Area Estimation with Subpixel Accuracy for Industrial Imaging Systems

M.S. klein Gebbinck, Theo E. Schouten

University of Nijmegen, Computing Science Institute
P.O. Box 9010, 6500 GL Nijmegen, The Netherlands
E-mail: mauricek@cs.kun.nl

June 30, 1997

Abstract

Industrial imaging systems used for product quality control often need to make an accurate estimation of an object's area. Usually the area is estimated by counting the number of pixels classified as belonging to the object. However, with a technique called decomposition, the number of pixels can be estimated more accurately at the same resolution, or as accurately at a lower resolution. Another way to reduce the hardware costs is to lower the spatial sampling rate. The negative effects on the area estimation accuracy can be counteracted by reconstructing the original image using interpolation. In this paper, we present the results of two experiments using HP ScanJet 4c images in which the capabilities and limitations of decomposition and interpolation were assessed quantitatively. It appeared that the resolution could be lowered with a factor 25 without decreasing the accuracy of the estimation. A reduction of the sampling rate, on the other hand, could not be compensated for as well. However, interpolation did seem to restrict the decrease in accuracy somewhat, although the results were not entirely conclusive due to hardware related reasons.

Keywords: industrial imaging system, HP ScanJet 4c scanner, area estimation, subpixel resolution, mixed pixels, linear mixture model, classification, decomposition, undersampling, linear interpolation, cubic convolution.

Classification: AMS-1991: 68U10; CR-1991: I.4.7, I.4.6, I.4.4, I.4.1, I.4.9.

1 Introduction

In the industry, imaging systems are often used for automated product quality control. In general these imaging systems must be highly cost effective out of competitive considerations. For a manufacturer with many production lines, each having several control points, it is important that the sensor unit in particular is as cheap as possible. If the sensor units—usually CCD[1]-cameras—operate under strenuous conditions and have to be replaced often, the requirement that these units should be low-cost is even stronger.

As the price of a CCD-camera goes down, the quality of the image it produces also deteriorates. One possibility to cut costs is to use less expensive optical components, e.g. lenses, which has a negative influence on the camera's resolution. Further savings can be achieved by using less expensive electronic circuitry, reducing the spatial and/or temporal sampling rate. For most applications the worse image characteristics decrease the accuracy of the estimation of the product properties. Fortunately, in some cases the negative effects introduced by inferior hardware can be compensated for in software.

[1] Charge Coupled Device

In this paper two methods for area estimation are described, which utilise the image grey level information to achieve more accurate results than could be expected given the image characteristics. In Section 3 a technique called *decomposition* is described, which estimates the fraction of a pixel covered by the object of interest. This method can be used to counteract the loss of accuracy due to an increased resolution cell size of the camera. The second method, *interpolation*, attempts to restore the originally continuous image, after which the object area is determined. Under certain conditions this technique, which is explained in Section 4, can be applied to undo the effects of a lowered spatial sampling rate. Both methods were applied to images captured with the Hewlett-Packard ScanJet 4c scanner (described in Section 2). The results of these experiments are presented in Section 5, while in Section 6 the conclusions are summarised.

2 The HP ScanJet 4C scanner

The operation of a flatbed color scanner is in many ways similar to that of a multispectral along-track scanner like the well-known HRV scanner aboard the SPOT satellite. Figure 1(a), adopted from Sabins [18], shows the operation of a typical multispectral along-track scanner. This scanner is equipped with three arrays of detectors, one for each spectral band, which are oriented perpendicularly to the flight direction. Associated with each detector is an instantaneous field of view (IFOV), subtending an area called ground resolution cell. All energy radiated or reflected from a resolution cell is focused onto the corresponding detector and transformed into an electrical signal. At regular intervals the signals of the all detectors are sampled and quantised in order to convert them into an array of digital numbers (DN's). As the satellite moves on, the IFOV of each detector sweeps other resolution cells lying further along the satellite's flight path.

One of the differences between an along-track scanner and a flatbed scanner is that the first relies on solar radiation while the latter needs an separate light source to expose the original. As is shown in Figure 1(b) (adopted from Bessner and Grant [1]), the radiation emitted by the light bar is reflected by the original, and the image is focused onto an array of CCD's using a lens and a couple of mirrors. After the output of the array has been recorded, the light bar moves on to scan the next line of the image. To capture color images, some flatbed scanners make three passes over the image, while applying a red, green, and blue filter. Other scanners only make a single pass and use a prism to disperse the light onto three different arrays of CCD's. An advantage of the latter approach is that the three images always match perfectly, while this might not be the case when doing three scans.

The Hewlett-Packard ScanJet 4c is a flatbed single-pass color scanner of around \$1000. The optical resolution of this scanner is 600dpi[2] in horizontal as well as in vertical direction. In fact, the term resolution is a bit confusing in this case. Sabins [18] defines spatial resolution as "the ability to distinguish between two closely spaced objects on an image." This definition shows that resolution is a subjective measure which depends on many factors, e.g. resolution cell size, spatial sampling rate, number of bits used for quantisation, etc. Specifying an optical resolution of 600dpi means that the maximum spatial sampling rate is 600 samples per inch. In horizontal direction this implies that 600 CCD's are placed on every inch of the sensor array, while in vertical direction the motor moving the light bar can take steps as small as $\frac{1}{600}$ inch (see Bone [2]). It is assumed that at the maximum sampling rate the consecutively sampled IFOV's are just contiguous, so there is no gap between neighbouring image pixels, nor do they overlap. According to Park *et al.* [14] this design is common for most sensor array systems, although to this date HP has not confirmed this.

Most scanners, including the ScanJet 4c, are able to scan at a different sampling rate than is specified by the optical resolution. Seemingly higher sampling rates are achieved through interpolation (see Section 4), which combines several truly sampled pixels to form an extra pixel in the midst of them. Because these new pixels are only pseudo-data, the spatial resolution is not improved. Furthermore, since the resolution cell size has not changed, pixels now have a significant overlap, as can be seen in Figure 2(a). Sampling at a lower rate in horizontal direction is usually

[2] dots per inch

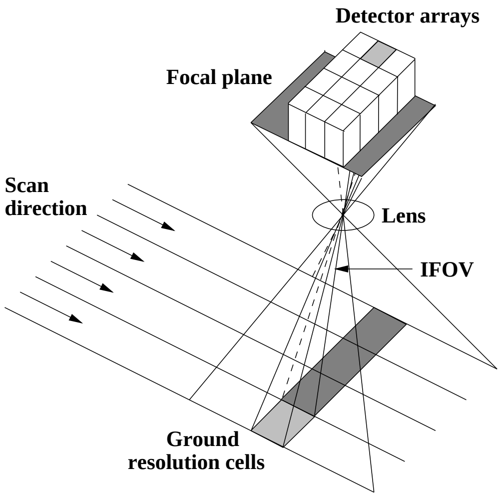

(a) Multispectral along-track scanner.

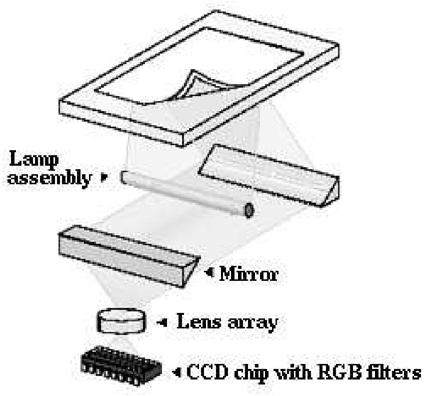

(b) Flatbed color scanner.

Figure 1: Multispectral along-track scanners and flatbed color scanners operate in a similar way. Radiation that is reflected within the scanner's IFOV is focused onto the corresponding detector in the array. After quantisation, the next line in the scan direction is sampled.

achieved by reading out only one out of every n CCD's. In vertical direction the same is accomplished by running the stepper motor moving the light bar at n times its minimum speed. This strategy is possible only if the desired sampling rate is an integral factor of the maximum sampling rate, and is applied solely by the lower-end scanners for which speed is of great importance (see Bone [3]). High-end scanners utilise a computationally more expensive technique, which must also be used for sampling rates that do not fit the maximum sampling rate an integral number of times. This technique, which is called resampling, consists of sampling the image at a higher rate than desired, after which it is converted to the correct frequency through interpolation[3]. A result of using a lower sampling rate is that neighbouring pixels are no longer contiguous, which is shown in Figure 2(b). For some applications like area estimation this may pose a problem.

