A New AOI Programming and Inspection Paradigm Based On Recent Studies in Neuroscience Reduces the Need for Human Intervention and Improves Program Stability and Quality

Pamela R. Lipson, Ph.D.
Imagen Incorporated and Landrex Technologies, Inc.

Abstract
The hidden cost of optical inspection systems is often in the programming time. In this paper we discuss a new AOI programming and inspection paradigm reduces the need for human intervention and improves program stability and quality. This paradigm is based on the principles of adaptability of the human brain. Our understanding of these principles has helped to facilitate the development of easy-to- program automated optical inspection machines that need little human intervention and yet exhibit a consistently high level of performance over extended periods.

Introduction
Everyday, we work with easy-to-program machines. They are called “humans.” On countless assembly lines across the world, for staggeringly varied industries, humans serve as quality inspectors. Their supervisors “program” them by simply showing them a couple of good samples and describing some of the faults to watch out for. That’s it. The human inspectors intelligently extrapolate the “programming” to generalize over other instances and are able to flag various kinds of problem instances even without having been shown their exact examples beforehand. They are rarely a burden on the supervisor and even get better with experience. Truly, in terms of ease of programming, humans are hard to beat. However, they come with some downsides. They are expensive. They are getting to be too slow to keep up with the frenetic beat rates of modern assembly lines. Their eyes are simply not good enough to match the output of ultra-fine manufacturing processes. They have difficulty keeping focused on very repetitive tasks. And, they need frequent tea/coffee/donut-breaks. The writing is on the wall. Human inspectors are giving way to automated optical inspection (AOI) systems on electronics lines the world over.

Unfortunately, programming these new systems is often a task so fraught with frustrations and tedium that managers are reluctant to adopt them. Programming time is one of the major hidden costs in deploying and maintaining an AOI system. AOI programs typically need constant modification to accommodate acceptable variation in the visual appearance of the PCB, including paste, components, and solder joints. Figure 1 illustrates some of these common variations.

Typically, a programmer must either show the AOI system the new appearances or modify the lights, algorithms, and camera views. A full-time programmer may be needed if conditions change rapidly, the system algorithms are particularly rigid, or changes in one part of the program cause unintentional deterioration in another part. An inspection program that exhibits low false failures and low false accepts at one moment may change its performance drastically and adversely over time, unless the program is constantly monitored and “tweaked” to keep its performance acceptable in the face of change and variation. All of this often adds up to a machine that is so difficult to program that is more trouble than it is worth. As a result, many such machines lie dormant shortly after purchase. What can be done to remedy this problem?

Clearly, if the goal is to make an easy-to-program inspection system, there is no better place to look than the human brain. What principles of data analysis does it use to achieve its generalization abilities? If we can answer this, we might be able to design a system that demands little and delivers a lot in inspection performance.

In this paper, we make a beginning towards this goal. Specifically, we enumerate five principles that are integral to brain function and its easy programmability. We describe how each of these principles translates into strategies for inspection system design. An inspection machine that embodied most, if not all, of these principles should be vastly easier to program, and better in performance, than the current state of the art. We illustrate how each of these may be used in a printed circuit board inspection system and at the end of the paper describe how well these concepts work in practice.

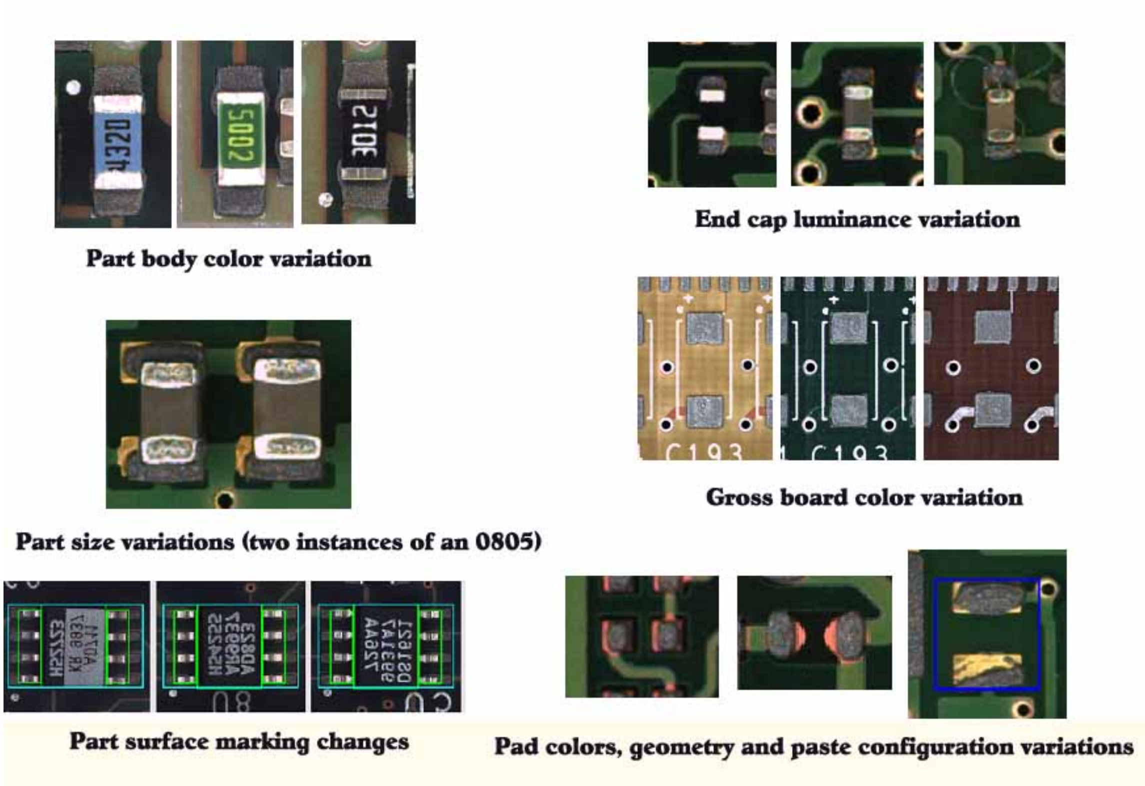

Figure 1 - Printed circuit boards exhibit a large amount of acceptable variability in how the boards, paste, parts, and solder joints look. It is difficult for an optical inspection system to tell a good board from a bad board when all the elements of a board can change their appearances.

Principle 1: It Takes a Village

The human brain is a modular information-processing system. Sensory information is sent to multiple areas of the cortex for analysis. Different areas of the brain process the same information in different ways, and then the results are pooled either in higher-order areas and/or via lateral connections (Zeki, 1993). One way to view this style of functioning is to compare it to a community of collaborating “agents” (Minsky, 1988), where each agent has its own area of expertise or domain. Individually, these domains might be fairly limited, and the agents might not even be the most powerful experts in those domains, but their collaborative interaction delivers superior results to the entire community.

Machine-learning scientists refer to this synergistic collaboration as “boosting.” Boosting is a general method for improving the accuracy of any given learning algorithm by combining multiple, often weak, classifiers to produce one that is surprisingly powerful (Drucker et al., 1994). This approach has several advantages. First, it simplifies system design because it eliminates the need to create a single, monolithic, all-knowing expert. Second, this approach enables the system to be far more adaptive. Agents that are confident participate in the decision, whereas less-confident agents’ opinions are either weighted less or dropped out of the ultimate decision. Third, a system based on a collection of experts is inherently more fault tolerant than a single one due to the built-in redundancy. Finally, it is easy to modify such a system to adapt to changing inspection requirements for PCB assemblies. Experts can be added or deleted without modifying the architecture of the system.

In the particular context of AOI, each member of the community corresponds to a software agent that performs a single type of inspection. Figure 2 shows some examples of these agents, whose decisions can be combined to determine if a part is present or absent. For instance, three agents might split up inspection duties this way, each specializing in one task: 1) examine the appearance of the object under test, 2) look at a component’s edges or transitions with the board, and 3) look for the board that is supposed to be underneath the part.

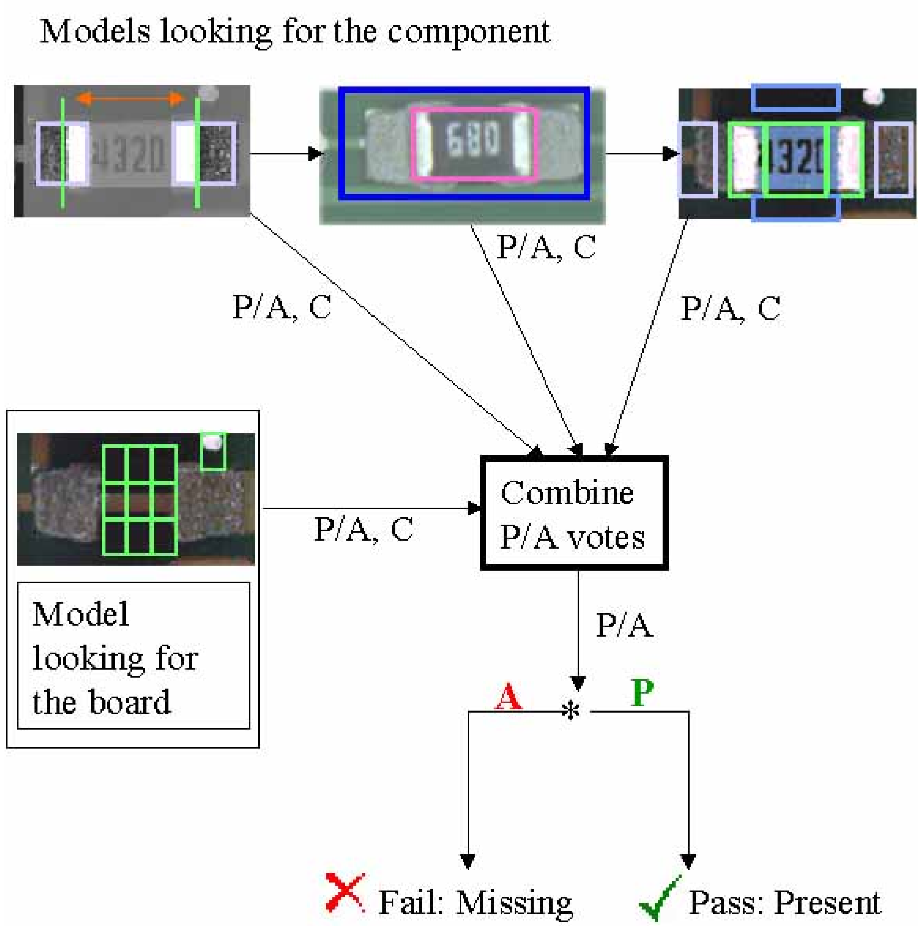

Figure 2 - Combining the use of multiple models or agents to determine if a part is present or absent. In this particular case, we have separate agents tuned to look for the part's appearance, the occlusion of the pads due to the presence of the endcaps, and the part's structure, (top). We can even add agents that look for the absence of the part or the presence of the board (left). Each agent makes an independent judgment about whether the part is present or absent and also provides a measure of confidence about that decision. By combining these independent votes and confidence measures, the system can make a final decision.

Principle 2: Learning Never Stops

The brain is always involved in learning. It does not switch from an exclusively "learning" phase to an exclusively "doing" phase. It learns from every new experience and constantly updates its world model, including modifying the model to define tolerances on allowable variations. Not only does this strategy create a more robust and dynamically adaptable internal representation, it also allows an assessment of the trajectory of changes in the world, i.e. the process of change. Knowing this trajectory allows the brain to anticipate variations in future inputs.

Translating this strategy to the AOI domain, agents can, during each inspection, dynamically sample the world and update its model of the visual aspects of the parts and board-under-test. Instead of relying on learned or programmed-in conditions that may be no longer valid, the agent uses sampling to handle changing conditions on an "inspection by inspection" basis. Having access to the continuous data acquired during each inspection has another advantage: it can enable the system to look for trends that signal changing conditions and thus perform process control.

One of the main sources of visual variation on printed circuit boards is color. Figure 3 shows an example of how an AOI system can adjust dynamically to changing board colors.

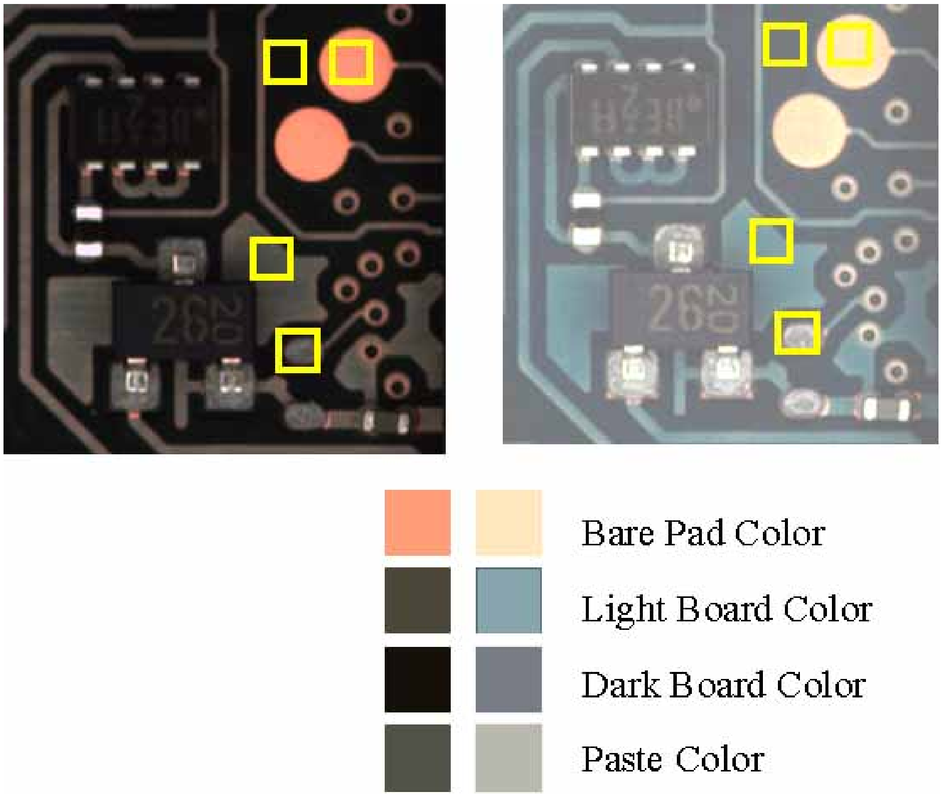

Figure 3 - A printed circuit board can be simply represented by 4-6 colors – bare pad, mask over copper (light board), mask over substrate (dark board), paste, silk screening, and flux. Rarely do any of these colors stay the same across boards. Thus, instead of optimizing for one board instance, we can dynamically sample these colors on every board by looking at predetermined locations that should contain the colors. This color data can be input into the algorithms that require this data; as a result, the correct color palette for each board in the inspection process is available to the algorithms, enabling them to handle the variation in color palette from board to board. This figure shows the color palette for two boards of the same type that differ in their colors.

