

Radiometric CCD Camera Calibration and Noise Estimation

Glenn E. Healey, *Member, IEEE*, and Raghava Kondepudy, *Student Member, IEEE*

Abstract—Changes in measured image irradiance have many physical causes and are the primary cue for several visual processes, such as edge detection and shape from shading. Using physical models for charged-coupled device (CCD) video cameras and material reflectance, we quantify the variation in digitized pixel values that is due to sensor noise and scene variation. This analysis forms the basis of algorithms for camera characterization and calibration and for scene description. Specifically, algorithms are developed for estimating the parameters of camera noise and for calibrating a camera to remove the effects of fixed pattern nonuniformity and spatial variation in dark current. While these techniques have many potential uses, we describe in particular how they can be used to estimate a measure of scene variation. This measure is independent of image irradiance and can be used to identify a surface from a single sensor band over a range of situations. Experimental results confirm that the models presented in this paper are useful for modeling the different sources of variation in real images obtained from video cameras.

Index Terms—CCD cameras, computer vision, camera calibration, noise estimation, reflectance variation, sensor modeling.

I. INTRODUCTION

SINCE the early and important work of Horn [10], machine vision researchers have derived algorithms from increasingly accurate models for image formation. Much of the recent work has concentrated on models for the interaction of light with matter [9], [17], [20], [24], [26]. Another aspect of image formation that has received less attention is the behavior of sensors. It is becoming clear, however, that accurate sensor models are important for many machine vision algorithms [12], [14], [21].

In this paper we examine the properties of charge-coupled device (CCD) image sensors from the perspective of machine vision. By studying the operation of a CCD camera imaging system, the various noise sources that corrupt digital pixel values can be quantified. The resulting camera model allows certain components of the sensor noise to be removed by a calibration procedure, while the remaining noise can be quantified accurately and accounted for by machine vision algorithms.

In addition to sensor noise, scene changes in reflectance and illumination can also cause variation in measured pixel values. By modeling reflectance and illumination variation, we define

the scene variation for a surface patch. Using the camera model developed in this paper, we show how scene variation can be estimated from single band images. This measure does not depend on image irradiance and can be used to identify a surface independent of the power of the scene illumination.

II. MODELING CCD VIDEO CAMERAS

The concept of the CCD was proposed [6] and experimentally verified [1] in 1970. Almost from the beginning, researchers recognized the potential advantages of CCD image sensors over vidicon tubes. In spite of the development of other solid-state imaging technologies [8], CCD video cameras are the most prominent sensing devices used in machine vision. This popularity is a result of their high resolution, high quantum efficiency, wide spectral response, low noise, linearity, geometric fidelity, fast response, small size, low power consumption, and durability. In addition, since a CCD is based on fixed sensing elements of equal size, the device provides precise spatial quantization that enables accurate spatial representation of images in a computer. In the 1980s, the consumer market for video cameras has led to the emergence of relatively inexpensive CCD cameras.

A. Overview of CCD Camera Operation

The operation of a CCD is often compared to measuring the spatial distribution of rainfall over a field by placing an array of buckets on the field [15]. Following a storm, the buckets are systematically transferred by conveyor belts to a metering station where the amount of water in each bucket is measured. Each measurement then represents the amount of rainfall at a particular location on the field.

Instead of measuring rainfall over a field, a CCD is used to measure the spatial distribution of light incident on a thin wafer of silicon. The measurement process relies on the fact that when a photon strikes silicon, an electron-hole pair is generated. The photon-generated electrons (photoelectrons) can then be collected in one of many discrete collection sites, each of which is associated with photons incident on a small area of the silicon surface.

Each collection site is formed by growing a thin layer of silicon dioxide on the silicon and depositing a conductive gate structure over the oxide. Applying a positive electrical potential to the gate creates the depletion region where photoelectrons may be stored. An electronic representation of the spatial distribution of the light incident on a CCD is formed by integrating photoelectrons in the individual collection sites (also called potential wells) over a fixed time interval.

Manuscript received July 9, 1992; revised March 1, 1993. This work was supported by the National Science Foundation under Grant IRI-9009811 and by the University of California Irvine Academic Senate Committee on Research. Recommended for acceptance by Associate Editor R. Bolle.

The authors are with the Department of Electrical and Computer Engineering, University of California, Irvine, CA 92717.
IEEE Log Number 9214430.

The process of charge coupling is used to move stored charge in a CCD. During charge coupling, charge packets at collection sites are transferred from site to site by manipulating gate potentials. The separation of individual charge packets is preserved during these transfers. The charge transfer efficiency (CTE) of a device quantifies the fraction of charge that can be effectively transferred from one collection site to an adjacent collection site.

An image is read out of the device by transferring the charge packets integrated at each collection site in parallel along electron conducting channels connecting columns of collection sites. A serial output register with one element for each column receives a new row of charge packets after each parallel row transfer. Before the next parallel transfer, the output register transfers each of its charge packets in sequence to an output amplifier that generates a signal proportional to the amount of charge. Once all of the charge packets in the serial register have been read out through the output amplifier, the next row can be transferred in parallel into the serial register. The process continues until the entire two-dimensional array of collection sites has been read out. Reliable techniques have been developed to prevent photoelectrons that might be generated during the read-out process from contaminating the charge packets being transferred.

In a video camera, the signal generated by the CCD is converted into an analog video signal. An RS-170 composite signal consists of interlaced lines of image and synchronization information and is updated at a rate of 30 frames/s. The video signal is low-pass filtered to prevent aliasing during subsequent sampling by a frame grabber that produces the final digital image.

B. The Camera Model

We begin by considering the signal generated for a single collection site. Ideally, an imaging system might produce a value I (number of electrons) at a collection site given by

$$I = T \int_{\lambda} \int_y \int_x B(x, y, \lambda) S_r(x, y) q(\lambda) dx dy d\lambda, \quad (1)$$

where T is integration time (in seconds), (x, y) are continuous coordinates on the sensor plane, $B(x, y, \lambda)$ is the incident spectral irradiance (Watts/unit area), and $q(\lambda)$ is defined as the ratio (electrons/Joule) of electrons collected per incident light energy for the device as a function of wavelength λ . $S_r(x, y)$ is the spatial response of the collection site. We note that $q(\lambda)$ is similar to device quantum efficiency that is defined as the ratio of electron flux to incident photon flux as a function of wavelength. As one might expect, there are several noise sources in CCD imaging systems that prevent us from measuring ideal pixel values. In what follows, we examine the effect of these noise sources.