3 Decomposition

The most likely measure to reduce the cost of a vision system or scanner is to use cheaper optical components which amount to a wider point spread function, in combination with a CCD-chip with fewer (non-contiguous) detector elements. Another option is to increase the size of the detector elements on the CCD-chip while reducing their number. In both cases the IFOV of each detector is increased, which leads to an image having less pixels, but where each pixel represents a larger area. As a result of the increased resolution cell size, smaller objects are harder to detect, and a larger fraction of the pixels of bigger objects is mixed. These *mixed pixels* occur at object boundaries and represent areas that partially belong to one object and partially to another object such as the background. Since the area of an object is usually estimated by summing the area of all its pixels, mixed pixels are a source of inaccuracy. For example, a pixel which is covered by a particular object for only 70% introduces an error of 30% of the total resolution cell area when being assigned to that object. However, based on the linear mixture model it is possible to decompose a mixed pixel using the spectral characteristics of its contributing objects, leading to a more accurate area estimation.

[3]For some sampling rates it is possible to sample at a lower rate and resample to the higher frequency without noticeably different results but an increase in speed (Bone [3]).

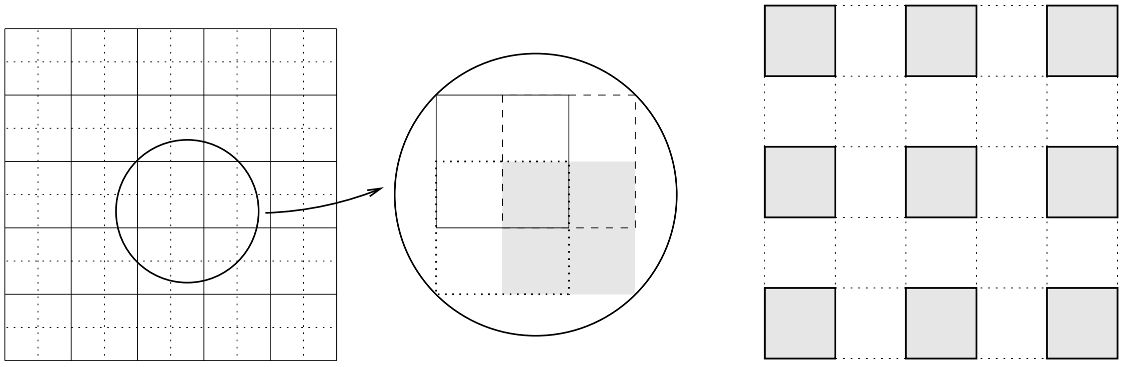

(a) The effects of using a higher sampling rate. (b) Using a lower sampling rate.

Figure 2: The consequence of using a different sampling rate. (a) Doubling the sampling rate results in quadrupling the number of pixels. Since the resolution cell size remains the same, neighbouring pixels now exhibit significant overlap. (b) Using half the maximum sampling rate reduces the number of pixels by $\frac{1}{4}$. Because the resolution cell size is unchanged, gaps between neighbouring pixels (grey squares) are introduced.

3.1 Linear mixture modelling

The image of an object that is large compared to the resolution cell size consists of a cluster of pure pixels, representing the interior of the object, surrounded by layer of mixed pixels, which are positioned at the object boundaries. With segmentation techniques such as region growing and edge detection these pure pixel regions can be identified, indicating which pixels may be mixed at the same time[4]. For each region that is big enough, statistics like the mean vector and the variance-covariance matrix can be calculated which are typical for the spectral reflectance of the corresponding object. Based on these characteristics or using class distributions from a database, mixed pixels can be decomposed and fractionally assigned to neighbouring regions, as has been demonstrated for satellite remote sensing images by Klein Gebbinck and Schouten [10, 11].

For remotely-sensed images the linear mixture model has shown to be a good abstraction for mixed pixels. Although there are many variants on this model (Pech *et al.* [16], Settle and Drake [19], Klein Gebbinck and Schouten [9]), the classical linear mixture model by Horwitz *et al.* [5] is still much used. Following the nomenclature of Settle and Drake, this model is given by

$$\mathbf{x} = \mathbf{M}\mathbf{f} + \mathbf{e}, \tag{1}$$

where vector $\mathbf{x}$ denotes the n-dimensional pixel, and vector $\mathbf{f}$ represents the proportions of the c objects that take part in the mixed pixel. The columns of the (nxc) matrix $\mathbf{M}$ list the mean vectors of the contributing objects, as calculated from the neighbouring regions of pure pixels. The variation in the reflectance of the objects, which are assumed to have multivariate normal distributions, is covered by the error vector $\mathbf{e}$. In theory $\mathbf{e}$ is also distributed multivariate normally, having a zero mean and a variance-covariance matrix $\mathbf{N}$ depending on the distributions as well as the fractions of the objects. However, it has been argued (Horwitz *et al.* [5]) that $\mathbf{N}$ can be approximated by the averaged variance-covariance matrices of the objects without much loss of accuracy, but resulting in a set of equations that is much more tractable.

If the number of components is at most equal to the dimensionality of the data ($c \leq n$), Equation (1) can be solved in infinitely many ways. From these solutions the one with the highest probability is chosen, which according to the maximum likelihood approach minimises

$$\mathbf{e}^T\mathbf{N}^{-1}\mathbf{e} + \ln|\mathbf{N}| = (\mathbf{x} - \mathbf{M}\mathbf{f})^T\mathbf{N}^{-1}(\mathbf{x} - \mathbf{M}\mathbf{f}) + \ln|\mathbf{N}| \tag{2}$$

[4]Isolated (pure) pixels constitute another class, representing objects having the size of a resolution cell.

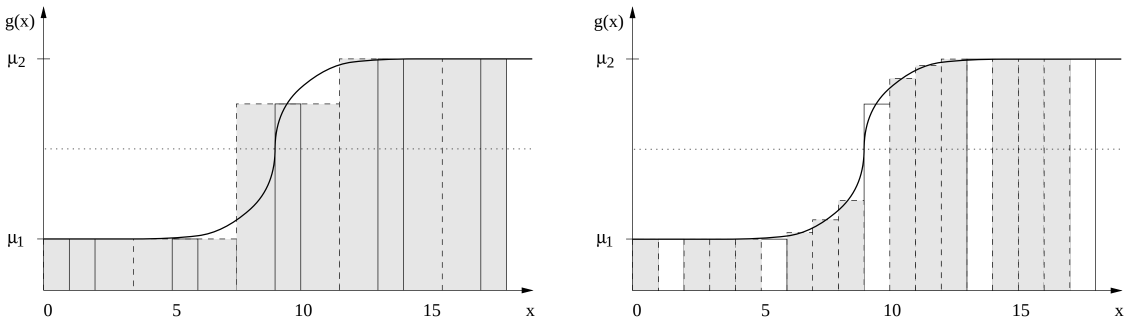

(a) Scale-up the area per sample.

(b) Interpolate between samples.

Figure 3: Handling a lower sampling rate. (a) The sampling rate of the one-dimensional signal (solid curve) has been reduced with $\frac{1}{4}$, so the length of each sample (solid rectangles) is multiplied by 4 (dashed rectangles). Classification of the samples—the dotted line is the decision boundary—allocates $\frac{15}{36}$ of the interval [0,18] to class 1 and $\frac{21}{36}$ to class 2, while according to the continuous signal both classes have the same length. (b) Reduction of the sampling rate with $\frac{1}{4}$ has been handled by interpolating three points (dashed, grey rectangles) between two successive samples (solid, white rectangles). Classification allocates the same length to both classes.

with $\ln|\mathbf{N}|$ representing the natural logarithm of the determinant of $\mathbf{N}$. In case that the variance-covariance matrix $\mathbf{N}$ is independent of the fractions vector $\mathbf{f}$, it can be derived analytically that the optimal solution is given by:

$$\hat{\mathbf{f}} = \mathbf{U}\mathbf{M}^T\mathbf{N}^{-1}\mathbf{x} \quad \text{with} \quad \mathbf{U}^{-1} = \mathbf{M}^T\mathbf{N}^{-1}\mathbf{M}. \tag{3}$$

Imposing the condition that the proportions should sum to one gives a slightly different solution, which can be found through standard Lagrangian analysis. The other constraint that should be satisfied is that all proportions should be positive (or zero), but this requires some additional processing afterwards. For further reading about both constraints is referred to Settle and Drake [19].

