DEEP LEARNING AT THE EDGE

Bring the Power of Deep Learning to Everyone—Beginners and Experts Alike

COGNEX

DEEP LEARNING AT THE EDGE
Bring the Power of Deep Learning to Everyone—Beginners and Experts Alike

TABLE OF CONTENTS

Introduction..3

What is edge learning?...4

Edge learning vs deep learning...6

How does edge learning work?..8

Using edge learning...9

Applications...10

Conclusion...11

INTRODUCTION

A common barrier to deploying deep learning in factory automation is the perceived level of difficulty involved. Today, advances in this technology are breaking the stigma.

Deep learning is now easier to use with the introduction of new technologies that process images at the “edge.” Deep learning at the edge, more informally referred to as “edge learning” is a subset of deep learning in which processing takes place directly on-device using a set of pre-trained algorithms. The technology is simple to setup, requiring less time and fewer images for training, compared to more traditional deep learning-based solutions. Requiring no domain expertise for deployment, edge learning is a viable automation solution for all—from machine vision beginners to experts.

Automated visual inspection is essential for improving manufacturing speed and accuracy. Line engineers are aware of these benefits and strive to automate visual inspection but are often held back by the additional complexities of machine vision, including the required level of technical expertise and extensive programming involved, much less those of more sophisticated deep learning applications. Even automation engineers experienced in the use of traditional machine vision recognize its limitations in cases of visual complexity, or subtle defects and variations. While deep learning is a great solution in these cases, its effective use requires advanced skills outside of an automation engineer’s expertise.

Edge learning is the solution for both engineers looking for an easy way to integrate automation into their lines and for expert automation engineers who regularly use rules-based machine vision tools but lack specific deep learning expertise.

WHAT IS EDGE LEARNING?

Edge learning embeds efficient rules-based machine vision within a set of pre-trained deep learning algorithms to create an integrated toolset optimized for factory automation. This technology does not require specialized knowledge of either machine vision or deep learning. Instead, line engineers can train edge learning using their existing knowledge of required tasks.

Using a single, smart camera-based solution, edge learning can be deployed on any line within minutes. This solution integrates high-quality vision hardware, machine vision tools that preprocess each image to reduce computational load, deep learning networks pre-trained to solve factory automation problems, and a straightforward user interface designed for industrial applications.

Edge learning differs from existing deep learning frameworks in that it is not general-purpose but specifically tailored for industrial automation. And it differs from other deep learning products in its focus on ease-of-use across all stages of application deployment. For instance, edge learning requires **fewer images** to achieve proof of concept, **less time** for image setup and acquisition, and **no specialized programming**.