Processing errors during CCD fabrication cause small variations in quantum efficiency and charge collection volume from collection site to collection site [5]. Thus, if a CCD array is uniformly illuminated, these variations will lead to a site-to-site nonuniformity in collected charge. This spatial nonuniformity is often referred to as fixed pattern noise. In

general, the site-to-site variation in quantum efficiency will exhibit some dependence on wavelength. In this paper, we will neglect this wavelength dependence to simplify camera calibration. We model the number of electrons collected at a site by KI , where I is as defined in (1) and K is a constant associated with the collection site that accounts for the product of a scaling of $q(\lambda)$ and a scaling of $S_r(x, y)$. We characterize K as having mean 1 and a spatial variance σ_K^2 over all of the collection sites. For a carefully manufactured device, there is little spatial variation in collection site response and σ_K^2 is small.

We note that we are assuming that the number of electrons collected at each site is independent of the number of electrons collected at other sites. This assumption can be violated in CCD sensors when a single site is illuminated with sufficient intensity to cause stored charge to overflow from a potential well and to mix with charge in other potential wells. This effect is called blooming and in severe cases can affect many collection sites in the vicinity of an overexposed site. In this work, we assume that our imaging system is configured to avoid overilluminating individual collection sites, and blooming effects are not considered.

Thermal energy in silicon generates free electrons. These free electrons, known as dark current, can be stored at collection sites and thereafter become indistinguishable from photoelectrons. The expected number of dark electrons generated is proportional to the integration time T and is highly temperature dependent. Cooling a sensor can reduce dark current generation to less than one electron per collection site per second [13]. For many devices, there are small fluctuations in the amount of dark current generated from collection site to collection site [8].

Shot noise is a result of the quantum nature of light and characterizes the uncertainty in the number of electrons stored at a collection site. This number of electrons follows a Poisson distribution so that its variance equals its mean. Shot noise is a fundamental limitation and cannot be eliminated. It is important to note that any dark current contributing to the charge stored at a collection site will increase the mean and therefore the variance in the number of electrons. The number of electrons integrated at a collection site is given by

$$KI + N_{DC} + N_S, \quad (2)$$

where N_{DC} is the number of electrons due to dark current and N_S is the zero mean Poisson shot noise with a variance that depends on the number of collected photoelectrons (KI) and the number of dark electrons (N_{DC}).

After charge is collected at each site, the CCD must transfer the charge to the output amplifier for readout. The charge transfer efficiency of real CCD's is less than 1. Charge that is not effectively transferred is either lost or deferred to subsequently transferred charge packets. The noise associated with charge transfer has been quantified [7]. Current buried-channel CCD's, however, achieve charge transfer efficiencies greater than 0.99999. It is reasonable, therefore, to neglect the effects of transfer inefficiency and to assume that all of the charge collected at each site is transferred to the output amplifier.

The on-chip output amplifier sequentially transforms the charge collected at each site into a measurable voltage. The amplifier generates zero mean read noise N_R that is independent of the number of collected electrons. Amplifier noise dominates shot noise at low signal levels and determines the read noise floor of the device. The voltage signal leaving the CCD is transformed into a video signal by the camera electronics. This signal is electrically low-pass filtered to prevent aliasing during subsequent sampling. This filtering also serves to remove high-frequency noise components. If we denote by A the combined gain of the output amplifier and the camera circuitry, then the magnitude of the video signal leaving the camera corresponding to each collection site is

$$V = (KI + N_{DC} + N_S + N_R)A. \quad (3)$$

To generate a digital image that can be stored in a computer, the analog signal from the camera is quantized both spatially and in magnitude using an analog-to-digital converter (ADC) on a frame grabber. For most systems, approximately one sample is taken for each collection site. The ADC approximates the analog voltage V using an integer multiple of a quantization step q so that each value of V satisfying

$$(n - \frac{1}{2})q < V \leq (n + \frac{1}{2})q \quad (4)$$

is rounded to a digital value $D = nq$, where n is an integer satisfying

$$0 \leq n \leq 2^b - 1 \quad (5)$$

and b is the number of bits used to represent D . To prevent clipping, q and b are chosen so that V does not exceed $(2^b - 0.5)q$. The quantization process can be modeled as the addition of a noise source N_Q so that

$$D = (KI + N_{DC} + N_S + N_R)A + N_Q. \quad (6)$$

Under reasonable assumptions [22], N_Q can be shown to be a zero mean random variable that is independent of V with a uniform probability distribution over the range $[-\frac{1}{2}q, \frac{1}{2}q]$ and a variance $\frac{q^2}{12}$.

A frame grabber that digitizes a composite video signal must separate synchronization information from image data. The uncertainty in the circuitry that performs this task can lead to linejitter that typically amounts to a small fraction of a pixel [4]. The effects of linejitter can largely be corrected and should eventually be eliminated [18]. In addition, nonvideo digital output cameras that are not susceptible to linejitter—such as the Cohu 4110, the Videk Megaplus, and the Photometrics Star 1—are gaining popularity. Since linejitter can be compensated for and causes primarily geometric distortion rather than radiometric distortion, we will not examine its effects further in this paper.

III. MODELING REFLECTANCE AND ILLUMINATION VARIATION

In this section, we examine how spatial variation in illumination and reflectance in a scene leads to spatial variation in the amount of charge collected across the sensor plane. Our goal is to model a nearly uniform reflectance card illuminated by

a nearly spatially uniform source. This model will be used to derive procedures for noise estimation and camera calibration in later sections. The problem of obtaining accurate geometric calibration is not addressed in this paper, but robust techniques have been developed [18], [25].

Define a three-dimensional coordinate system (x, y, z) with z along the optical axis of the imaging system and the (x, y) plane coinciding with the sensor plane. Consider a planar matte surface oriented normal to the optical axis of the system and illuminated by a distant source. For this configuration, the imaging geometry is accurately modeled by an orthographic projection. We assume a magnification of 1. Let $[\lambda_1, \lambda_2]$ be the range over which $q(\lambda)$ is nonzero. We can model the spatially varying spectral reflectance of the surface by

$$R(x, y, \lambda) = \bar{R}(\lambda) + r(x, y, \lambda), \quad (7)$$

where $\bar{R}(\lambda)$ is the mean reflectance for each λ and $r(x, y, \lambda)$ represents the random spatial variation of the reflectance as a function of wavelength. Thus,

$$E[R(x, y, \lambda)] = \bar{R}(\lambda) \quad \lambda \in [\lambda_1, \lambda_2] \quad (8)$$

and

$$E[r(x, y, \lambda)] = 0 \quad \lambda \in [\lambda_1, \lambda_2], \quad (9)$$

where E is an expected value operator applied over (x, y) .

