NBSIR 79-1764

Image Pattern Recognition in Industrial Inspection

Gordon J. Vanderbrug
Roger N. Nagel

Automation Technology
Mechanical Processes Division
Center for Mechanical Engineering
and Process Technology
National Engineering Laboratory
National Bureau of Standards
Washington, DC 20234

Issued September, 1979

U.S. DEPARTMENT OF COMMERCE
NATIONAL BUREAU OF STANDARDS

NBSIR 79-1764

IMAGE PATTERN RECOGNITION IN INDUSTRIAL INSPECTION

Gordon J. Vanderbrug
Roger N. Nagel

Automation Technology
Mechanical Processes Division
Center for Mechanical Engineering
and Process Technology
National Engineering Laboratory
National Bureau of Standards
Washington, DC 20234

Issued September, 1979

U.S. DEPARTMENT OF COMMERCE, Juanita M. Kreps, ***Secretary***
Luther H. Hodges, Jr., ***Under Secretary***
Jordan J. Baruch, ***Assistant Secretary for Science and Technology***
NATIONAL BUREAU OF STANDARDS, Ernest Ambler, ***Director***

TABLE OF CONTENTS

1. INTRODUCTION

Image analysis and pattern recognition can play a large role in the industrial arena. In this report we present a state of the art survey illustrating the application of these technologies to the problem of visual inspection. The important aspects of these technologies for visual inspection are:

> Digitization and Quantization - the image must be converted to a form where it can be processed digitally.
>
> Compression - techniques such as contour coding which reduce the amount of data.
>
> Enhancement - techniques which manipulate gray levels or modify histograms so as to transform the image into a form which is more suitable for processing.
>
> Segmentation and Feature Extraction - techniques such as thresholding, edge detection, and template matching which are used to segment the image and produce features.
>
> Classification-techniques such as discriminate functions, decision trees, and inspection procedures which are employed to make decisions.

References 1 and 2 are good introductions to image analysis and pattern recognition. Reference 3 contains a discussion of the visual inspection problem, while reference 4 discusses the use of sensors in industrial automation applications.

The use of visual inspection in industry is relatively new, thus the applications we have surveyed include production systems (installed on the factory floor), prototype systems (designed for use in production), and feasibility studies (performed as a research effort). In feasibility studies the needs for computational simplicity and cost have been de-emphasized and further development would be needed to produce a production system.

The motivating factors in applying visual inspection systems vary. In some cases the task is simple, but the speed required (several decisions per second) precludes the use of human inspectors. In other cases the job is boring and human inspectors do not perform well (and don't want the job). In addition to the human factors mentioned above, there are a host of economic and environmental motivations for the use of automated visual inspection. A further and increasingly important need for visual inspection, arises in the need to provide sensory feedback for robot manipulators. In the case of a robot, a human is by definition not available, yet often there is a need to perform visual tasks. The lack of a vision system on commercially available robots has both limited the range of robot applications, and increased the cost of many robot installations.

Section Two of this report surveys a variety of stand alone visual inspection systems developed by industry. The third section surveys the work on robot "eyes" developed for inspection problems. In the fourth section, a list of other applications for automated inspection is presented. Finally, in the last section we present our conclusions, and an assessment of the future role for visual inspection in industry.

2. Visual Inspection Systems

The explicit inspection of parts and assemblies in the manufacturing process has been estimated to represent 10% of the total labor cost. This 10% figure is second only to the cost of assembly, which is estimated at 22% of the labor cost [4]. Thus there is a high potential payoff in cost when an automated visual inspection system can be installed.

In this section we present a brief overview of several visual inspection systems. In general these systems match a set of predefined features found by image analysis and pattern recognition techniques with a stored range of values for these features. The end result is to label the part as acceptable and allow it to pass, or to reject the part as defective. The vision system is usually connected to some part handling equipment, such as a transfer line, which removes rejected parts from production.

For each of the applications, the problem will be stated by describing the manual process, or by indicating the difference between the acceptable and the defective parts. We will indicate the hardware configuration used in the system, and describe the techniques used in the inspection procedure. If possible, we will provide information on performance.

2.1 General Motors Chip Alignment

Figure 1 is an assembly consisting of a large heat sink, a weld cup on the corner of the heat sink, and a Darlington IC chip which is bonded to the surface of the heat sink [5]. The assembly is a component in the High Energy Ignition system in General Motors automobiles. The manual inspection process of these assemblies is illustrated in Figure 2 and consists of the following steps. The assemblies are automatically fed into a rotating table where an operator 1) visually (through a microscope) determines the position and orientation of the chip, 2) inspects for any gross defects in the chip, 3) manually manipulates a set of electrical test probes over the base area of the chip using joy sticks, and 4) initiates the lowering of the probes onto the chip contacts and the automatic electrical test process. GM Research Laboratories developed the vision programs and GM Manufacturing Development produced a production line system to automate this process.

The processing is illustrated in Figure 3. It begins by histogramming the orientations of an edge detector. The maximum value is taken as the approximate orientation of the chip on the heat sink. Chips that have too much tilt for the automatic test equipment are identified. At the same time that the edge detectors are applied, corner detectors are used to locate possible corners of the chip. A square template, equal to the size of the chip, is used to search the space of possible corners, to determine the actual position and orientation of the chip. At this point, chips that are off the edge of the heat sink or too close to the weld cup are identified.

Simple contrast detection schemes are used to verify that the chip is not broken, undersized, or fractured (see Figure 4). Thus the system determines the position and orientation of the chip, rejects chips that are structurally defective, and directs the alignment of the test probes over the chip contacts. A schematic of the system is shown in Figure 5.

The prototype system used a 50 X 50 pixel camera, with the field of view twice the width of the chip. Sixteen gray levels were used. The system was equipped with an infrared illumination source. A production line system shown in Figure 6, was installed in a Delco Electronics plant in January of 1977.

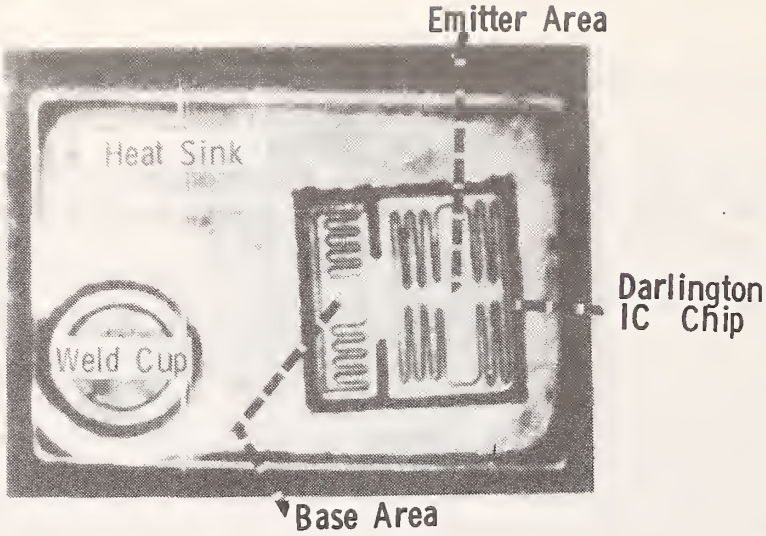

Figure 1

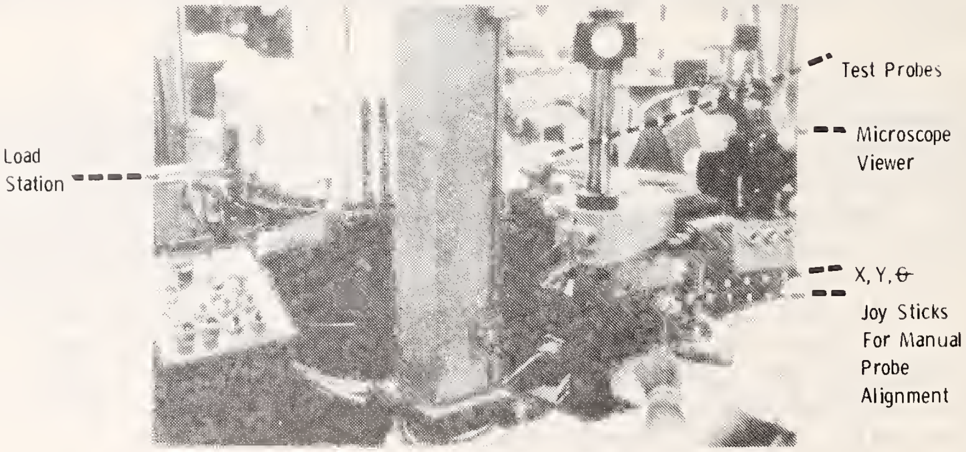

Figure 2
Darlington Electrical Test Probe Station

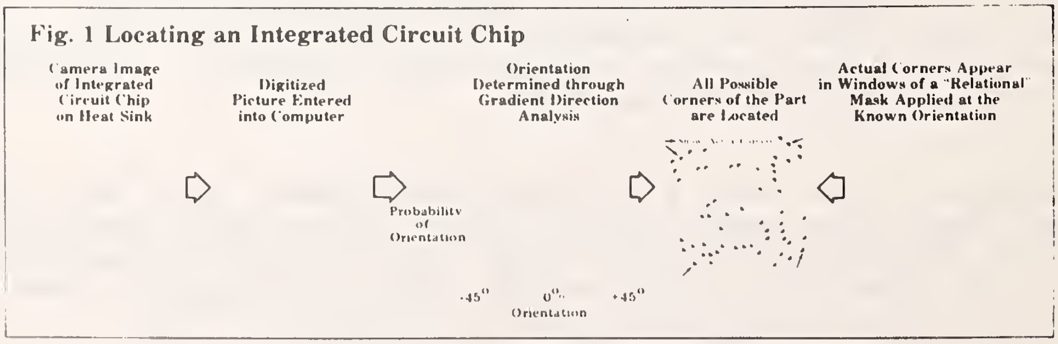

Figure 3

Partial Chip

Inspection for chip integrity rejected part

Figure 4

Figure 6

SIGHT-I installation in use at Delco Electronics' Kokomo, Indiana plant. The operator is loading a supply tray for the machine.

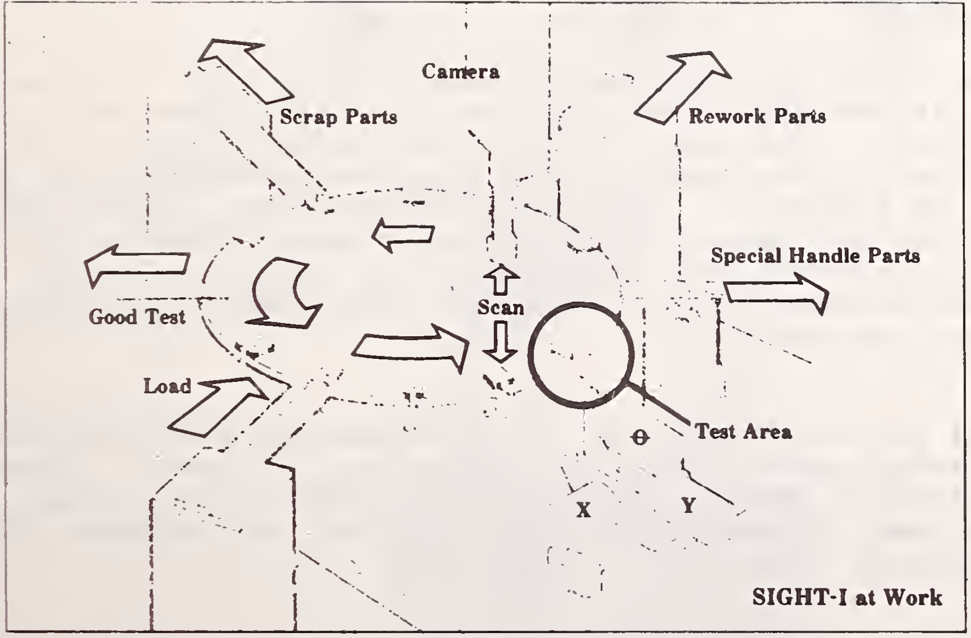

Figure 5

2.2 Western Electric Series 700 Connector

Figure 7 shows the Western Electric Series 700 Connector and the internal contact [6]. They are used for making quick, reliable, water resistant splices of two or three insulated conductors. The connector consists of a slotted "U" shaped contact, preassembled into a two piece sealant-filled plastic housing. A splice is made by inserting the wire into the housing and compressing it to make the contact cut through the wire insulation, and thus to make the connection.

