Automating Tolerance to Process Variation

Pamela Lipson
Imagen, Inc. and Landrex Technologies
Cambridge, MA and Santa Clara, CA

Abstract
No two printed circuit boards look exactly alike. Even across two adjacent boards on an assembly line, one can find significant differences arising out of normal process variation – the components and the boards can change color, size and surface markings. A key challenge for inspection systems is to automatically handle such allowable variation, and distinguish it from other variations that constitute defects.

Automated optical inspection (AOI) systems have emerged as an important test strategy in printed circuit board manufacturing to detect defects. Typical AOI systems depend in large measure on heuristics-based (trial and error experimentation) data as the means to establish typical conditions, the degree of normal variation, the thresholds for nominal pass/fail conditions, the lighting/camera conditions to best view the object, and the parameters for providing variable measurement data. The use of empirical processes is a sound basis to make decisions where the sample population being employed is large enough to mimic the whole. However, user assessments of heuristics require skill and experience of programmers. As a result, the competency of empirical methods is built up over time and over volume by basing 'goodness' criteria on historical values and historical volumes, as seen through the filter of user judgment. Where time and volume are insufficient to establish stable norms, where user judgment of good and bad are questionable, or where variation of the elements of the printed circuit board is significant, it becomes difficult to effectively deploy heuristic-based programs.

In this paper, we present a new technology called 'Configural Recognition' that provides built-in tolerance to normal process variation. The technology was initially developed at the MIT Artificial Intelligence Laboratory for applications such as natural scene classification, face recognition, and trademark logo search. In these applications, normal variation is a significant and challenging problem for standard computer-vision systems. The technology is grounded in studies of how humans visually recognize objects.

Over the past six years the technology has been employed for the task of printed circuit board inspection and process control. The benefits of the technology in the PCB domain are its native ability to recognize PCB artifacts without re-sensitizing the objects under test, thus eliminating or reducing the necessity of establishing test norms A second benefit is that as variation is a known entity and accepted as an inevitable but compensated activity, far fewer examples need be used to generate a PCB test program. Finally, the technology allows for the optical-inspection system to make the leap from finding defects to providing reliable and repeatable variable measurement data with the same ease of use.

Introduction
Variation in images is truly the crux of the image processing or pattern-recognition problem. Computers are very good at comparing two images and making an assessment if they are exactly the same. Do they have the same layout? Do they have the same colors? Do they have the same textures? However, when one takes several pictures of an object, rarely do the images come out exactly the same. The reasons for this are:

- The camera position/focus/angle has changed
- The lighting has changed
- The object has moved
- The object has changed
- The object in the new picture is another instance of the object the original picture
- There are multiple objects in the image causing a range of confusing changes

The inability of computer vision systems to handle "acceptable" or "natural" variations of an object image has prevented such vision systems from being widely deployed in the marketplace and has constricted their applications to only the most controlled environments. (See Figure 1.)

Figure 1 - All the eight images above depict the same "object concept." Each of the four images on the left is easily recognized by a computer vision system as the same object or as being from the same object class. However, the images on the right depict some common variations, including lighting, pose, and resolution changes. These images are very difficult for a computer vision system to recognize as belonging to the same class as the original four images.

Vision systems that are the most successful are the ones deployed in the area of industrial inspection. However, even in these areas, most systems on the market give the programmer the responsibility of creating a program that can handle the normal process variation that the system is likely to see. They do this by providing the user with a tool kit that consists of a number of cameras of different views, variable lighting, a suite of image processing algorithms, and a set of thresholds that can be tuned. In these systems, the user is asked to find the best combination of cameras, lights, algorithms, and thresholds that best account for all the natural variation and possible substitutions. As a result:

- They generally require lengthy programming times to validate program stability
- They generally require many board examples to validate that good and bad thresholds are set properly
- They generally are intolerant of substitutions and normal variations in the process
- The program quality is highly dependent on the skill and experience of the operator who crafted the program
- They generate substantial numbers of false calls, requiring continuous adjustment to review and correct over time.