4 Interpolation

Another way to reduce the cost of an imaging system and to increase the speed of image acquisition at the same time is lowering the spatial sampling rate. As has been described in Section 2, most scanners are able to do so, scanning images with a sampling rate that is below the optical resolution. Although this strategy requires less memory and is considerably faster, it results in an image where neighbouring pixels no longer represent contiguous areas (see Figure 2(b)). For area estimation in particular this is a big problem, since the area of an object is usually calculated by summing the area of all pixels belonging to that object. One possibility is to multiply the object area by a factor equal to the square of the reciprocal of the factor by which the sampling rate has been reduced, e.g. lowering the sampling rate by $\frac{1}{2}$ gives a multiplication factor of 4. As can be seen in Figure 3(a), this usually leads to an inaccurate estimation. An approach which may give more accurate results can be found in the field of image registration, where the objective is to register a pair of images with subpixel accuracy. The intensity interpolation method (Tian and Huhns [21]) attempts to reconstruct one of the images as closely as possible before the displacement of the other image is calculated. If the same reconstruction technique is used before area estimation, the accuracy of the estimation may be higher as well—see Figure 3(b). Depending on the image characteristics and the sampling rate used, perfect reconstruction and thus exact area estimation is possible.

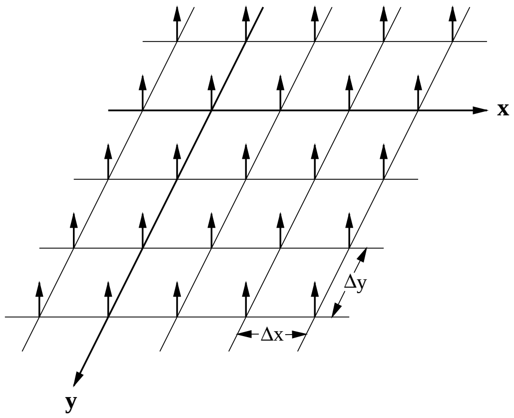

(a) Dirac delta function sampling array.

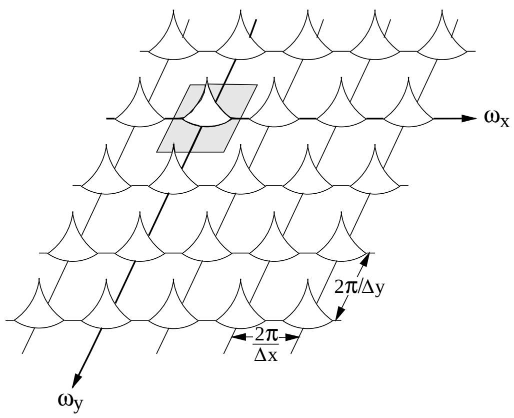

(b) Spectrum of a sampled, band-limited image.

Figure 4: Sampling of an image. (a) A continuous image is sampled by multiplying it by a two-dimensional sampling function. (b) The spectrum of the sampled image contains many copies of the spectrum of the continuous image. If the ideal image is band-limited and not undersampled, it can be reconstructed by taking the inverse Fourier transform of the spectrum in the fundamental interval (grey rectangle).

4.1 Image sampling and reconstruction

In a perfect image sampling system, a continuous, ideal image field is sampled by multiplying it by a spatial sampling function(Gonzalez and Woods [4], Pratt [17]). The sampling function is composed of an infinite array of Dirac delta functions, which are arranged in a grid of spacing $(\Delta x, \Delta y)$ as is shown in Figure 4(a). After taking the two-dimensional Fourier transform, it can be observed (Figure 4(b)) that the spectrum of the sampled image consists of the spectrum of the ideal image infinitely repeated over the frequency plane in a grid of spacing $(\frac{2\pi}{\Delta x}, \frac{2\pi}{\Delta y})$. As long as the different copies of the spectrum do not overlap, the continuous image can be reconstructed by selecting the spectrum in the fundamental interval using an ideal low-pass filter, and taking the inverse Fourier transform.

Unfortunately, perfect reconstruction of the continuous image can hardly ever be achieved in practice. One reason is that digital images are always spatially limited but rarely periodic, leading to visible distortions at the image boundaries. However, much more serious artifacts are introduced by overlapping spectra. A first cause of spectral overlap is that the ideal image is not band-limited, i.e. at infinitely high frequencies the spectrum of the image still displays some power. For example, if an image contains a sharp edge, its spectrum contains high frequency components. Another reason for spectral overlap is that the rate at which the continuous image was sampled is lower than the Nyquist frequency[5]. As a result, the grid in the frequency plane becomes too closely spaced to fit the spectra without foldover. Taking the inverse Fourier transform of the spectrum in the fundamental interval will not yield the original image, because the spectrum is corrupted by spurious frequency components from adjacent spectra. This type of distortion, which is usually called aliasing, can be avoided by applying a low-pass filter to the image prior to Fourier transformation. However, removal of the higher frequencies may blur the image considerably.

Image reconstruction can be done in the spatial domain just as well as in the frequency domain. One advantage of working in the spatial domain is that interpolation, as this type of reconstruction is often called, does not require computationally expensive Fourier transforms to be calculated. According to the Fourier convolution theorem, multiplying two Fourier transformed functions is

[5]The Nyquist frequency is determined by the highest frequencies in the image. If these cutoff frequencies are ω_{xc} and ω_{yc}, the sampling intervals should satisfy $\Delta x \leq \frac{\pi}{\omega_{xc}}$ and $\Delta y \leq \frac{\pi}{\omega_{yc}}$.

equal to convolving the corresponding spatial functions and taking the Fourier transformation of the result. Since ideal low-pass filtering is nothing more than multiplying the image spectrum by a rectangular window, perfect interpolation in the spatial domain is achieved by convolving the image with the inverse Fourier transform of that window function. Unfortunately, the inverse Fourier transform of a (one-dimensional) rectangular window is the sinc function, which is defined as $\mathrm{sinc}(x) = \frac{\sin(\pi x)}{\pi x}$. The sinc is not spatially limited, resulting in the multiplication of every pixel by all others when performing the convolution. Since this is undesirable for practical reasons, other interpolation kernels with a small width must be used. Well-known examples of this approach are linear and cubic convolution interpolation.

4.2 Linear interpolation

Nearest neighbour, linear, and cubic B-spline interpolation all belong to the same family of interpolation methods. Nearest neighbour interpolation (Parker *et al.* [15]) gives the point to be interpolated the value of the sample to which it is closest. For a one-dimensional signal this is equivalent to convolution with a rectangular kernel that has a width of one sampling interval. Linear interpolation (Parker *et al.* [15]) assigns the intermediate point a value that is a linear combination of the two samples on both sides of it, where each weight factor is inversely proportional to the distance of corresponding sample. The interpolation kernel that is needed is found by convolving the nearest neighbour kernel with itself. The cubic B-spline interpolation kernel (Hou and Andrews [6]), finally, is constructed by four convolutions of the rectangular kernel. As a result, the kernel spans four samples, two on each side of the point to be interpolated. Although cubic B-spline interpolation gives the best results, it is also computationally expensive because for every four consecutive samples a matrix problem has to be solved (Hou andAndrews [6]). Nearest neighbour interpolation, on the other hand, does not suffice either, because samples are merely replicated. For the problem of area estimation this boils down to exact same situation as is reached by scaling the area per sample—see Figure 3. Linear interpolation seems to combine the best of both worlds: it has a computational complexity that is just above the complexity of nearest neighbour interpolation, and has an accuracy that is only slightly less than that of cubic B-spline interpolation.

In the previous paragraph it was stated that the linear interpolation kernel equals the convolution of a rectangle function with itself. As can be seen in Figure 5(a), the result is a triangular kernel whose base is two samples wide. The Fourier transform of this kernel is displayed in Figure 5(b). A comparison between the spectra of the linear interpolation kernel and the ideal low-pass filter shows some loss of power below the Nyquist frequency, resulting in an image that is a bit blurred. Furthermore, some spurious frequencies are present above the Nyquist frequency, possibly leading to aliasing effects. However, these effects will probably have only a minor influence on the accuracy of the final area estimation.