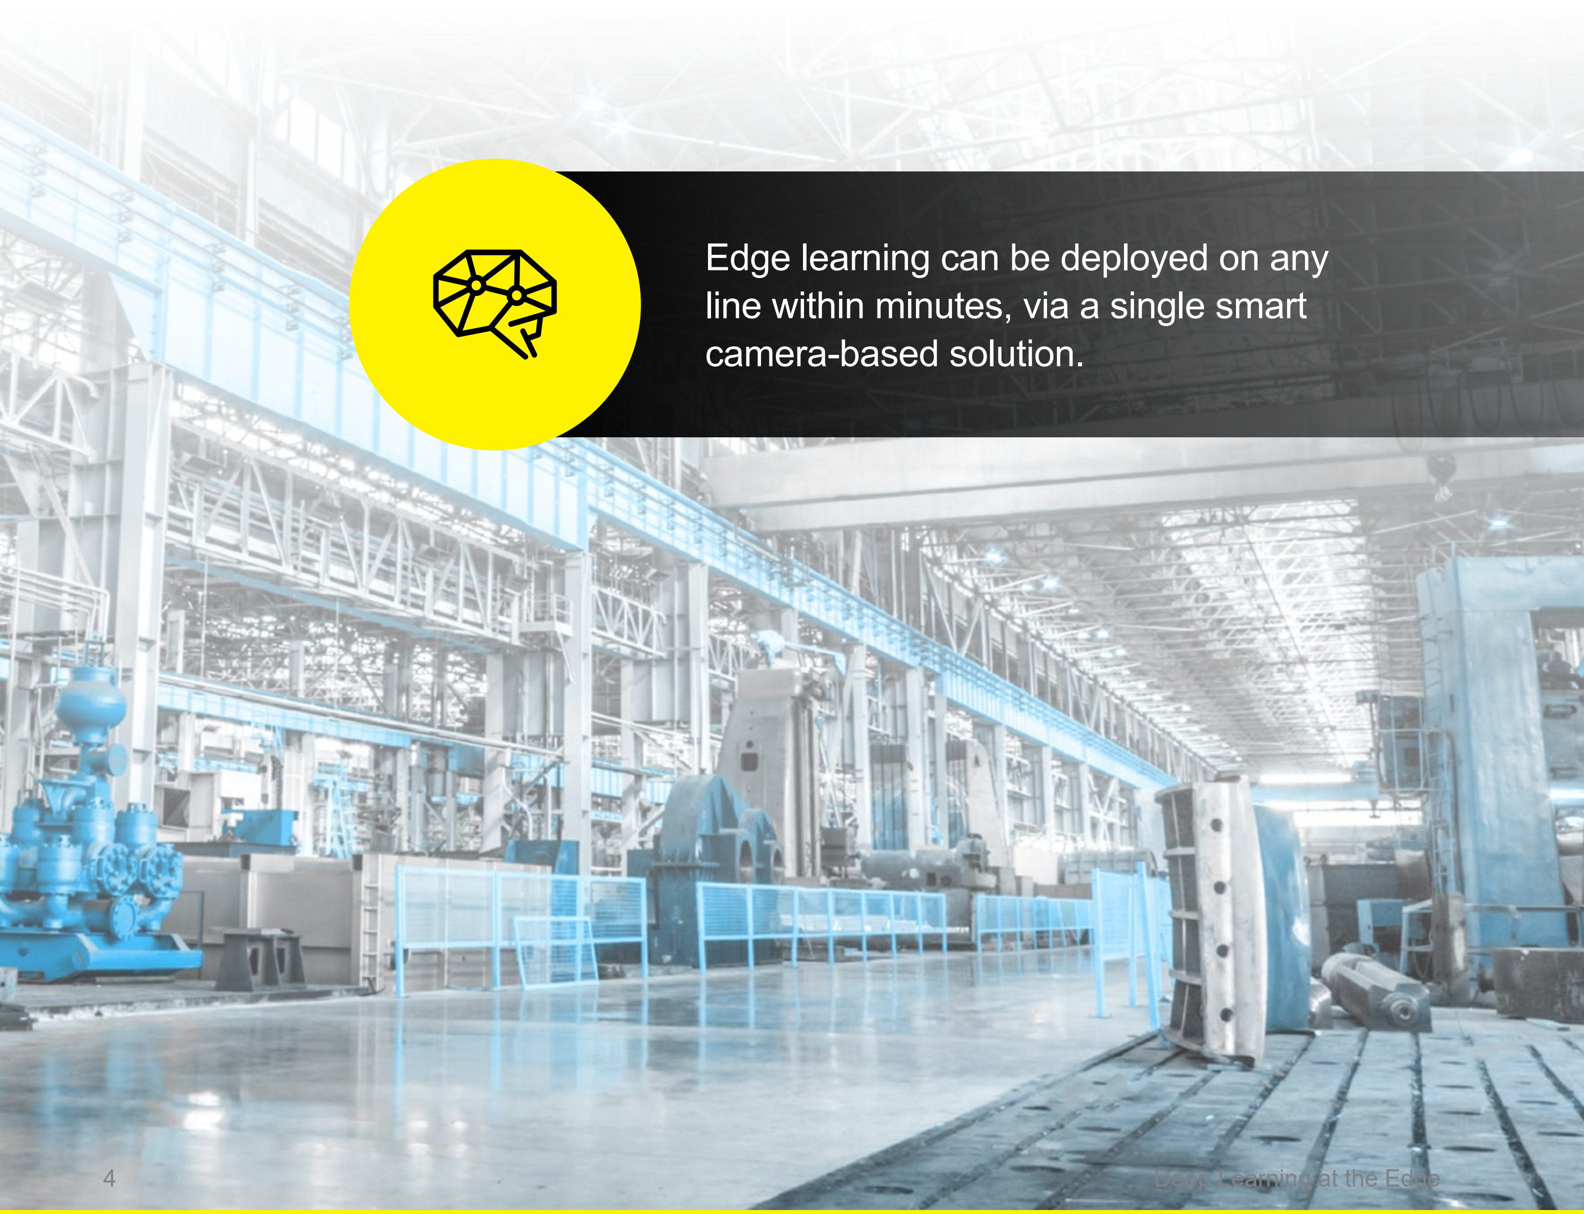

Edge learning can be deployed on any line within minutes, via a single smart camera-based solution.