In this paper we discuss a fundamentally different approach to the problem of handling normal variation automatically. Our technique, called Configural Recognition™, is designed so that it does not depend on modulated lighting or camera conditions to elicit defects. In the particular embodiment of our system, we require only static, ordinary white light and single vertical camera, using these to collect 'flat', well-lit color images. Rather than using hardware-intensive methods to set up artificial imaging conditions to amplify the defects (and decrease the variation) present, our technique utilizes a new language to describe visual objects that can automatically handle the specific types and characteristics of objects to be inspected in the image.

As a result of using this new language, we can test the key elements of the image automatically, and relieve the user from qualifying, identifying and compensating for variant elements. A system equipped with this technique can provide:

- Single board programming (as in-circuit test does today)
- Outstanding operation in low-volume, high-mix operations
- Low false flag rates (approaching 5-10 ppm in practice)
- Freedom from perpetual tweaking to compensate variations over time
- Tolerance to user component and PCB substitutions
- Tolerance to normal variation

We have implemented this technique in the Landrex Optima 7200 post-place inspection system. At the end of this paper we will discuss some of our real results that show these benefits.

Process Variation in Printed Circuit Boards

When we first started working in the printed circuit board inspection field, we felt sure that the problem of finding components on printed circuit boards would be fairly straightforward. After all, from the CAD data and fiducial positions, we knew where all the parts were on the board and what type of parts they were. In addition, as the designers of the AOI machine, we were in charge of controlling the image-capture system. As described earlier, typical vision systems work very well when they are asked to make judgments about objects that look very similar and when these objects are viewed under very similar lighting conditions. We knew there would be some amount of natural variation. However, what we didn't realize

is that most of the visual attributes of a printed circuit board and its components are not controlled and, thus, are highly variable. There are an enormous number of visual changes in the printed circuit board, the paste and its components.
After starting the project, we began to collect images of printed circuit boards from a variety of different manufacturers. We also visited manufacturers of components, paste and bare boards. What we found was that:

1) The bare board appearance is dependent on the underlying metal substrate (e.g. copper) and non-metal substrate and the mask used on the board. Mask color is dependent on what the manufacturer has on hand and can range from red to blue to green. The speed at which the mask is placed over the board gives different shades of color on the board and across boards. (See figure 2(a) for a range of board colors.)

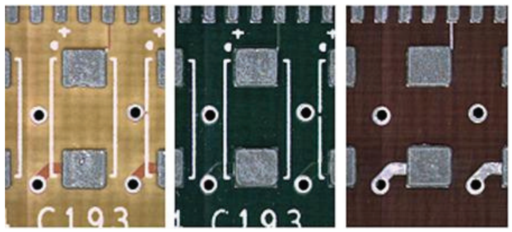

Figure 2(a) - Examples of printed circuit board color variations. In this case, the same board can be manufactured in colors such as yellow, green and red.

2) The component color, marking, and to some extent size are dependent on the manufacturer and also the environment. Humid conditions can make the leads oxidize and look, in some cases, as dull as the paste. (See figure 2(b).)

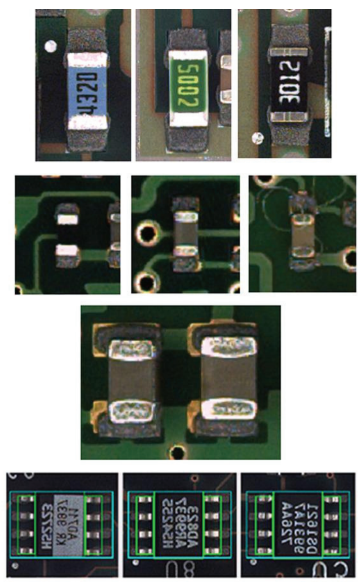

Figure 2(b) - Examples of component variations.
The top row shows components with three different colors. The second row shows components with different levels of oxidation on the end caps. The third row shows two 0805s which have very different sizes placed next to each other on the board. The third row shows examples of different markings on components. A printed circuit board inspection machine must be able to automatically handle this range of acceptable variations.

3. Each manufacturer can decide how much paste they want to put on the pads. In today's manufacturing environment, most pads are not fully pasted to avoid bridging problems. Regions from two semi-pasted pads can look very much like two shiny ends of a component such as a resistor. (See figure 2(c).)