For spatially varying illumination, we model the spectral irradiance incident on the planar surface by

$$L(x, y, \lambda) = \bar{L}(\lambda) + l(x, y, \lambda), \quad (10)$$

where

$$E[L(x, y, \lambda)] = \bar{L}(\lambda) \quad \lambda \in [\lambda_1, \lambda_2] \quad (11)$$

and

$$E[l(x, y, \lambda)] = 0 \quad \lambda \in [\lambda_1, \lambda_2], \quad (12)$$

where, as before, E is an expected value operator over (x, y) . The reflected radiance of the surface in the direction of the camera is then

$$R(x, y, \lambda)L(x, y, \lambda) = \bar{R}(\lambda)\bar{L}(\lambda) + \epsilon(x, y, \lambda), \quad (13)$$

where

$$\begin{aligned} \epsilon(x, y, \lambda) = & r(x, y, \lambda)\bar{L}(\lambda) + \bar{R}(\lambda)l(x, y, \lambda) \\ & + r(x, y, \lambda)l(x, y, \lambda). \end{aligned} \quad (14)$$

If we assume that the variation in surface reflectance $r(x, y, \lambda)$ and the variation in the illumination $l(x, y, \lambda)$ are uncorrelated, then $E[\epsilon(x, y, \lambda)] = 0$ for $\lambda \in [\lambda_1, \lambda_2]$.

For the image of such a surface, the spectral irradiance pattern incident on the sensor is given by

$$B(x, y, \lambda) = [R(x, y, \lambda)L(x, y, \lambda) * p(x, y, \lambda)]t(\lambda), \quad (15)$$

where $*$ denotes spatial convolution, $p(x, y, \lambda)$ is the point spread function of the optics, and $t(\lambda)$ is the spectral transmission of the optics.

From (1), the ideal number of electrons collected at a site (a, b) is given by

$$I(a, b) = T \int_{\lambda} \int_y \int_x B(x, y, \lambda) S_r(x - X_a, y - Y_b) \cdot q(\lambda) dx dy d\lambda, \quad (16)$$

where $\int_y \int_x$ denotes integration over the area of the collection site, (X_a, Y_b) are the coordinates of the center of the collection site, and $S_r(x, y)$ is defined as the response of a collection site that is centered at $(0, 0)$. Note that while in Section II we considered I at a single collection site, we now consider $I(a, b)$ defined over all the collection sites. We can break $I(a, b)$ into a constant component S and a spatially varying component $E(a, b)$ according to

$$I(a, b) = S + E(a, b), \quad (17)$$

where

$$S = T \int_{\lambda} \int_y \int_x (\bar{R}(\lambda) \bar{L}(\lambda) * p(x, y, \lambda)) t(\lambda) \cdot S_r(x - X_a, y - Y_b) q(\lambda) dx dy d\lambda \quad (18)$$

$$E(a, b) = T \int_{\lambda} \int_y \int_x (\epsilon(x, y, \lambda) * p(x, y, \lambda)) t(\lambda) \cdot S_r(x - X_a, y - Y_b) q(\lambda) dx dy d\lambda. \quad (19)$$

Since $E[\epsilon(x, y, \lambda)] = 0$ for $\lambda \in [\lambda_1, \lambda_2]$, $E(a, b)$ has zero mean and a variance that depends on the spatial variation of the reflectance of the planar surface and on the spatial variation of the illumination. The illumination $L(x, y, \lambda)$ can often be controlled in applications to achieve acceptable levels for $I(a, b)$. In addition, the overall system transmission can be scaled using neutral density filters or by adjusting the camera lens aperture.

The blurring by $p(x, y, \lambda)$ and the spatial integration by the collection sites serve to spatially low-pass filter the reflected spectral irradiance pattern $R(x, y, \lambda)L(x, y, \lambda)$. Ideally, this low-pass filtering would be sufficient to prevent aliasing during the sampling by the CCD collection sites. For radiance patterns containing high spatial frequencies, however, the filtering can be insufficient and significant aliasing can occur [11], [16]. When images are taken of patterns containing high spatial frequencies, steps must be taken to account for the aliasing of the input pattern that occurs during sampling at the collection sites.

IV. ESTIMATING SENSOR NOISE

In many situations, image sensors can be characterized and calibrated before they are used in applications. From the CCD camera model described in Section II-B, we can estimate the dependence of the noise variance on the signal level using controlled images. In this section, we describe a method for estimating this dependence for the part of the noise that does not vary from collection site to collection site. In Section V, a calibration algorithm is presented that removes

the effects of the spatially varying noise sources that are not considered in this section. Other algorithms for estimating noise variance have been applied to arbitrary images [2], [19]. These algorithms, however, assume that noise variance is independent of signal level, which as shown in II-B is not the case for images sensed by CCD cameras.

Generalizing the individual pixel model of (6), a digitized two-dimensional image of a planar surface of a material is described by

$$D(a, b) = (K(a, b)I(a, b) + N_{DC}(a, b) + N_S(a, b) + N_R(a, b))A + N_Q(a, b), \quad (20)$$

where $I(a, b)$ is modeled using (17). Quantities with properties that depend on pixel location (a, b) are indicated explicitly.

In what follows, we describe a method for estimating the variance of the noise sources that are intrinsic to the CCD. Consider a pixel (a, b) with corresponding spectral irradiance $B(x, y, \lambda)$. The observed $D(a, b)$ is a random variable

$$D(a, b) = \mu(a, b) + N(a, b), \quad (21)$$

where the expected value of $D(a, b)$ is given by

$$\mu(a, b) = K(a, b)I(a, b)A + E_{DC}(a, b)A, \quad (22)$$

where $E_{DC}(a, b)$ is the expected value of $N_{DC}(a, b)$. $N(a, b)$ is the zero mean noise given by

$$N(a, b) = N_I(a, b) + N_C(a, b). \quad (23)$$

$N_I(a, b)$ is the part of the noise that depends on the number of collected electrons

$$N_I(a, b) = N_S(a, b)A \quad (24)$$

and $N_C(a, b)$ is the part of the noise that does not depend on the number of collected electrons