Principle 3: Don't Ignore the Forest When Looking at a Tree

Local information is often ambiguous. When we look at an impressionistic painting, it is not clear what any particular blob stands for, until we expand our window of analysis to include context. A blob of green paint looks like a green smudge until we examine it as part of its context and then easily identify it as a tree. This observation applies not just to paintings, but also to all manners of real-world image analysis. The brain relies heavily on context in order to interpret images accurately and discriminate relevant visual information from noise and other imaging imperfections (Cox et al, 2004).

Figure 1 shows there is a great deal of variation in appearance on a PCB. However, there is even a greater level of variation at the very local level. Shiny parts have range of luminances and colors and tend to reflect the image-capture system. Colors change in the mask across the board and over different substrates. These local level variations can create confusion for an AOI system.

These problems can be addressed by having the AOI system consider not just the local data but also the context within which the local data appears. The contextual cues help constrain and make sense out of the information at the local level. For instance, an 0201 when viewed on its own appears to be a random assortment of different colored pixels. However, when it is examined in the context of the pads and board around and the expected board underneath it, then it can be correctly classified. The same 0201 can be analyzed in terms of how it looks relative to other 0201s to determine good from bad. Figure 4 shows an example of analyzing an 0201's orientation by referring to its context.

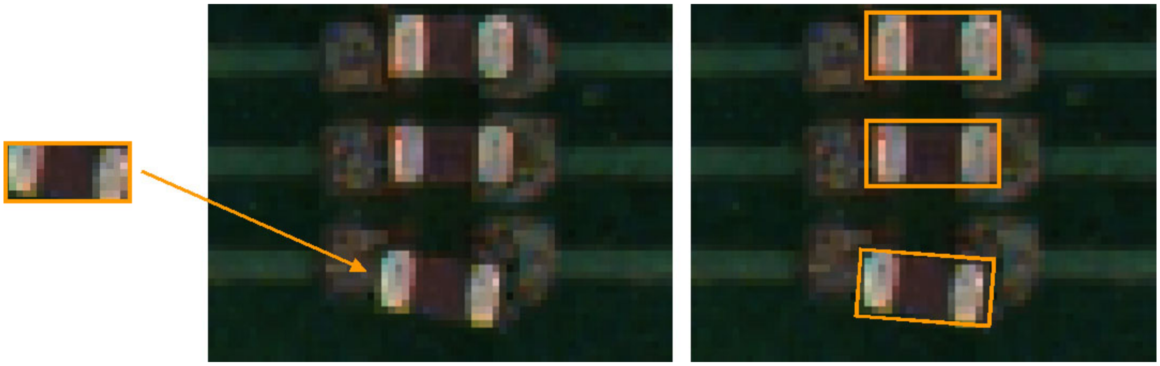

Figure 4 - Looking at only local regions of an image can result in misleading interpretations. Often these pixels are not sufficiently constrained for an image-processing system to make an accurate diagnosis. This figure shows an 0201 on the left. Without context, this part looks correctly oriented. However, when looking at it with respect to other properly placed components and by taking into account where the pads and board pixels should be when the part is correctly placed, it becomes apparent that this part is actually rotated.

Principle 4: Qualitative Complements Quantitative

In an inspection and measurement domain, it is tempting to make each image measurement as information rich and as detailed as possible. Often such systems interpret the images literally by taking viewing the absolute color or luminance of each pixel as an important piece of information. Others look at fine details such as the precise transition between an object and its background or interior transitions between luminance or color regions. These transitions are known as edges. It is appealing for system designers to try to reduce the world to these transitions and interpret them based on their precise location, size, orientation, and color. However, new studies on how humans process images show something quite different. Neurophysiologic studies of cells in the visual cortex, the seeing part of the brain, suggest that many neurons might encode image information far more coarsely. A single neuron might not actually encode precise information such as RGB values or size and orientation of edges. In fact, we know that humans are terrible at providing information about absolute attributes of world objects. Instead researchers have found that these neurons might be responsible for indicating simply *the direction of contrast* over large regions, such as whether the left side is brighter or darker than the right.

The question is why would the brain choose this coarse relative metric over something more quantitative and precise? We now understand that this is an optimal strategy for understanding and tolerating acceptable variations in how things look. Relative information is more adaptable to image changes than metric measurements and, therefore, more robust to imaging variations. A face seen under different illumination conditions or from different view points has the same relative qualities, but very different absolute qualities (Sinha, 2002). The bottom line is that it is the same face irrespective of these conditions, and the brain needs to recognize it as such.

Understanding this principle can be very helpful to an AOI system designer. Applying this principle means we can create representations for elements on a PCB board (either faults or good examples) that are based on qualitative relations between regions. Because this representation gets at the underlying structure of the element, it is likely to be more invariant than those that include more detailed information. (Figure 5 shows how a qualitative representation can be used to determine part presence/absence.) Thus, this is one instance where disregarding information yields positive dividends. We should note that some quantitative information may be important, especially when color encodes a value of a component. However, we can combine both qualitative and quantitative representations by having agents that attend to each type of information as described in principle 1.

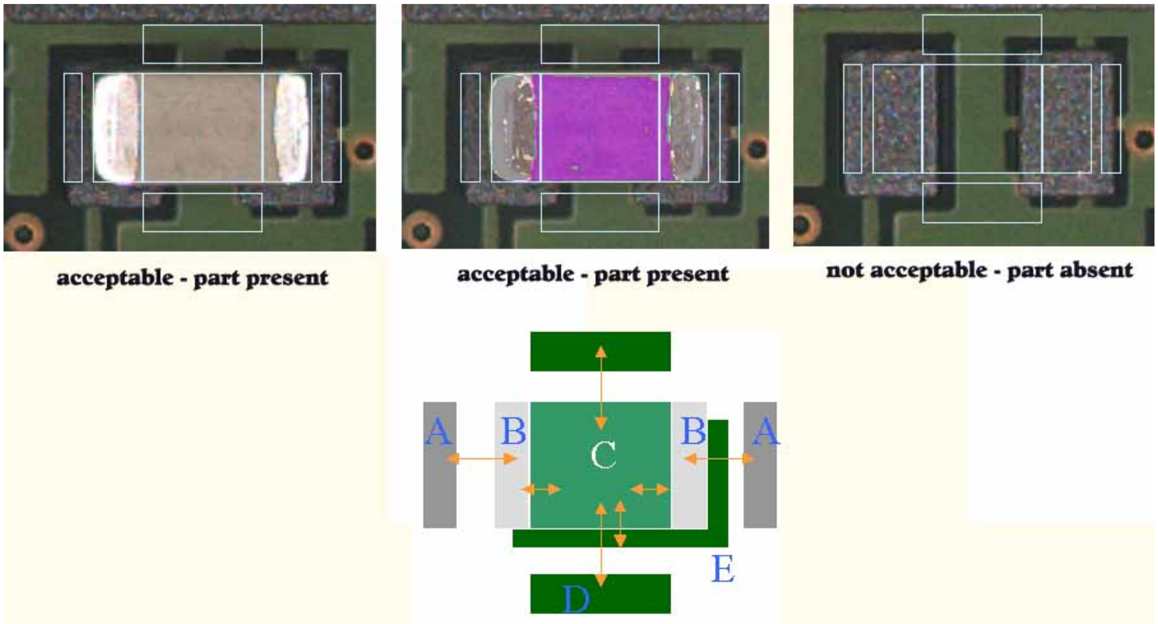

Figure 5 - We can represent a capacitor by a set of regions and their relative relationships. As a result, we can now encode a large range of acceptable variations while still distinguishing unacceptable variation. For instance, this figure shows a qualitative model of a capacitor on a pre-reflow board (bottom) which includes regions for the endcaps, body, paste, background around the part, and background underneath the part. The representation encodes the generalities that the endcaps are often different from the body, the body is often different from the endcaps, the body is often different from the background (on the sides or below the part), and the endcaps are often brighter than the paste. As a result, this simple representation is satisfied when the endcaps are bright or dark, the body is brown or pink, and the board color changes. However, when the part is absent, the representation fails for the reasons that the endcaps regions are not brighter than the paste and the body is not different from the background, leading to a judgment of unacceptable variation.

Principle 5: Prior Knowledge Bootstraps and Experience Refines

One of the big questions in neuroscience concerns how the brain achieves its impressive face processing abilities so rapidly. We believe that even at birth or shortly after a human infant is proficient at detecting faces. The ability to detect and even recognize faces even in complex scenes improves quickly afterwards. What underlies such rapid learning? A prominent proposal suggests that the brain does not start out with a blank slate (Morton and Johnson, 1991). Evolution has equipped even the brain of a newborn with a coarse representation of a human face. This representation allows and encourages the infant to look for patterns that are more likely to be faces and to learn specific faces. This is a beautiful example of how the brain uses prior knowledge to bootstrap its learning, and real-world experience to refine its concepts. The use of domain-specific knowledge right at the outset not only gives the brain a significant level of functionality right away, it also greatly expedites learning through experience.

We can use the same principle in the AOI domain by giving our agents built-in knowledge about PCB inspection and then combining this knowledge with learned data about the board-under-test. The use of application-specific knowledge provides a base level of understanding of a PCB board from the start. This is helpful to the programmer because the knowledge is already built into the system. If the system is shown a live example of the board, parts, joints, etc. to be inspected, the system can tailor its built-in biases for the board in question. Finally, with the addition of dynamic sampling, the system can make an informed decision based on different layers of knowledge (dynamic, learned, built-in) regarding how to handle unexpected or highly varying conditions. We would expect such a system provide consistently good inspections over long periods of time without significant human modification of the program. In addition, if the system were to make a mistake, one or two labeled examples would allow it to adjust its decision- making ability fairly easily. Figure 6 shows an example how an AOI system might combine built-in knowledge and learned knowledge in order to generate a high-quality inspection program quickly.

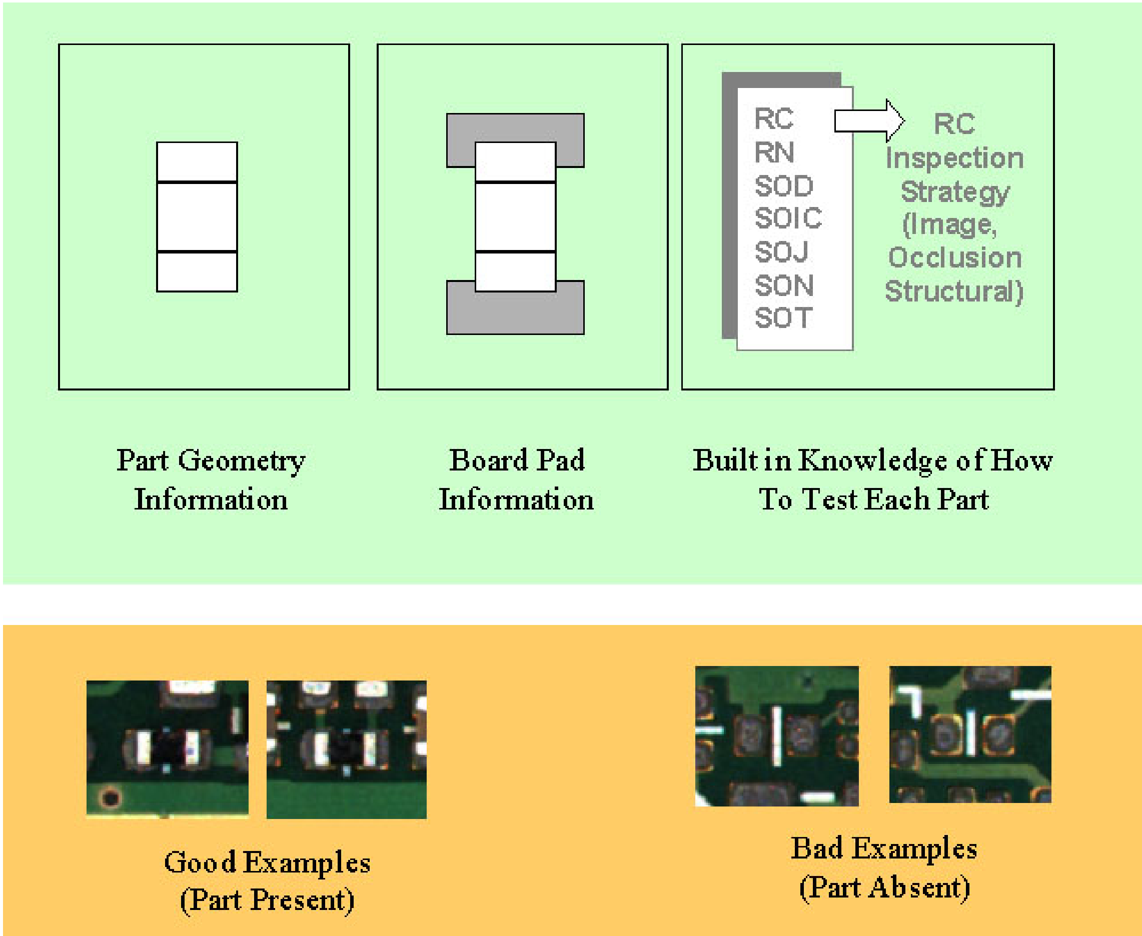

Figure 6 - An optical inspection system can be given an intrinsic ability to inspect parts on a printed circuit board, without ever having seen any real-world examples, by building in information such as part geometry, board pad data, and strategies for inspecting each known part type. By integrating built-in knowledge with real examples from the board-under-test, the inspection system can generate very quickly an inspection program that achieves excellent results after only a few sample boards. Additionally, good or bad examples can be fed back into the system at any time to modify the system's behavior.