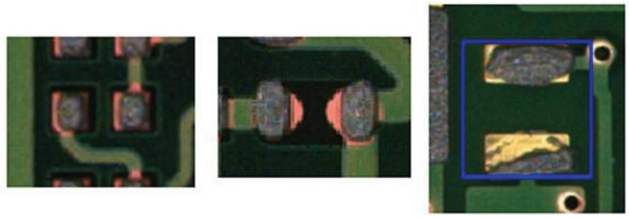

Figure 2(c) - Examples of partially pasted pads.
Pairs of bright metal patches can easily be confused for pairs of endcaps.

The larger question is if the board colors can change, the pad appearance can change, and the part appearance can change, how can an optical inspection system determine if it is looking at the board without a part present, a board with the correct part present in the right place, a board with the correct part present in the wrong location, or a board with the wrong part present?

Typical Methodologies Used By Optical Inspection Systems
Typical optical inspection systems perform image-analysis in one of three ways. However, none of these methods is really well suited to handle the large range of acceptable variations as described above.

1. Image-Correlation Method
The first method is called “image correlation.” This technique compares a stored image of the part with the image under test. This technique has the effect of superimposing the stored image on the image under test, taking the difference of each pixel in the stored image with its corresponding pixel in the image under test. *This technique looks for a “fixed spatial pattern” of “fixed quantitative colors.”* If the sum of the absolute differences is small, then the image under test is declared to be “the same” as the stored image. However, if the differences are great, then the image under test must either not contain the part or contain the wrong part. One main question to ask is how much “difference” is required before the diagnosis changes from “the same part” to “wrong part” or “missing part”. This question is difficult to answer.

The image-correlation technique works very well if the image under test looks exactly like the stored image. Figure 3(a) shows an example of a black resistor. If the image under test has a resistor that has the same size, same colorings and same markings, this technique works very well. See figure 3(a) for an example of similar parts that look the same. However, if the part has acceptable variations in size, color, or markings, then the system risks a false failure. See figure 3(b) showing a similar part that has changed from black to blue. In addition, if the board changes from a color that is *different* from the part to a color that is *similar* to the part, then the system risks a false accept. Most systems let the user decide how much “difference” in the way patterns look is acceptable or not acceptable. In addition, they often let the user change the lights and possibly even the angle of the camera to help reduce the visible effects of these variations.

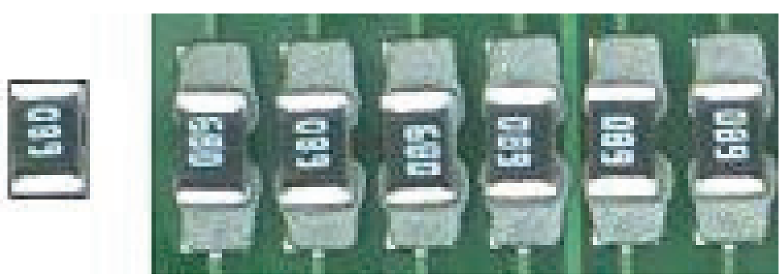

Figure 3(a) - A template image is shown on the left. Image-based correlation systems work well when all the examples look very similar to the template image, as those shown on the right.

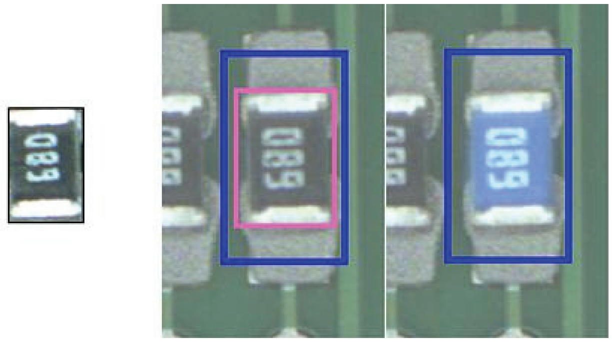

Figure 3(b) - Image-based correlation systems do not work well when any attribute of the part appearance changes. The template in this example is shown on the left. When applied to the two images on the right, it is easy for a correlation-based system to identify the black part. However, when trying to identify the blue part, because the differences between the images are so great, the system might diagnose the part as missing.