$$N_C(a, b) = N_R(a, b)A + N_Q(a, b). \quad (25)$$

Since $N_S(a, b)$ is a Poisson noise source, the variance of $N_I(a, b)$ is

$$\sigma_I^2(a, b) = A^2(K(a, b)I(a, b) + E_{DC}(a, b)) \quad (26)$$

and the variance of $N_C(a, b)$ is

$$\sigma_C^2 = A^2\sigma_R^2 + \frac{q^2}{12}, \quad (27)$$

where σ_R^2 is the variance of $N_R(a, b)$ and q is explained in Section II-B. From the independence of $N_I(a, b)$ and $N_C(a, b)$, the total variance of the noise $N(a, b)$ is given by the sum of the variance of the noise component $\sigma_I^2(a, b)$ that depends on the signal level $I(a, b)$ and the variance of the constant component σ_C^2

$$\sigma_N^2(a, b) = \sigma_I^2(a, b) + \sigma_C^2. \quad (28)$$

A convenient way to estimate the total noise variance is to digitize two images D_1 and D_2 using the same nearly spatially invariant irradiance field $B(x, y, \lambda)$ for each image. Such a

field can be obtained using a nearly uniform calibration card illuminated by a single source. We then have

$$D_1(a, b) = \mu(a, b) + N_1(a, b) \quad (29)$$

$$D_2(a, b) = \mu(a, b) + N_2(a, b). \quad (30)$$

If we form the pixel-by-pixel difference image $D_\Delta(a, b)$, we obtain

$$D_\Delta(a, b) = D_1(a, b) - D_2(a, b) = N_1(a, b) - N_2(a, b). \quad (31)$$

Note that while in $D_1(a, b)$ and $D_2(a, b)$ each pixel (a, b) has a characteristic mean, each pixel of $D_\Delta(a, b)$ has a zero mean. From the independence of $N_1(a, b)$ and $N_2(a, b)$, each pixel of $D_\Delta(a, b)$ has a variance of $2\sigma_N^2(a, b)$. From (26) and (28), we can write

$$\sigma_N^2(a, b) = A^2(I(a, b) + \bar{E}_{DC}) + A^2(K(a, b) - 1)I(a, b) + \sigma_C^2 + A^2(E_{DC}(a, b) - \bar{E}_{DC}), \quad (32)$$

where \bar{E}_{DC} is the mean value of $E_{DC}(a, b)$ over the array. Using the properties that

$$|K(a, b) - 1| \ll 1 \quad (33)$$

and

$$|E_{DC}(a, b) - \bar{E}_{DC}| \ll \bar{E}_{DC} \quad (34)$$

and assuming that the spatial variation in $I(a, b)$ is small so that we can replace $I(a, b)$ by its spatial mean \bar{I} , we can approximate $\sigma_N^2(a, b)$ by

$$\sigma_N^2 = A^2(\bar{I} + \bar{E}_{DC}) + \sigma_C^2. \quad (35)$$

Using this approximation, each pixel in $D_\Delta(a, b)$ can be considered to be a sample of a zero mean random variable with a variance of $2\sigma_N^2$.

Given the digital images $D_1(a, b)$ and $D_2(a, b)$, the expected value of $\mu(a, b)$ over all of the pixels is

$$\mu = \bar{I}A + \bar{E}_{DC}A. \quad (36)$$

If each image has n rows ($1 \leq b \leq n$) and m columns ($1 \leq a \leq m$), then we can estimate μ by

$$\hat{\mu} = \frac{1}{2M} \sum_{1 \leq a \leq m} \sum_{1 \leq b \leq n} D_1(a, b) + \frac{1}{2M} \sum_{1 \leq a \leq m} \sum_{1 \leq b \leq n} D_2(a, b) \quad (37)$$

and σ_N^2 by

$$\hat{\sigma}_N^2 = \frac{1}{2(M-1)} \sum_{1 \leq a \leq m} \sum_{1 \leq b \leq n} (D_\Delta(a, b) - \hat{\mu}_\Delta)^2, \quad (38)$$

where $\hat{\mu}_\Delta$ is the sample mean of the pixels in $D_\Delta(a, b)$ and $M = nm$. $\hat{\mu}$ and $\hat{\sigma}_N^2$ are consistent estimates of μ and σ_N^2 . In general, the variance of the estimator $\hat{\sigma}_N^2$ depends on the fourth moment of the pixel values of $D_\Delta(a, b)$ about the mean. If we assume that the pixels of $D_\Delta(a, b)$ are normally distributed [23], then the variance of $\hat{\sigma}_N^2$ is

$$\text{var}[\hat{\sigma}_N^2] = \frac{2(\sigma_N^2)^2}{(M-1)} \quad (39)$$

and can be approximated by

$$\text{var}[\hat{\sigma}_N^2] \approx \frac{2(\hat{\sigma}_N^2)^2}{(M-1)}. \quad (40)$$

The normal assumption is reasonable for high levels of \bar{I} where $N_I(a, b)$ dominates $N_C(a, b)$ and the Poisson distribution of the shot noise approaches a normal distribution.

If we now consider P pairs of images $D_1(a, b)_i$ and $D_2(a, b)_i$ ($1 \leq i \leq P$) where each pair is obtained at a different level of \bar{I} , then we can compute $\hat{\mu}_i$, $\hat{\sigma}_{N_i}^2$, and $\text{var}[\hat{\sigma}_{N_i}^2]$ for each pair using (37), (38), and (40). From this set of estimates, we can estimate A and σ_C^2 .

From (35) and (36) we have that

$$\sigma_N^2 = A\mu + \sigma_C^2. \quad (41)$$

We will compute estimates \hat{A} and $\hat{\sigma}_C^2$ for the parameters A and σ_C^2 using a line fitting technique that minimizes the weighted sum of squares

$$\chi^2 = \sum_{1 \leq i \leq P} \frac{(\hat{\sigma}_{N_i}^2 - (\hat{A}\hat{\mu}_i + \hat{\sigma}_C^2))^2}{\text{var}[\hat{\sigma}_{N_i}^2]}. \quad (42)$$

This minimization produces the maximum likelihood estimates for A and σ_C^2 under the assumption that each of the $\hat{\sigma}_{N_i}^2$ is normally distributed [3]. This assumption is justifiable since $\frac{(M-1)\hat{\sigma}_{N_i}^2}{\sigma_N^2}$ has a chi-square distribution in $M-1$ degrees of freedom [23] so that for large M , $\hat{\sigma}_{N_i}^2$ will be approximately normal. Details of the minimization are given in an appendix.

