AUTOMATED INSPECTION USING GRAY-SCALE STATISTICS

Stephen T. Barnard

SRI International, Menlo Park, California

ABSTRACT

A method for using gray-scale statistics for the inspection of assemblies is described. A test image of an assembly under inspection is registered with a model image of a nondefective assembly and the two images are compared on the basis of two statistical tests: a χ^2 test of the two marginal gray-level distributions and the correlation coefficient of the joint distribution. These tests are made in local subareas that correspond to important structure, such as parts and subassemblies. The subareas are compiled in an off-line training phase. The χ^2 measure is most sensitive to missing or damaged parts, whereas the correlation coefficient is most sensitive to mispositioned parts. It is also possible to detect overall lighting changes and misregistration with these measures. Two examples are presented that show how the tests detect two types of defects.

I INTRODUCTION

Binary machine-vision techniques have received a great deal of attention for industrial inspection [1,2,3,4]. High-contrast lighting and thresholding may be used to obtain an accurate silhouette that can be processed at video rates to yield useful features, such as area, perimeter, centroid, and higher moments. In addition, structural information is available in the geometric relationships between the local features of the outline (holes, corners, and so on). This kind of information is sometimes sufficient for some industrial automation (IA) tasks, such as part identification and acquisition. Other tasks, however, are not so easily approached. Although many simple parts can be adequately represented by a mere outline, most assemblies cannot because they are typically composites of several overlapping parts or subassemblies. Binary techniques will not be effective in such cases because thresholding will not, in general, separate the important components. Thorough binary inspection of even simple parts may not be feasible if one wishes to find defects in surface finish or other types of defects not limited to the outline.

Gray-scale techniques have lately received more attention [5,6,7]. Compared to binary methods, there is a great variety of ways to use gray-scale information. This paper describes an approach for exploiting gray-scale information for inspection in a very basic form. Statistical tests of marginal and joint intensity distributions are used to compare test assemblies with an ideal model assembly. These tests are very efficiently computed and they are sensitive to defects and discrepancies that are not easily handled in the binary domain.

II REPRESENTATION

We use a representation that directly specifies the expected appearance of the assembly. Important structures of the assembly are represented as subareas that are tagged for special consideration. Using this representation, an image of a test assembly is more-or-less directly compared to a model with a minimum of preprocessing. It is necessary to do some sort of photometric and geometric normalization to the test image to bring it into correspondence with the model.

At the lowest level, an assembly is represented by a gray-scale image. At the highest level, an assembly is represented by a list of named subareas, each of which corresponds to a particularly meaningful segment. Attached to each of these subareas are properties that specify the important characteristics; these characteristics identify the corresponding segment as "good" or "defective." Ideally, these subareas could have arbitrary shapes and sizes, but, for now, think of them as rectangular windows. Each has a specific location and size in the normal reference frame. The inspection system begins with a representation of an ideal assembly, called the model. This may be constructed interactively in a training phase. Given a low-level representation of a test case (i.e., an image), the system proceeds to build a high-level representation by comparing segments of the test image to segments of the model.

The first step is to bring the test image into geometric registration with the model image. This is not strictly necessary. We could directly compare regions in the normal reference frame

(i.e., in the model image) with regions in the translated and shifted reference frame. Nevertheless, geometric registration simplifies further processing by establishing a one-to-one correspondence between model pixels and test pixels.

We have assumed that the positional variation of the test assemblies is restricted to translation and rotation. We must therefore determine three parameters--Δx, Δy, and θ. There are several ways to approach this problem. Binary techniques may be adequate to determine the normalization parameters to subpixel accuracy with a system such as the SRI vision module [3]. In the gray-scale domain, one may search for maximal cross-correlation, although this will probably be very time-consuming without special hardware. A potentially more efficient method is to find distinguishing local features in the test image and then match them to their counterparts in the model [8]. Once the translation and rotation has been determined it is easy to register the test image using a linear interpolation [9].

III STATISTICAL COMPARISON

Two statistical measures that are useful for comparing model and test subareas are the χ^2 test and the correlation coefficient. They are both extremely efficient and simple to implement, and they are sufficiently different to distinguish two broad classes of defects.

A. The χ^2 Test

The χ^2 test measures the difference between two frequency distributions. Let $h_m(k)$ be the frequency distribution of gray-scale intensities in a model window. Let $h_t(k)$ be the frequency distribution of a test window. We can consider h_m to be a hypothetical ideal distribution. The χ^2 test gives a measure of how far h_t deviates from the hypothetical distribution h_m. The significance of the test depends on the number of samples.

$$\chi^2 = \sum_k \frac{((h_m(k)-h_t(k))^2}{h_t(k)} .$$

This yields a measure of difference, but, to be consistent with what follows, we want a measure of similarity. Let

$$r = e^{\chi^2/c}$$

where c is some positive constant. r is a measure of the similarity of two distributions (in the χ^2 sense). If the distributions are identical, then r will be unity; if they are very different, r will be close to zero.