Real Results: The Principles Applied to a Pre-reflow PCB Inspection System

The Landrex Optima 7200 ™ optical pre-reflow inspection system and process-control tool utilizes each of these principles. The system was born out of collaboration between MIT brain scientists, computer scientists and engineers with the addition of novel technology developed at the Artificial Intelligence Laboratory and licensed by Imagen Incorporated. The goal of the collaboration was to use these principles of brain function to make a machine that would be easy to use. The system has been tested extensively in different geographical locations (US, Europe, Asia), in low and high volume manufacturers and over a variety of product types. It achieves its ease of use through the following five core features, which correspond to the principles enumerated above:

- combines multiple types of image processing in a dynamic way
- dynamically samples visual attributes of the boards, parts, and paste
- models both the parts *and* the board underneath and around the parts
- uses the patented novel technique of Configural Recognition ™ which is qualitative in nature
- incorporates prior knowledge of the structure of PCB boards and parts.

Using these principles in combination improved performance, in terms of false accepts/false rejects, over traditional methods such as image correlation and edge matching by three orders of magnitude. In terms of programming time, a good test program can be created in as little 10 minutes with performance of 10-100 ppm. Real-work examples bear out the ease of use of the system. During a short-term evaluation, where 5 million components were inspected, the program required only 15 discrete and unique adjustments. Once the adjustments were made, there was no reoccurrence of the issues. In a longer-term study, a customer inspected approximately 500,000 boards over a period of 2.5 months. In this case, after the initial creation of the program and a short debugging session, the program ran *without* user intervention and was able to *consistently and reliably* inspect the boards. The injection of the system in the process improved product quality by an order of magnitude.

The Wave of the Future: Test Systems Based on Brain Principles

One might ask why these and other brain-inspired principles haven't previously been incorporated into AOI systems. The first reason is that up until recently, the building blocks of the AOI systems were not sufficient to attempt this kind of processing. More specifically, the capabilities of hardware, computers, camera, memory, and lighting were so rudimentary that only simple images could be captured and only simple image processing algorithms could be employed. In the past 5-8 years we have seen a revolution in the hardware arena, which has resulted in faster computers, higher-resolution cameras, color capacity, cheaper memory, and the introduction of the white LED. Thus, AOI machines now have the capacity to "see better" and the basic platform to analyze what the systems are seeing.

A similar hardware revolution has happened in brain science. Hardware advances from the last 30 years have produced MRI and functional MRI that allow us to image the brain and to view its activation patterns at a detailed level. As a result, our knowledge of how the brain recognizes objects has increased tremendously.

We are now well poised to exploit the synergies between hardware advances in the industrial domain and the basic understanding of the brain in the scientific domain. What I have presented here is one example of this synergy related to making machines easier to program. We can expect that many more such insights for the design of robust and effective test systems will flow from a deep understanding of brain science in the coming years.

Bibliography

Cox, D., Meyers, E. and Sinha, P. (2004). Contextually evoked object-specific responses in human visual cortex. Science, Vol. 303, No. 5667, pp 115-117.

Drucker, H., Cortes, C., Jackel, L.D., LeCun, Y., and Vapnik, V. (1994). Boosting and other ensemble methods. *Neural Computation*, 6(6):1289-1301.

Minsky, M. (1988). The Society of Mind. Simon & Schuster.

Morton, J., & Johnson, MH (1991). CONSPEC and CONLERN: A two-process theory of infant face recognition. Psychological Review, 98, 164-181.

Sinha, P. (2002). Qualitative representations for recognition. In Lecture Notes in Computer Science, Springer-Verlag, LNCS 2525, pp 249-262.

Zeki, S. (1993). A Vision of the Brain. Blackwell Publishing.

From brain science to AOI

A new AOI programming and inspection paradigm reduces the need for human intervention and improves program stability and quality.

Pamela Lipson

Imagen Incorporated and Landrex Technologies

Post-Place/Pre-Reflow

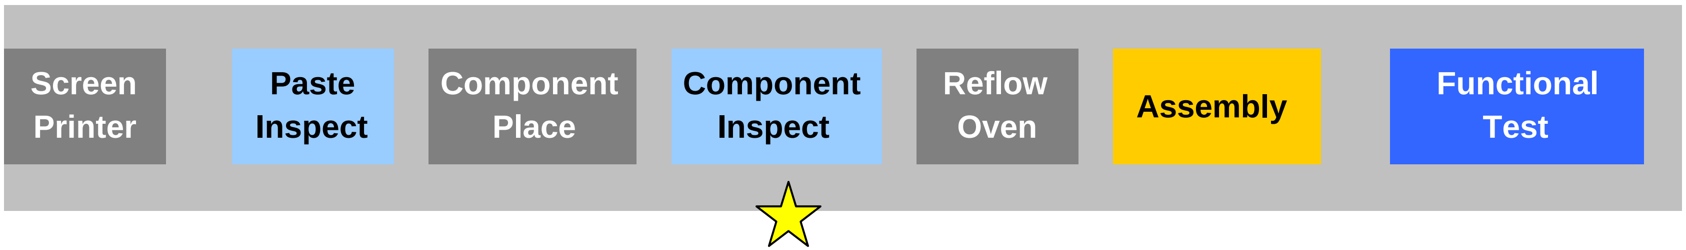

In this talk we give examples of post-place, pre-reflow defects.

However, all principles can apply to paste inspection, post-reflow, x-ray etc.

Defect Detection & Process Monitoring

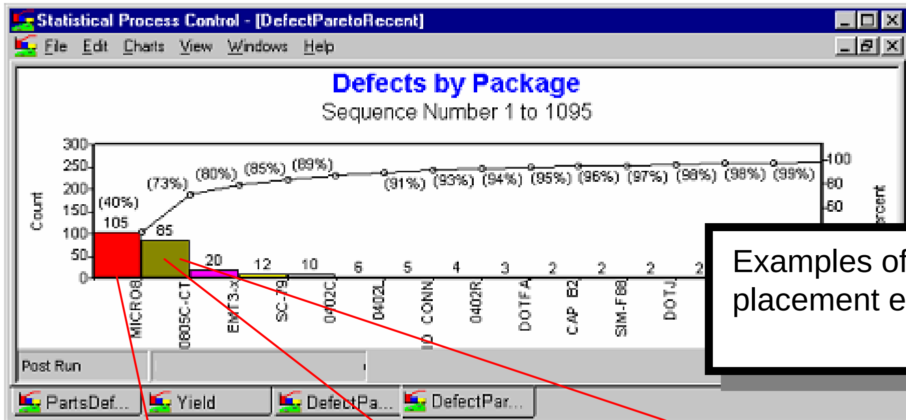

Examples of 'random' placement errors

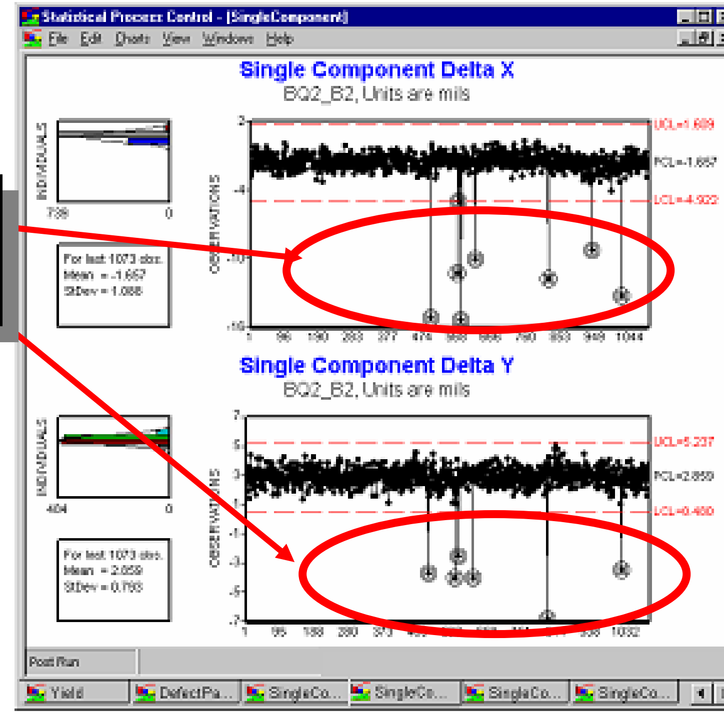

U5 Mic

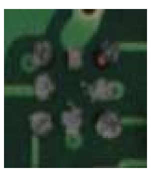

Missing

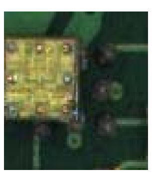

Upside Down

Billboard

0805C

Billboard

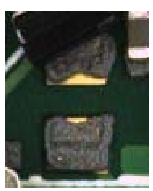

Missing

SOT

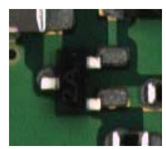

Off Pad

Problem:

- **Parts not cleanly placed by nozzles**
- **Parts being lifted, dropped**

Key AOI challenge:
Handling normal process variation

Part colors vary

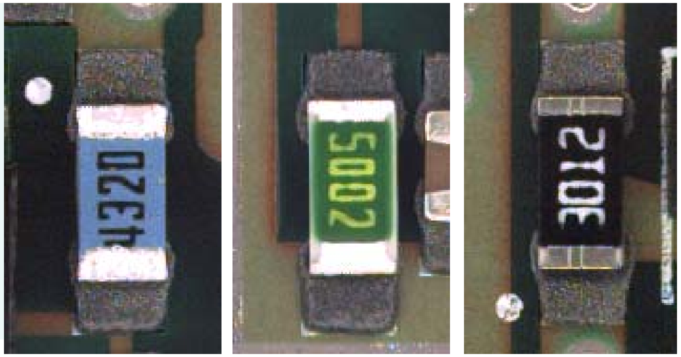

Board colors vary

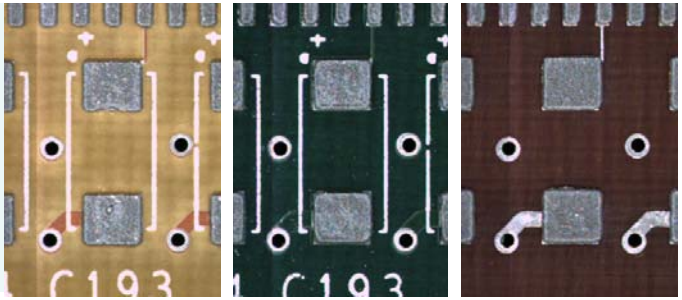

Part markings vary

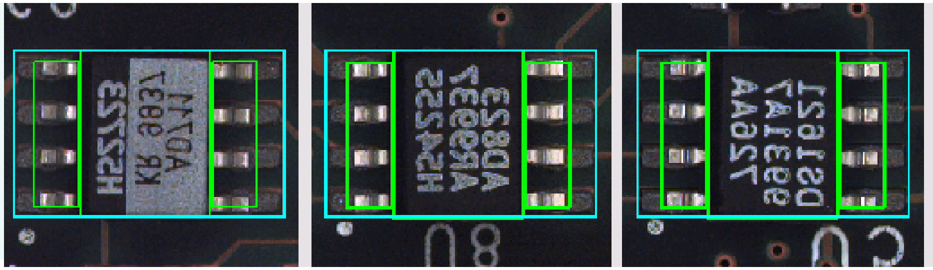

Key AOI challenge:
Handling normal process variation

End cap luminances vary

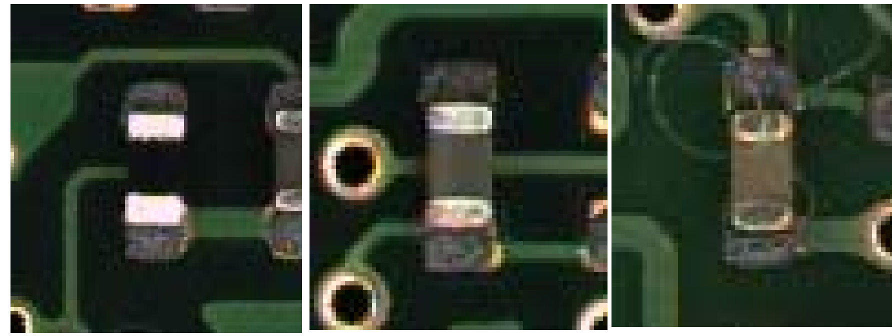

Part sizes vary

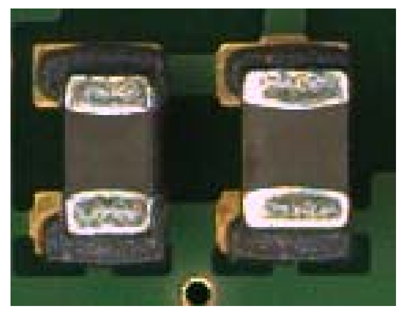

Pad colors and paste configuration vary

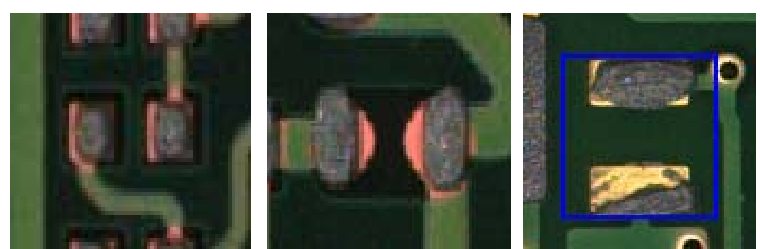

A typical AOI process:

A perpetual sequence of ‘tweaks’