V. CAMERA CALIBRATION

In this section, we describe and analyze a calibration procedure that corrects for spatial nonuniformities due to fixed pattern noise and variations in dark current. An important benefit of this procedure is that the noise in a resulting corrected image is dominated by shot noise at high levels of I rather than by fixed pattern noise.

A. Estimating Variation in Dark Current

Suppose that we digitize an image in a dark environment so that $I(a, b) \equiv 0$. From (20), each pixel in such an image is modeled by

$$D(a, b) = (N_{DC}(a, b) + N_S(a, b) + N_R(a, b))A + N_Q(a, b) \quad (43)$$

and has mean $E_{DC}(a, b)A$ and variance $\sigma_N^2(a, b)$. If we average s of these dark images, we obtain an estimate of the mean $E_{DC}(a, b)A$ that has a variance $\sigma_N^2(a, b)/s$. Let $\widehat{D}_D(a, b)$ denote a dark reference image obtained in this way.

B. Estimating Fixed Pattern Variation

In order to estimate the fixed pattern noise $K(a, b)$, we consider a series of images of a nearly uniform reflectance card illuminated by a nearly spatially uniform source. Combining (17) and (20), each image is described by

$$D(a, b) = (K(a, b)[S + E(a, b)] + N_{DC}(a, b))A + N(a, b), \quad (44)$$

where $N(a, b)$ is the zero mean noise of (23). To estimate $K(a, b)$ using $D(a, b)$, we must be able to distinguish the fixed pattern noise $K(a, b)$ from spatial variation $E(a, b)$ caused by nonuniform illumination and surface reflectance and also from spatial variation due to the other noise sources $N_{DC}(a, b)$ and $N(a, b)$.

The effects of spatially nonuniform illumination and surface reflectance can be overcome by considering n_1 different imaging configurations. Each new configuration is selected such that for each collection site, a new patch on the card illuminated by a new part of the source is imaged for each new configuration. A new configuration, therefore, can be obtained by moving both the calibration card and the light source.

Suppose that for each imaging configuration we digitize and average n_2 frames. Since $N(a, b)$ has zero mean, this averaging will produce an estimate of

$$(K(a, b)I_i(a, b) + E_{DC}(a, b))A, \quad (45)$$

where $I_i(a, b) = S + E_i(a, b)$ is the number of electrons that would be collected at pixel (a, b) by an ideal sensor for the i th imaging configuration. The variance of this estimate is $\sigma_N^2(a, b)/n_2$. By subtracting the dark reference image $\bar{D}_D(a, b)$ we can obtain for each configuration an estimate

$$e_i(a, b) \approx K(a, b)I_i(a, b)A \quad 1 \leq i \leq n_1 \quad (46)$$

with spatial variation due only to fixed pattern noise and spatial variation in illumination and surface reflectance.

To estimate $I_i(a, b)A$, we consider $m \times m$ windows centered at each (a, b) . The window size is chosen large enough so that the mean of $K(a, b)$ over the window is nearly 1 and small enough so that variation in illumination and surface reflectance is small except for the possibility of isolated outliers due, for example, to a small spot on the calibration card. For our setup, a window size of 9×9 provides these desired properties. We use the mean $\bar{e}_i(a, b)$ of $e_i(a, b)$ over an $m \times m$ window centered at (a, b) to estimate $I_i(a, b)A$.

After considering n_1 imaging configurations, for each pixel (a, b) we have n_1 points $(\bar{e}_i(a, b), e_i(a, b))$. From (46), each of these points lie approximately on a line through $(0, 0)$ with slope $K(a, b)$.

To estimate the slope $K(a, b)$, we define

$$m_i(a, b) = \frac{e_i(a, b)}{\bar{e}_i(a, b)} \quad 1 \leq i \leq n_1. \quad (47)$$

We compute the mean of the slopes of the n_1 points

$$\bar{m}(a, b) = \frac{1}{n_1} \sum_{1 \leq i \leq n_1} \frac{e_i(a, b)}{\bar{e}_i(a, b)}. \quad (48)$$

Using this line as an initial estimate, we remove outlier points from the original point set $(\bar{e}_i(a, b), e_i(a, b))$ that have $m_i(a, b)$ differing from $\bar{m}(a, b)$ by more than a fixed amount. To identify outliers, we first consider the distribution of the slopes. Variation in slopes is due mostly to small variations in card reflectance. For a good calibration card this variation is approximately Gaussian, with a small variance and with some




Fig. 1. $(\bar{e}_i(a, b), e_i(a, b))$ for fixed (a, b) , $n_1 = 20$.

isolated outlier points of much higher or lower reflectance. Fig. 1 shows a point set with one outlier point.

Our goal is to obtain an estimate of the slope that is not influenced by the outliers. We retain each point with a slope in $\bar{m}(a, b) \pm p\hat{\sigma}$, where $\hat{\sigma}$ is an estimate of the standard deviation σ in $m_i(a, b)$ over different configurations. The estimate $\hat{\sigma}$ is obtained by averaging the sample standard deviation in $m_i(a, b)$ over many different locations (a, b) . After removing these outliers, the points are averaged again and a new slope is computed to produce an estimate $\hat{K}(a, b)$ of $K(a, b)$. The estimate of $K(a, b)$ is not sensitive to the choice of p because outliers tend to be far from the mean. We used $p = 1.5$ for our experiments. Thus,

$$\hat{K}(a, b) = \frac{1}{n'_1(a, b)} \sum_{1 \leq i \leq n'_1(a, b)} \frac{e_i(a, b)}{\bar{e}_i(a, b)}, \quad (49)$$

where $(\bar{e}_i(a, b), e_i(a, b))$ are the $n'_1(a, b)$ samples that remain after removing outliers.

Since the points outside the cone defined by $\bar{m}(a, b) \pm p\hat{\sigma}$ have been removed, the variance in the remaining samples of $m_i(a, b)$ is less than $(p\hat{\sigma})^2$. Thus,

$$\text{var}[\hat{K}(a, b)] \leq \frac{1}{\hat{n}_1} (p\hat{\sigma})^2, \quad (50)$$

where var denotes variance and \hat{n}_1 is the smallest $n'_1(a, b)$ over the image. Since the quantization step of the imaging system is q and the maximum intensity representable is about $2^b q$, choosing \hat{n}_1 so that $\text{var}[\hat{K}(a, b)]$ is less than $\frac{q^2}{10(2^b q)^2}$ ensures sufficient accuracy in $\hat{K}(a, b)$. Thus, we use $\hat{n}_1 = 10(2^{2b})(p\hat{\sigma})^2$. The worst-case ratio $\frac{n_1}{\hat{n}_1}$ can be measured, allowing us to determine n_1 .

