

# Deep Learning Delivers Fast, Accurate Solutions for Object Detection in the Automated Optical Inspection of Electronic Assemblies

By Charlie Zhu, CyberOptics Corporation

When automated optical inspection (AOI) works, it is almost always preferable to human visual inspection. It can be faster, more accurate, more consistent, less expensive, and it never gets tired. But there are some challenging applications. Some tasks that are very simple for humans are quite difficult for machines. Object detection is an example. Given an image containing a cat, a dog and a duck, a human can instantly confirm the objects' presence, even when they overlap, and tell you exactly what points in the image are included in each one. This seemingly simple task can be very challenging for AOI. In electronic assembly, manufacturers may want to confirm the presence or absence of a component that varies little in position or appearance from assembly to assembly. This is a relatively simple task. A more difficult problem, and the focus of the work described here, is the detection of corner fill used to secure integrated circuits (IC) to a substrate. Though the location of the fill is relatively constant, its shape and size may vary from instance to instance. This variability makes detection much more complicated.

## Background and terminology

Recent advances in artificial intelligence (AI) in the area known as machine learning (ML), and especially deep learning (DL) with artificial neural networks (ANN), have made significant progress in solving problems like object detection. Artificial intelligence can be described as the study and development of intelligent systems that perceive their environments and take actions to achieve goals. Machine learning focuses on software algorithms that automatically learn to improve a machine's performance in a certain task based on feedback from past performance. Deep learning is a subset of machine learning. Deep learning runs on artificial neural networks, which learn using processes modeled on biological networks, i.e., brains. For example, in a brain, the more often a connection is used, the stronger the connection becomes – “neurons that fire together, wire together”. A neural network would implement that behavior in software. The “deep” in deep learning refers to the many layered architecture of the neural networks it uses. Finally, learning may be supervised – the system is trained using labeled data (explicit examples of the concept being learned), or unsupervised – the system must itself discover patterns occurring the training data. Supervised learning is faster, but unsupervised learning opens the possibility that the system may find previously unknown relationships within the data.

The machine vision problems that the DLs are trying to solve can be broken down into three sub-types of increasing complexity (Fig 1): classification, detection, and segmentation. Presented with an image containing an object, classification attempts to match the object to a known class – “Do I recognize an object in this image?” Given an image containing an object, no object, or many different objects, detection tries to detect the objects, classify them, and determine their size and position – “Do I recognize objects in this image? What are they? Where are they? How big are they?” Finally, segmentation tries to associate each point (pixel) in the image with an object. Segmentation tries to draw a closed border around each detected object. Any pixel inside the border is part of that object. Any pixel not in an object is background.

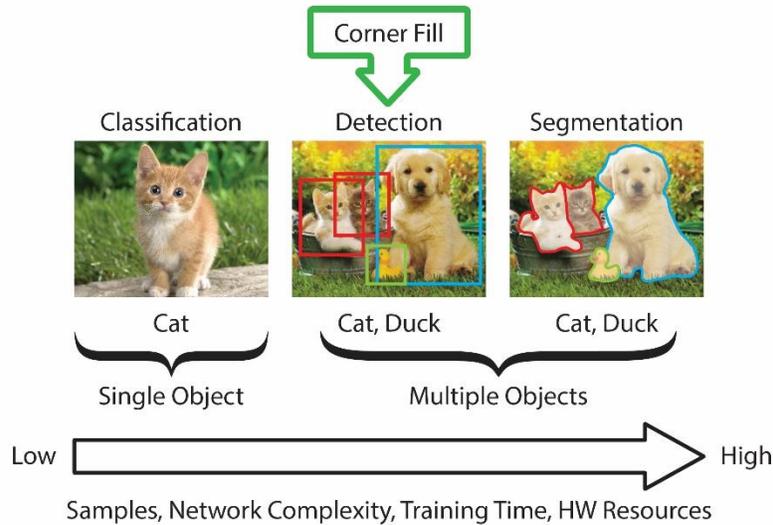


Figure 1. Object detection problems can be divided into 3 sub-types of increasing complexity: classification, detection, and segmentation.

### Corner fill inspection

A memory manufacturer uses fill under the corners of ICs to bond the package to the substrate. They needed to inspect for the presence or absence of fill and ensure that there is neither too much, nor too little. They required a solution that could measure the length of the corner fill and compare it to specifications.

Traditional methods of corner fill inspection, such as blob analysis, are challenged by the lack of gray level specificity in this application. Blob analysis attempts to find a continuous blob within a certain intensity or contrast range and sometimes breaks larger blobs into separate smaller blobs, all too short to meet specification, rather than a single good one, and reported a false negative result. This is an example of a problem that is challenging for AOI but relatively easy for humans, who can readily see that the fill is continuous in the example. Results from inspection with blob analysis were inconsistent and unreliable, with many false negatives. The manufacturer sought a more reliable approach.