In order to reconstruct a continuous function $g(x)$, the sampled function $g_s(x_n)$ is convolved with a scaled version of the kernel $r(x)$. As is shown in Figure 6(a), this means that each sample is overlaid with a kernel, which has been scaled by the sampling interval in horizontal direction and by the sampled value in vertical direction. The value of the reconstruction $g_r(x)$ is determined by summing all contributions at the x that is to be interpolated. Since the width of the kernel is only two samples, each interpolated point can be calculated simply by taking

$$g_r(x) = g_s(x_i)r(d) + g_s(x_{i+1})r(d-1), \tag{4}$$

where x lies between the consecutive samples x_i and x_{i+1} at a distance $d\Delta x$ $(0 < d < 1)$ from x_i. Although Equation (4) only describes how to interpolate in one dimension, the same principle can be followed when dealing with higher dimensional data. Figure 6(b) shows how two-dimensional image data can be interpolated by first performing two one-dimensional interpolations in one direction, followed by a single interpolation in the other direction. This way of linearly interpolating points along two separable orthogonal axes is called bilinear interpolation.

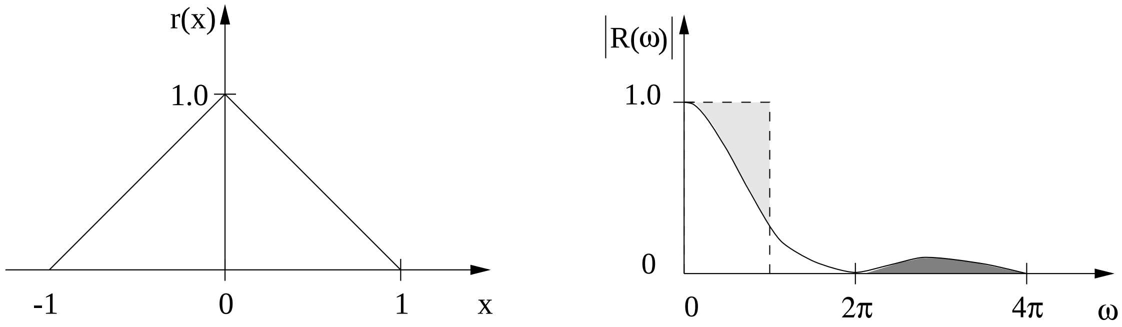

(a) Spatial domain representation.

(b) Frequency domain representation.

Figure 5: The linear interpolation kernel. (a) While the sinc function has an infinite width, the linear interpolation kernel is only two samples wide. (b) Spectra of the sinc (dashed) and linear (solid) interpolation kernels. The latter shows some loss of ideal frequencies (light area) which leads to a blurred image, as well as some spurious higher frequencies (dark area) resulting in aliasing effects.

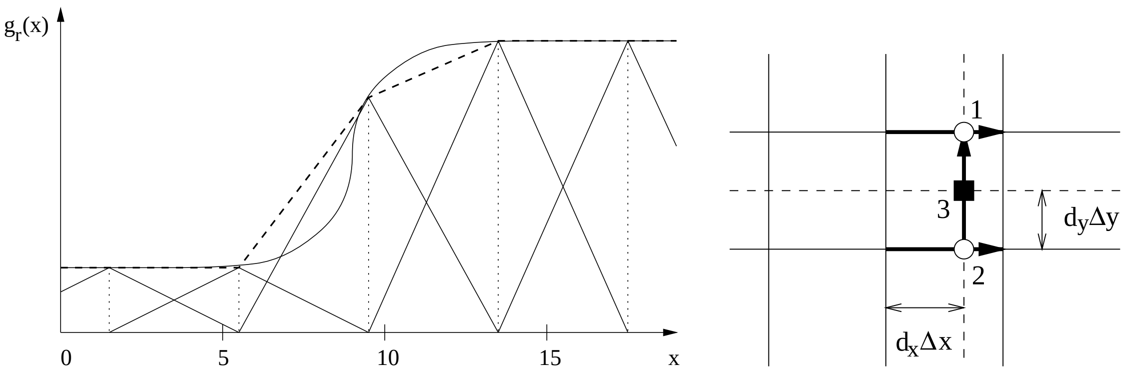

(a) Interpolation in one dimension.

(b) Bilinear interpolation.

Figure 6: Linear interpolation. (a) The continuous signal is reconstructed by overlaying each sample with a scaled version of the triangular kernel. (b) An image can be reconstructed by interpolating two intermediate points (white circles) in x-direction, followed by the desired point (black square) in y-direction.

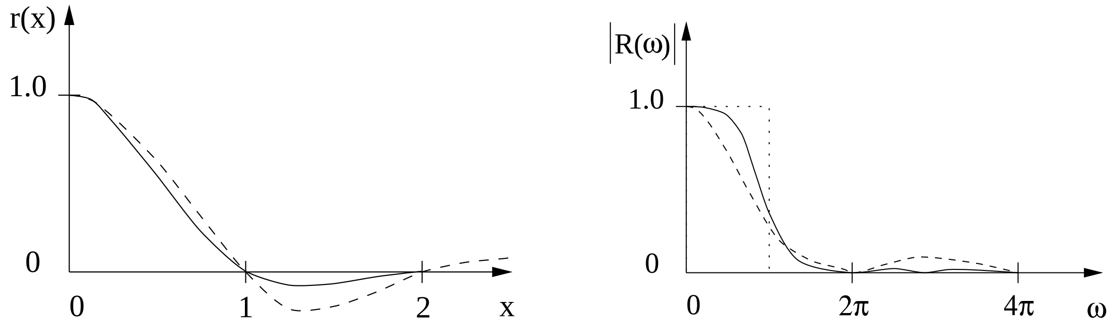

(a) Spatial domain representation.

(b) Frequency domain representation.

Figure 7: The cubic convolution interpolation kernel. (a) Right half of the symmetric sinc (dashed) and cubic convolution (solid) kernels. The sinc has an infinite width, the cubic convolution kernel is only four samples wide. (b) Spectra of the sinc (dotted), linear (dashed), and cubic convolution kernel (solid). The ideal low-pass filter is approximated better by cubic convolution than by linear interpolation.

4.3 Cubic convolution

As has been described in Section 4.1, exact reconstruction in the spatial domain can only be achieved by convolution with the sinc function. But since this kernel is infinitely wide, other kernels that are spatially limited, e.g. the linear interpolation kernel, have been suggested. Another option is truncating $\mathrm{sinc}(x)$ to a practical width, say $|x| \leq 2$. However, just like cutting off higher frequencies of the spectrum causes ringing in the image, truncating the sinc function introduces ripples in the spectrum(Park and Schowengerdt [13]). A way to suppress these ripples is to replace the truncated sinc by an approximating function having a continuous slope at $x = \pm 2$. A well-known implementation of this idea is cubic convolution (Keys [8]), which is also known as parametric cubic convolution (Park and Schowengerdt [13]), (bi)cubic interpolation (MATLAB [12]), and high-resolution cubic spline interpolation (Parker *et al.* [15]). The last name may give the reader the impression that cubic convolution is the same as cubic B-spline interpolation (Hou and Andrews [6]), but this is not the case. What they have in common is that they both use a cubic spline interpolation kernel, which by definition is a piecewise continuous function composed of third-order polynomials; the constraints imposed on these polynomials, however, are different.

The one-dimensional cubic convolution kernel $r(x)$ must be (1) symmetric, (2) zero outside the interval $[-2, 2]$, (3) continuous with a continuous first derivative, and (4) satisfy $r(0) = 1$, $r(x) = 0$ for $x = \pm 1, \pm 2$. As is shown by Keys [8], these four constraints restrict the general cubic spline kernel to

$$r(x) = \begin{cases} (\alpha+2)|x|^3 - (\alpha+3)|x|^2 + 1 & |x| < 1 \\ \alpha|x|^3 - 5\alpha|x|^2 + 8\alpha|x| - 4\alpha & 1 \leq |x| < 2 \\ 0 & \text{otherwise.} \end{cases} \tag{5}$$

The parameter α is equal to the slope of $r(x)$ at $x = 1$, and used to be set at $\alpha = -1$, which is the slope of the sinc function. However, it has been proved that in general $\alpha = -\frac{1}{2}$ is a better choice (see Park and Schowengerdt [13]). Figure 7(a) shows the cubic convolution kernel for $\alpha = -\frac{1}{2}$, while in Figure 7(b) the corresponding spectrum is displayed. From the spectra it is clear that cubic convolution approximates the ideal low-pass filter much better than linear interpolation does. However, since the kernel's spectrum is different from the ideal low-pass filter both below and above the Nyquist frequency, the interpolated image will still exhibit some blurring as well as aliasing effects.