A total of n_2 images are averaged to reduce the variance of $N(a, b)$ to obtain the estimate $e_i(a, b)$ in (46). A method for estimating the variance of noise sources has been presented in Section IV. Using (41) and the estimates \hat{A} and $\hat{\sigma}_C^2$, we

can estimate the noise σ_N^2 as a function of the mean image intensity. The variance of the estimates $e_i(a, b)$ is given by σ_N^2/n_2 . Thus if we allow an upper limit of $\frac{2}{10}$ on the variance of the estimates, we can determine the required n_2 .

Define a corrected version $D_C(a, b)$ of an image $D(a, b)$ by

$$D_C(a, b) = \frac{D(a, b) - \widehat{D_D}(a, b)}{\widehat{K}(a, b)}. \quad (51)$$

From (20) and (51),

$$D_C(a, b) = \left(I(a, b) + \frac{N_S(a, b)}{\widehat{K}(a, b)} + \frac{N_R(a, b)}{\widehat{K}(a, b)} \right) A + \frac{N_Q(a, b)}{\widehat{K}(a, b)}, \quad (52)$$

where we are assuming that the errors in the estimates $\widehat{D_D}(a, b)$ and $\widehat{K}(a, b)$ are small compared to the variance of the remaining noise. At high levels of $I(a, b)$, the noise in the corrected image is dominated by the shot noise $N_S(a, b)$, which has a variance that is proportional to $I(a, b)$. This is the best we can achieve with an imaging system. By contrast, at high levels of $I(a, b)$ the noise in the uncorrected image $D(a, b)$ is dominated by the fixed pattern variation $K(a, b)$ and has a variance that is proportional to $I^2(a, b)$.

This correction procedure requires that $D(a, b)$ is a linear function of $I(a, b)$. Fortunately, the photoelectric effect is a very linear process, and CCD's exhibit excellent linearity. One must be cautious, however, because many video cameras are adjusted to produce a nonlinear response. This nonlinearity must be corrected (see [21]) if the calibration procedure of this section is used for such a camera.

VI. ESTIMATING SCENE VARIATION

Using the models of Section III, we examine how spatially varying illumination and reflectance in the scene combine with camera noise to produce variation in measured pixel values. From our camera model, we show that it is possible to produce separate estimates of the image variance due to camera noise and to scene variation. In addition, this analysis allows us to compute a measure of scene variation from an image that is independent of the image irradiance level.

As in Section III, we consider the case of planar surfaces of a single material. This case is of particular interest because the amount of scene variation measured in images of such surfaces is often comparable to the level of image variation due to camera noise. From (17) and (52), the corrected image of such a surface is

$$D_C(a, b) = SA + E(a, b)A + \frac{1}{\widehat{K}(a, b)}(N_I(a, b) + N_C(a, b)), \quad (53)$$

where $N_I(a, b)$ and $N_C(a, b)$ are the noise components defined by (24) and (25). The mean of $D_C(a, b)$ is

$$\mu_D = SA, \quad (54)$$

and the spatial variance of $D_C(a, b)$ is

$$\sigma_D^2 = \sigma_E^2 A^2 + \sigma_{KN}^2, \quad (55)$$

where σ_E^2 is the spatial variance of $E(a, b)$ and

$$\sigma_{KN}^2 = E \left[\frac{1}{\widehat{K}^2(a, b)} \right] \sigma_N^2, \quad (56)$$

where σ_N^2 is defined by (35). From (54) and (55), we have that

$$\sigma_D^2 - \sigma_{KN}^2 = \frac{\sigma_E^2}{S^2} \mu_D^2. \quad (57)$$

We will refer to $V_S = \frac{\sigma_E^2}{S^2}$ as the scene variation. As shown in Section III, V_S depends on the surface, the illumination, and the imaging configuration, but it is independent of scalings of either the illumination or the imaging system response.

Using (57), V_S can be approximated from estimates of σ_D^2 , σ_{KN}^2 , and μ_D . We can estimate μ_D and σ_D^2 from $D_C(a, b)$ using the sample mean $\widehat{\mu}_D$ and sample variance $\widehat{\sigma}_D^2$. To estimate σ_{KN}^2 in (56), we must estimate each factor. The expected value in (56) can be estimated from the values of $\widehat{K}(a, b)$ corresponding to $D_C(a, b)$. Equation (41) can be used to estimate σ_N^2 by

$$\widehat{\sigma}_N^2 = \widehat{A}(\widehat{\mu}_D + \overline{D_D}) + \widehat{\sigma}_C^2, \quad (58)$$

where $\overline{D_D}$ is the spatial mean of the dark reference image $\overline{D_D}(a, b)$.

Since the scene variation V_S is independent of the level of image irradiance, it is useful for identifying a surface in situations where the imaging configuration is fixed but where μ_D may change over time. Such an environment frequently occurs in visual inspection tasks where the imaging geometry is tightly controlled, but where the observed image irradiance values may change due, for example, to drift in the illumination intensity or a change in the camera integration time. For these cases, the spatial variance of the measured pixel values will not be a stable scene descriptor since it changes with image irradiance. The scene variation V_S , on the other hand, will not change with changes in μ_D and therefore can be used for scene description and surface identification over a range of image irradiance levels.

VII. EXPERIMENTAL RESULTS

The methods described in this paper were used to characterize and calibrate a monochrome video CCD camera imaging system based on a Sony XC-77 camera and a RasterOps TC-PIP frame grabber. The camera has a linear response and is equipped with an infrared blocking filter. In this section, experimental results are presented to verify the accuracy of the models for sensor noise and scene variation.