The craftsman normally does not make a visual or electrical test of a splice. Hundreds or thousands of individual connector splices make up a single cable splice, which is typically buried after completion. Currently, 30% of the parts are inspected manually, a percentage based on the cost of inspection and product quality. One hundred million of these parts are manufactured each year. Figure 8 shows various categories of defective parts, such as a missing or misoriented contact, and insufficient sealant.

The inspection process is based on identifying features which characterize good connectors. The connector is inspected by viewing it from the side using transmitted light. The contact appears as a dark upside down "U", and the presence of the sealant can be determined by the amount of transmitted light in the appropriate areas.

Nine features are used to characterize a connection. Average intensities in the areas "B", "C", and "F" (see Figure 9) are used for the sealant. The area N is used for normalization. The contact position is determined by the five measurements shown in the lower half of Figure 9. These measurements are made using a median filtering operation on three adjacent lines of columns. Inspection is carried out by an algorithm involving these nine features and thresholds that have been determined to separate the good and the bad parts.

A prototype system (Figure 10) was developed at Bell Laboratories using a solid state, 100 X 100 resolution TV camera. Since the amount of data to measure the sealant parameters is minimal, this resolution was determined by the accuracy requirements for the contact position.

Presently, Western Electric is integrating the inspection and parts handling components. The inspection component is being implemented on an LSI-11, and will operate two GE - TN 2200 cameras. The performance goal is one part per sec. for each camera.

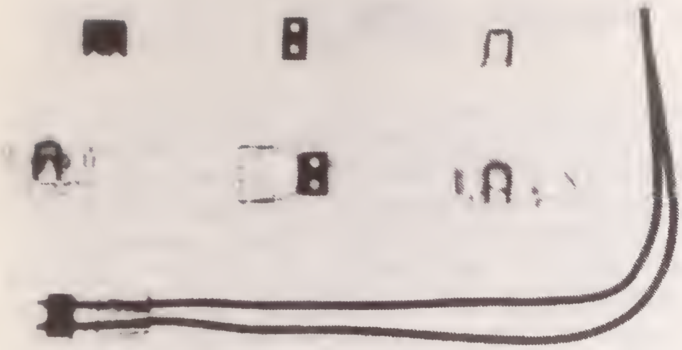

Figure 7
Western Electric
Series 700 Connector

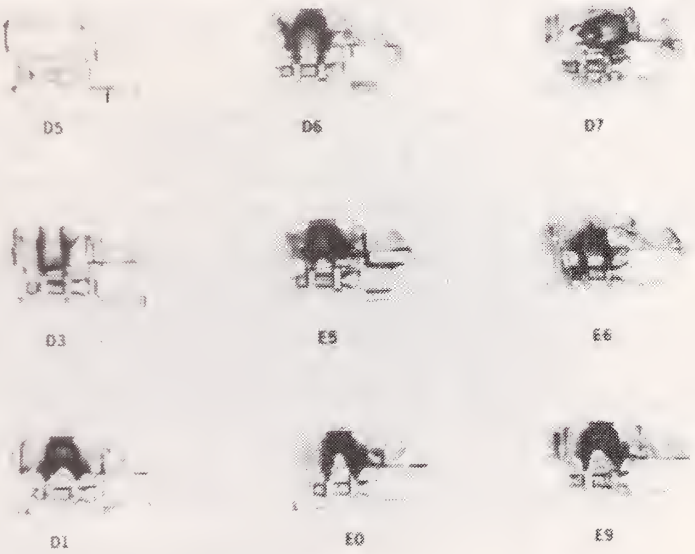

Figure 8
Defective Parts

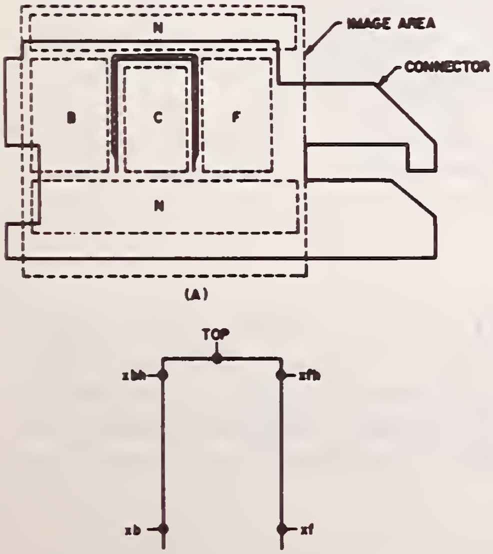

Figure 9

Figure 10

2.3 General Electric Syntactic Analysis System

General Electric has developed a system for inspecting industrial parts which uses syntactic pattern recognition [7]. The system uses a finite state grammar to analyze the median curves of parts such as screw threads or springs. The system, shown in Figure 11, is notable for its degree of generality and the amount of processing that it does in hardware. It consists of a rotating glass disk which back lights the object as it passes in front of a RETICON solid state linear imager. The rotation of the disk is synchronized with the scan clock of the imager. The output is thresholded to produce a binary signal corresponding to the part silhouette, such as that shown in Figure 12.

The outline of the part silhouette is obtained in special purpose hardware by locating edge transitions. The presence of dust and dirt on the glass and gaps within the silhouette are handled by filtering, and by requiring adjacent edge transitions on successive lines. The edge silhouette determines the median curve.

Straight line approximations are done sequentially by an algorithm which looks for the next data points inside an envelope. The envelope is allowed to change slightly, based on the length of the line and the average slope. The algorithm has been implemented with a standard 181 arithmetic logic unit chip, and is capable of line fitting to data points at 2 X 20,000 points/sec.

The next step is to perform the grammatical analysis using syntactic pattern recognition. The terminal symbols of the grammar are numerals chosen to represent straight line segments of various lengths and orientations, as shown in Figure 13. For example, the symbol 9 represents a short segment oriented in a northeast direction. The grammar determines a finite state machine which is used to parse the median axis, and thus, to accept or reject the part.

An example sequence of non-terminal symbols (i.e. numerals) that represent one down slope and one up slope is shown in the left part of Figure 14. The state transition network for this example is in the right part of the figure. The down slope contains a small noise burst which is indicated in the illustration and in the network.

The grammar is extensive - it contains 1360 productions. It is finite state, but has been implemented with a stack for the last three symbols, so that limited backtracking can be performed. The prototype system for the syntactic analysis has been implemented on an Interdata 7/16 minicomputer, and can process 15 parts/sec. Currently the hardware preprocessor is capable of processing 5 parts/sec.

The finite state grammar approach is sufficiently flexible to be a candidate approach for inspecting any manufactured part where the median curve is a good representation. One difficulty in the syntactic approach is in determining the grammar. There has been some work in grammatical inference, and this work may some day be useful for syntactic pattern recognition applications to industrial inspection.

Figure 11

Figure 12
(Top Left)

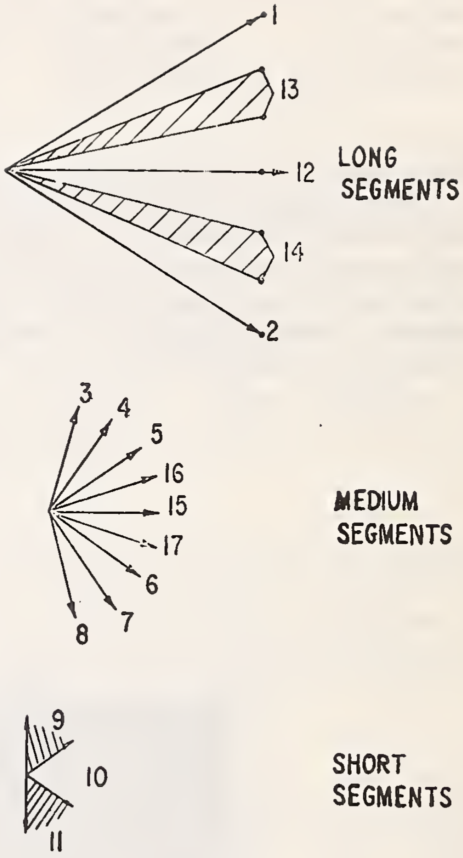

Figure 13

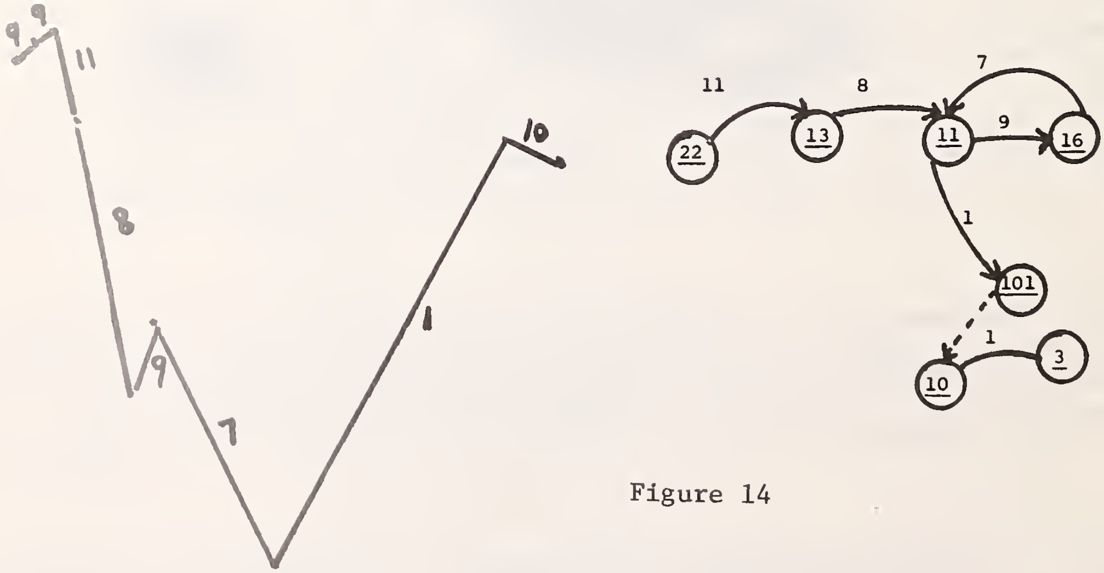

Figure 14

2.4 Lockheed Shell Inspection

Lockheed Palo Alto Research Laboratory performed a study which demonstrated the feasibility of digital image analysis techniques for the automatic inspection of 105mm projectiles using x-rays [8]. The defects which appear on the x-rays are cracks, cavities, porosity, and base separation. The steps taken in the feasibility study are as follows.

The original digital image of the projectile is shown in the upper left of Figure 15. The image has a 500 micron resolutiion and nine bit intensity quantization. The upper right image of Figure 15 shows markings of the shell boundary used to adjust the shell orientation. The lower left image of Figure 15 outlines the region which will be searched for defects. In the final image of this figure, this region has been transformed by a "flat field" correction algorithm. This algorithm fits a second order polynomial to the raster-scanned intensity values, subtracts the fit values from the actual intensity values, and scales the result. Some hints of possible defects begin to be visible in this image.