Reconstruction of the continuous function $g(x)$ is done by convolving the sampled function $g_s(x_n)$ with the kernel $r(x)$, which is scaled by the sampling interval horizontally and by the corresponding sample vertically. Analogous to linear interpolation, the value of the intermediate point $g_r(x)$ is found by the summation of all non-zero kernels at x. Since the cubic convolution kernel has a width of four samples, the sum to be calculated is

$$g_r(x) = g_s(x_{i-1})r(d+1) + g_s(x_i)r(d) + g_s(x_{i+1})r(d-1) + g_s(x_{i+2})r(d-2), \tag{6}$$

assuming that x lies between samples i and $i+1$ at $d\Delta x$ $(0 < d < 1)$ from the i-th sample. Two-dimensional data is handled by doing four one-dimensional interpolations along one axis, followed by a single interpolation along the other axis, just like is done in case of bilinear interpolation (see Figure 6(b)).

5 Experiments

Since an industrial imaging system was unavailable, all images had to be captured using a flatbed scanner. A flatbed scanner like the HP ScanJet 4c, however, can take images of two-dimensional scenes only. The design of the required two-dimensional object is described in Section 5.1. With the scanned images, experiments were carried out to investigate the potential of the decomposition and interpolation methods, see Sections 5.2 and 5.3 respectively.

5.1 Image design

Although any object could have been used for the evaluation of the area estimation algorithms, an image of a very simple integrated circuit (IC) was chosen to simulate an industrial imaging application. At first, some experiments were carried out with an image drawn by hand using colored pencils. The results, however, were biased because the statistical distribution of the pixels near the object boundaries deviated from that of the pixels of the object's interior. This difference was caused by the desire to clear sharp edges and the resulting fear to cross them when filling the outline. Since the area estimation algorithms are judged on their performance on these border pixels, this was clearly unacceptable. Another approach that failed was to cut the elements of the IC out of colored paper and paste them on a white sheet of paper. Because the object was no longer two-dimensional, complicated shadowing effects appeared which made an accurate area estimation impossible.

The solution to all of the above problems was to construct a PostScript image of the IC and make a hardcopy using a color printer. The printer used was the Tektronix-IIIPXi, which has a resolution of 300dpi. Since the HP ScanJet 4c has an optical resolution of 600dpi, the image had to satisfy a few constraints. The first limitation was that the object had to be constructed out of rectangles only to prevent the edges from being jagged. The second restriction was that the only colors that could be used were black, cyan, magenta, and yellow, which are available from the ink-cartridge. All other colors are realised by dithering, which results in noticeably speckled object with an inpracticle statistical distribution[6]. In combination with a white background, the IC's body was assigned the color black, while the leads were given the color cyan; the result is shown in Figure 8.

Before the image was scanned, first the scanner's sweetspot was determined. The sweetspot is the part of the imaging area that is least affected by the minor inconsistencies in the scanning mechanism. It can be determined by making a low-resolution scan—about 75dpi is high enough—of a clean white surface. After equalising the histogram of the resulting image, spots that are too dark can be easily recognised (Sullivan [20]). The sweetspot of the scanner that was used is shown in Figure 9.

To create a reference segmentation, the image of the IC was placed on the scanner's sweetspot and scanned at the maximum resolution of 600dpi[7]. Careful inspection of the image showed

[6] For the primary colors red and green (but not blue) no such pattern was observed either.

[7] Turn the option Pixel Averaging off.

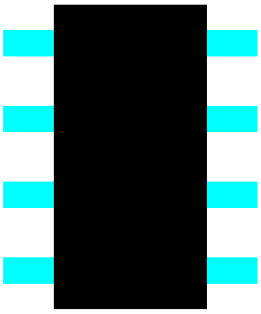

Figure 8: Design of a simple IC. The size of the black body is 450x900 pixels, while each of the cyan leads measures 150x75 pixels.

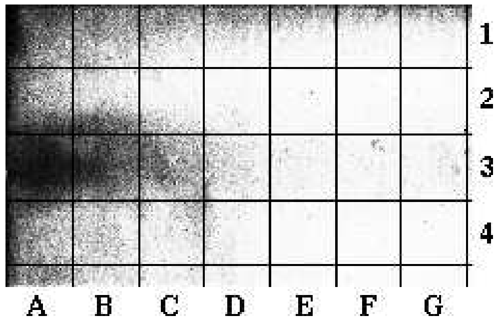

Figure 9: Image of the scanning area after histogram equalisation. The IC is placed on grid-cell F4, which is part of the scanner's sweetspot.

that the object boundaries were somewhat blurred, indicating that neighbouring pixels had some overlap. Although this is a violation of the assumption that the consecutively sampled IFOV's are just contiguous, the overall effect is assumed to be negligible for all images except those of high resolution. After calculation of the statistical distribution of the white, black, and cyan regions, it was decided to classify the image using least squares instead of maximum likelihood. The advantage of using the least squares approach, which in fact is equal to maximum likelihood classification with identity variance-covariance matrices, is that the blurred pixels are divided between the surrounding regions without favoring the region with the most variation. After classification it was discovered that, due to spectral confusion, several layers of pixels on the boundary between the background and the body were assigned to the leads class. Therefore, the image was classified once more using only the classes background and body, and the results were used to correct the original classification manually. Furthermore, noisy edges were somewhat smoothed to reconstruct the originally rectangular regions.

Although the IC had been designed to have a body of 450x900 pixels and leads of 150x75 pixels each, the reference segmentation displayed a body of 460x909 pixels and an average lead size of 152x83 pixels. This aberration is too high to be explained by the difference in resolution between the printer and the scanner alone. An explanation could be that because the colored regions were plotted and the background was not, the body and leads were oversized due to the ink running out. Another possibility is that the CCD-chip had some sort of non-linear mapping between the intensity and the output digital number (DN). If, for instance, the response curve equals the gamma function with $\gamma = \frac{1}{2}$ (see Figure 10), an intensity of 0.5 is mapped to a DN of 64, while an intensity of 1.0 is assigned a DN of 255. Therefore, blurred pixels that have intensities between these limits get DN's that are closer to 64 than to 255. For example, an intensity of 0.75 gets a DN of 144 instead of 192, and is classified to the region of intensity 0.5 instead of to the region having an intensity of 1.0. This non-linear transformation may be corrected by a similar transformation with $\gamma = 2$, but this was not done because the exact transformation was unknown and only one of the possible causes. Furthermore it was assumed that overestimation of the colored regions was sufficiently compensated for by using the reference segmentation itself, which had larger regions than was specified in the design.

5.2 Decomposition

In this experiment it was examined how accurately the area of the body, leads, and background could be estimated as the IFOV of each detector element increased. Since only one flatbed scanner was available and its detector elements could not be changed, images captured with CCD's with larger IFOV's had to be simulated. The usual way to construct these images is to apply spatial degradation of high-resolution images of the same scene. Justice *et al.* [7] describe a technique for spatial degradation that takes several other sensor characteristics (e.g. sensor blur) into account

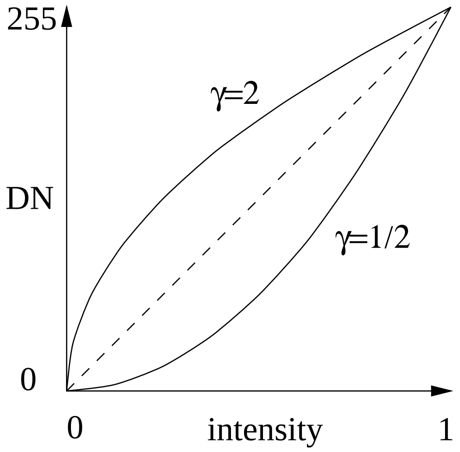

Figure 10: The gamma function. If $0 < \gamma < 1$, low intensity regions are assigned more blurred pixels than high intensity regions; in case $\gamma > 1$, high intensity regions are favored.

as well. Although this technique produces very realistic images, a model of the sensor's (optical) transfer function ((O)TF)[8], which is the Fourier transform of the point-spread function (PSF), is needed. Since the OTF of the ScanJet 4c was unknown, the low-resolution images were constructed by simple averaging.

