact of the above compensation technique on the performance of a subpixel edge localizator, we have embedded the method into a camera calibration procedure [22, 20] and compared the results with and without compensation. Camera calibration consists in estimating *intrinsic* and *extrinsic* parameters of an image acquisition system through the analysis of the views of a *calibration pattern*. In the adopted camera model, the intrinsic parameters are the *optical center* (intersection between the optical axis and the image plane), the *focal length* and two parameters that describe the radial distortion of the optical lens. The extrinsic parameters are represented by the relative position and orientation of the camera with respect to the target. It is quite evident that the reliability of the calibration procedure critically depends on how accurately certain *fiducial marks* of the calibration pattern are localized.

The calibration pattern used in the experiment is planar and exhibits a set of regularly spaced black squares on a white background, as shown in Fig. 9. The position of the fiducial marks, i.e. the corner points of the squares, is known [10] with a precision of ± 5 μm. In order to perform an accurate camera calibration, it is necessary to localize the fiducial marks of the test image with the best achievable precision. Being the fiducial marks corner points of squares, they can be localized by intersecting edges detected with subpixel accuracy.

The adopted calibration procedure estimates the camera parameters and provides us with a measure of the estimate accuracy, based on the standard deviation of the error between the detected position of fiducial marks on the image plane, and their position computed through the camera model. The accuracy measurement has been used as an evaluation of the performance of the edge localization algorithm, and a comparative evaluation of the results with and without ELE compensation has been done.

As already mentioned in Section 4, the ELE compensation requires the determination of offset parameters. In fact, the offsets have been added to the list of intrinsic parameters of the CCD camera and estimated by the calibration procedure. By doing so, the estimated offsets can be used in other applications and for the ELE correction.

The ELE statistics associated to the test image of Fig. 9 can be assumed uniform with good approximation, therefore the calibration target is suitable also for the estimation of the compensation curve.

Fig. 10 shows the edge points of the test image, obtained with a technique based on cubic interpolation and flex point search. From such edges it is quite straightforward to visualize the ELE associated to the adopted subpixel technique. In fact, by magnifying all horizontal (vertical) edges of one row (column) of squares, we obtain the curves of Fig. 11, whose

[10] The calibration pattern has been measured with a monocomparator beforehand.

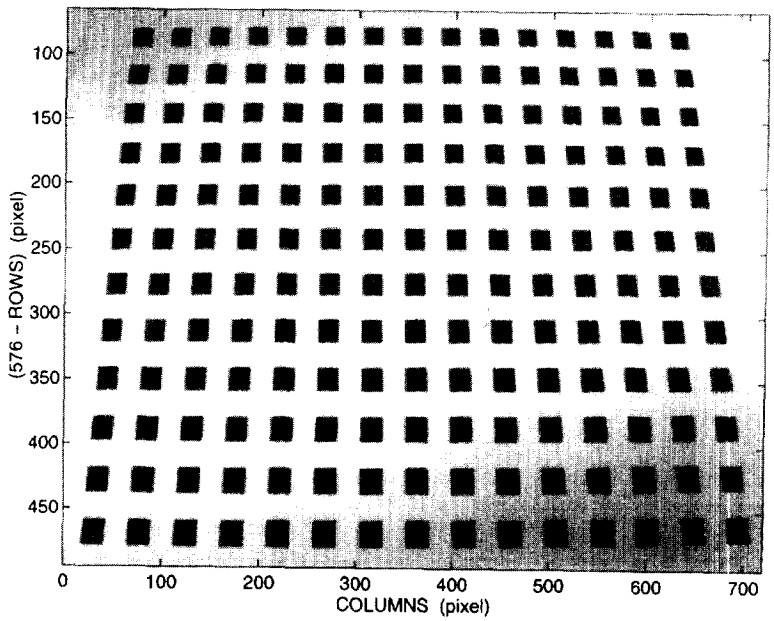

Fig. 9. Calibration pattern used to obtain the calibration results of Table 1.