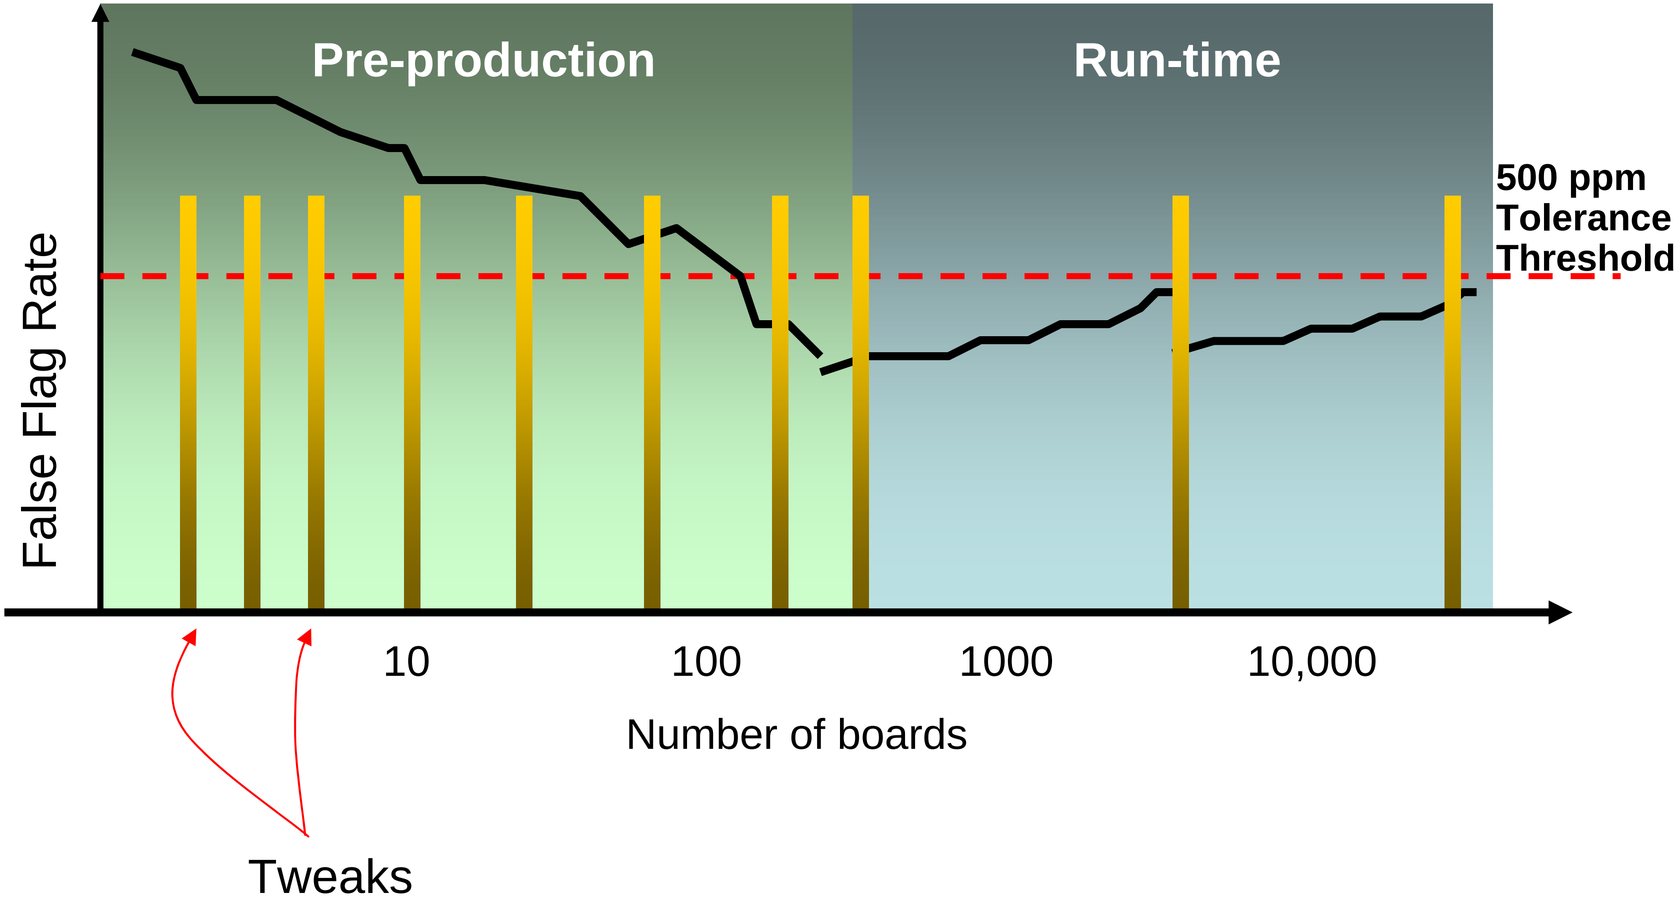

Why is this so?

Why is this so?

Because common pattern-matching techniques, such as correlation and edge-matching etc. are brittle

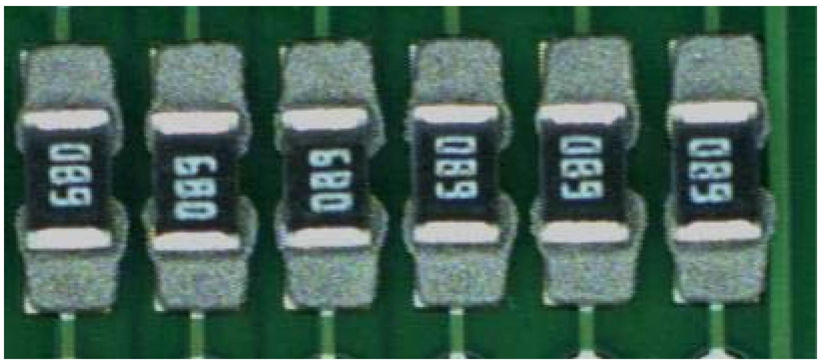

Only accept black parts

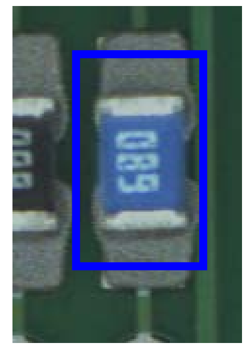

Rejects other colored parts

False positives

False negatives

Why is this so?

Because common pattern-matching techniques, such as correlation and edge-matching etc. are brittle

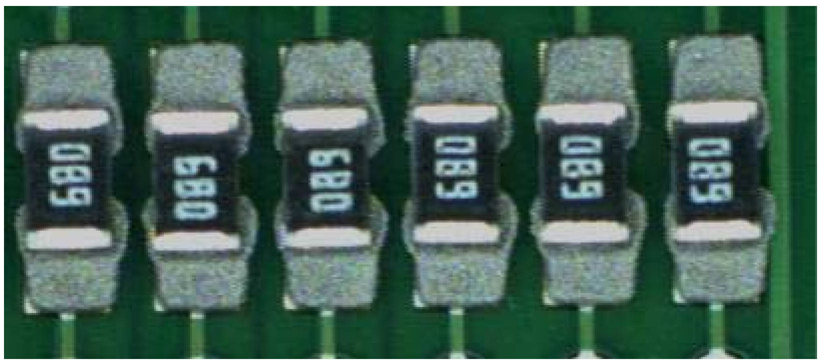

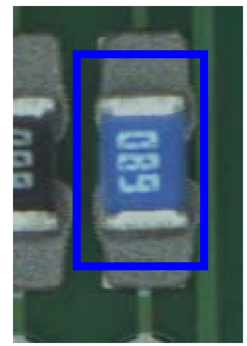

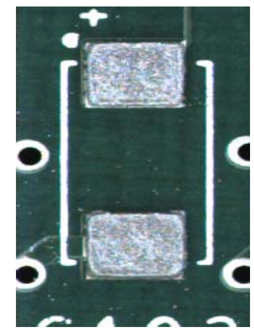

Accepts multiple colored parts, may also accept the board when the part is absent

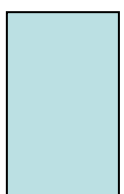

False positives

False negatives

How can we devise a less brittle, more robust matching strategy?

What can we turn to for ideas?

How can we devise a less brittle, more robust matching strategy?

What can we turn to for ideas?

The very best example of a system capable of robust visual processing is the human brain

The most impressive aspect of the brain's visual processing:
Ability to generalize across variations

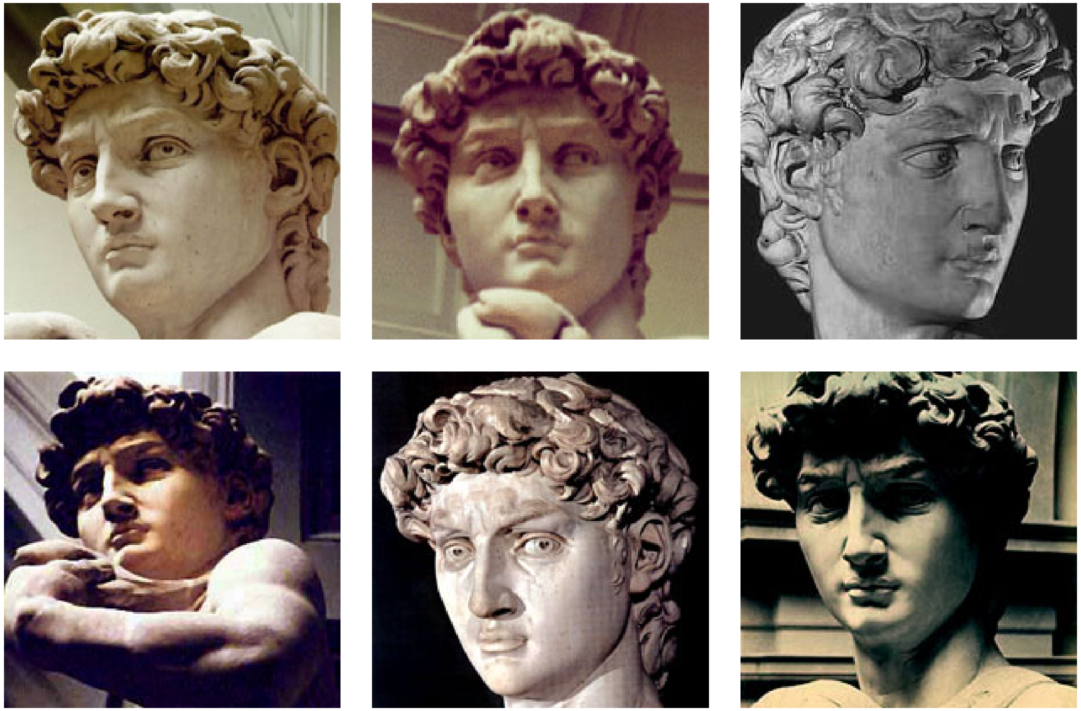

If we can get clues about how the brain accomplishes this, we might be able to design better AOI systems.

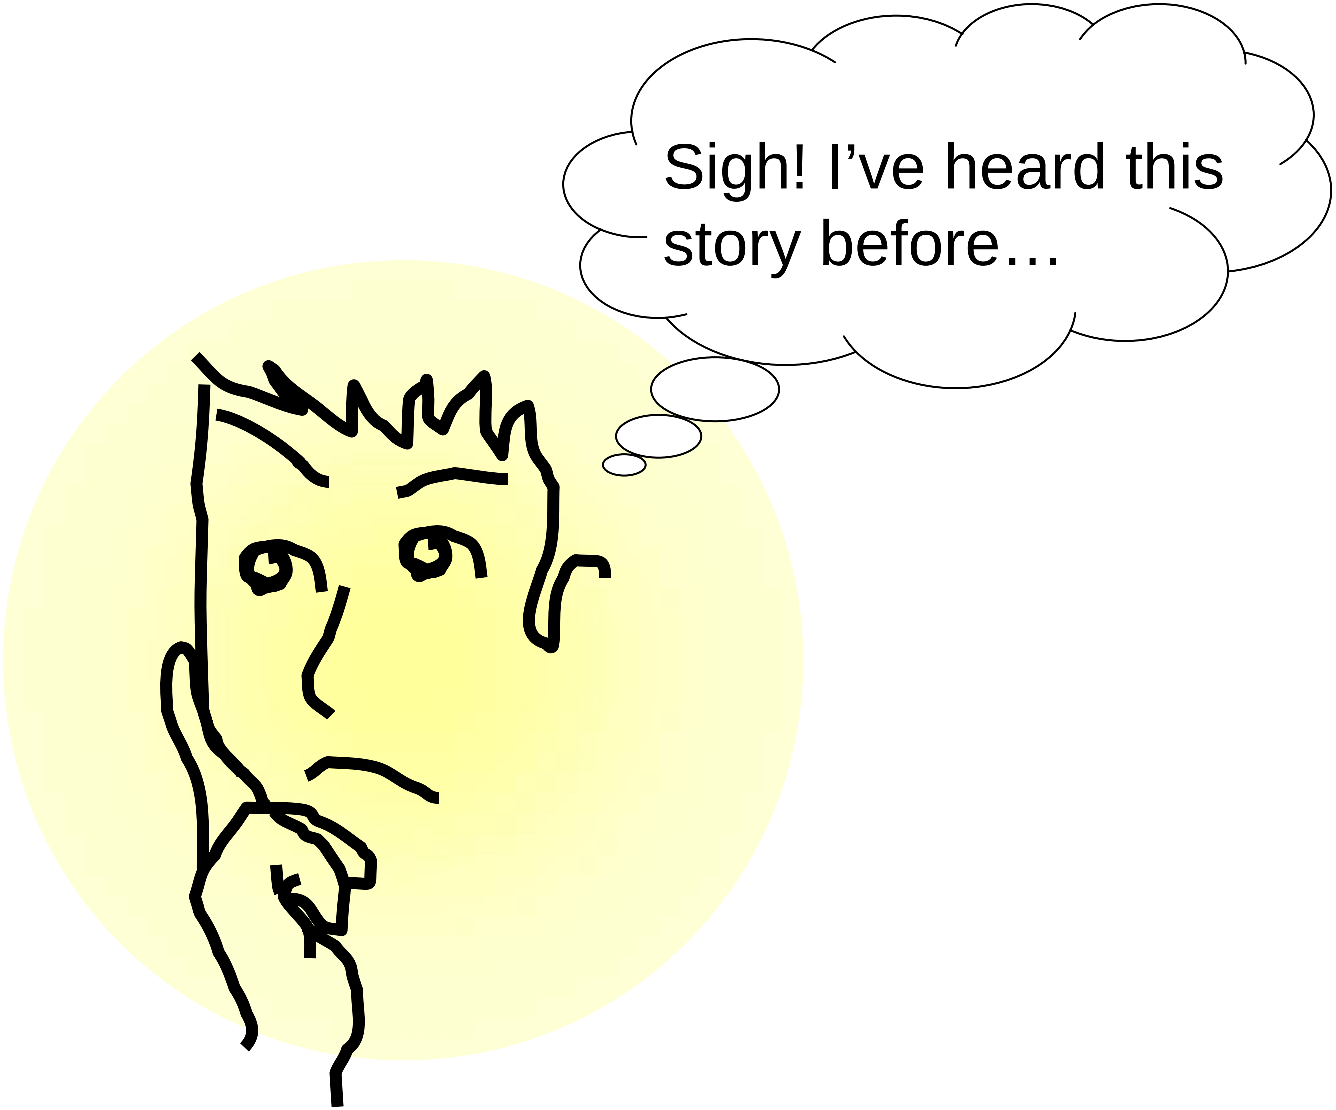

The story so far:
Big dreams, small results!