Individual defective picture points are shown in the left of Figure 16 as bright points. They are determined on a point by point basis. This determination is made by comparing certain statistics on the intensities inside a small neighborhood to the statistics on the intensities inside a larger neighborhood.

In the final image (Figure 16 RIGHT) the individual defective picture points have been associated into distinct defects. This process is a simple region growing algorithm with a threshold (in this case 25 picture points) for the minimum allowable defect size. Features such as size, length, width, and orientation are computed and used to classify the defects. The details of the classification have not been reported.

The 500 micron resolution was sufficient to locate the above defects, but was not sufficient to check for base separation of the projectile. For this reason, the base region was rescanned at a higher resolution for that part of the inspection process.

The feasibility study was considered a success and a prototype system is under development.

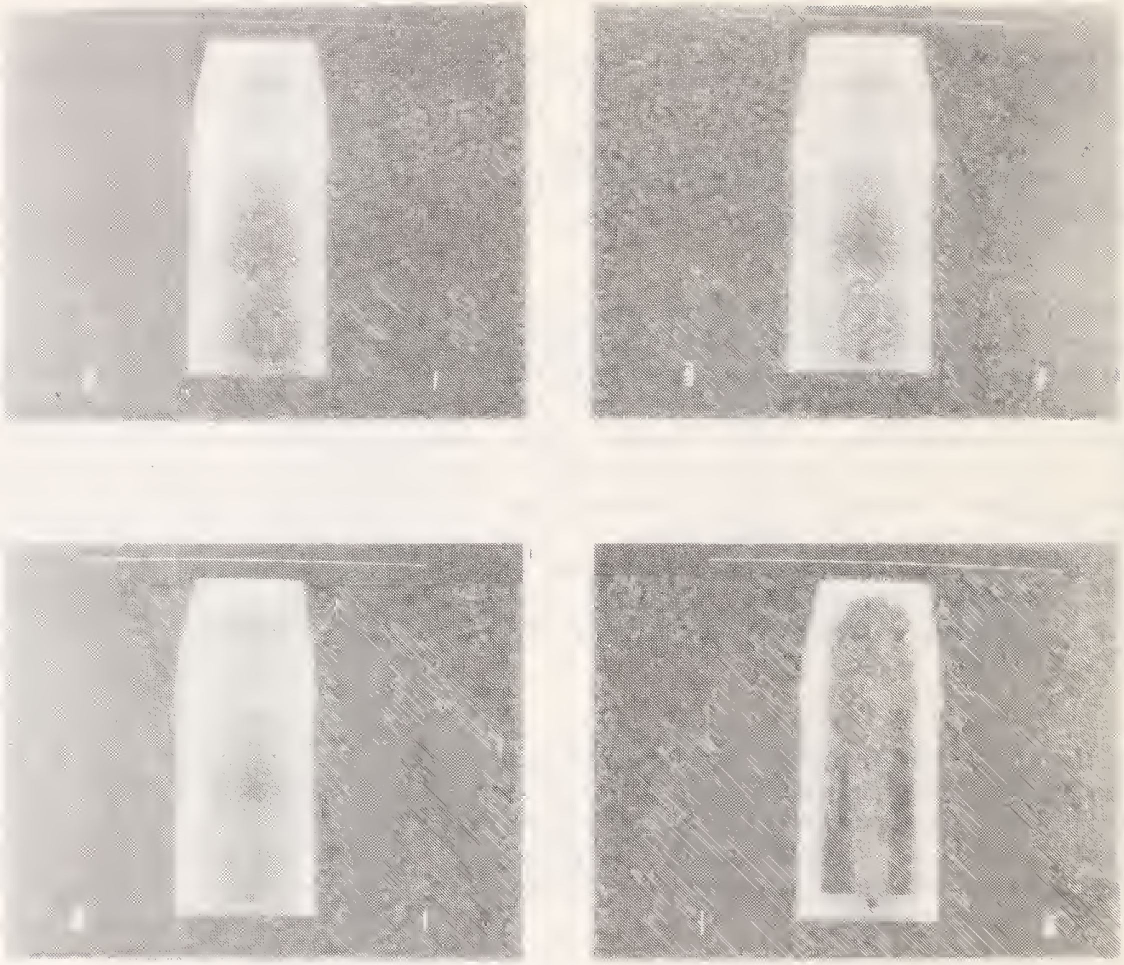

Figure 15 (Above)

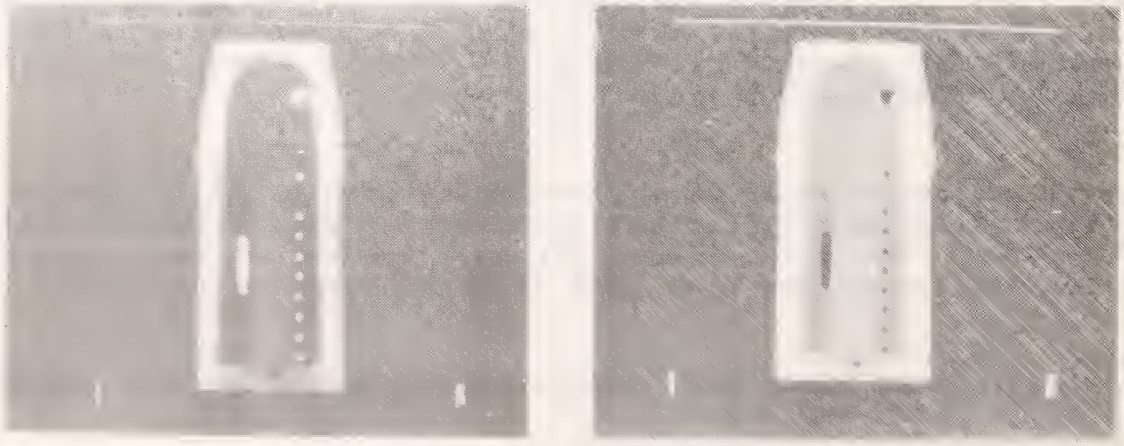

Figure 16

2.5 Recognition Systems, Inc. - Optical Power Spectrum Analysis

Recognition Systems, Inc. has adopted an approach toward automatic inspection based on optical power spectrum analysis [9]. An optical computer takes the Fourier Transform of a part. The important features of the resulting power spectrum are sensed by a solid state photodetector, and fed into a digital computer for analysis. A block diagram is shown in Figure 17.

One of the applications that Recognition Systems has demonstrated is the inspection of hypodermic needles (Figure 18). A hypodermic needle diffracts a laser beam through a lens to produce the power spectrum. The needle is placed in the front focal plane of the lens, and the detector is placed in the back focal plane, where the power spectrum is formed. The detector consists of rings and wedges to measure energy in various regions of the power spectrum. A solid state electronics package consisting of an amplifier, a multiplexor, and an A/D convertor delivers the signal to a computer.

This approach is based on certain properties of the Fourier Transform that make it possible to predict the structure of the transformations of various objects. In particular, circles are transformed into circles, and lines are transformed into lines that pass through the origin at 90 degrees to the object line. In a smooth picture the falloff in moving out from the origin is very gradual, while in a busy picture it is much more rapid. Directional biases in the picture show up as directional differences in the rate of the falloff. This is why the detector consists of a semicircle of annular rings and a semicircle of wedges.

Figure 19 shows both a good and a defective needle point, and their respective power spectra. The defective needle contains a horizontal edge that is not present in the case of the good needle. This edge results in a vertical component of diffracted energy in the detector. By examining the pattern of the wedge intensities, the presence of the horizontal edge, and hence the defect, is detected.

The computer executes the inspection algorithm and controls the material handling subsystem. The overall throughput of the system is 10 needles per second.

It is reported that the savings over manual inspection cost is sufficient to pay for the machine in less than one year. The needle point system was installed in a factory approximately two years ago, and is the first industrial inspection application by Recognition Systems. Presently they are working on applications to photomask inspection, paper printability determination, and fabric defect inspection. The fabric inspection (Figure 20) will produce a quality grading of the product.

The power spectrum of a high quality fabric exhibits a regularity that is not present in the power spectrum of the low quality fabric. The individual lobes of the pattern for the low quality fabric are more deformed. Note the effect of a double-thread in the woven cloth on the non-control lobes in the power spectrum of the weave pattern in Figure 20.

Figure 21 is a prototype drawing for the inspection system. The plans are for 100% inspection of webs of fabric up to 48 inches wide at the rate of 100 yards per minute. Based on individual decisions across the width of the web, point counts for various defects will be assigned, and an overall grade for each piece of cloth passing through the system will be determined.

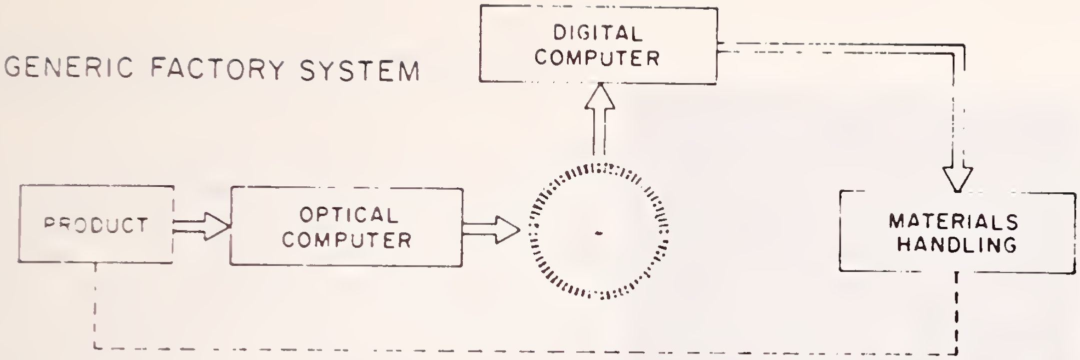

Figure 17 (Above)

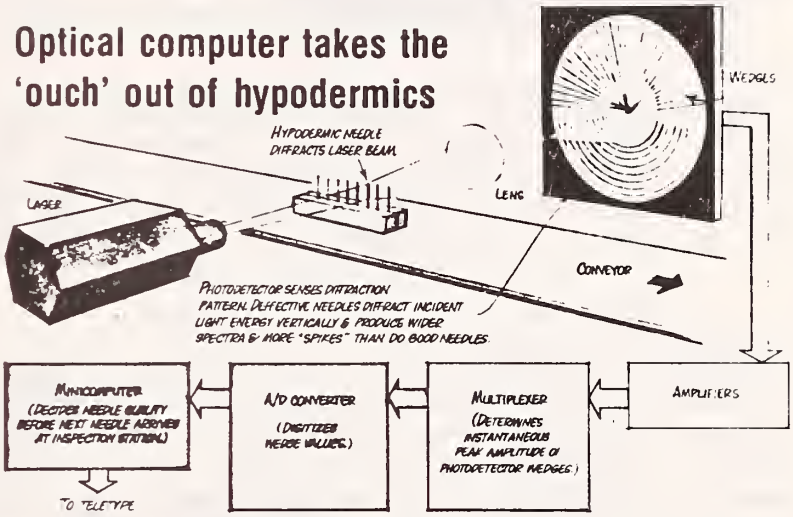

Optical computer takes the 'ouch' out of hypodermics

Figure 18

[illegible] Good Needle Point [illegible]