A. Estimating Sensor Noise

The parameters A and σ_C^2 that describe the noise variance as a function of the signal level μ have been estimated for our imaging system using the algorithm described in Section IV. The estimates $\widehat{\mu}$, $\widehat{\sigma}_N^2$, and $\text{var}[\widehat{\sigma}_N^2]$ were computed using (37), (38), and (40) from a fixed 100×100 -pixel region in images of a nearly uniform reflectance card (Applied Image MT-6) illuminated by a nearly spatially uniform source (Newport Model 765). Different levels of μ were obtained by placing




Fig. 2. The sensor noise estimate.

different neutral density filters in front of the camera. The results of the estimation are plotted in Fig. 2. Each "o" indicates a point $(\hat{\mu}_i, \hat{\sigma}_{N_i}^2)$ that has been estimated from a pair of images $D_1(a, b)_i$ and $D_2(a, b)_i$, as described in Section IV. The line in Fig. 2 is the best fit to the points in the sense of (42). The estimated parameters are $\hat{A} = 0.003384$ and $\hat{\sigma}_C^2 = 1.1196$, giving a χ^2 of 4.3496. This value of χ^2 corresponds to a goodness of fit probability $Q(\chi^2/4) = 0.3607$ (see Appendix). Thus it is reasonable to accept the model of (41) and the fitted parameters.

The camera calibration method of Section V determined n_1 to be 11 and n_2 to be 18 for our imaging system. Keeping a safety margin, we used $n_1 = 20$ and $n_2 = 20$ to determine $\widehat{D}_D(a, b)$ and $\widehat{K}(a, b)$. The mean of $\widehat{K}(a, b)$ over the array was 0.99998, with a standard deviation of 3.72×10^{-3} . The resulting $\widehat{K}(a, b)$ applied to an image of maximum intensity causes 27% of the pixels to be modified by an amount exceeding one quantization step, with 3% of these modifications exceeding two quantization steps. The maximum modification over the sensor is five quantization steps at maximum intensity. This was repeatable over several experiments. The mean of $\widehat{D}_D(a, b)$ over the array was 10.74, with a standard deviation of 0.555.

B. Estimating Scene Variation

The relationship defining scene variation in terms of image statistics and sensor noise (see (57)) was tested on images of painted squares on the Macbeth color chart using 100×100 -pixel regions within a painted square. The parameters μ_D and σ_D^2 are estimated using the sample mean and sample variance of the pixels in the region. The value of σ_N^2 in (56) is estimated using (58) with the fitted values of \hat{A} and $\hat{\sigma}_C^2$ from the noise estimation procedure. The expected value in (56) and \widehat{D}_D in (58) are estimated by the appropriate means over the area in the calibration images $\widehat{K}(a, b)$ and $\widehat{D}_D(a, b)$ that corresponds to the 100×100 region of interest.

Fig. 3 shows the results for four different painted squares. Each of the seven o's indicate a point $(\hat{\mu}_D, \hat{\sigma}_D^2 - \hat{\sigma}_{KN}^2)$ estimated from a fixed 100×100 image region within the




Fig. 3. Experimental verification of (57).

square. Different levels of μ_D are generated by placing neutral density filters in front of the camera. We recall from Section VI that the scene variation V_S does not change when μ_D is changed in this way. The curves in Fig. 3 are the best least squares fit to the points of the form given in (57) and show that the dependence is predicted accurately by the model for this data.

C. Relative Contributions to Image Variance

The variance over a region of an image of a single surface is due to variation in the illumination and surface reflectance, spatial variation in the dark current, variation in $K(a, b)$, and the sensor noise components. The spatial variance in dark current and σ_C^2 are independent of intensity level. The shot noise variance is linearly dependent on intensity, and scene variance is proportional to the square of intensity.

Fig. 4 shows the fractional contribution of each of these factors to the total variance for a square on the color chart. Note that the total variance increases as μ increases. We observe that the scene variance is the major contributor to the total variance at higher intensities, while σ_C^2 is the dominant contributor at lower intensities.

VIII. CONCLUSION

Image sensors play the important role in machine vision of transforming the signal of interest in the 3-D world into the digital image that is processed by algorithms. Given the underconstrained character of the problems addressed by machine vision, accurate sensor models can often be used to improve the capability of an algorithm. In this paper, we have analyzed the noise properties of CCD camera imaging systems. This analysis was used to develop techniques for noise characterization and camera calibration. Experiments with a CCD video camera suggest that these models can be used to quantify accurately the noise properties of a CCD imaging system.

Using a physical reflectance model, we have examined the effect of small variations in material reflectance on digitized image irradiance values. Using our sensor noise model and




Fig. 4. Parts of total noise versus intensity.

calibration scheme, the image variance due to sensor noise can be separated from the variance due to effects in the scene. This allows us to estimate scene variation for a surface patch from its image. This quantity is independent of image irradiance and can be used for surface identification in the presence of changes in illumination.

APPENDIX

The parameters \hat{A} and $\hat{\sigma}_C^2$ that minimize χ^2 of (42) are given by [3]

$$\hat{A} = \frac{S_1 S_5 - S_2 S_3}{S_1 S_4 - (S_2)^2} \quad (59)$$

$$\hat{\sigma}_C^2 = \frac{S_3 S_4 - S_2 S_5}{S_1 S_4 - (S_2)^2} \quad (60)$$

where

$$S_1 = \sum_{1 \leq i \leq P} \frac{1}{\text{var}[\hat{\sigma}_{N_i}^2]} \quad (61)$$

$$S_2 = \sum_{1 \leq i \leq P} \frac{\hat{\mu}_i}{\text{var}[\hat{\sigma}_{N_i}^2]} \quad (62)$$

$$S_3 = \sum_{1 \leq i \leq P} \frac{\hat{\sigma}_{N_i}^2}{\text{var}[\hat{\sigma}_{N_i}^2]} \quad (63)$$

$$S_4 = \sum_{1 \leq i \leq P} \frac{\hat{\mu}_i^2}{\text{var}[\hat{\sigma}_{N_i}^2]} \quad (64)$$

$$S_5 = \sum_{1 \leq i \leq P} \frac{\hat{\mu}_i \hat{\sigma}_{N_i}^2}{\text{var}[\hat{\sigma}_{N_i}^2]} \quad (65)$$

The covariance matrix that quantifies the uncertainty in the estimate $(\hat{A}, \hat{\sigma}_C^2)$ is given by

$$C = \frac{1}{S_1 S_4 - S_2^2} \begin{bmatrix} S_1 & -S_2 \\ -S_2 & S_4 \end{bmatrix}. \quad (66)$$

The value of χ^2 for the fit has a chi-square distribution with $P - 2$ degrees of freedom. If χ_M^2 is the value of χ^2 corresponding to the estimate $(\hat{A}, \hat{\sigma}_C^2)$, then it is useful to consider the goodness of fit measure

$$Q(\chi_M^2/P - 2) = \int_{\chi_M^2}^{\infty} f(x) dx, \quad (67)$$

where $f(x)$ is the chi-square density with $P - 2$ degrees of freedom. $Q(\chi_M^2/P - 2)$ is then the probability that a value of χ^2 obtained by the fitting procedure will exceed χ_M^2 given that the model is correct. Thus, very small values of Q are evidence for an incorrect model.