Prior theories have either aimed too high...

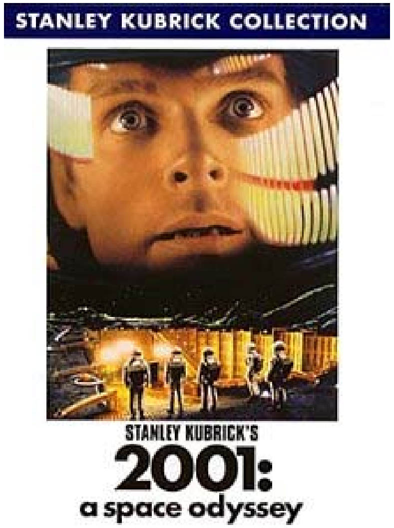

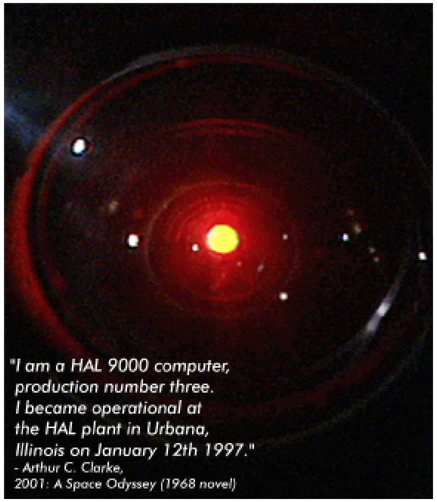

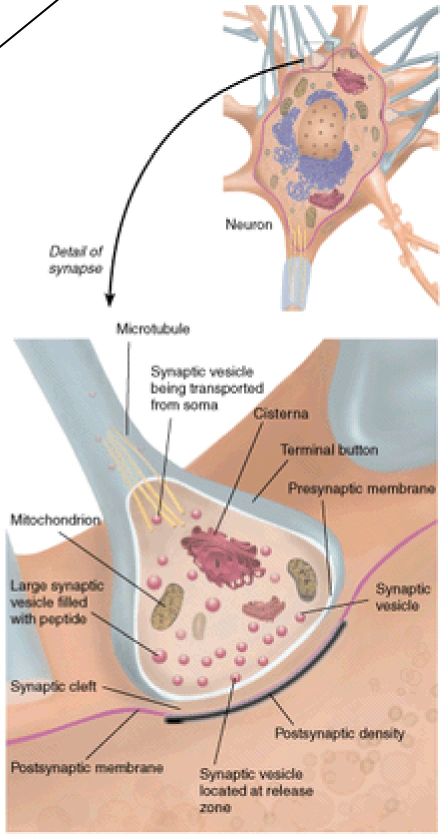

... or too low

What is different in this talk?

We use mid-level principles of the brain

1. It takes a Village
2. Learning never stops
3. Don’t ignore the forest when looking for the trees
4. Think relative
5. Incorporate prior Knowledge & refine with experience

1. It takes a Village

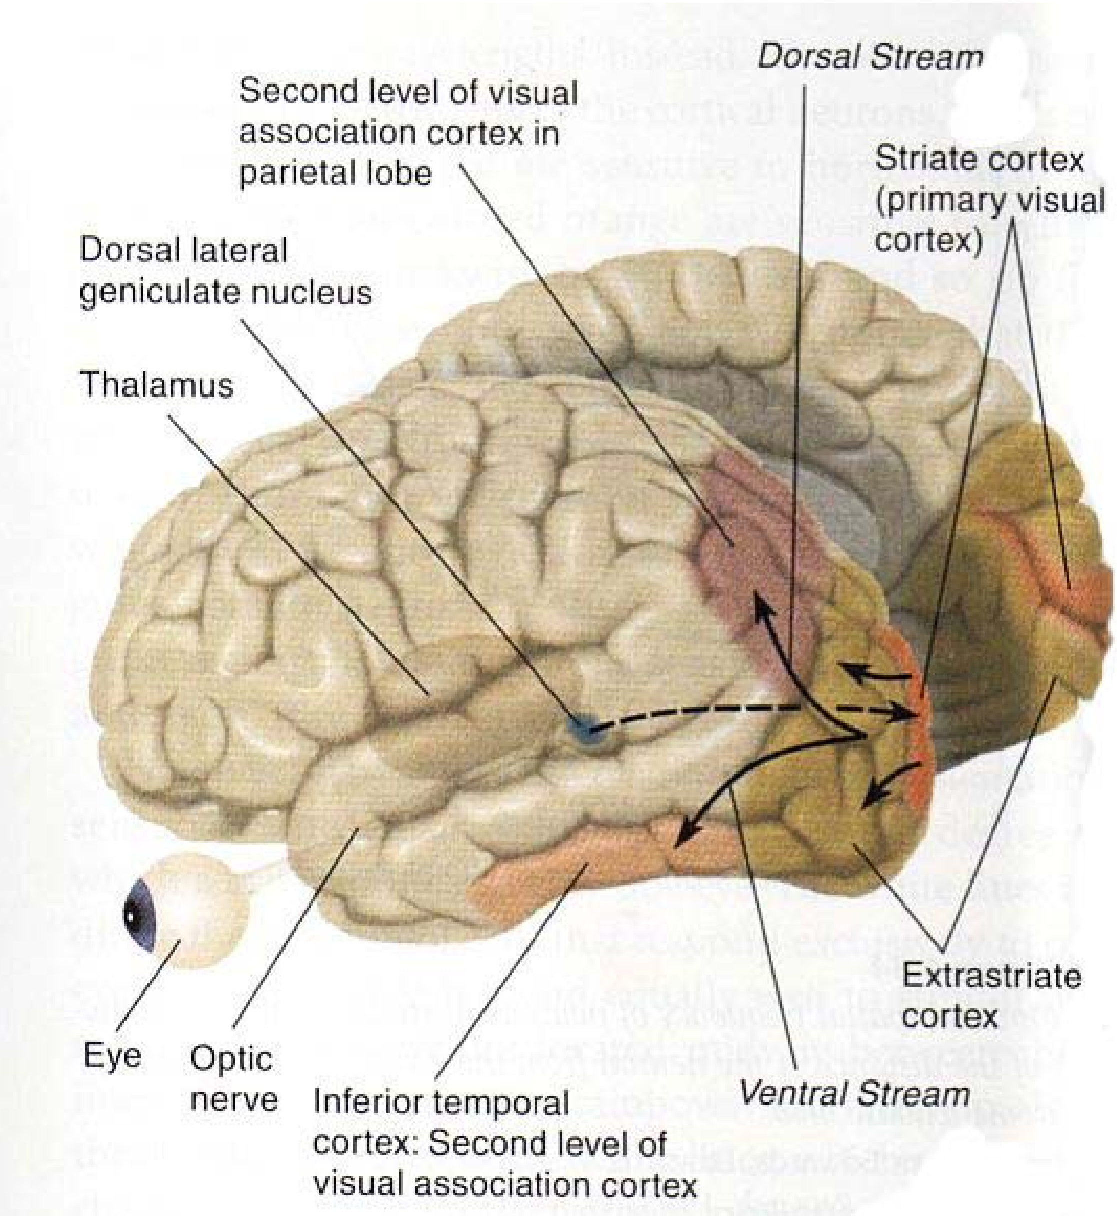

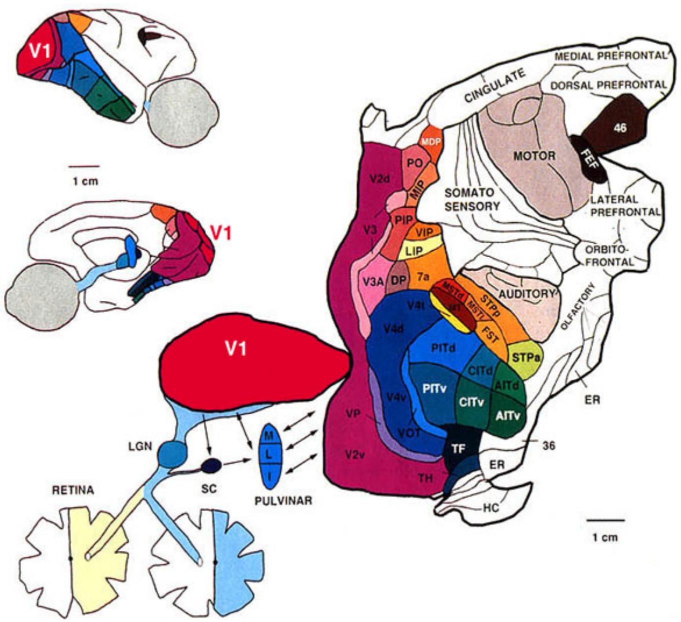

The visual part of the brain has multiple partitions, each concerned with a particular kind of image analysis.

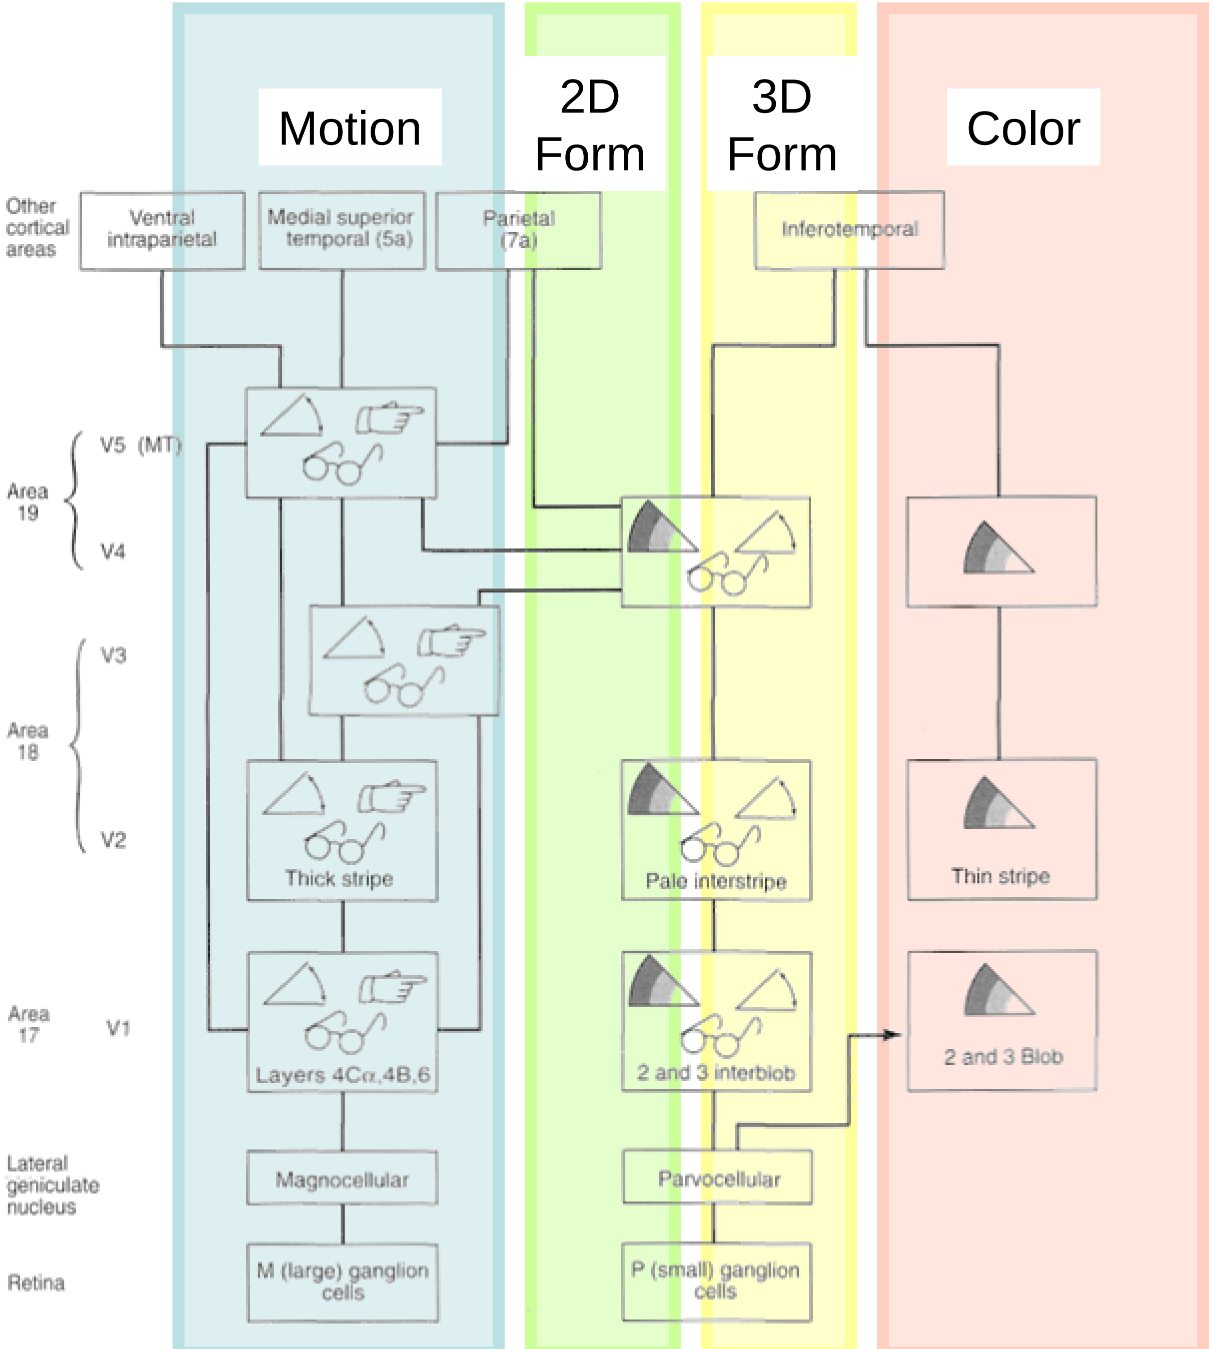

Thus, every visual stimulus is multiply analyzed and represented.

This redundancy enhances robustness since deficiencies in one stream can be compensated by another.