[illegible] Diffraction Pattern of Good Needle Point

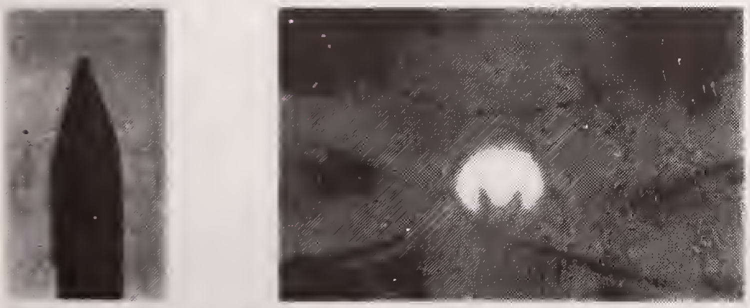

[illegible]

[illegible] Diffraction Pattern of Defective Needle Point

Figure 19

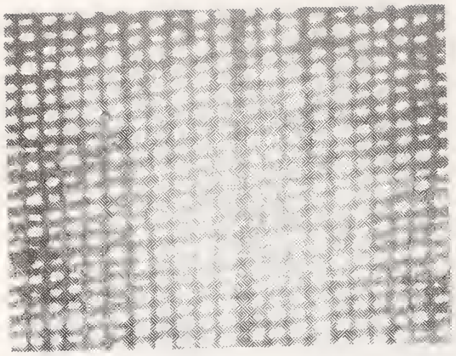

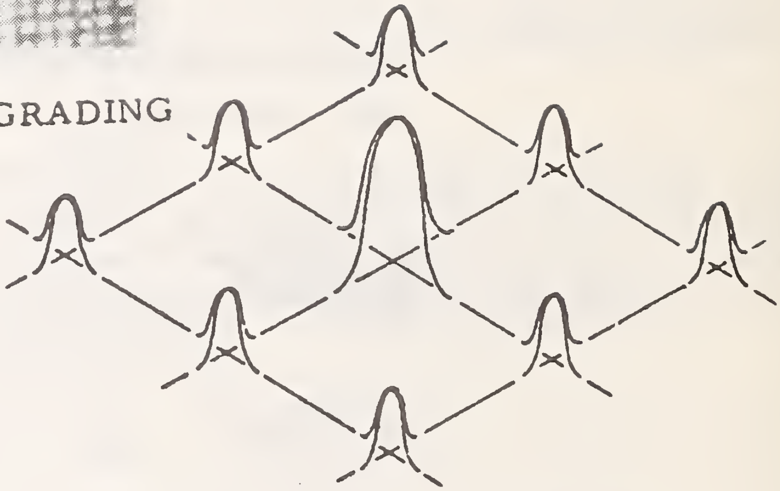

Figure 20
Fabric Inspection

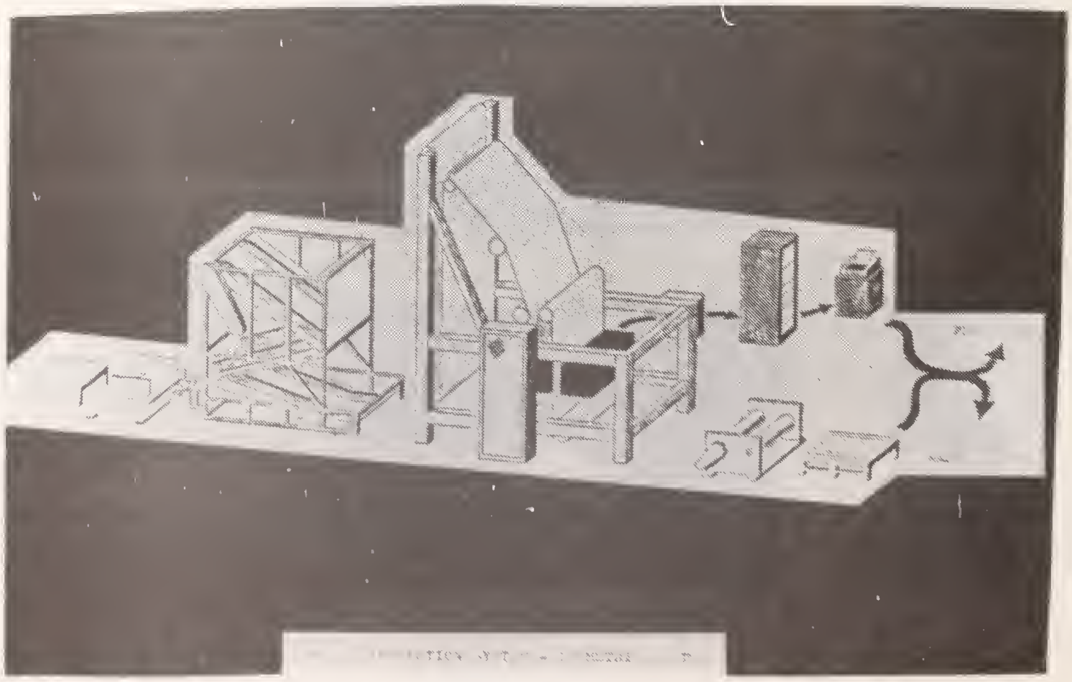

Figure 21

2.6 Hitachi Shadow Mask Inspection

A color picture tube has a shadow mask at the back of its front panel. It is made of a thin iron plate on the surface of which is a grid of several hundred thousand small rectangular holes. The holes are formed by contact-exposure printing and chemical etching. The patterns of the master plate are frequently damaged in production or in the printing process, hence the master plate must be inspected. Manual inspection of a master plate is shown in Figure 22. An automatic visual inspection system for shadow mask master plates has been developed by Hitachi, Ltd. [10].

A master pattern is a 610 x 800 mm plate with rectangular patterns approximately .10 x .60 mm in size and approximately .10 to .20 mm apart. Figure 23 shows the kinds of defects and the dimensional tolerances in the inspection process. Defects include chips, projections, and pin holes in the patterns, as well as dimensional errors.

The image acquisition component of the experimental equipment (Figure 24) consisted of a solid state linear image sensor, a light source, a sliding table with driving motors and control circuits, and a microscope. The table is moved to form a raster scan of the master plate. The image sensor is a RETICON linear array, which provides a resolution of 7 micrometers under the conditions that the clock frequency is 5 MHz and the magnification of the object is 10 times.

The video signal is thresholded and processed in a minicomputer. The processing is done by comparing successive patterns, and by examining features of individual patterns. Features used include area comparison, maximum width comparison, distance between patterns, and width change within a single pattern. Each field is 128 x 512 picture points, and contain 3 pattern stripes.

A production line machine (Figure 25) for the automatic inspection of master plates has been constructed. It is reported that the machine operates at 200 stripes/sec. and the performance is equal to that of skilled inspectors.

Figure 22

No.	Name	Feature	Limit size
1	chip-1	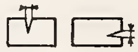	10μm
2	chip-2	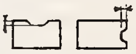	10μm
3	projection		10μm
4	pinhole		10μm
5	missed		10μm
6	too large	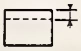	10μm
7	too small	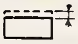	10μm
8	dot	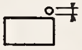	10μm
9	bridge width		10μm
10	position error	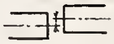	10μm

Figure 23 Types of Defects

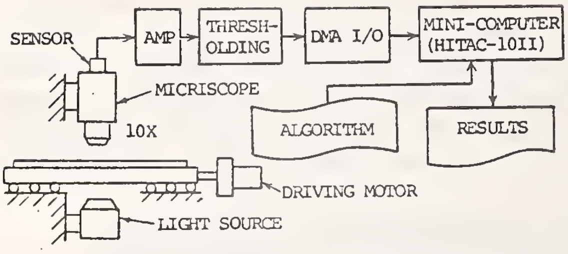

Figure 24 (Left)

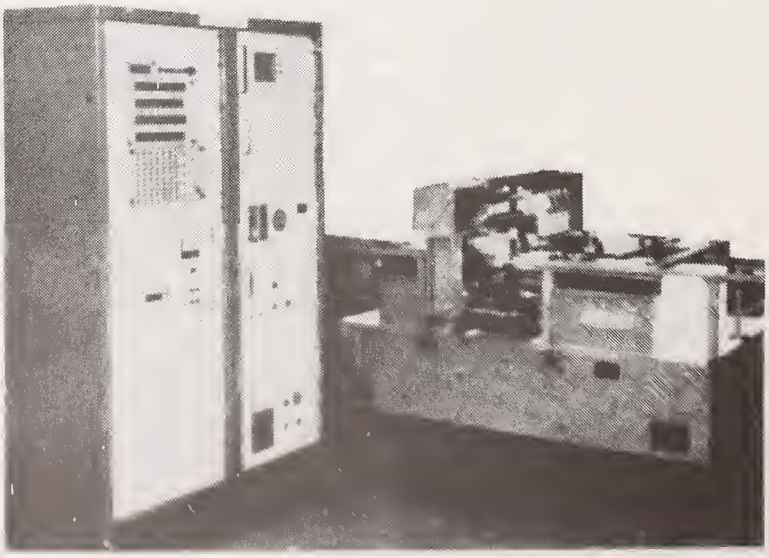

Figure 25

2.7 Hitachi Relay Switch Inspection

Hitachi, Ltd. has built a machine for the automatic inspection of contact points of relay switches [11]. Defects (Figure 26) such as displacement of the contact point, scratches, and welding splashes hinder the performance of the relay switch. Visual inspection criteria have been established for each of these defects. For example, the displacement should be less than 0.1mm along the X and the Y axis.

The system configuration (Figure 27) consists of a TV camera, a vertical reflected light source, and an objective lens for magnification, and a movable table. The inspection process takes place when the table is stationary.

The image processing for a single contact point is done in a sequence of windows of the entire image. Within each window a threshold is chosen, so that processing is done on a binary image. Thresholds are determined on the basis of the maximum and minimum gray levels within the window.

Figure 28 shows a window (labeled "inside frame") which corresponds to the smallest acceptable contact size. Previous processing has determined the contact center, which is used to fix the location of the window. The search for scratches and the adherence of an alien substance is restricted to this region. The search itself consists of looking in 7 X 7 neighborhoods for sets of above threshold picture points which form a chain across the neighborhood. In a similar manner, examining the contact point for a welding splash is done by searching with a different size neighborhood along the perimeter of the larger window in Figure 28.

Figure 29 shows the appearance of the production line machine. It employs two cameras, and is capable of inspecting 32 contact points in 5.4 sec.