2. Histogramming Method

A second common technique recognizes that the absolute spatial appearance of the pixels may change, but assumes the quantitative color/luminance of the pixels will stay mainly the same. This technique is often called “histogramming” or “binning.” In this technique, the AOI system looks in a region and counts the number of pixels that have a certain value or property. If enough of the pixels in the image under test satisfy the expected count, then the system declares the part or object present. *This technique looks for a “non-fixed spatial pattern” of “fixed quantitative colors.”* Thus, if the part is black, the system looks for a sufficient count of black pixels to declare that the part is present. This technique helps to reduce some amount of sensitivity to part-size variations or marking variations. However, this approach can increase false positives because the technique has deliberatively ignored all the spatial information that is important in the signature of a part.

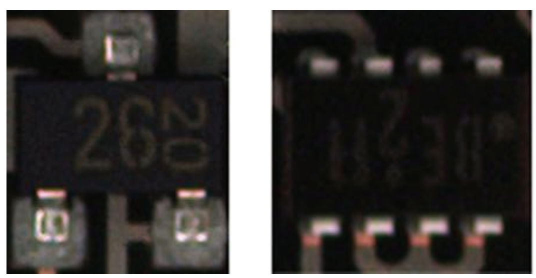

Figure 3(c) - Two images that have very similar distributions of colors and luminances. The histogramming approach, which does not take into account spatial configuration, might incorrectly identify these as the same part type.

3. Edge-Detection Method

There is a third approach, which looks for the transitions between the part and the board. This technique is called “edge detection.” If the system can identify a pattern of transitions or edges that are consistent with the part, then the system will declare the part present. *This approach looks for a “fixed spatial pattern of edges”, but ignores the “color/luminance attributes”.* The difficulty of this approach is that these transitions are difficult to identify because there are either too many distracting features or too little contrast between the board and the part to see any transition at all. These transitions represent what is typically known as “high frequency” features and are highly prone to noise.

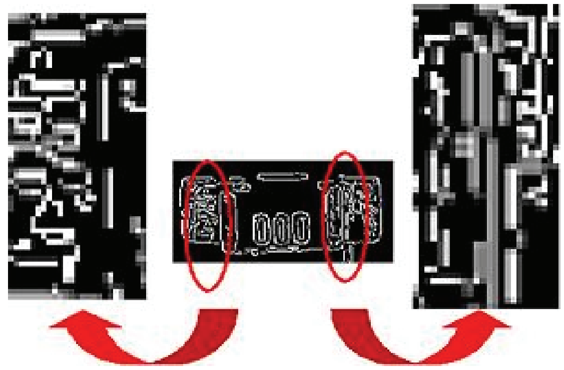

Figure 3(d) - This figure shows a processed image of a part. The white lines indicate strong edges or transitions in the image. The two images on the sides are enlargements of the end of each end cap and also paste regions. Looking at the enlarged images, it is difficult for a person and a computer vision system to tell which edges represents the true transition from the end cap to the paste.

The Basics of Our Approach

If we look at the strengths and weaknesses of these approaches it seems like an optimal approach would keep some information about the spatial layout or structure of a part, but wouldn't be so sensitive to the absolute color or luminances of the parts or the boards or the high-frequency noise-prone transitions. We, thus, developed an approach called "configural recognition" that looks for part structure, but measures part properties in a *qualitative* rather than a *quantitative* way. The configural recognition approach is essentially a new language to describe visual objects. It was originally developed at the Massachusetts Institute of Technology's Artificial Intelligence Laboratory and has been applied to a variety of visual-recognition tasks such as natural scene classification, face detection, and logo trademark matching.

The configural recognition scheme encodes class models as a set of salient low-resolution image regions and salient qualitative relationships between those regions. The configural recognition system constructs class models for parts using a wide vocabulary of relative relationships, including both spatial and photometric. In the current system, part models are described using several types of relative relationships between image regions. Each of these relationships can have the following values: less than, greater than, equal to, or different than. The spatial relationships used are relative horizontal and vertical descriptions. The photometric relationships, for example, can include luminance relationships, and color (e.g. R, G, B) relationships. We also encode relative size, where the size of the patch is described by how many pixels it covers. Intra-patch relative chromatic relationships may also be used in the model description.