B. The Correlation Coefficient

Let h_{mt} be the joint frequency distribution of the model and test windows. That is, $h_{mt}(u,v)$ is the frequency with which a model pixel has gray value u and its corresponding test pixel has gray value v.

Let m_1 be the mean of h_m and m_2 be the mean of h_t.

Let σ_1 be the standard deviation of h_m, and σ_2 be the standard deviation of h_t.

The central moments of the joint distribution h_{mt} are

$$\mu(i,j) = \frac{1}{n} \sum_k ((X_m(k)-m_1)^i * (X_t(k)-m_2)^j)$$

where $X_m(k)$ and $X_t(k)$ are gray values of the kth pixel in the model image and and the test image, respectively. The correlation coefficient, ρ, is

$$\rho \quad \frac{\mu(1,1)}{\sigma_1 \sigma_2}$$

ρ is in the interval $[-1,1]$. If it has the value +1 or -1 the total "mass" in the joint frequency distribution must lie on a straight line. This will be the case when the test image and the model image are identical, and ρ will be +1. In general, if there is a linear functional dependence between the test and model windows, ρ will be +1 (or, in the extremely unlikely case that one window is a "negative" of the other, ρ will be -1). If the windows are independent distributions, however, ρ will be 0. We can reasonably expect that intermediate values will measure the degree of dependence between the two windows.

C. r vs. ρ

r is not sensitive to the location of pixels. It simply measures the degree of similarity between two marginal distributions. ρ, on the other hand, measures the degree to which the model pixels agree with their corresponding test pixels; therefore, it is sensitive to location. This implies that r is a good test for missing and severely damaged parts (for they are likely to change the distribution of the test pixels compared to the distribution of the model pixels), while ρ is a good test for the proper location of parts (for misalignment will change the joint distribution).

A systematic change in lighting can also be detected. r would be small because the lighting change would change the intensity distributions in the test windows, but ρ would be large because the test and model distributions would still be well-correlated. If this observation were made for only one window, it would not be very meaningful. However, if we make the reasonable assumption that most of the windows locate regions that are not defective, this leads to the observation that a systematic pattern of small r and large ρ indicates a lighting change.

Small misregistration errors are also detectable. Small misregistration would produce

large r because the marginal distributions of the test windows would not be much different from the model windows. On the other hand, ρ would be smaller than if there were poor registration because the windows would not correlate as well. The same result for a single window would be caused by a misplaced part, but, again using the assumption that most of the windows locate non-defective regions, a systematic pattern of large r and small ρ over many windows would indicate an overall misregistration.

These relationships are summarized in Table 1.

Table 1

Defect Pattern vs. r and ρ

	OK	Missing	Misplaced	Lighting	Registration
r	LARGE	SMALL	LARGE	SMALL	LARGE
				(SYSTEMATIC)	
ρ	LARGE	SMALL	SMALL	LARGE	SMALL

D. A Two-Stage System

The gray-scale statistics discussed above provide a basis for an integrated minimal system for inspection that is composed of two modules--a training module that is used off-line and that allows a human operator to build a high-level description of the assembly, and an inspection module that matches images of test assemblies to the model built in the training phase. In the training phase the operator works with an image of an ideal assembly, identifying the important parts that are likely to be defective and describing the modes of defects that may be relevant. For example, the location of a particular part may not be precisely fixed, but rather permitted to range over a rather large area. In this case the operator might indicate that r (the location insensitive measure) is a relevant test, but not ρ. In another case there may exist areas that have extremely variable appearance, perhaps because of individual part identification markings, and these areas might be excluded from testing altogether. In the normal case, however, a part will be fixed in one location and orientation, perhaps with some tolerance, and the operator will merely specify allowable limits for r and ρ.

The on-line inspection phase is entirely automatic. The knowledge about the assembly collected in the training phase is applied to specific test assemblies and a judgment is made as to how well what is known fits what is seen. The r and ρ measures are computed for each area and are used to derive probability estimates of the various types of defects.

IV EXPERIMENTS

We have tried the r and ρ tests on two assemblies.

Figure 1 and Table 2 show the results for a water pump. The upper left portion of Figure 1 is the "model" image of a nondefective pump in a normal position and orientation. Several windows representing important parts of the assembly are also shown. The upper right portion of Figure 1 is a "test" image of a defective assembly. The round, dark pulley in the center is missing. In the lower left part of Figure 1 the test image has been registered with the model. The lower right image is the difference between the model image and the registered test image, and has been included to indicate how close the registration is. Table 2 shows the r and ρ values for the various windows. Note that r and ρ are both very small for the (missing) pulley compared to the other (nondefective) parts, just as predicted. r is also small for the window representing the total assembly because this includes the defective pulley.