The same four area estimation algorithms the authors had used before when dealing with satellite images (see [10, 11]) were applied. These algorithms are data- and model-driven decomposition, and data- and model-driven classification[9]. The decomposition methods divide each pixel between the neighbouring regions (data-driven) or the classes in the database (model-driven), while the classification methods, which are based on the maximum likelihood principle, assign each pixel to a single neighbouring region (data-driven) or a single class (model-driven). Both data-driven methods need the image to be segmented into pure regions prior to decomposing or classifying the remaining pixels. Because the segmentation process may introduce errors which have nothing to do with neither decomposition nor classification, this step was omitted by extracting the pure regions from the reference segmentation. All pixels that did not belong to a single object were considered to be mixed pixels, and were processed by the area estimation algorithms.

After the area estimation algorithms had processed all mixed pixels, the results were quantitatively evaluated using several error measures. To determine how accurately the method had estimated the components of a mixed pixel, for each pixel p

$$e_p = \frac{1}{2} \sum_i \left| \hat{\mathrm{f}}_{p,i} - \mathrm{f}_{p,i} \right| \tag{7}$$

was calculated. In this equation $\hat{\mathbf{f}}_\mathbf{p}$ and $\mathbf{f}_\mathbf{p}$ are the estimated and real fractions vector, respectively, and i ranges over all classes. The average error in accuracy of all pixels that were processed is denoted by $< e_p >$.

Another error measure that was calculated is based on the deviation between the real and estimated total area per class i. If these figures are denoted by $A_i = \sum_p \mathrm{f}_{p,i}$ and $\hat{A}_i = \sum_p \hat{\mathrm{f}}_{p,i}$, respectively, where p ranges over all processed pixels, the following error measure was computed:

$$e_A = \frac{1}{2} \sum_i \left| \hat{A}_i - A_i \right| . \tag{8}$$

This error measure has the area of a single pixel as a unit of measurement. But because the area of a pixel differs for all of the simulated images, e_A is also expressed in square millimeters to show

[8] The better-known modulation transfer function or MTF is the absolute value of the OTF.

[9] Data-driven is sometimes called region-based or field-based, model-driven is called pixel-based

the absolute error. A third expression is e_A as the percentage of the total area occupied by the pixels that were processed. This last expression was introduced because, although the number of pixels to be processed decreased as the IFOV increased, the area they covered became larger. As a result, it was expected that the error given in pixels would decrease, and the error expressed in mm^2 would increase; the error given as a percentage of the total area, however, was unpredictable.

The results of the experiment are presented in Table 1. For all area estimation algorithms it can indeed be observed that e_A expressed in pixels decreased and e_A expressed in mm^2 increased as the IFOV got larger. At the same time, the error given as a percentage of the total area tended to decrease as well. The reason for this is probably that, as the IFOV expanded, the proportion of blurred pixels was reduced, while more and more pixels that were relatively easy to process were incorporated. This way the resulting mixed pixel satisfied the linear mixture model much better, making decomposition in particular more reliable. The same reason can be used to explain the decreasing error in accuracy as well. However, the lower e_p was also the result of the fact that an increasing number of pixels was averaged to make a low-resolution pixel. Since according to the theory of statistics the variance of the estimation of the mean value decreases as the number of averaged samples is increased, again the linear mixture model was better satisfied, leading to a more accurate decomposition and classification. The fact that at the lowest resolution the results deteriorated can be explained by the absence of the pure pixel regions for the IC leads, which hampered the correct operation of the data-driven methods. The surprisingly good results achieved by the decomposition methods on the 300dpi image on the other hand may be a coincidence, caused by the large variation in pixel values at high resolutions. However, the fact that the resolution of the color printer used is 300dpi as well may have something to do with that.

A comparison of the four area estimation algorithms clearly demonstrates the superiority of the data-driven decomposition method, both in terms of a more accurate estimation per pixel as well as a better estimation of the total area per class. A bit less good results were achieved by model-driven decomposition, which at some distance is followed by the classification methods. Of these last two techniques, in general the data-driven approach seems to give a little bit better results than model-driven classification, though for some resolutions their order is reversed. As was to be expected, the error in accuracy for the classification methods was quite high—25%–60% vs. 1%–16% for decomposition—because pixels were assigned to a single region or class only. However, the areal error e_A was also significantly higher for all resolutions. A detailed analysis of the results revealed that this was caused by the spectral resemblance of the combination black and white with cyan, which had already been noted during the creation of the reference segmentation. The spectral confusion led to an underestimation of the background area (white) and the IC's body (black), and an overestimation of the area of the leads (cyan). In contrast to the classification methods, which are limited to a single class, model-driven decomposition is able to divide a black-and-white pixel into black and white, but in some cases the cyan alternative is just more attractive. Since no pure cyan pixels were processed that might have been decomposed into black and white, errors due to spectral confusion were made in one direction only, leading to a structural overestimation of the area belonging to the class cyan.

The conclusion of this experiment is that data-driven decomposition is able to estimate the area of an object from a low-resolution image just as accurately as model-driven classification can do from an image of higher resolution. At higher resolutions the estimation by data-driven decomposition is even more accurate, but only by a relatively small factor.

5.3 Interpolation

The following experiment had been devised to investigate the accuracy of area estimation when decreasing the spatial sampling rate. For this purpose the image of the IC used in the previous experiment was scanned at a number of different resolutions. As was described in Section 2, lower-end scanners like the HP ScanJet 4c can sample images at a sub-maximum rate without resampling as long as the chosen sampling rate is an integral factor of the maximum sampling rate. Because resampling by the scanner might bias the results of the area estimation methods, apart from the maximum sampling rate of 600dpi also the rates 300dpi, 200dpi, 150dpi, 120dpi, 100dpi,

resolution		data-driven decomposition				model-driven decomposition				data-driven classification				model-driven classification			
		$<e_p>$	e_A			$<e_p>$	e_A			$<e_p>$	e_A			$<e_p>$	e_A		
scale	dpi		pixels	mm^2	% of total		pixels	mm^2	% of total		pixels	mm^2	% of total		pixels	mm^2	% of total
2	300	0.074	48.743	0.349	3.0	0.100	74.700	0.535	4.6	0.511	746.250	5.349	46.0	0.596	452.250	3.242	27.8
3	200	0.139	79.588	1.284	5.6	0.161	79.838	1.288	5.7	0.435	328.476	5.298	23.3	0.533	332.476	5.363	23.6
4	150	0.106	35.806	1.027	2.8	0.121	72.841	2.089	5.6	0.452	412.317	11.823	31.8	0.438	194.317	5.572	15.0
5	120	0.072	41.025	1.838	4.2	0.079	42.160	1.889	4.3	0.364	187.920	8.419	19.2	0.422	198.400	8.889	20.3
6	100	0.098	29.150	1.881	3.6	0.121	45.772	2.953	5.6	0.357	121.618	7.846	14.9	0.460	163.183	10.528	20.0
8	75	0.062	13.361	1.532	1.9	0.073	25.963	2.978	3.6	0.238	53.838	6.175	7.5	0.267	71.958	8.253	10.0
10	60	0.046	5.549	0.994	1.2	0.057	12.414	2.225	2.6	0.307	70.850	12.697	14.7	0.385	121.850	21.837	25.3
12	50	0.049	9.321	2.405	2.0	0.069	20.706	5.343	4.5	0.305	39.011	10.067	8.5	0.347	71.011	18.325	15.5
15	40	0.030	2.669	1.076	0.7	0.055	14.610	5.891	3.8	0.319	19.374	7.812	5.0	0.396	12.846	5.180	3.3
20	30	0.026	1.490	1.068	0.6	0.039	7.636	5.474	2.9	0.343	38.216	27.395	14.3	0.365	50.429	36.150	18.9
24	25	0.026	1.888	1.949	0.8	0.040	7.894	8.149	3.3	0.347	34.235	35.339	14.3	0.363	41.235	42.565	17.3
25	24	0.029	2.225	2.492	1.0	0.030	4.402	4.931	2.0	0.316	47.920	53.674	21.6	0.333	45.920	51.434	20.7
30	20	0.021	1.117	1.802	0.6	0.033	4.865	7.847	4.2	0.317	24.846	40.074	13.4	0.334	34.846	56.203	18.7
40	15	0.016	0.675	1.935	0.5	0.026	2.982	8.551	2.2	0.261	7.356	21.092	5.3	0.261	12.356	35.429	9.0
50	12	0.023	0.769	3.445	0.7	0.023	2.133	9.556	2.0	0.369	14.229	63.750	13.4	0.394	14.229	63.750	13.4

Table 1: Results of four area estimation algorithms on images of different resolution. For data-driven decomposition the threshold is set to the standard value of 12, for data-driven classification it is set to an optimal value that differs with the resolution (low for images of low resolution, high for high-resolution images).