The brain's visual hardware is really a 'village' of interacting processors.

Applying “It takes a Village” to AOI

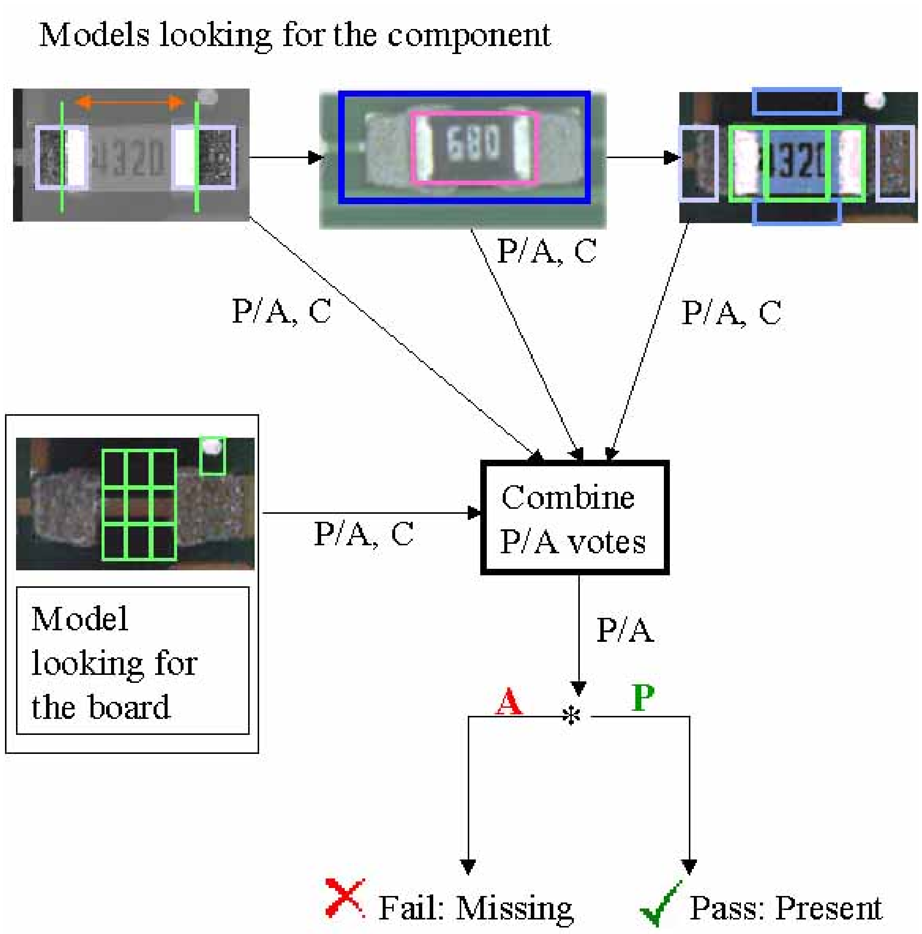

2. Learning never stops

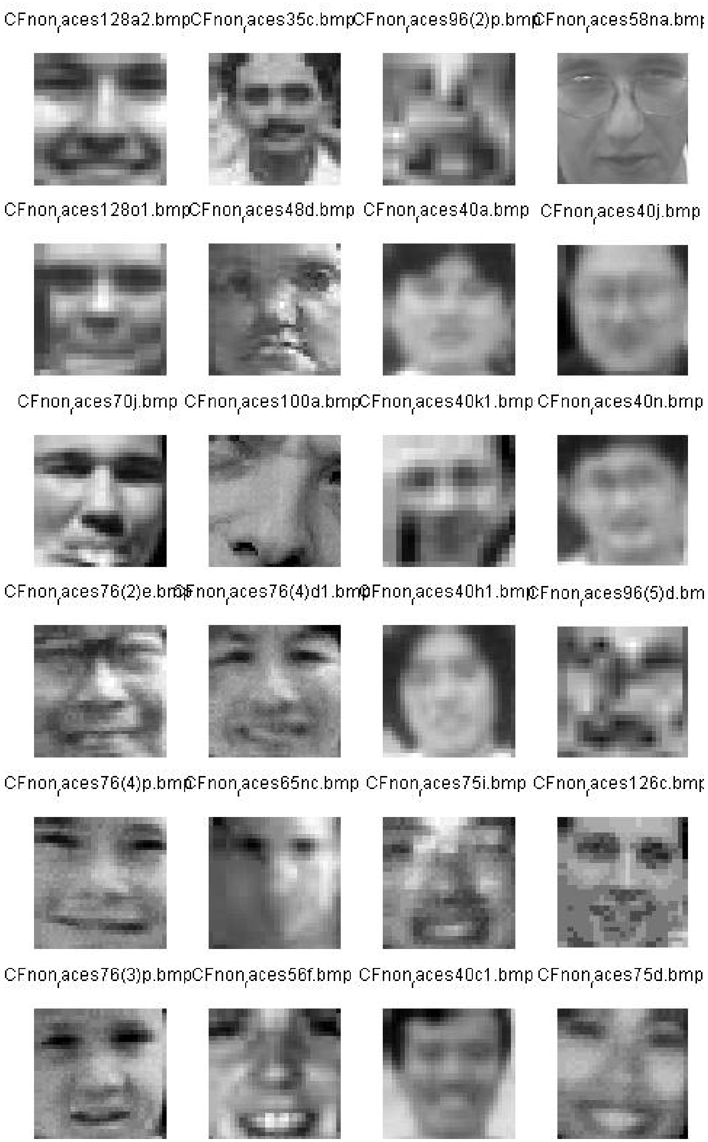

Faces

Non-faces

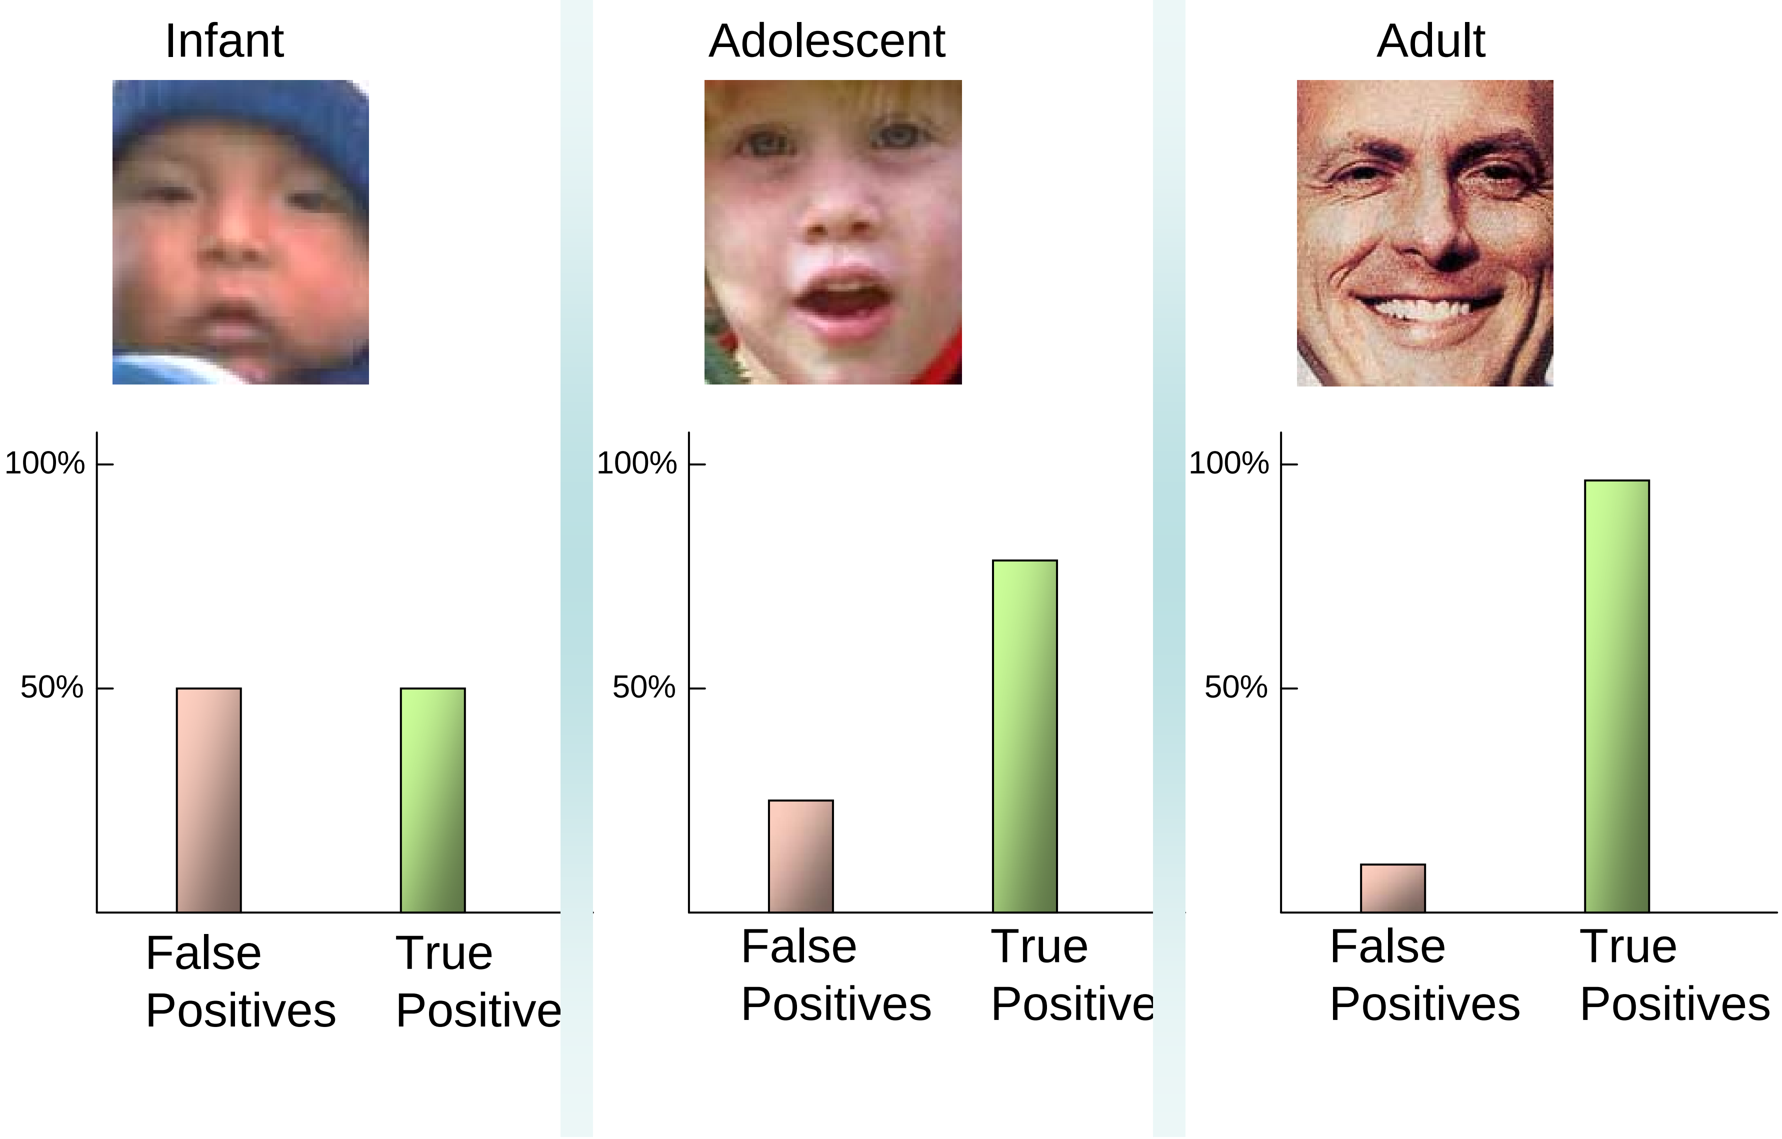

Performance keeps improving through on-line learning throughout the life-span

Bouvrie and Sinha, 2007, Neurocomputing (in press)
Sinha, Project Prakash (in progress)

Applying “Learning never stops” to AOI

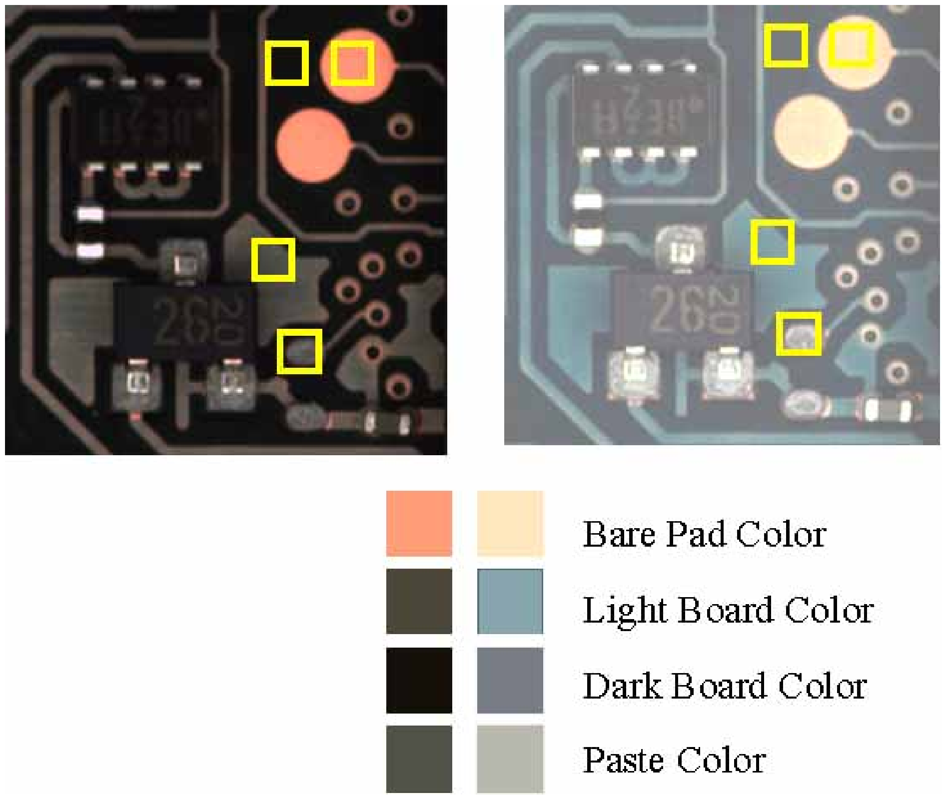

3. Don't ignore the forest when looking for the trees

Using contextual cues to improve diagnostics

a

* Cox, Meyers and Sinha, 2004, Science

3. Don’t ignore the forest when looking for the trees

Using contextual cues to improve diagnostics

** Cox, Meyers and Sinha, 2004, Science*

Contextually driven activations are spatially coincident with intrinsically driven ones