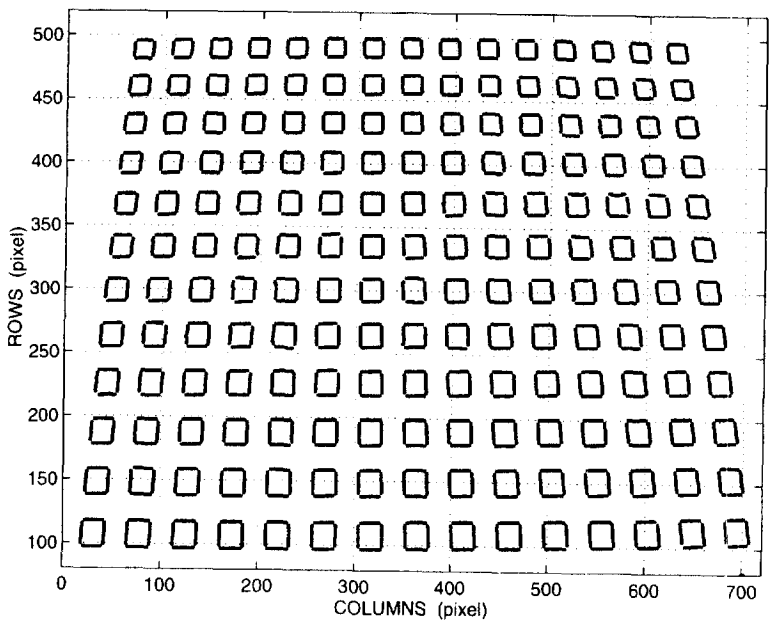

Fig. 10. Edge points detected in the image of the calibration pattern.

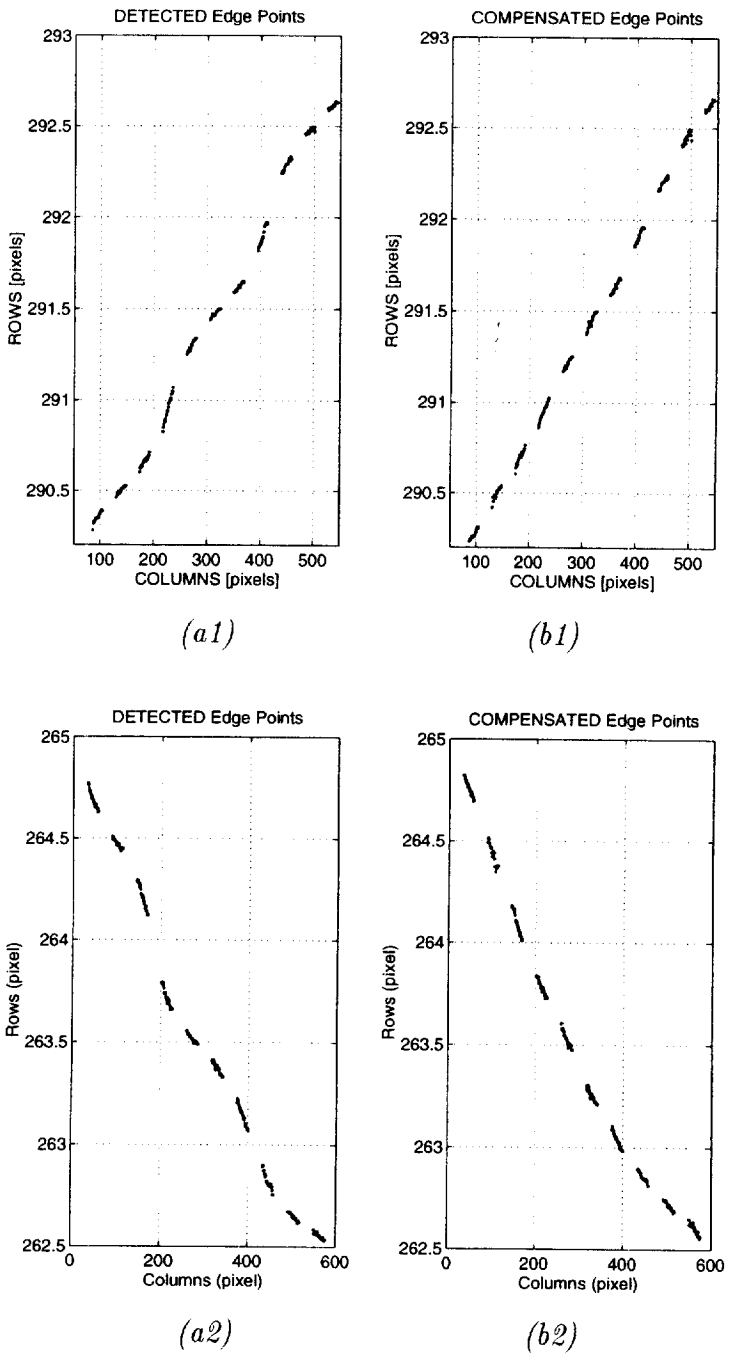

Fig. 11. Magnification of (a1, b1) rows 290–293 and (a2, b2) rows 262–265 of Fig. 13. (a) Linear interpolation and threshold crossing; (b) same as (a) with error compensation.