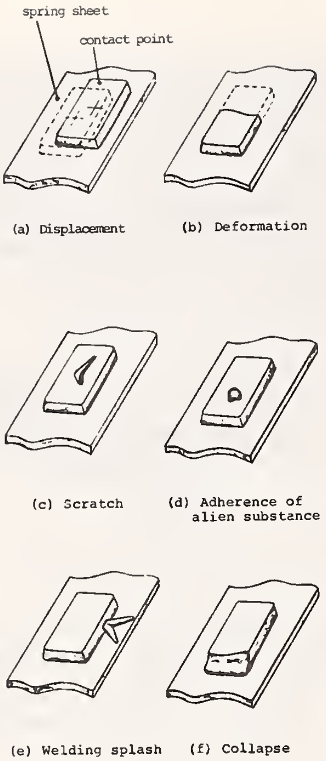

(a) Displacement (b) Deformation

(c) Scratch (d) Adherence of alien substance

(e) Welding splash (f) Collapse

Figure 26 (Above)

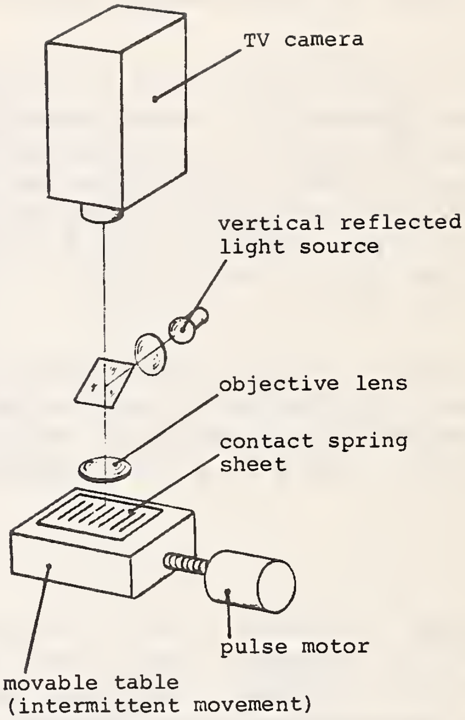

Figure 27 (Above)

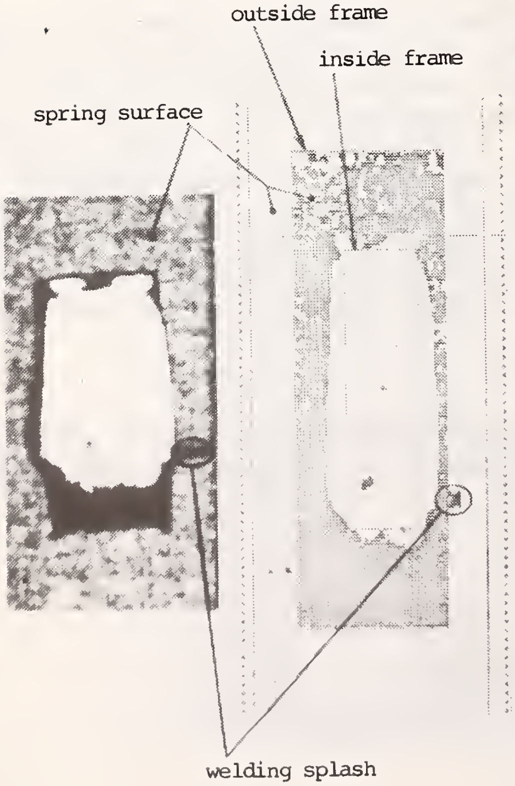

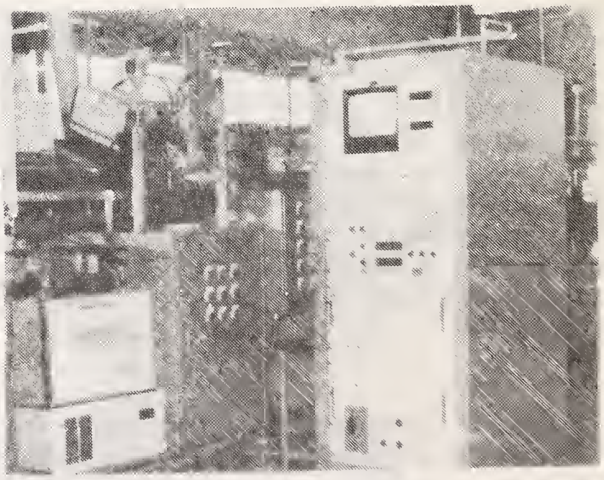

Figure 29 (Above)

Figure 28 (Left)

3. Visual Inspection for Robots

A visual inspection capability is a potentially important component of a programmable manipulator or robot. Typically, a human factory worker performs several implicit inspection tasks in a materials handling job. It is reasonable to expect the worker to inspect for: the presence or absence of a part; completeness of a part; presence of major defects such as cracks, burrs, and a variety of simple defects. If a robot is to perform this implicit inspection aspect of materials handling it needs sensory feedback. Vision sensors using the techniques of image analysis and pattern recognition are now appearing on prototype robots and in feasibility studies.

Even when the robot is not performing an inspection task, sensory feedback is important. Sometimes the identity of the part is not known, such as in a foundry where a number of different castings could be on a conveyor belt.

Often, however, the identity of the part in a material handling situation is known, but its position and orientation are not known. A vision system provides the information necessary for a robot to acquire the parts.

There is still another category where sensory feedback for robots is important, where the approximate position and orientation of a part is known. Here the choice is to achieve higher accuracy in part location through the conveyor, assembly line and special purpose machinery, or to equip the robot with sensory feedback. The latter is potentially more economical.

3.1 General Motors Part Recognition

General Motors Research Laboratories has developed a vision system that can determine the position and orientation of complex curved objects in gray-level noisy scenes [12,13].

The system is able to recognize parts such as connecting rods and steering knuckles that are lying flat on a conveyor belt which is at a known distance from the camera. No special illumination is used. The system stores a model of each part that it is able to recognize. Additional models for parts are added to the system by showing the part under favorable conditions.

Figure 30 is a 256 X 256, 16 gray level image of a connecting rod. The Hueckel edge operator produces the 400 edge points shown in Figure 31. Edge points are linked together using a local proximity and orientation criteria to form chains (Figure 32). The program puts an asterisk at the end of open chains and shades the interior regions of closed chains. Curve fitting methods are used to replace the chains by what are called concurves, which are sequences of straight lines and arcs. Figure 33 shows that nine chains have been replaced by four concurves, which consist of 19 straight lines and arcs. A program uses this data to obtain model concurves (Figure 34), which are stored in the data base so that future connecting rods can be detected.

Figure 35 shows the digitized picture of six parts on a conveyor belt, the edge points, the concurves, and the models superimposed on the gray level picture. The problem of matching the concurves of a part to the right model amounts to finding the translation and rotation which maps the model into the concurves of the part. There are two steps in the matching process. First, one model concurve is matched against one image concurve to determine a tentative transformation. Second, the transformation is checked in a global manner by matching a representation of the total model data with the total image data. The search aspect of the matching process is aided by a heuristic which measures the matching likelihood for each model/image concurve pair that is possible. This heuristic is based on general characteristics of the concurves. The system can work when the parts are occluding, such as shown in Figure 36. The effort was a feasibility study. Presently work is going on to reduce the computational and storage requirements.

Digitized picture of
connecting rod

Figure 30 (Above)

Figure 31 (Below)

Result of applying Hueckel edge operator on Fig. 30 Output: 400 edge points containing location, direction, and average intensity on both sides.

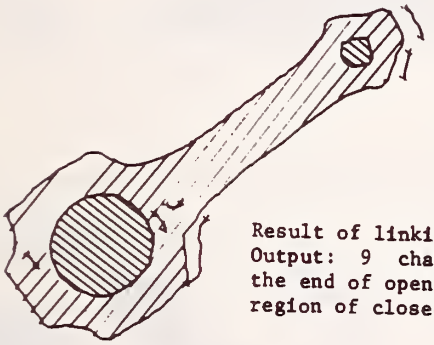

Result of linking edge points to form chains. Output: 9 chains. The program puts a "*" at the end of open chains and shades the interior region of closed chains.

Figure 32 (Left)

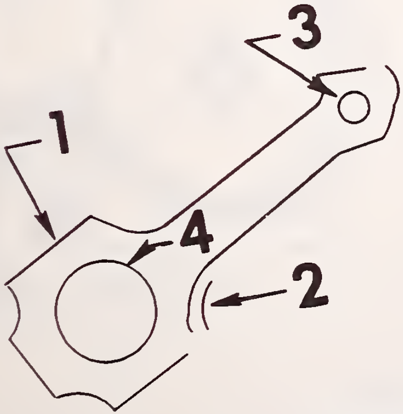

Figure 33

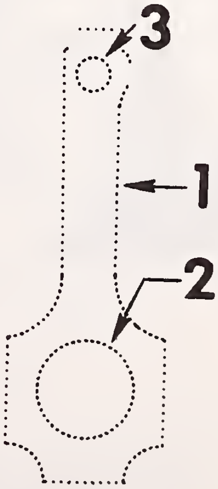

Figure 34

Result of curve fitting to form concurves. Output: 4 concurves which consist of 19 straight lines and arcs.

Model concurves obtained by program from table of input geometric data.

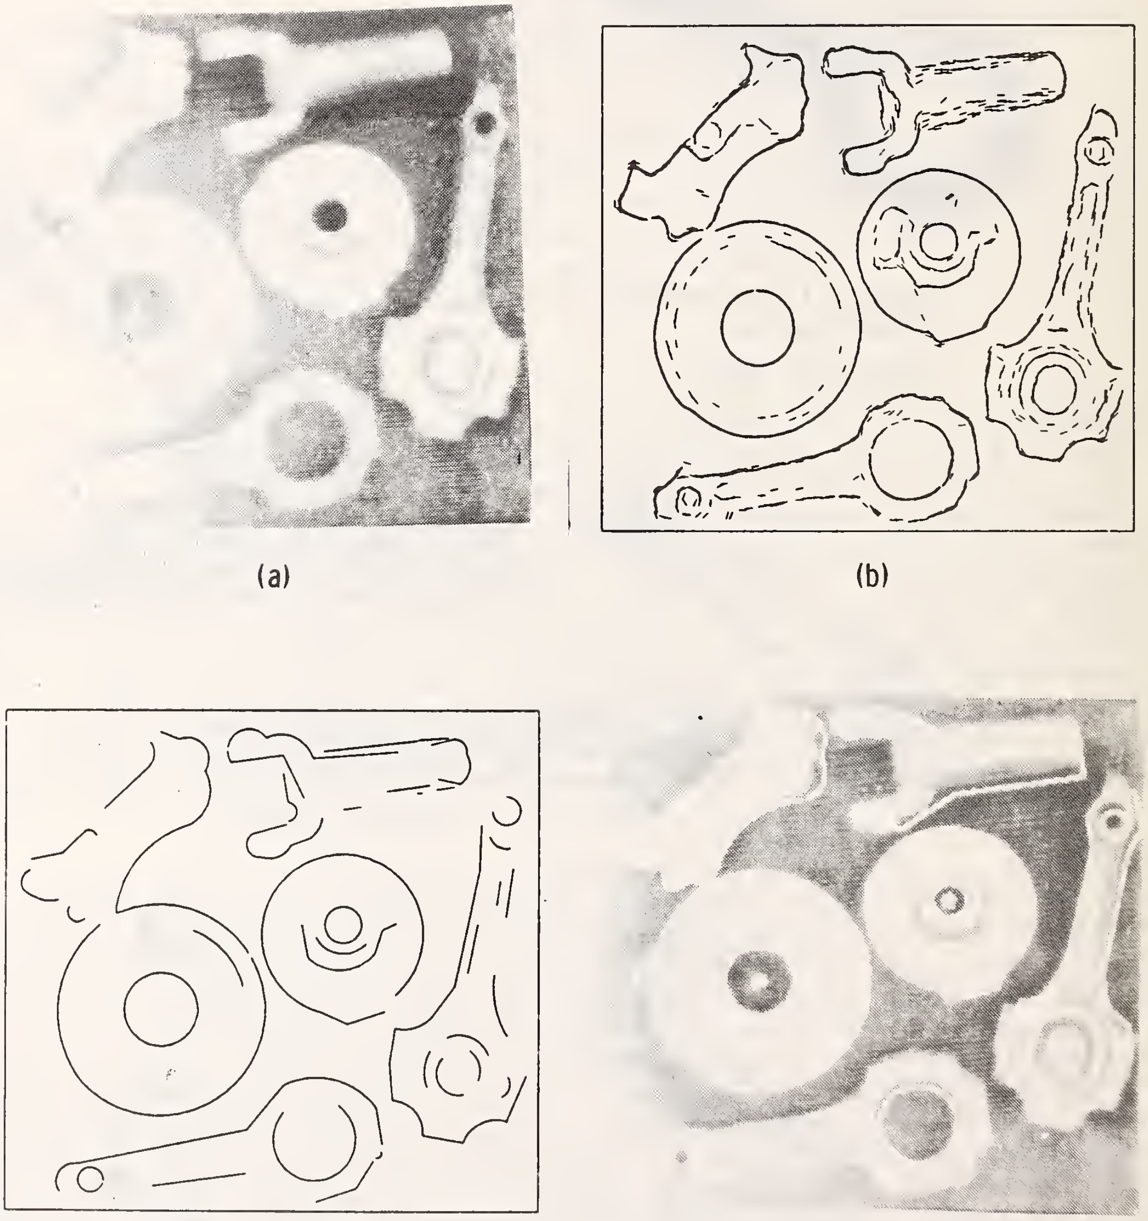

(a) (b)

(c) (d)

Figure 35

Six parts on a conveyor belt. (a) Digitized Picture. (b) Edge Points. (c) Concurves. (d) Models superimposed on gray-level picture.