ACKNOWLEDGMENT

The authors thank J. Dadlani and C. Lee for dedicated programming support. We also thank Referee A, whose careful review and thoughtful comments greatly improved the clarity of this paper.

REFERENCES


- [1] G. Amelio, M. Tompsett, and G. Smith, "Experimental verification of the charge coupled device concept," *Bell Syst. Tech. J.*, vol. 49, pp. 593-600, Apr. 1970.
- [2] P. Besl and R. Jain, "Segmentation through variable-order surface fitting," *IEEE Trans. Pattern Anal. Machine Intell.*, vol. 10, no. 2, pp. 167-192, Mar. 1988.
- [3] P. Bevington, *Data Reduction and Error Analysis for the Physical Sciences*. New York: McGraw-Hill, 1969.
- [4] H. Beyer, "Linejitter and geometric calibration of CCD cameras," *ISPRS J. Photogrammetry and Remote Sensing*, vol. 45, pp. 17-32, 1990.
- [5] M. Blouke and J. Janesick, *Arms Control Verification: The Technologies That Make It Possible*. 1987.
- [6] W. Boyle and G. Smith, "Charge coupled semiconductor devices," *Bell Syst. Tech. J.*, vol. 49, pp. 587-593, Apr. 1970.
- [7] J. Carnes and W. Kosonocky, "Noise sources in charge-coupled devices," *RCA Rev.*, vol. 33, pp. 327-343, June 1972.
- [8] M. Collet, "Solid-state image sensors," *Sensors and Actuators*, vol. 10, pp. 287-302, 1986.
- [9] G. Healey, "Using color for geometry insensitive segmentation," *J. Opt. Soc. Am. A*, vol. 6, no. 6, pp. 920-937, June 1989.
- [10] B. K. P. Horn, "Shape from shading: A method for obtaining the shape of a smooth opaque object from one view," MIT AI Lab., Tech. Rep. 232, Nov. 1970.
- [11] F. Huck, N. Halyo, and S. Park, "Aliasing and blurring in 2-D sampled imagery," *Appl. Opt.*, vol. 19, no. 23, pp. 2174-2181, July 1980.
- [12] K. Ikeuchi and T. Kanade, "Modeling sensors: Toward automatic generation of object recognition program," *Comp. Vision, Graphics, Image Processing*, vol. 48, pp. 50-79, 1989.
- [13] J. Janesick, T. Elliott, S. Collins, M. Blouke, and J. Freeman, "Scientific charge-coupled devices," *Opt. Eng.*, vol. 26, no. 8, pp. 692-715, Aug. 1987.
- [14] G.J. Klinker, S.A. Shafer, and T. Kanade, "A physical approach to color image understanding," *Int. J. Comput. Vision*, vol. 4, pp. 7-38, 1990.
- [15] J. Kristian and M. Blouke, "Charge-coupled devices in astronomy," *Scientific American*, vol. 247, pp. 67-74, Oct. 1982.
- [16] J. Krumm and S.A. Shafer, "Sampled-grating and crossed-grating models of moire patterns from digital imaging," *Opt. Eng.*, vol. 30, no. 2, pp. 195-206, Feb. 1991.
- [17] H.-C. Lee, E. Breneman, and C. Schulte, "Modeling light reflection for computer color vision," *IEEE Trans. Pattern Anal. Machine Intell.*, vol. 12, no. 4, pp. 402-409, Apr. 1990.
- [18] R. Lenz and D. Fritsch, "Accuracy of videometry with CCD sensors," *ISPRS J. Photogrammetry Remote Sensing*, vol. 45, pp. 90-110, 1990.

- [19] P. Meer, J.-M. Jolion, and A. Rosenfeld, "A fast parallel algorithm for blind estimation of noise variance," *IEEE Trans. Pattern Anal. Machine Intell.*, vol. 12, no. 2, pp. 216-223, Feb. 1990.
- [20] S.K. Nayar, K. Ikeuchi, and T. Kanade, "Surface reflection: Physical and geometrical perspectives," *IEEE Trans. Pattern Anal. Machine Intell.*, vol. 13, no. 7, pp. 611-634, July 1991.
- [21] C. Novak, S. Shafer, and R. Wilson, "Obtaining accurate color images for machine vision research," in *Proc. SPIE Conf. Perceiving, Measuring, and Using Color*, Santa Clara, CA, Feb. 1990.
- [22] A. Oppenheim and R. Schaffer, *Digital Signal Processing*. Englewood Cliffs, NJ: Prentice, 1975.
- [23] A. Papoulis, *Probability and Statistics*. Englewood Cliffs, NJ: Prentice, 1990.
- [24] S.A. Shafer, "Using color to separate reflection components," *COLOR Research and Application*, vol. 10, no. 4, pp. 210-218, 1985.
- [25] R. Tsai, "A versatile camera calibration technique for high accuracy 3D machine vision metrology using off-the-shelf TV cameras and lenses," *IEEE J. Robotics and Automat.*, vol. 3, no. 4, pp. 323-344, Aug. 1987.
- [26] L. Wolff, "Polarization-based material classification from specular reflection," *IEEE Trans. Pattern Anal. Machine Intell.*, vol. 12, no. 11, pp. 1059-1071, 1990.



Glenn E. Healey (S'85-M'89) received the B.S.E. degree in computer engineering from the University of Michigan, Ann Arbor, in 1984, and the M.S. degree in computer science in 1985, the M.S. degree in mathematics in 1986, and the Ph.D. degree in computer science from Stanford University, Stanford, CA in 1988.

From 1984 to 1989 he was affiliated with the Computer Science Department, IBM Almaden Research Center as a Research Student Associate and Visiting Scientist. In 1989 he joined the Electrical and Computer Engineering Department at the University of California, Irvine, where he is currently an Assistant Professor. His research interests in machine vision and image processing include physical modeling of visible and infrared image formation, algorithm development and characterization, and high-performance architectures.



Raghava Kondepudy (S'89) received the B.Tech. (Hons.) degree in electronics and electrical communication from the Indian Institute of Technology, Kharagpur, India, in 1990, and the M.S. degree from the University of California, Irvine, where he is currently a doctoral student.

His interests include model development for applications in computer vision and image processing.