75dpi, 60dpi, 50dpi, 40dpi, 30dpi, 25dpi, 24dpi, 20dpi, 15dpi, and 12dpi were selected for image acquisition[10]. All images were interpolated using both linear and cubic convolution interpolation to reconstruct the original image sampled at 600dpi. Since interpolation can only be done between samples, first each image was expanded by duplicating the top row at the bottom and the left column at the right to make sure all interpolated images would get the same size. Next, a large number of intermediate points were calculated depending on the original sampling rate, e.g. for a sampling rate of 300dpi the number of new points to be calculated equalled $\left(\frac{600}{300}\right)^2 - 1 = 3$ times the original number of samples (90,000 for every square inch scanned). After interpolation the extra row and column were removed again. Although this procedure might seem a bit arbitrary, it is of little concern because the region of interest was not located near the image boundaries.

The area estimation methods that were applied are the simple scaling approach shown in Figure 3(a), and the interpolating techniques linear interpolation and cubic convolution. It had already been observed in the previous experiment that, due to spectral confusion between the class combination body/background and the leads class, the standard maximum likelihood approach resulted in many classification errors. Since through interpolation the number of pixels that are subject to spectral confusion was increased considerably, this method for area estimation was clearly unsuitable. Instead, the model-driven decomposition technique was used together with a post-processing step in which each pixel was assigned to the class that constituted its largest fraction. Verification of this method was done using the scanned 600dpi image and showed only minor deviations from the reference segmentation.

A quick look at the original images revealed that the scanned area was not identical for all sampling rates. A reason for these displacements might be that the coordinate system of the scanner is recalibrated before each scan. As a consequence, the reference segmentation could not be used to find out the exact composition of each pixel, making a quantitative evaluation of the results on a pixel basis impossible. However, the areal error measure e_A could still be determined by counting the number of pixels per class in both the estimated and the reference segmentations to compute $\hat{A}_i$ and A_i respectively, and calculating Equation (8) next. Since for all sampling rates the images had either been scaled or interpolated to simulate a sampling rate of 600dpi, converting e_A from pixels to square millimeters was not really necessary. However, this conversion was done anyway to make a comparison with the results of the previous experiment possible. Furthermore, because for all sampling rates the entire image was processed without filtering-out the pure pixels first, there was no need to express the areal error as a percentage of the total area.

The results in Table 2 show that the area estimates were not very accurate, especially in case the sampling rate was reduced considerably. However, an explanation for this phenomenon is not very difficult. Due to the sharp edges present in the IC image and the high frequencies they are associated with, the Nyquist frequency of the image was rather high. As the sampling rate was further decreased, the image became more and more undersampled, making reconstruction increasingly less reliable. Looking at Figure 11, it is obvious that the resulting area estimation was not very accurate.

In general, the interpolation methods were more accurate in estimating the areas body, leads, and background than the simple scaling method. Linear interpolation seemed to give a slightly better estimation than cubic convolution, although the results are not very conclusive. An explanation for this observation might be that in spite of the fact that cubic convolution can reconstruct an image more exactly in theory, linear interpolation precedes in this case because the image intensity at the object boundaries was—assumingly—subject to linear mixing. Because of the match between the imaging model and the reconstruction method, better results were achieved. Surprisingly, the scaling approach was more accurate than either interpolation method at the sampling rates 200dpi and 50dpi. A detailed analysis of these results showed that at 200dpi the area of the body and leads was severely underestimated, while at 50dpi their area was heavily overestimated. A first reason could be that at these rates the positions at which the image was sampled were unfavorable with regards to the reference segmentation. However, as is shown in Figure 12 for the 200dpi image, the number of sampling lines intersecting with an object at 200dpi is at most one

[10] A sampling rate of less than 12dpi is not supported.

resolution		scaling		linear interpolation		cubic convolution	
dpi	scale	pixels	mm^2	pixels	mm^2	pixels	mm^2
300	2	547	0.980	271	0.486	308	0.552
200	3	5464	9.792	6059	10.858	6075	10.887
150	4	755	1.353	222	0.398	216	0.387
120	5	2673	4.790	1551	2.780	1598	2.864
100	6	1176	2.108	607	1.088	675	1.210
75	8	5128	9.190	1317	2.360	1307	2.342
60	10	2873	5.149	1760	3.154	1920	3.441
50	12	7613	13.643	9392	16.832	9456	16.946
40	15	2965	5.314	1981	3.550	1869	3.349
30	20	5027	9.009	4226	7.573	4226	7.573
25	24	9667	17.324	9232	16.545	8711	15.611
24	25	20838	37.344	14602	26.168	15352	27.512
20	30	16073	28.805	14996	26.874	15844	28.394
15	40	23688	42.452	20404	36.566	20031	35.898
12	50	58073	104.073	45503	81.546	49424	88.573

Table 2: Areal error e_A for three area estimation methods applied to images sampled at different rates. The column scale gives the expansion of the image in each direction. For all sampling rates and methods e_A is expressed in 600dpi pixels.

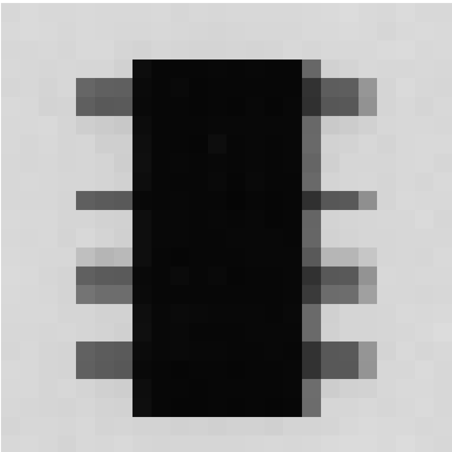

Figure 11: Green band of the IC image scanned at 12dpi. The leads were designed to have the same size, but the sampling rate was too low to capture this.

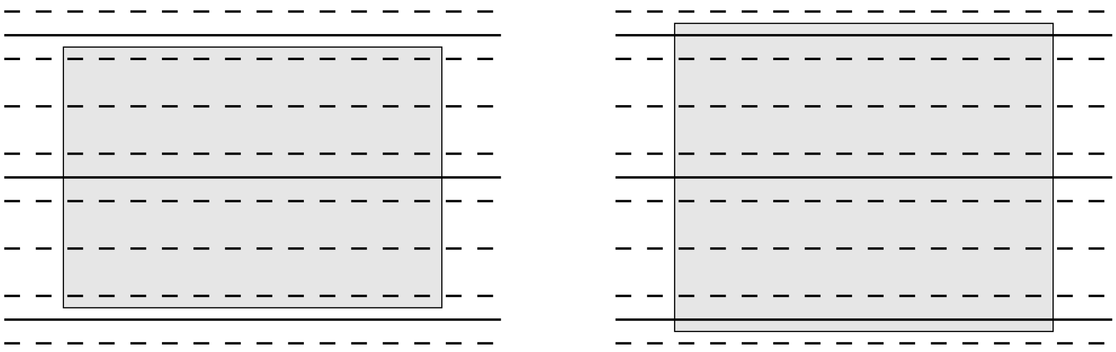

(a) Lower sampling rate leads to a smaller area. (b) Lower sampling rate leads to a larger area.