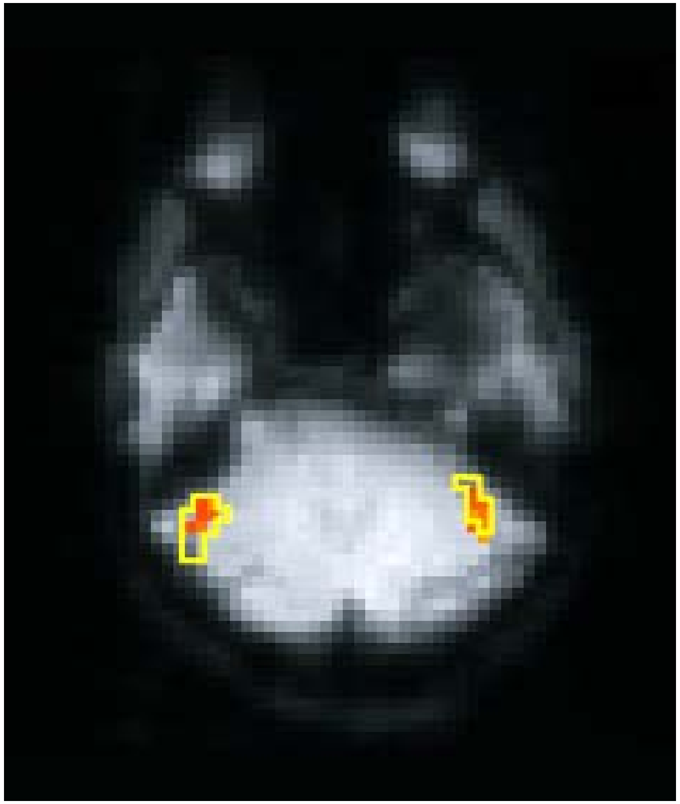

Activation with intrinsically defined faces

Activation with contextually defined faces

Cox, Meyers and Sinha, 2004, Science

Applying "Don't ignore the forest when looking for the trees" to AOI

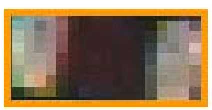

Applying “Don’t ignore the forest when looking for the trees” to AOI

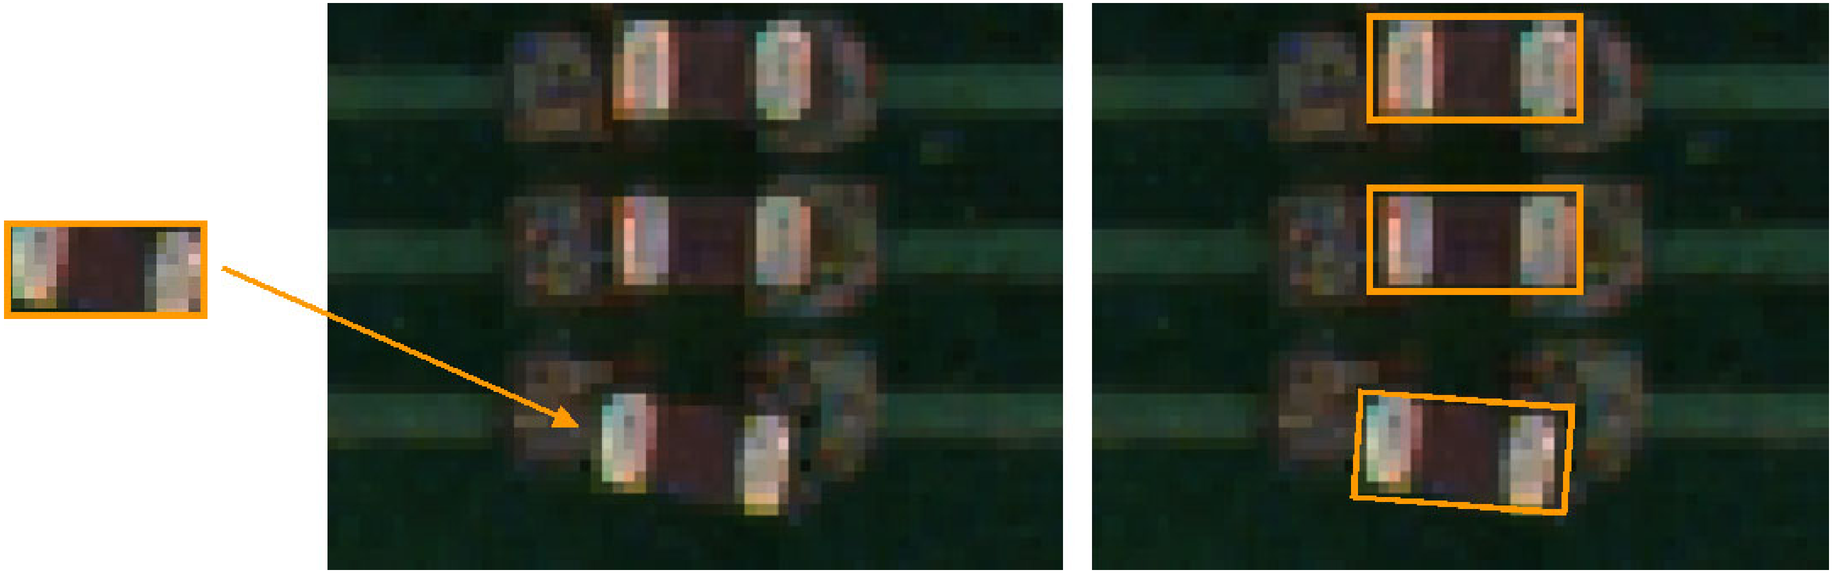

4. Qualitative complements quantitative

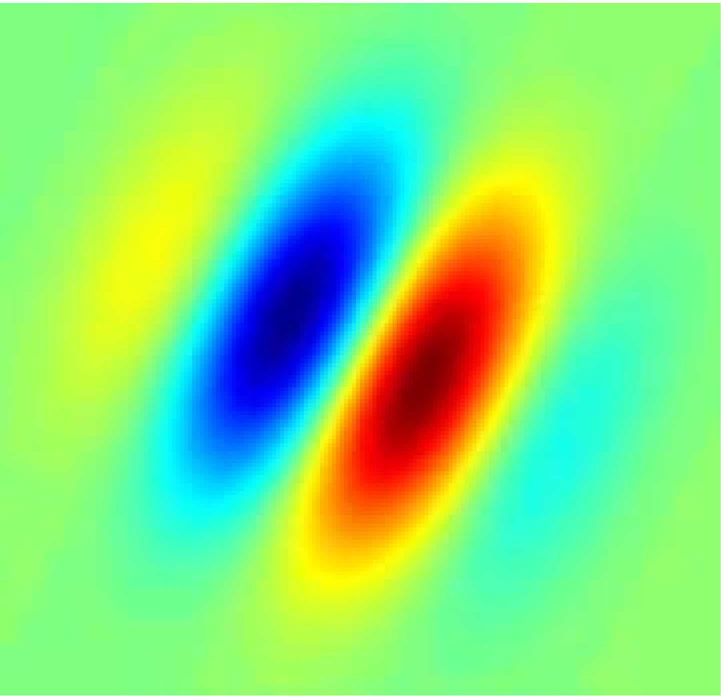

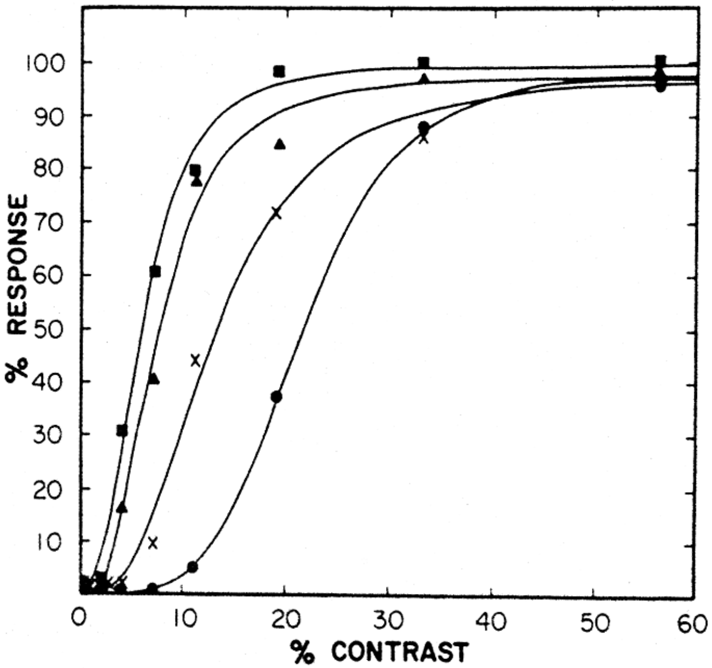

Most neurons in the early stages of the visual cortex compare image attributes in adjacent regions. Many of these comparisons are 'qualitative'. Many of the fields are quite large (3 degrees of visual angle)

Hubel Wiesel, 1959-1972
(Nobel Prize 1981)

Albrecht and Hamilton, 1989, J. Neurophys.
Sinha, 2002, Springer-Verlag LNCS

Qualitative image comparisons

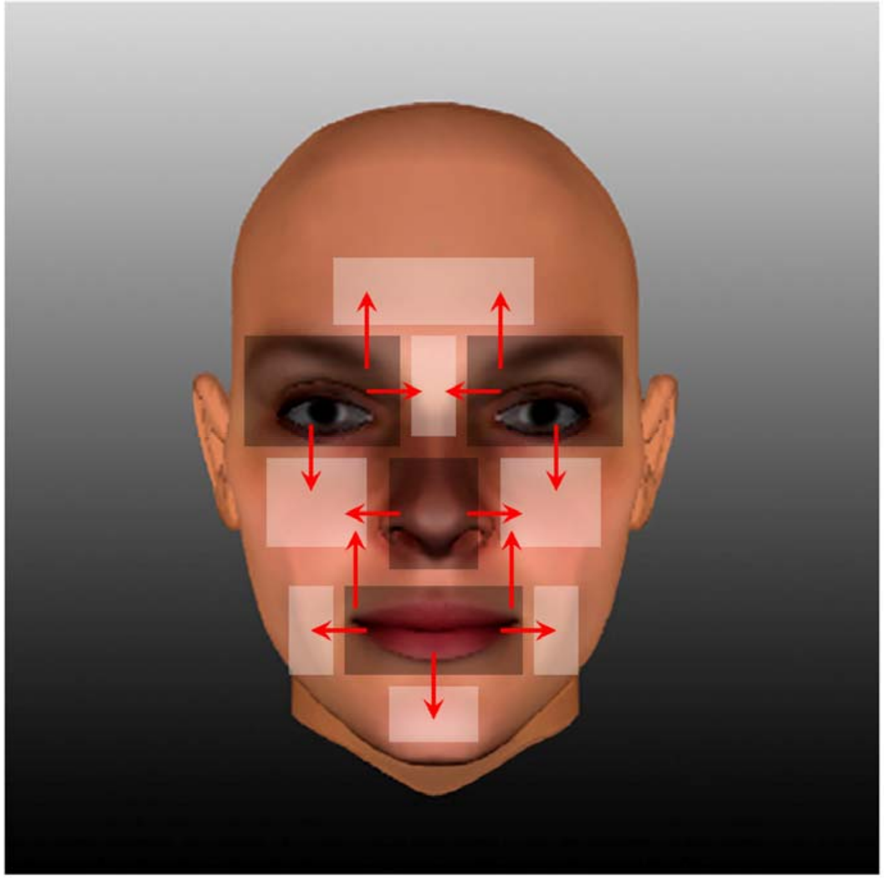

13 February, 2007, New York Times

Results of using a qualitative representation to detect faces

Sinha, 2002, Springer-Verlag LNCS

Applying "Qualitative complements quantitative" to AOI

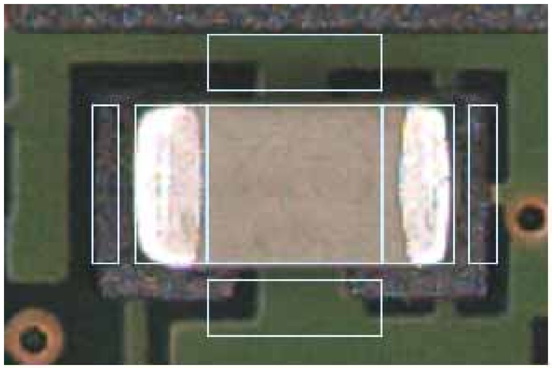

acceptable - part present

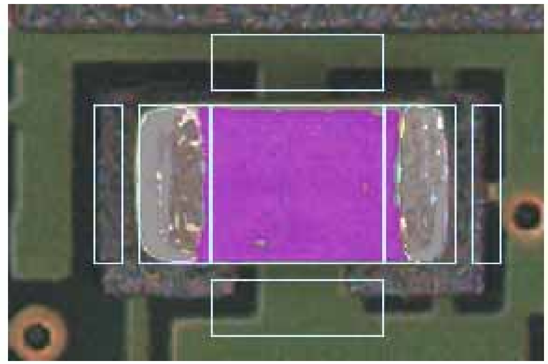

acceptable - part present

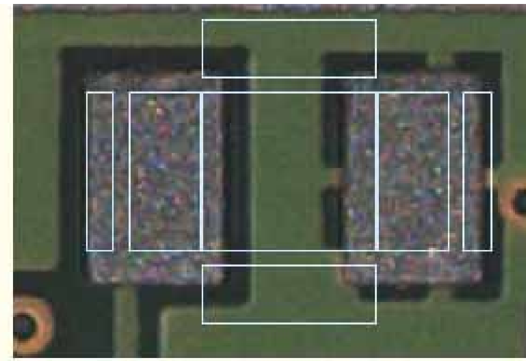

not acceptable - part absent

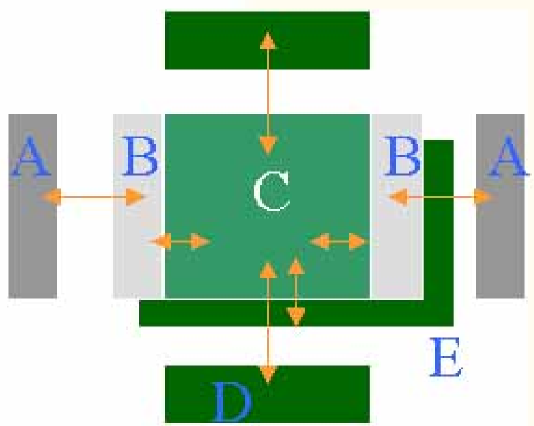

5. Prior Knowledge Bootstraps and Experience Refines

The greater the familiarity with a face, the more the degradation that can be tolerated. Prior knowledge helps in filling-in missing information.