The basic idea is that we can construct part models by these patches or regions and then encode relative relationships between the regions that will help to distinguish the part a) when it is present, b) from the board when the part is absent, or c) from a wrong part. Having the regions helps to encode the structure of the part. For instance, a capacitor always has two end caps and a body. The capacitor sits on two pads (pasted or partially pasted). The capacitor also has a relationship to the board. There should be board regions visible to the sides of the capacitor. There is also a board region below the capacitor that is visible when the capacitor is absent. Figure 4(a) shows these regions.

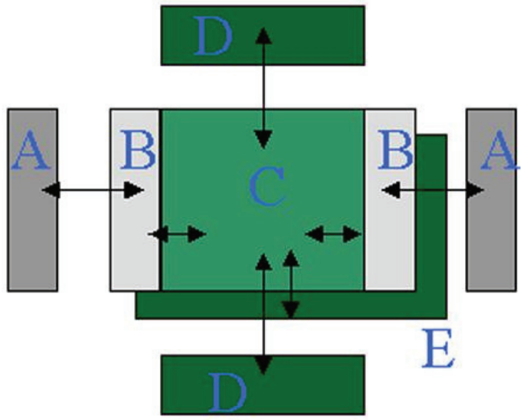

Figure 4(a) - Example of a model of capacitor (and the board around it) written in the language of regions and their relative relationships.

Once the spatial relationships that encode the parts structure have been determined, we can add in the photometric relationships. There are several photometric relationships that almost always remain common across all capacitors. The first is that the end caps are generally lighter than the paste. The second is the body is generally a different color from the end caps. Finally, the body is usually a different color from the background both below and to the sides of the parts. With these simple relationships we can encode both part structure and some aspects of part appearance that ideally will account for the large amount of acceptable variation that exists.

Let's review these regions and relationships on three images–two where the part is present and the one where the part is not present.

In the first image of Figure 4(b), there is a brown capacitor on a green board. Looking at the regions and their relationships, we see the body is different from the end caps. The end caps are brighter than the paste. The body is different from the background. All of our relationships from our model are satisfied and the system would declare the part present.

In the second image of Figure 4(b), there is a pink capacitor with oxidized end caps on a green background. Both the body color and end cap luminance have changed. Ideally, we'd like the system, *without any changes to our representation,* to be able to identify the part if present. Looking at these regions and their relationships, we see that they are still the same. The body is a different color from the end caps. The end caps, while oxidized, are still brighter than the paste. The body is a different color from the background. Once again the relationships from our model are satisfied and the system would declare the part present.

Now let's take a look at the case where the part is not present in the third image in Figure 4(b). We put our regions down on the image under test where we expected the part. In evaluating our regions and relationships, we find the body is a different color than the end caps. However, the end caps are not distinctly brighter than the paste, and the body is not different from the background. Thus, the results of these relationships indicate that the part is not present.

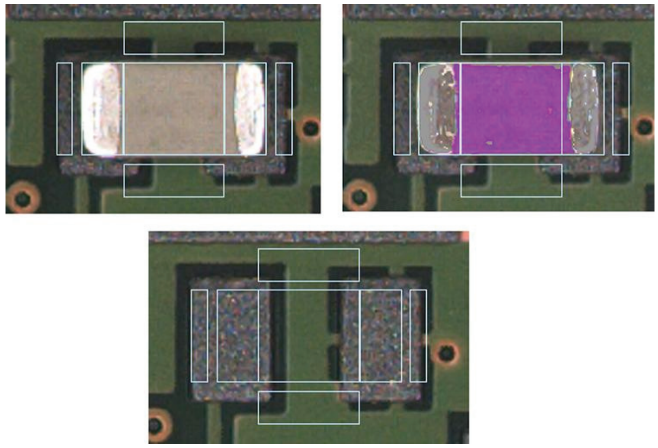

Figure 4(b) - Three examples of how the regions and relationships can determine part presence or absence. The first two images show natural variations of a capacitor. The third image shows a board when the capacitor is not present. We have overlaid the regions from our model on each of the images to illustrate our analysis process.