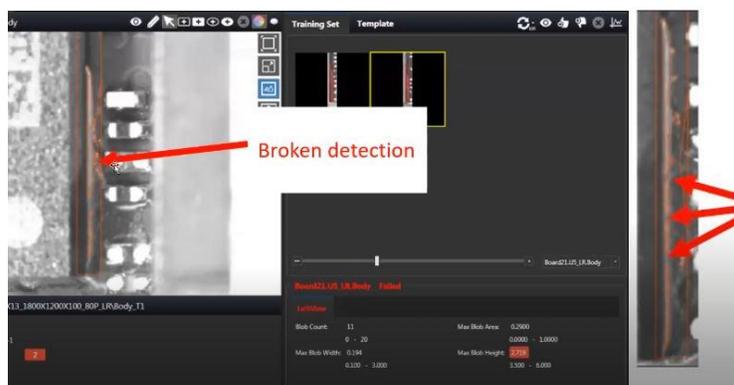


Figure 2. Blob analysis incorrectly divided this good corner fill into several smaller pieces and reported a false negative result.

The manufacturer desired a solution that could confirm the presence of absence of the corner fill and measure its size – a classic example of object detection, mid-level in complexity among classification, detection and segmentation as described above. Research in the use of deep learning for object detection has made dramatic progress in recent years, driven by demand across a variety of applications, including facial recognition and autonomous driving. Autonomous driving shares some requirements with the corner fill application. It needs to be fast, i.e., it needs to detect objects, such as pedestrians and other cars in nearly real-time. It needs to determine how big the object is, and where it is in the field of view, with enough accuracy to avoid a collision. It does not need to precisely define the edges of the object.

A deep learning algorithm for object detection that has gained wide acceptance, is used in autonomous driving applications. Earlier approaches to object detection repurpose classifiers to perform detection, while this approach frames object detection as a regression problem to spatially separated bounding boxes and associated class probabilities. A single neural network predicts bounding boxes and class probabilities directly from full images in one evaluation. Since the whole detection pipeline is a single network, it can be optimized end-to-end directly on detection performance. The approach applies a single neural network to an image, divides the image into regions, and predicts bounding boxes that are weighted by class probabilities. It shares the same network architecture across all classes, which simplifies programming and speeds inferencing. The network can be trained on a personal computer with a single GPU (graphics processor). Once trained, inferencing can run on a device as simple as a mobile phone.

Corner fill is located below the corners of a relatively flat, rectangular package where it is not easy to see from any one point of view. Some corner fill inspection systems use a top-down camera and a mirror to view all sides of the package as it rotates – an approach that adds time and complexity to the data acquisition process. The system used in this work (CyberOptics SQ3000™ Multi-Function System for AOI, SPI and CMM powered by Multi-Reflection Suppression™ (MRS™) sensor technology) incorporates a unique optical sensor originally designed for three-dimensional inspection and metrology using phase shift profilometry. The sensor views the inspection target simultaneously through four side-view cameras positioned off the normal axis at azimuths of 0°, 90°, 180°, and 270°. For corner fill inspection, the side-view cameras can instantly acquire images of all four sides, without mirrors or rotating the sample (Figure 3).

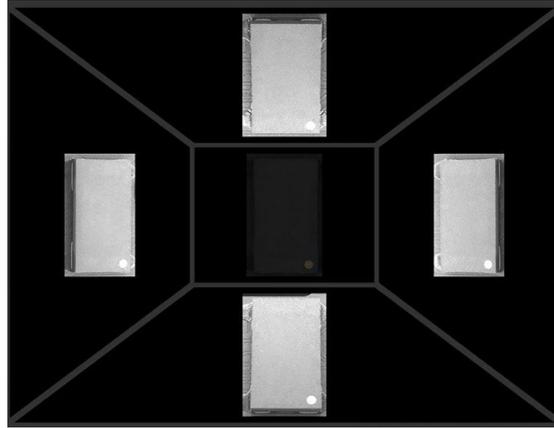


Figure 3 The MRS sensor in the inspection system uses 4 cameras arranged around the sample to provide detailed views of all sides. The corner fill is located under the corners of the package where it cannot be seen from any single point of view.

The system was trained and tested with a set of 72 images. 62 images were used for training and 10 for validation. Corner fill was labeled in the training images and the training was run on a standard personal computer with dual GPUs. One simplified model was able to perform all necessary tasks, powered by deep learning. The bounding box proved to be sufficiently accurate for corner fill measurements (Figure 4). The program runs smoothly and delivers a simplified user experience.

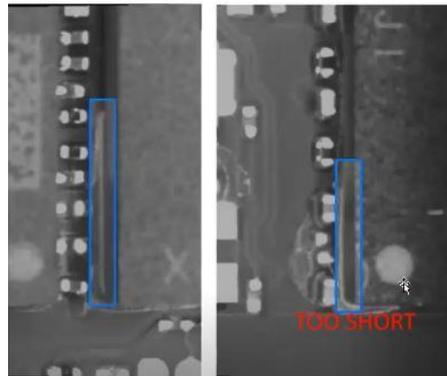


Figure 4. Examples of corner fill measurements using the bounding box reported by the deep learning object detection algorithm. The left image shows good corner fill meeting the length specification. The right image shows corner fill too short.

## Conclusion

We have described the use of an AOI system (CyberOptics MRS-Enabled SQ3000™) and an associated deep learning algorithm to inspect corner fill in an electronics assembly application. The results confirm robust performance in detecting the presence or absence of corner fill and measuring its length. The system was easy to train and runs readily on a standard PC. Integration of the deep learning algorithm into standard system software would allow factory engineers to train networks that could then be run locally for inspection. There are many more potentially valuable applications for deep learning object detection in SMT and semiconductor applications that we are actively pursuing.