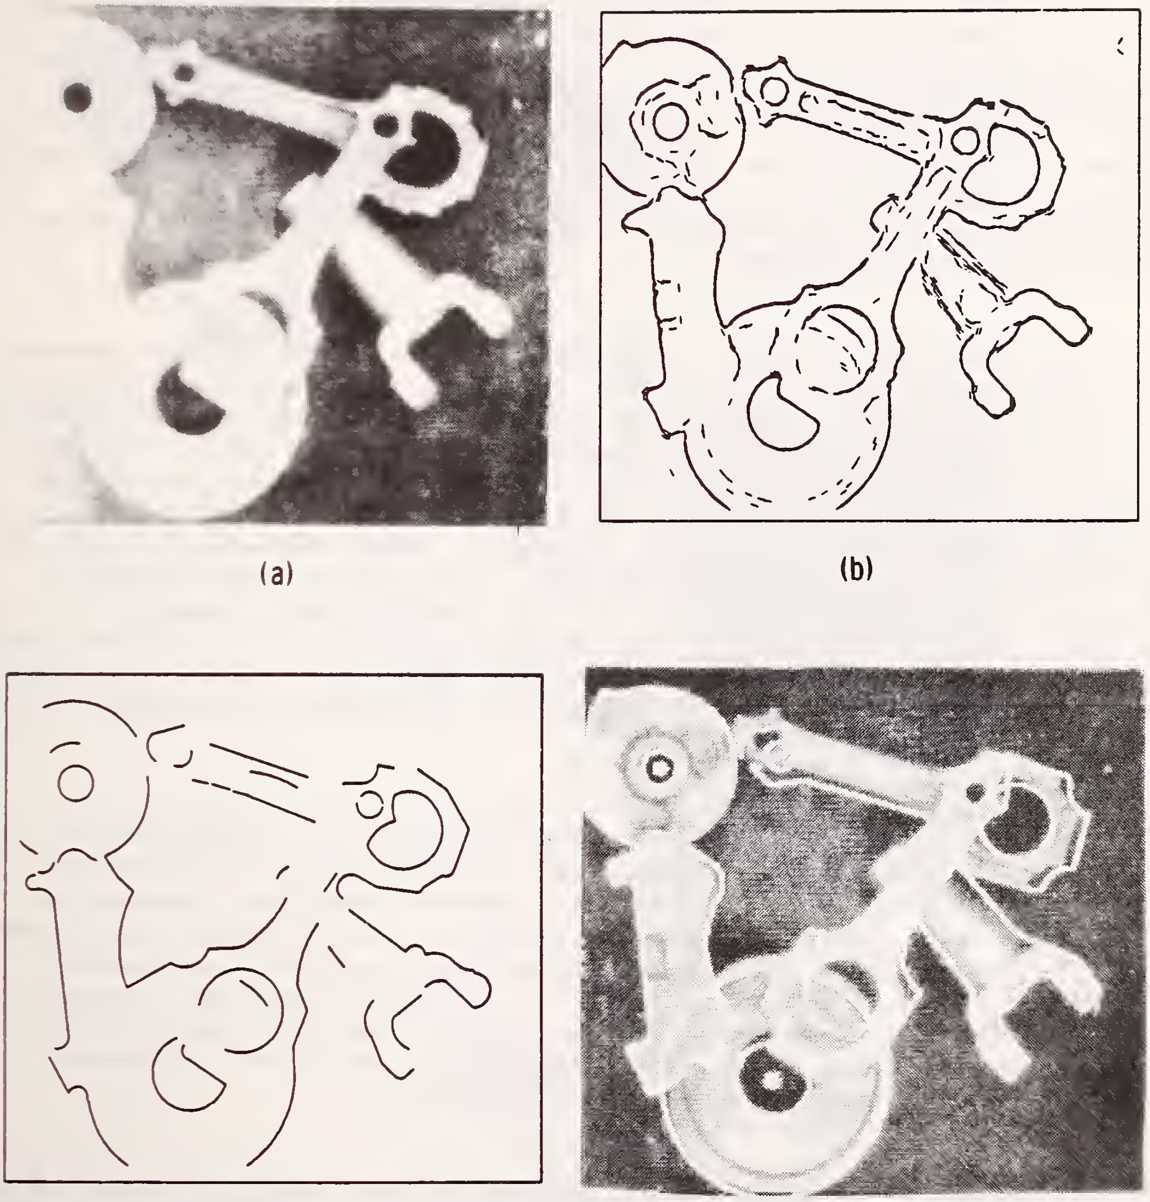

(a) (b)

(c) (d)

Figure 36

Six parts on conveyor belt with considerable occlusion.
(a) Digitized Picture. (b) Edge Points. (c) Concurves.
(d) Models superimposed on gray-level picture.

3.2 SRI International

SRI International has been working in the area of vision systems for inspection and for materials handling [14,15]. Their system begins operation by thresholding the image of a part, such as one of the castings shown in Figure 37. A connectivity analyzer determines the outline of the part, and a set of features is computed. The features used include: perimeter of the outline, square root of the area, total hole area, minimum radius, maximum radius, average radius, and a ratio of perimeter to area. These features were chosen because they are invariant with respect to translation and rotation.

The four castings shown in Figure 37 have seven stable states, as shown in Figure 38. All of the features are computed for each of the stable states. The system automaticaly computes a binary decision tree, such as the one shown in Figure 39. This tree is then used to recognize individual parts. The feature at the top of the tree is taken to be the one which has the largest gap between the classes. In this example, that feature is called x3, and happens to be the minimum radius in the outline. It separates the connecting rod, the piston sleeve, and the brake caliper from the rest of the stable positions. One can readily see that the minimum radius feature performs this separation. Recursively repeating this process generates the entire binary decision tree.

SRI has also used their system for inspection applications such as for the inspection of the lamp bases shown in Figure 40. The system provides the user with the capability to determine the edges of the bases and the contacts, and provides programming tools to assist him in implementing an inspection procedure. The procedure must determine that: 1) there are two contacts, 2) that the spacing between them is correct, and 3) that the position of the two contacts with respect to the base is correct. In Figure 41, only one part passes inspection.

SRI has mounted a solid state camera in the hand of their Unimate robot (Figure 42). The camera is mounted parallel to the axis of the hand. The mode of operation to precisely locate an object is to: 1) position the camera over the object, 2) use the previously described image processing software to find the object, 3) move over a fixed distance equal to the separation of the hand and the camera, and 4) move straight down to the object.