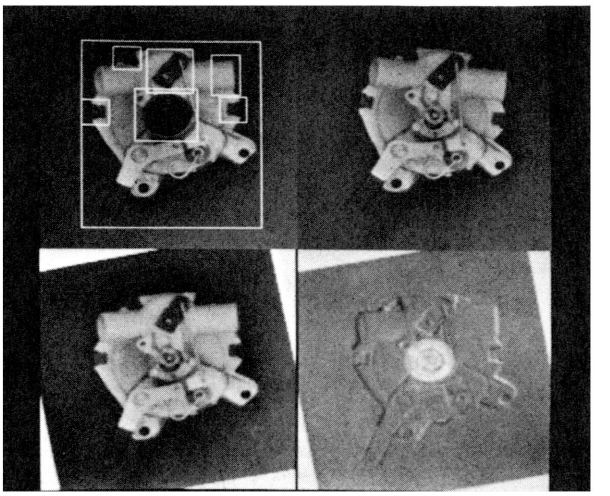

Figure 1. Pump

Table 2

Pump Statistics

	r (c=800)	ρ	
Total	.492	.801	
Pulley	.236	.354	<= Defect
Link	.981	.824	
Spout	.919	.904	
Clip1	.862	.879	
Clip2	.978	.780	
Clip3	.949	.898	

Figure 2 and Table 3 show the results for a hardcopy computer terminal. In this case the write head is not positioned in the correct place. Note that the window for the "head" includes the entire area where it might be located. As predicted, the r value is high, while the ρ value is small. In practice this might not be considered a defect because it may be permissible for the head to be located anywhere along the track. If this were the case, the window could be tagged to indicate that ρ is not a relevant test.

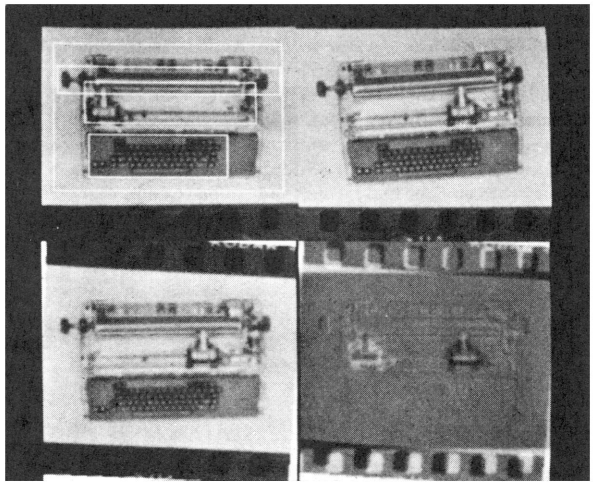

Figure 2. Terminal

Table 3

Terminal Statistics

	* r (c=800) *	ρ *	
Total	.746	.910	
Platen	.674	.868	
Head	.890	.458	<= Defect
Keys	.740	.923	

REFERENCES

1. M. Ejiri, T. Uno, M. Mese, and S. Ikeda, "A Process for Detecting Defects in Complicated Patterns," Computer Graphics and Image Processing, Vol. 2, pp. 326-339 (1973).

2. W .C. Lin and C. Chan, "Feasibility Study of Automatic Assembly and Inspection of Light Bulb Filaments," Proc. IEEE, Vol.63, pp. 1437-1445 (October 1975).

3. G. J. Gleason and G. J. Agin, "A Modular Vision System for Sensor Controlled Manipulation and Inspection," A.I. Center Technical Note 178, SRI International, Menlo Park, California (March 1979).

4. G. J. Gleason, A. E. Brain, and D. McGhie, "An Optical Method for Measuring Fastener Edge Distances," A.I. Center Technical Note 190, SRI International, Menlo Park, California (July 1979).

5. M. L. Baird, "SIGHT-I: A Computer Vision System for Automated IC Chip Manufacture," IEEE Trans. on Syst., Man, Cybern., Vol. SMC-8, No. 2 (February 1978).

6. W. Perkins, "A Model Based Vision System for Industrial Parts," IEEE Trans. Comput., Vol. C-27, pp. 126-143 (1978).

7. R. J. Woodham, "Reflectance Map Techniques for Analyzing Surface Defects in Metal Castings," Ph.D. Dissertation, Massachusetts Institute of Technology, Cambridge, Massachusetts (September 1977).

8. S. T. Barnard, "The Image Correspondence Problem," Ph.D. Dissertation, University of Minnesota, Minneapolis, Minnesota (December 1979).

9. A. Rosenfeld and A. Kak, Digital Picture Processing (New York: Academic Press, 1976).