Figure 12: Sampling an image at a different rate can both decrease and increase the estimated area of an object. (a) When going from a high sampling rate (dashed lines) to a lower rate (solid lines), the number of intersecting sampling lines divided by the scaling factor of three is reduced by one (1 vs. 6/3=2). (b) Lowering the sampling rate when scanning a slightly bigger object raises the number of intersecting sampling lines by one (3 vs. 6/3=2).

line more or less than the number of lines at 600dpi divided by the reduction in sampling rate. Since after interpolation of the 200dpi image the number of sampling lines intersecting with the IC's body was 11 less than according to the reference segmentation, while the maximum deviation could only be 3, this reason alone is not enough to explain the deviation. The same holds for the 50dpi image, where the extra number of lines equaled 15, while a number of 12 was the maximum possible. A second explanation could be that, due to recalibration of the scanner's light bar before each scan, the distributions of the classes in the 200dpi image did not match the class distributions in the database. As a result, the pixels of the sampled image could not be classified correctly, leading to overestimation (underestimation) of the background if its mean vector had significantly higher (lower) values than its counterpart in the database. However, comparison of the distributions of the 200dpi and 50dpi images with the database did not reveal any large deviations. The third reason that might explain why the area of the body and leads was largely underestimated at 200dpi is that the speed of the stepper motor moving the light bar was not scaled by a constant factor. If at 200dpi the light bar was moved at more than three times the speed at which it was moved at 600dpi, the body and leads would appear to be smaller. This explanation is consistent with the fact that the deviation between the classified image and the reference segmentation was found in vertical direction mainly, while in horizontal direction the deviation was small. Since the reference segmentation in this case would never produce a correct estimation of the areal error, the scaling method could easily give better results than the interpolation methods. If the scaling method overestimated the area of an object because of its inaccuracy, while the reference segmentation erroneously indicated that the area should be even larger, the possibly correct estimation of the interpolation method would result in a higher areal error.

It can be concluded that estimation of the area of objects in an image that was sampled at a low(er) rate while keeping the IFOV constant is not very accurate. Especially if the image contains sharp edges, the risk of undersampling is big, resulting in a poor estimation. Interpolation techniques give somewhat better results than the scaling method, although they do not differ very much. The difference in performance is probably more pronounced if the image contains less high frequencies, making a more exact reconstruction possible. Increasing the number of boundary pixels will make the difference larger as well, since only at the object boundaries the area estimation methods will show a different behaviour.

6 Conclusions

To reduce the costs of an industrial imaging system, two options can be considered which both deteriorate the image characteristics. The first possibility is to use cheaper optical components, which results in an increased IFOV. The second way to save money is to use a lower sampling rate, which requires less expensive electronic circuitry. For both strategies several methods were examined to determine their capabilities and limitations with regard to an area estimation application.

To combat the negative effects of an increased IFOV, two methods based on the decomposition of mixed boundary pixels were described. The image of an IC scanned by the HP ScanJet 4c at its optical resolution was used to mimic an industrial imaging system. Next, the image was spatially degraded by several factors using a simple averaging method to simulate an increased IFOV. Experiments in which the degraded images were processed by the two decomposition methods and two traditional methods based on classification clearly showed that the data-driven decomposition technique gave the best results. The worse performance of the classification methods was caused by the fact that they can assign a mixed pixel to the class having the largest fraction at best, while the inaccuracy of the model-driven methods is due to spectral confusion. At 12dpi the accuracy of data-driven decomposition proved to be at least as good as the accuracy of the classification methods at a much higher resolution of 300dpi.

To deal with a reduced sampling rate, which results in pixels that only partially cover an object, two interpolation methods were described. These methods attempt to reconstruct the original scene before the area of an object is estimated. Using images acquired by the HP ScanJet 4c scanner at different sampling rates, the results of the interpolation methods were compared to the results of a simple scaling method. To avoid inaccurate estimations due to spectral confusion, classification of the images was executed by decomposing each pixel and assigning it to the class having the largest fraction. Because the different images could not easily be registered, the methods were evaluated based on the difference between the estimated area of each object and the corresponding area according to the reference segmentation only. Although the results were rather inaccurate because the image contained some high frequencies making it susceptible to undersampling, the interpolation method seemed to be somewhat more accurate than the scaling approach. For an image that contains edges that are less sharp but higher in number this difference will be probably more pronounced.

References

[1] H. Besser and J. Trant. Introduction to imaging: issues in constructing an image database. WWW-page http://www.ahip.getty.edu/intro_imaging/Tbl.html, 1995.

[2] J. Bone. The scanning FAQ. WWW-page http://www.infomedia.net/scan/The-Scan-FAQ.html, 1995.

[3] J. Bone. Re: reduced scanner resolution. E-mail message, May 8, 1996.

[4] R.C. Gonzalez and R.E. Woods. *Digital Image Processing.* Addison-Wesley Publishing Company, Inc., 1992.

[5] H.M. Horwitz, R.F. Nalepka, P.D. Hyde, and J.P. Morgenstern. Estimating the proportions of objects within a single resolution element of a multispectral scanner. In *Proceedings of the 7th International Symposium on Remote Sensing of Environment*, pages 1307–1320, 1971. Ann Arbor, Michigan.

[6] H.S. Hou and H.C. Andrews. Cubic splines for image interpolation and digital filtering. *IEEE Transactions on Acoustics, Speech, and Signal Processing*, ASSP-26(6):508–517, 1978.

[7] C.O. Justice, B.L. Markham, J.R.G. Townshend, and R.L. Kennard. Spatial degradation of satellite data. *International Journal of Remote Sensing*, 10(9):1539–1561, 1989.

[8] R.G. Keys. Cubic convolution interpolation for digital image processing. *IEEE Transactions on Acoustics, Speech, and Signal Processing*, ASSP-29(6):1153–1160, 1981.

[9] M.S. klein Gebbinck and T.E. Schouten. Decomposition of mixed pixels. In J. Desachy, editor, *Image and Signal Processing for Remote Sensing II*, pages 104–115. SPIE 2579, 1995. Paris.

[10] M.S. klein Gebbinck and T.E. Schouten. Accurate area estimation by data-driven decomposition of mixed pixels. Technical Report CSI-R9622, Computing Science Institute, University of Nijmegen, 1996.

[11] M.S. klein Gebbinck and T.E. Schouten. Application of data-driven decomposition to Landsat-TM images for crop area estimation. Technical Report CSI-R9626, Computing Science Institute, University of Nijmegen, 1996.

[12] The Mathworks, Inc., 24 Prime Park Way, Natick, MA 01760. *MATLAB Reference Guide*, October 1992.

[13] S.K. Park and R. Schowengerdt. Image reconstruction by parametric cubic convolution. *Computer Vision, Graphics, and Image Processing*, 23:258–272, 1983.

[14] S.K. Park, R. Schowengerdt, and M.-A. Kaczynski. Modulation-transfer-function analysis for sampled image systems. *Applied Optics*, 23(15):2572–2582, 1984.

[15] J.A. Parker, R.V. Kenyon, and D.E. Troxel. Comparison of interpolating methods for image resampling. *IEEE Transactions on Medical Imaging*, MI-2(1):31–39, 1983.

[16] R.P. Pech, A.W. Davies, R.R. Lamcroft, and R.D. Graetz. Calibration of LANDSAT data for sparsely vegetated semi-arid rangelands. *International Journal of Remote Sensing*, 7(12):1729–1750, 1986.

[17] W.K. Pratt. *Digital Image Processing.* John Wiley & Sons, Inc., 1978.

[18] F.F. Sabins. *Remote sensing.* Remote Sensing Enterprises, Inc., 2nd edition, 1987.

[19] J.J. Settle and N.A. Drake. Linear mixing and the estimation of ground cover proportions. *International Journal of Remote Sensing*, 14(6):1159–1177, 1993.

[20] M.J. Sullivan. *Make Your Scanner a Great Design and Production Tool.* North Light Books, 1507 Dana Avenue, Cincinnati, OH 45207, 1995. WWW-page http://www.hsdesign.com/scanning/toc.html.

[21] Q.I. Tian and M.N. Huhns. Algorithms for subpixel registration. *Computer Vision, Graphics, and Image Processing*, 35:220–233, 1986.