Figure 37 (Left)
Four Foundry Castings

Figure 38 (Right)
Seven Stable States

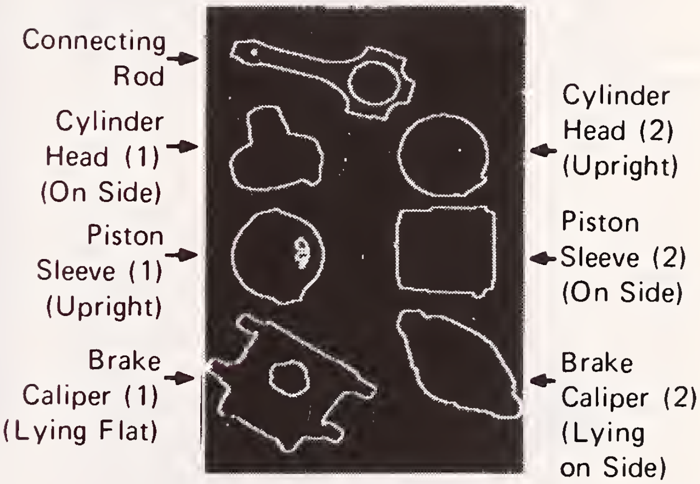

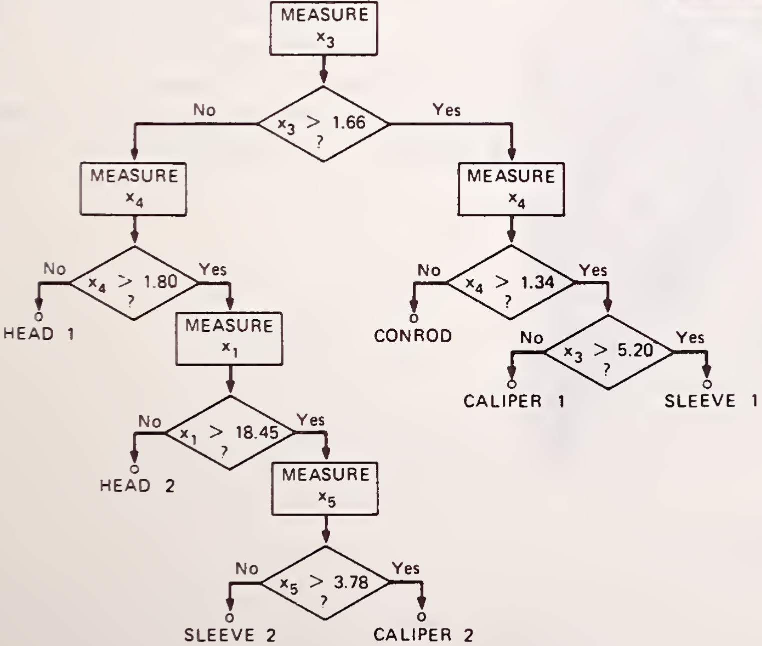

Figure 39 (Left)
Decision Tree

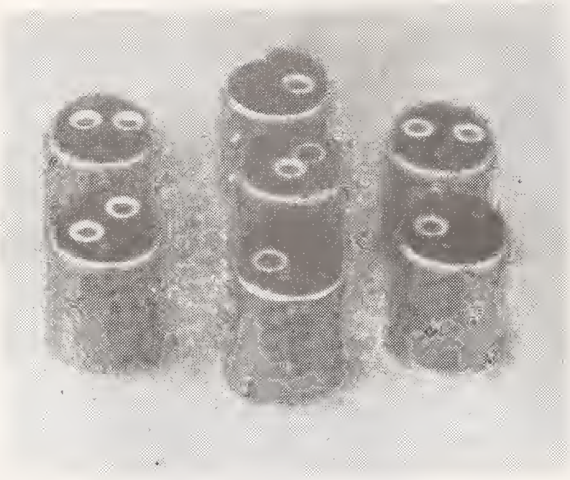

Figure 40 (Left)
Lamp Bases

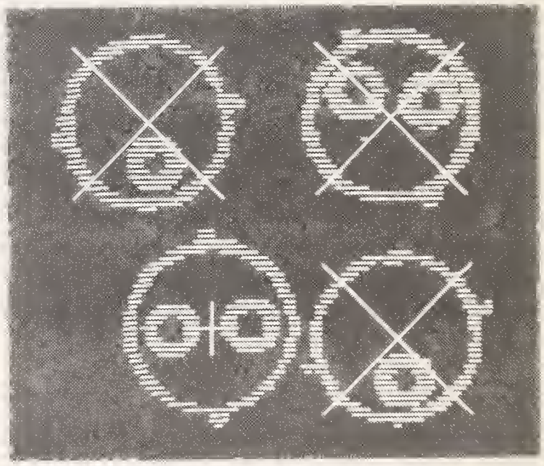

Figure 41 (Above)
Inspected Lamp Bases

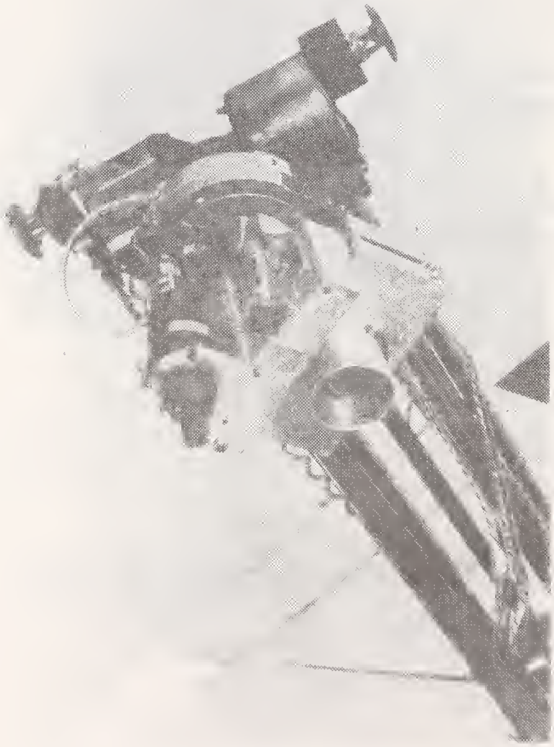

Figure 42
Camera on Unimate

3.3 Auto-Place Robot

Figure 43 shows an industrial robot made by Auto-Place Inc. At the Robots II Conference in 1977 they presented a demonstration that showed how their robot could modify its behavior from sensory input. The sensory input comes from a solid state camera mounted horizontally on top of the robot. Light is deflected from a horizontal work table to the camera. In the demonstration, the robot picked up a pair of dice that had been randomly placed on the work table. It also determined the throw of the dice, by simply adding up the black points (obtained by a simple threshold algorithm) which correspond to the dots on the face of the dice. In an earlier demonstration Auto-Place used two robots working together with a camera to sort a deck of cards into suits (see Figure 44). The first robot would pick up a card from a card tray with a suction grip and hold it up to a camera mounted separately from the robot. The camera and computer system determine the suit by looking at one of the suit symbols, thresholding it, and using a simple shape criterion. If the card did not fall into the proper catagories (club, diamond, heart or spade) the robot was instructed to turn the card over. If again the suit could not be determined, the robot dropped the card on the floor. When a suit was determined, the robot which picked it up, placed it into the sorted pile of hearts or diamonds as appropriate. When a club or spade was found, the card was passed to the second robot and appropriately placed in the proper sorted pile. Auto-Place has not yet reported an industrial application of their robot with a camera. They are, however, currently working on an inspection application of vision for the Jeep Corporation. This inspection system will consist of ten cameras linked to a microcomputer. The system will analyze the body-mounting pad heights on an automative frame. This analysis determines the amount of shimming required at each body- mounting location, in order to insure a flat mounting plane for the body.

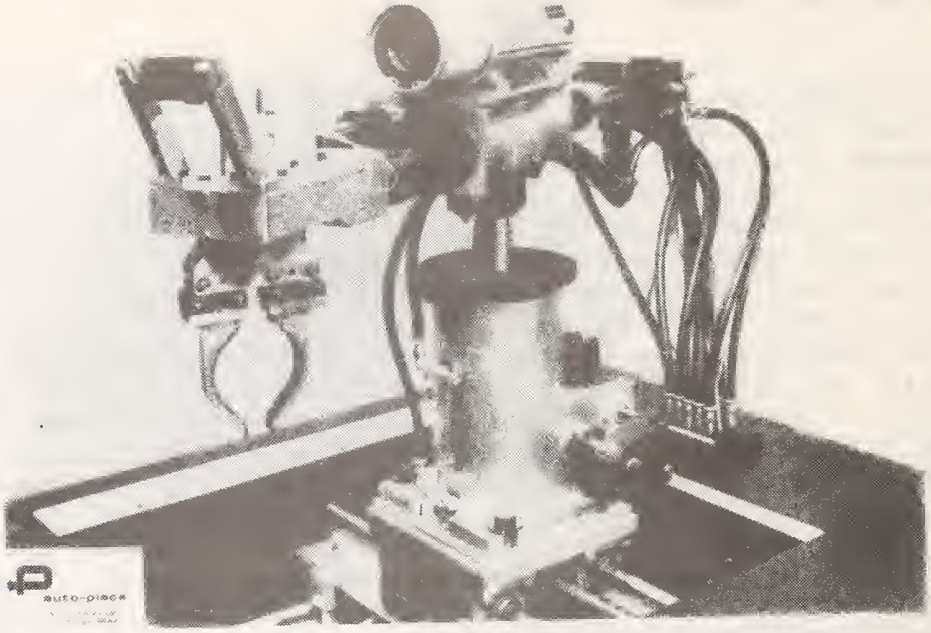

Figure 43

Figure 44

3.4 National Bureau of Standards

The Automation Technology Program at the National Bureau of Standards is developing a vision system for its robot [17]. The vision system consists of a solid state camera, a strobographic light source, and an 8-bit microprocessor. The camera is mounted obliquely at the wrist of the robot, so that its field of view covers a region extending from inside the finger tips out to a distance of one meter. The light source flashes a plane of light parallel to the wrist of the robot into the region. Figure 45 shows the configuration of the camera, flash unit, and the field of view. The resulting image sensed by the camera has a variable scale (shown in Figure 46) due to the angulation of the camera. The net result is that objects far away from the wrist appear in the coarsely calibrated region of the image, whereas objects close to the wrist appear in the finely calibrated region. The plane of light produces an image with bright line segments. The computer system computes a run length encoding representation of the line segments. The interpretation algorithm uses only the line segment data, and based on the fact that triangulation gives range data, the slope of the lines indicate the orientation of the object, and the end points of the lines provide information on the edges to be grasped. Figure 47 shows the camera and flash unit mounted on the robot, and Figure 48 illustrates the line segments obtained from some simple objects. The computer can control the duration of the flash, and therefore the contrast in the image. The camera system provides an image with 128 by 128 pixels, and is quantized to eight bits.

The software component of the system is still being developed. A typical application is in materials handling, such as a three step sequence for the task of acquiring an object which has been randomly placed on a table. First the table is scanned in a plane which is approximately parallel to the table. The object will appear in the image, and range values for its locations can be computed. In general, these values will only be estimates of the objects' location, since the computation will be done in the course resolution part of the scale. Next, the estimate of the location can be used to move closer and take a second view of the object. This more precise determination of the location can be used to signal the control to move the arm above the object in preparation for grasping it. Finally, a third look can be used to verify conditions prior to grasping. Figure 48 shows a set of simulated images which illustrate this sequence.

The intent of the system is to demonstrate the concepts for a highly flexible low cost, vision system. By using the plane of light and a dedicated microprocessor, a hard wired data reduction scheme contained in the vision sensor has been achieved. This permits the analysis of vision data to take place in real time as the robot moves about in its environment.