What we have shown in these three examples is that a very simple language that combines the structure of the part and board in combination with qualitative measures can be an effective and robust tool for determining part presence/absence even in the face of extreme acceptable variations. Beyond the qualitative encoding of regions and relationships, one of the benefits of the approach is that we are aggregating information over large regions; this gives us tolerance to high frequency noise.

The Science behind Configural Recognition

The configural recognition approach is based on three main principles derived from studies of human perception. These principles apply to any class of visual objects including faces, natural scenes, and printed circuit board components. We describe below the three primary principles using the example of natural scenes:

1. The importance of global scene configuration. The image on the right in Figure 5(a) has been derived from the one on the left by dividing the latter into pieces and altering their positions. Both images share identical chromatic and (to a large extent) textural statistics. However, these images do not belong to the same perceptual class because they have different overall configurations. Several systematic psychological studies have demonstrated that a visual stimulus presented with a correct spatial configuration allows viewers to more accurately and rapidly identify the stimulus and its parts than the same stimulus presented with incorrect spatial relationships [1][2][3]. Our conclusion, we derived is that the overall organization of a scene's or object's parts strongly influences a viewer's interpretation and perception of an image.

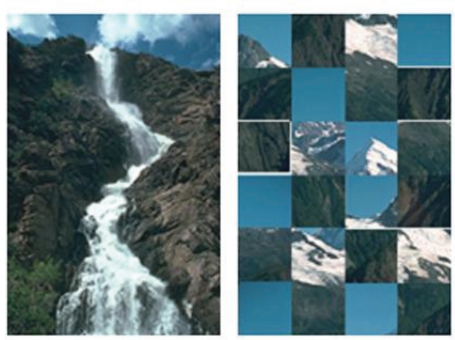

Figure 5(a) - A mountain picture and its scrambled counterpart.
Although both images contain the same color and textural characteristics, perceptually we would not classify the second image in the same category as the first.

2. The use of qualitative measurements. Figure 5(b) shows three snow-capped mountain scenes that are examples of a single class of objects. These three unique scenes have in common three relevant regions: a blue region (A), a white region (B), and a grey region (C). Region A is always located above region B that is above region C. Although these three regions are located at diverse absolute locations and vary in absolute size, one constant is that all the regions across the images have the same relative spatial layout.

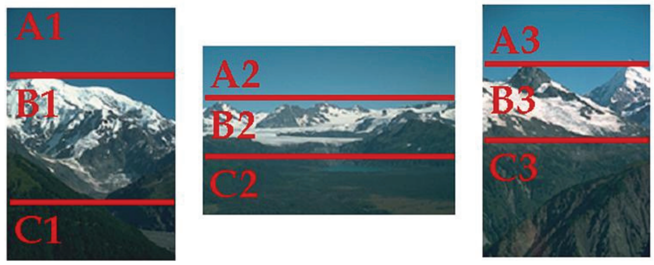

Figure 5(b) - Three snow-capped mountains are shown. Each is divided into three regions (A, B, C). Even though the corresponding regions in the three images differ in their absolute sizes, positions, and colors, their relative spatial and photometric relationships are largely the same.

Just as relative spatial relationships may be used to encode the overall configuration of scene content, relative color or luminance relationships between image regions may also be important for perceptual classification of scenes. The corresponding image regions in figure 5(b) may not have the same absolute color, but the relative photometric relationships between the regions are the same. For instance, in all the images region A is bluer and brighter than region C. This suggests that the classification of a scene may remain valid as long as the *relative* relationships between the image regions remain the same, even though the *absolute* values of the regions may change.

However, when the relative relationships are violated, often the percept and therefore the classification of that image is greatly altered. The difficulty that observers experience in recognizing photographic negatives is a case in point. Figure 5(c) shows an example of three images. The left image shows a rocky coastal scene. The second image shows the same scene with the color range stretched. Here the human percept of the image is still of a rocky coastal scene. Thus, even though the absolute colors have changed, the relative relationships of the regions remain the same (e.g. the surf is whiter and brighter than the coastline). In the right most image, the relative relationships between the regions are inverted, yielding an image that more closely resembles an iceberg than a rocky coast. Our hypothesis is that the relative relationships between the colors and luminances in the image are the key to human perception an image's content.