Table 1
Standard deviation of the final residual error in the calibration procedures, with and without error compensation

Std. deviation σ (pixel)		
Without comp.	With comp.	Improvement
0.0813	0.0454	44.2 %

oscillations are mainly caused by the ELE. By comparing the curves of Fig. 11, obtained with and without compensation, we observe a substantial reduction of the ELE.

A comparison between the accuracy of the calibration experiments that we have obtained with and without error compensation is reported in Table 1. The standard deviation of the calibration points with error compensation results as being 0.045 pixel, which is approximately 44% less than what we obtained without compensation.

The improvement in the performance of the calibration procedure shows that the impact of the ELE compensation technique can be significant in certain applications where precision is crucial. Furthermore, they confirm the fact that the systematic space-invariant component of the ELE is quite relevant.

Notice that the proposed ELE compensation method is one-dimensional, as it can be applied to either nearly horizontal or nearly vertical edges. Extensions to the more general two-dimensional case are possible by taking into account the fact that we would have to construct an approximation of two EL surfaces (two EL functions of two parameters). In fact, we would need to express position and orientation of a detected edge as a function of position and orientation of the ideal edge that generated it. It is worth noticing, however, that edges that are either nearly horizontal or nearly vertical are the most sensitive to ELE. In fact, with reference to Fig. 11, it is not difficult to realize that the density of ripples due to ELE increases with the edge angle with respect to horizontality or verticality. In particular, a nearly horizontal edge gives rise to a very *slow* ripple. In this case we can correctly estimate the actual edge location only when a large number of edge points is available. ELE compensation allows us to dramatically reduce such a number. As the edge slope increases, the ripple periodicity increases as well, which makes ELE compensation progressively less effective.

6. Conclusions

In this article we have proposed and analyzed a method for improving the performance of sub-pixel edge localization techniques, which is based on the compensation of their edge localization error (ELE). In particular, we have shown how to estimate the EL function and how to derive the ELE compensator from it. We have also evaluated the performance of the ELE compensation method in a concrete situation, by determining its impact on the accuracy of a camera calibration procedure.

The improvement in the calibration accuracy due to ELE compensation has been shown to be quite significant (44%), which can be crucial especially in applications of low-cost photogrammetry and 3-D reconstruction from multiple views, and justifies its adoption whenever it is important to maximize the precision of the edge localization without significantly affecting the total cost of the acquisition system.

Acknowledgements

The authors are grateful to Prof. H.G. Musmann of Universität Hannover, Germany, and Prof. F. Rocca of Politecnico di Milano, Italy, for their precious help in the development of this article.

References

[1] D.F. Barbe, Imaging devices using the charge-coupled concept, Proc. IEEE 63 (1) (1975).
[2] H.A. Beyer, Geometric and radiometric analysis of a CCD-camera based photogrammetric close-range system, Ph.D. Thesis, Institut für Geodäsie und Photogrammetrie, ETH, Zurich, 1992.
[3] M. Born, E. Wolf, Principles of Optics, Pergamon Press, Oxford, 1959.
[4] R.W. Brodersen, D.D. Buss, A.F. Tasch, Experimental characterization of transfer inefficiency in charge-coupled devices, IEEE Trans. Electron Devices ED-22 (2) (1975).
[5] J. Canny, A computational approach to edge detection, IEEE Trans. Pattern Anal. Mach. Intelligence 8 (6) (1986).
[6] S.G. Chamberlain, D.H. Harper, MTF simulation including transmittance effects and experimental results of charge-coupled imagers, IEEE Trans. Electron Devices Vol. ED-25 (2) (1978).
[7] R. Deriche, Fast algorithms for low-level vision, IEEE Trans. Pattern Anal. Mach. Intelligence 12 (1) (1990).