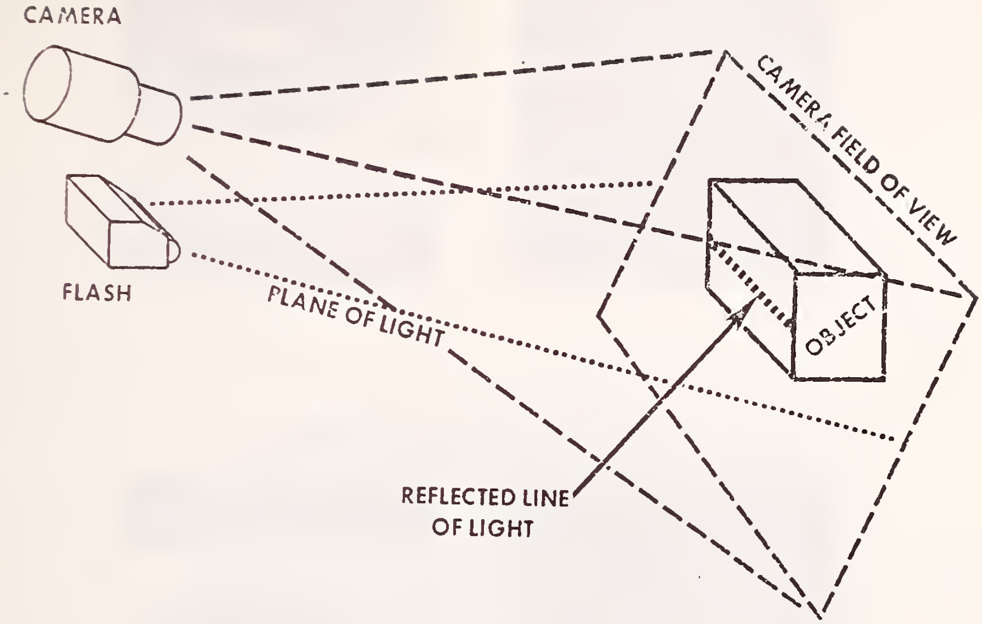

Figure 45

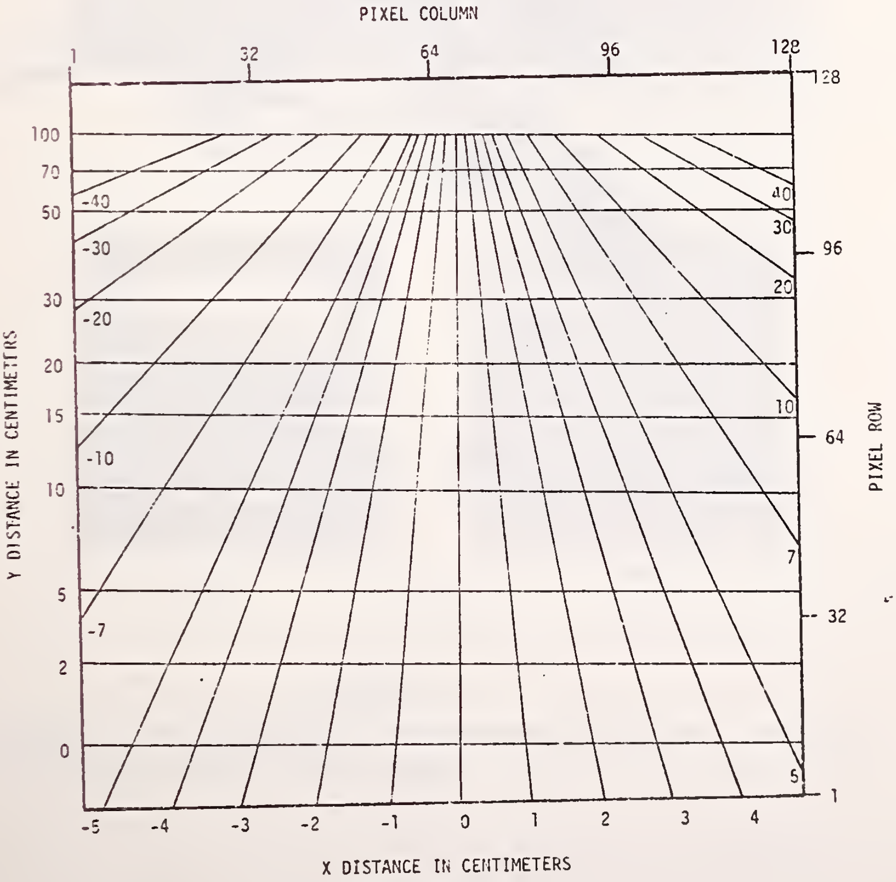

Figure 46

Figure 47

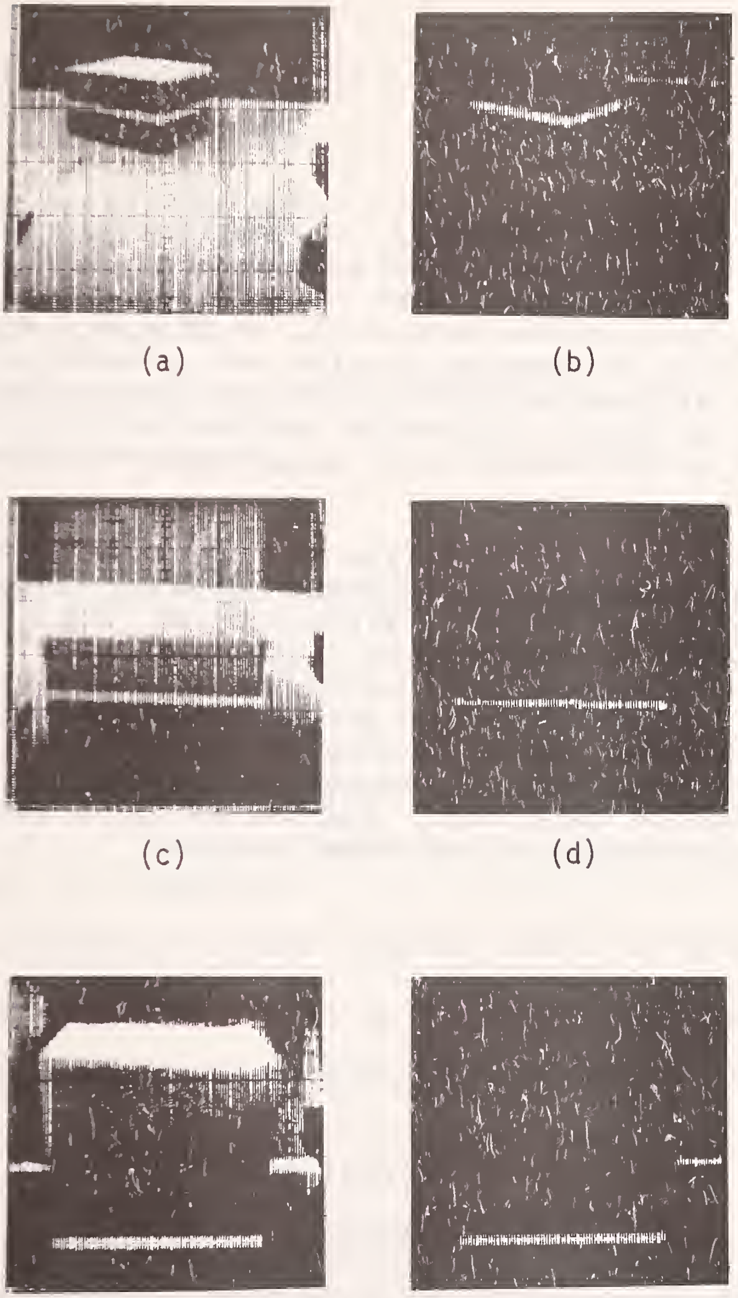

(a) (b)

(c) (d)

(e) (f)

Figure 48

A series of pictures which illustrates the three step part acquisition sequence for a box.

4.0 Other Applications

In section 2 and 3 we presented a sampling of application in visual inspection systems. There are other applications at other institutions. For example, work on printed circuit board inspection has been done at the University of Missouri and at Princeton University. Work on computer vision for industrial applications has gone on at Stanford University, at Hitachi Central Research Laboratories, at Osaka University, and at the University of Rhode Island. This is not meant to be an exhaustive enumeration.

In this section we present a list of other applications of visual inspection to industrial problems. Information on applications by private companies is often difficult to obtain. Often the very fact that a company is considering an application is proprietary. Nonetheless, we do have a list of industrial applications to present. There are no details of the applications available. However, the diversity of the list is itself thought provoking. Kenneth White of Proctor and Gamble collaborated in the compilation of this list. Additional applications were obtained from, reference 18 and other sources.

1. grasshopper contamination of string beans
2. frozen food quantity count variations
3. dust detection on LSI substrates
4. bottle fatigue detection, to prevent explosions
5. thermometer inspection
6. label inspection on plastic bottles
7. golf ball label inspection
8. IC chip wire bonding
9. measurement of the effectiveness of a window defogger over time
10. automatic debarking, initial timber cutting, sorting of lumber
11. measurement of biscuit height so that packaging equipment doesn't jam
12. unthreaded nut detection
13. button inspection for the correct number of holes
14. sorting of porcelain seals by shape and configuration
15. automatic closed loop control of extruded gelatin sausage casings
16. hot roll steel web width control
17. particle inspection of pharmaceutical liquids
18. closure integrity inspection for pharmaceuticals
    - . is the rubber seal inside the cap provided by the supplier?
    - . is the cap tightened?
    - . is it not cross threaded?
    - . is the seal over the cap properly in place?
19. blister pack inspection - is there one pill or blister?

20. china plate inspection for pinholes and bubbles on the plate

21. automated separation of whole almond nut meat from debris, shells, and damaged meat.

5.0 Conclusions and Future Trends

What are some common characteristics of the inspection problems that we have discussed in detail? A first observation is that the majority of the applications were hard automation, in the sense that they were specific solutions to specific problems.

One exception to this is the syntactic pattern recognition approach to inspecting parts, such as springs and screws, where the median axis is a good representation of the part. The primary task in applying this technique to a new application is the development of a new grammar. The optical power spectrum approach is also somewhat of an exception to the hard automation observation. Here the user must develop a classification procedure based on the 32 wedge and the 32 ring features of the photodetection.

The parts recognition work at General Motors and at SRI International are also exceptions in that new part systems can be added to these systems. These systems are primarily intended for material handling applications as opposed to visual inspection. The barrier of hard automation for visual inspection has been addressed at SRI, in that they have developed interactive programming aids to assist the user in developing new inspection procedures.

The systems that are in (or are nearly in) the factory are for applications with high production volumes. This tends to require high inspection rates: one per second for the Series 700 Connector; 5-15/sec. for the screws and springs.

High inspection rates result in a big emphasis on data reduction. Image processing applications inherently begin with large amounts of data. One guideline to keeping a handle on the amount of data is to use the coarsest acceptable resolution. This can be determined in part by the smallest dimension that must be measured. In part it seems that some experimentation is required to determine the coarsest acceptable resolution.

Similar things are true for quantization. Many of the applications discussed used binary images. Only the shell inspection application used a fine quantization.

Probably the best parameter in estimating what inspection rate may be achievable is the complexity of the inspection process itself. How difficult will it be to identify the features that are necessary to distinguish between the good and the bad parts? How complex will the segmentation process be? What will the computational complexity of the decision procedure be?

One heuristic in estimating the complexity of the inspection process is the amount of difficulty that a technically competent person, who has no previous knowledge of the component being inspected, experiences in understanding the manual procedure. Ideally, the manual procedure should be rather obvious.

A good inspection application will make effective use of techniques for controlling the environment. If possible, the part should be viewed from a fixed location, with a known orientation. Using an array detector under these conditions results in processing only a single image. Sometimes, such as with the spring and the screw application, the second dimension can be effectively provided by moving the part. Sometimes such as with the shadow mask inspection, moving the part is the only option.

Control of the illumination, such as with lighting techniques, can reduce the complexity of the task. The same is true for control of the optical parameters (such as color and reflectivity) of the part being inspected, although often these parameters are not in the domain of control of the inspection designer.

What is the future for automated visual inspection systems? The future is good! The future is good because there exists a substantial, fundamental understanding of image pattern recognition, and basic research in this area continues. The future is good because the semiconductor manufacturers have provided a variety of linear and array sensors, and functioning cameras using these devices. These solid state cameras are potentially rugged and reliable, and only cost in the neighborhood of one to three thousand dollars. The future is good because the microprocessor is substantially reducing the cost of computational power. Also the cost of memory is coming down.

So the future is good. Nonetheless, the future is not assured. The production engineer and the factory manager must be brought to the point were they view automated inspection systems as one more tool to accomplish their jobs. This means that adequate attention must be paid to the details of the production engineer's concerns for performance and cost. A track record of successful applications must be built up. That track record is only beginning to appear.

'A (REV. 9-78)

DEPT. COMM.
BIBLIOGRAPHIC DATA SHEET

1. PUBLICATION OR REPORT NO.
79-1764 (NBS)

2. Gov't. Accession No.

3. Recipient's Accession No.

4. TITLE AND SUBTITLE
Image Pattern Recognition in Industrial Inspection

5. Publication Date

6. Performing Organization Code

7. AUTHOR(S)
Gordon J. VanderBrug and Roger N. Nagel

8. Performing Organ. Report No.

9. PERFORMING ORGANIZATION NAME AND ADDRESS

NATIONAL BUREAU OF STANDARDS
DEPARTMENT OF COMMERCE
WASHINGTON, DC 20234

10. Project/Task/Work Unit No.

11. Contract/Grant No.

12. SPONSORING ORGANIZATION NAME AND COMPLETE ADDRESS *(Street, City, State, ZIP)*
Same as 9

13. Type of Report & Period Covered

14. Sponsoring Agency Code

15. SUPPLEMENTARY NOTES

☐ Document describes a computer program; SF-185, FIPS Software Summary, is attached.

16. ABSTRACT *(A 200-word or less factual summary of most significant information. If document includes a significant bibliography or literature survey, mention it here.)*

Most manufacturing processes require a visual inspection of some aspect of the process. In the most straight forward applications, a vision system is used to inspect for part completeness, or to check for flaws in the manufacturing product. In more sophisticated use of vision systems the inspection task may be secondary to tasks such as part location, identification, or determining part orientation. In general these tasks are important when the vision system is used in conjunction with a robot manipulator. In robot systems, vision is needed to allow the robot to acquire, manipulate, and inspect parts without the need for elaborate fixturing, or complex part delivery systems. The labor cost of manual inspection, and the high cost of special purpose part delivery systems for robots have led many manufacturers to investigate vision systems for use in manufacturing. Digital image processing and pattern recognition are providing the basis for a growing number of attempts to achieve an automated vision system.

This paper begins with a brief historical perspective on image processing and pattern recognition. Next a series of state of the art examples of visual inspection systems, and then robot vision systems is presented. The paper contains a list of other areas in manufacturing for the application of vision systems, and concludes with an assessment of the future of vision systems in manufacturing.

17. KEY WORDS *(six to twelve entries; alphabetical order; capitalize only the first letter of the first key word unless a proper name; separated by semicolons)*

Automation, inspection, image processing, manufacturing, pattern recognition, robotics, vision systems

18. AVAILABILITY ☒ Unlimited

☐ For Official Distribution. Do Not Release to NTIS

☒ Order From Sup. of Doc., U.S. Government Printing Office, Washington, DC 20402, SD Stock No. SN003-003-

☐ Order From National Technical Information Service (NTIS), Springfield, VA. 22161

19. SECURITY CLASS (THIS REPORT)
UNCLASSIFIED

20. SECURITY CLASS (THIS PAGE)
UNCLASSIFIED

21. NO. OF PRINTED PAGES
45

22. Price
$4.50