Sinha, 2002, Nature Neuroscience

Applying "Prior Knowledge Bootstraps and Experience Refines" to AOI

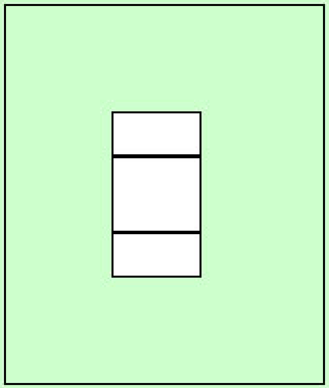

Part Geometry Information

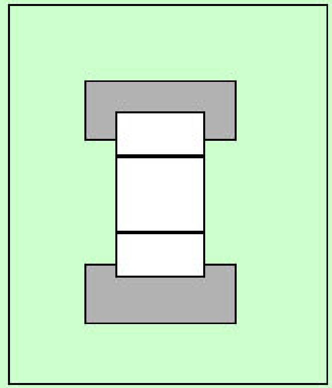

Board Pad Information

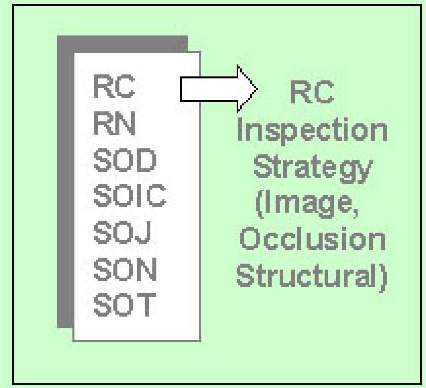

Built in Knowledge of How To Test Each Part

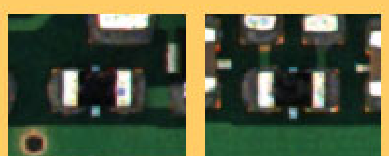

Good Examples (Part Present)

Bad Examples (Part Absent)

Putting it all together…

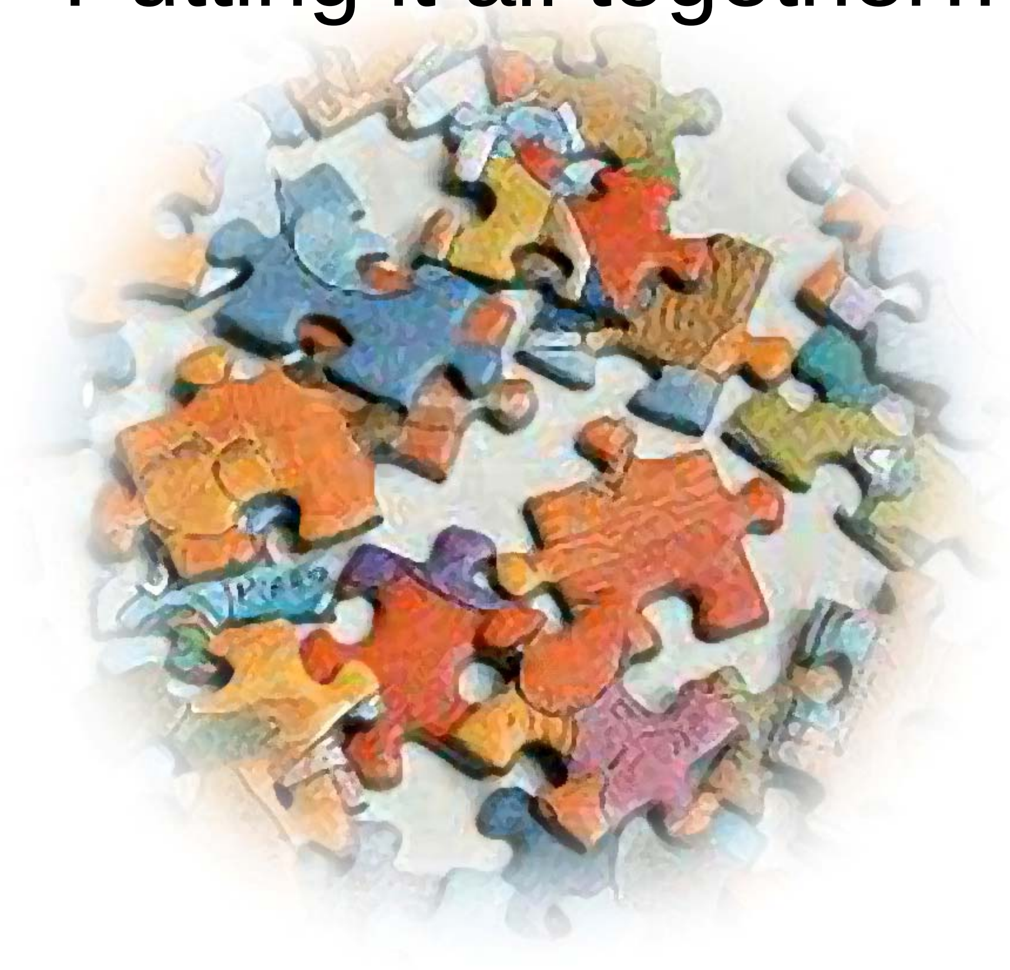

Applying the five principles to a pre-reflow, process monitoring machine

A real case study

Applying these principles to a pre-reflow, process monitoring machine

Landrex Optima7210 ™

Screen Printer	Paste Inspect	Component Place	Component Inspect	Reflow Oven	Assembly	Functional Test

Post Place defect detection and measurement system:

- Attribute data (such as "missing part," "wrong part," "wrong label")
- Variable data (e.g., precise and repeatable measurement of part position).

How does the 7200 incorporate the 5 principles?

Principle 1 (collective action): For instance, if the part color changes, some models may fail the part or be uncertain about whether it is present. Other models, however, that look for structure or occlusion of the board will indicate that the part is present.

Principle 2 (Learning never stops): Allows the system to know the board color has changed from the sample board and to distinguish this from a part color or luminance change.

Principle 3 (using context): Allows the system to analyze the part in the context of the board and parts around it, thus, adding more information to the process.

Principle 4 (qualitative encoding): Models use qualitative encoding to look for part structure and thereby gain invariance to color or luminance changes.

Principle 5 (Prior Knowledge Bootstraps and Experience Refines): Allows the system to tolerate variations by prior knowledge of part-type and by learning the specific examples on the board;

the engineer does not need to develop a complicated program to teach the system all possible variations and combinations.

The consequences

From 1000 ppm to 10 ppm

Initial false-call rate

False-call rate after use of the five principles

PLUS

Reduced need to tweak program

A Study (Long Term)

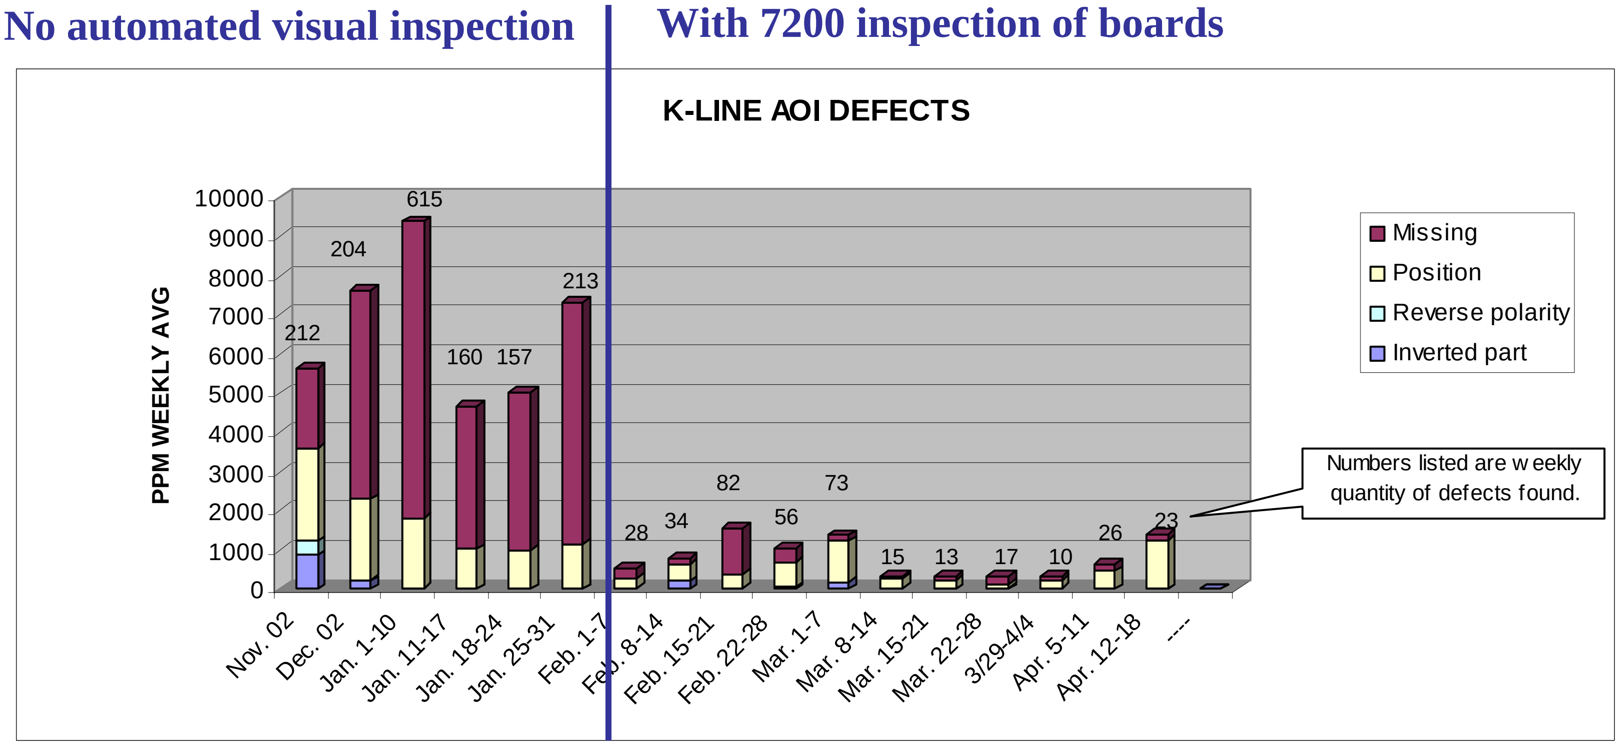

500,000 boards – no change to the program

The road ahead

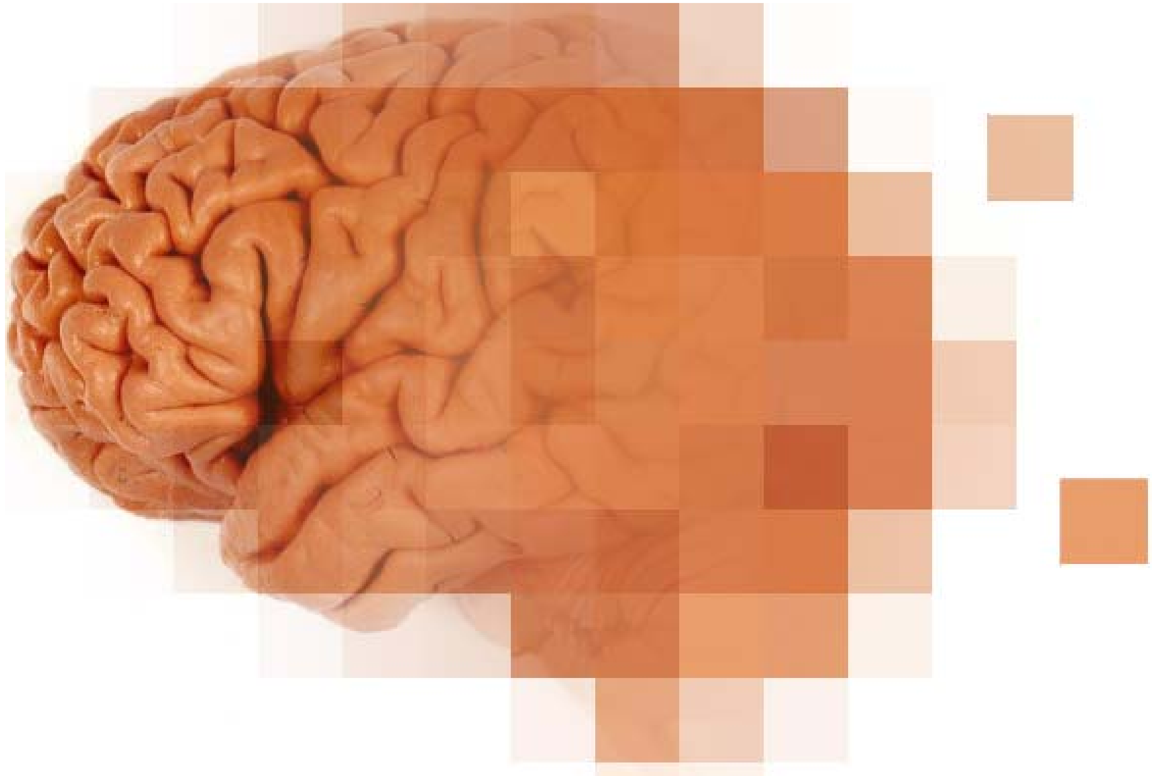

- A hardware revolution has occurred: We now have unprecedented computational resources and imaging capabilities.

- A software revolution is on the way: Brain science provides a roadmap for how best to make use of these resources for optimizing performance of AOI systems.

lipson@imagen-inc.com