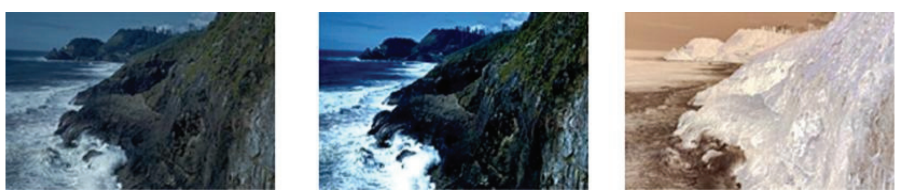

Figure 5(c) - Three images that demonstrate the importance of maintaining the qualitative or relative relationships between scene regions. The first two images are perceived as belonging to the same perceptual class–rocky coastlines– even though they differ in their absolute colors. The third image, an inverted image, is seen as belonging to a different class of images – e.g. snowy icebergs, because the color/luminance relationships between the regions are inverted from the original image.

3. The sufficiency of low spatial frequency information for scene classification. Figure 5(d) shows several thumbnails of readily recognizable images. An arrangement of low-frequency photometric regions is the only information retained in these small images, yet the image is easily recognizable. This observation leads us to the conclusion that humans need little detailed information to recognize many objects and scenes. As a result, as designers of an AOI system, our assumption is that we can base our classification algorithm on an image's low-frequency information.

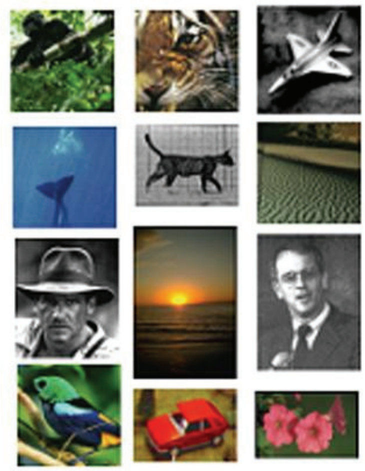

Figure 5(d) - Low-resolution images may be sufficient for classification. These images are identifiable Despite their extremely poor resolutions.

Deploying Configural Recognition in An AOI System
Our approach is deployed in the Landrex Optima 7200. It is the main technique that is used to determine presence/absence of parts. In the terminology of the machine, the technique is known as the "structural model." In addition to determining presence/absence, the technique can also be used to find part position. Because of its use of low- frequency information, the resolution of measurement of part position is at the pixel level. We utilize a geometric (edge) model after the structural model to more precisely localize the part. Because we have already found the part with the structural model at a pixel level, the geometric model is only asked to make a very fine adjustment and therefore is much less prone to making a mistake than when this technique is used as the full workhorse to find the part and determine part presence/absence. (See Figure 6(a).)

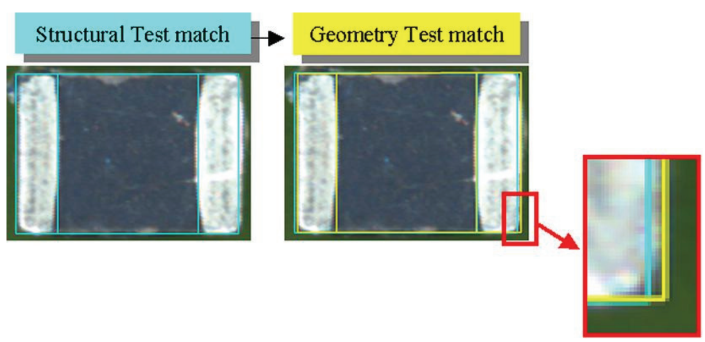

Figure 6(a) - This figure shows localization of the part by the structural model followed by a fine adjustment by the geometry model. The light blue lines in the image on the left denote the part's position, as determined by the structural model. Once the structural model has found the part's position, the geometric model is asked to make a very fine adjustment to adjust each edge to represent the true transition from the edge of the part to the board. The yellow lines show the localization of the part edges by the geometric model. The window on the right shows a blow up of the corner of the part to illustrate how fine the adjustment is.

The structural model has built in-knowledge about part structure. This comes from the library the user enters with body sizes and shapes and lead size shape and spacing and also a one word part description such as "CC" or "SOT". From this very generic part outline, the model is able to create a set of regions. Based on the part type, it then populates pairs of regions with relationships. The system uses built-in knowledge, for instance, that metal leads will be generally brighter than the paste and a different color from the body to bootstrap the process.

After the model is populated with regions and relations based on built-in knowledge, it then validates and emphasizes these relationships based on two sets of boards, a pasted board without parts and a pasted board with parts. At any time in the future, if an example comes along that is outside of the boundaries of acceptable variations that the model can handle, then the new example can be imported into the system and the model will attempt to adjust to the new example. The model has a measure of "goodness of inspection." At any time, if the model finds the examples are within too extreme a range, it can let the programmer know that its inspection quality is compromised by trying to encompass this extreme range. While the situation is unlikely, it is easy within the system to generate two or more structural models to handle this extra wide range of variation.

The structural model is not only used to determine whether the part is present; we also create a separate structural model, known as the "negative structural model," to determine whether the board is present. This negative structural model has encoded regions and relationships that represent the board when the part is not present. The benefit of this negative structural model is that it is not sensitive to variations that occur in color within a board or across boards. It is also not sensitive to the movement of board features such as silk-screening or other board markings.

Each type of structural model outputs a probability. In the case of the regular part structural model, the probability indicates how certain the model thinks the part is present. In the case of the "negative" structural model, the probability indicates how certain this model thinks the part is absent. The probabilities of both models can be combined to determine an overall probability of the part as present or absent.

Results

The results of using the structural model, as compared to using other techniques described earlier, are dramatic. Programming time ranges from anywhere from 4 hours to 10 minutes, depending on whether the CAD data is in the right format and the shapes in the library cover the majority of the parts on the board. Usually, the false failure and false accept rate start at about 200 ppm; by adding a few examples to the system, this rate can be easily brought down to 50-5 ppm.

In addition to shorter programming times and low false fail/false accept rates, this approach eliminates the need to constantly tweak machines. From observing the results of the machine as it runs 24/7 in production environments all over the world, we have seen that the machine does not require constant tweaking. Except for an occasional example or set of examples that are outside of the expected range, the machine is able to maintain its ppm level without operator/programmer intervention. In one evaluation over a two-week period, only 20 discrete "debugging" adjustments were made to a program for a board. Once the adjustments were made, the system never exhibited the same problem again. After those adjustments the system maintained its stable performance. At another customer site, after initial programming and minor debugging the system ran over a million boards *without any program adjustments* while maintaining an extremely low false failure/false accept rate.

Conclusion

Vision systems are routinely asked to handle very challenging situations. They are asked to deal flawlessly with normal variation even though they may have been trained on only a few examples of such variation. We have found by many years of observation that variation is especially large in the field of printed circuit boards. A vision system must be able to automatically handle the range of variation that occurs such as part color change, board color change, and paste configuration changes. In this paper, we have described a novel language for representing visual objects based on qualitative rather than quantitative aspects of images. Our approach is motivated by studies of human perception. The approach is featured in the inner workings of a post-place printed circuit board inspection system. The approach allows the AOI system unparalleled ability to automatically handle normal variation in post-place printed circuit boards and to adapt to new variations. As a result the inspection system frees the programmer from constant tweaking of the program while providing exceptional inspection capabilities.

End Notes

[1] M. Bar and S. Ullman, "Spatial context in recognition," *Perception*, Vol. 25, pp. 342-352, 1996.
[2] I. Biederman, "Perceiving real world scenes," *Science,* Vol. 177, pp. 77-80, 1972.
[3] C. Cave and S. Kosslyn, "The role of parts and spatial relations in object identification," *Perception*, Vol. 22, pp. 229-248,1993.

Notes and Acknowledgments:

The ideas described in this paper form the basis of multiple US and international patents and patent applications.