Quick classification: an example

Edge learning is powerful enough to analyze multiple regions of interest (ROIs) in its field of view and classify each of those regions into multiple categories. This enables it to do sophisticated assembly verification.

For example, edge learning can verify and sort four sections of a frozen meal tray, packed by robots on a high-speed line. In each tray, the bottom center section contains the protein, the top left the vegetable, the top middle the dessert or side dish, and the top right the starch. Each of the sections can contain multiple SKUs, say chicken, turkey, or meatloaf in the protein section, and rice, potatoes, or pasta in the starch section.

With a simple click and drag, each region can be defined for edge learning and locked to invariant features on the meal tray. After that, it is trained to classify each section of the tray with only a handful of images, often as few as two for each possible class. Within minutes of training, edge learning will accurately classify the different sections at high speeds. If more variation is introduced, for instance either a new class or new variety within the same class, the tool can be updated with a few images of the new category.

What works for frozen meal trays also works for printed circuit boards (PCBs) and other complex assemblies.

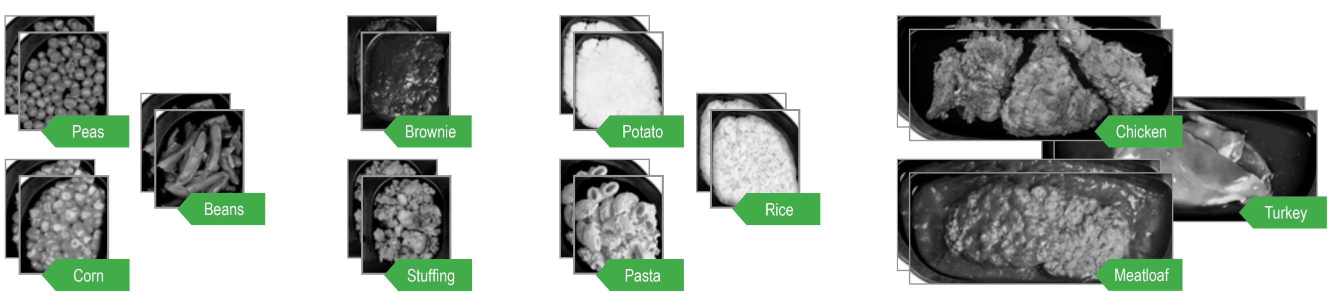

Training each section of the lunch box requires only a handful of images, often as few as two for each possible class.

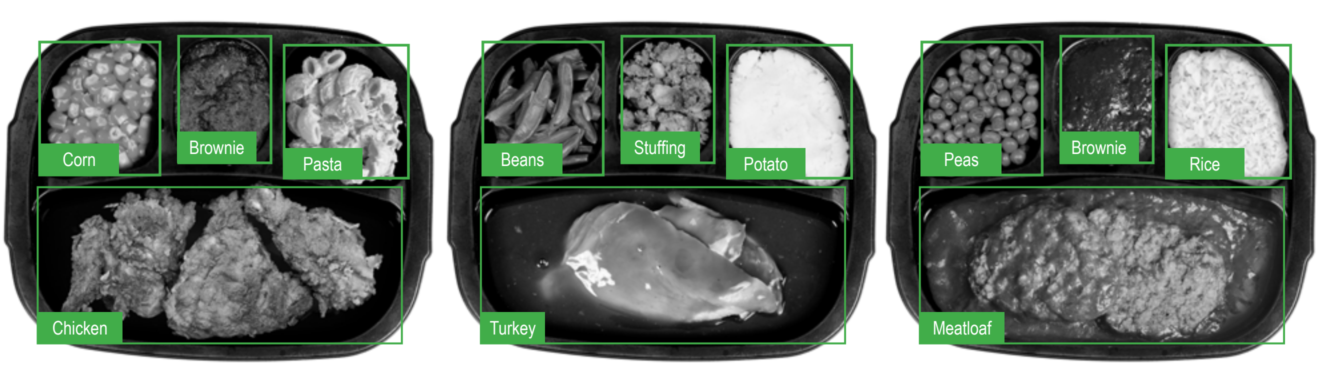

Within minutes, edge learning can accurately classify the lunch boxes at high speeds.

EDGE LEARNING VS. DEEP LEARNING

Deep learning simulates the way interconnected neurons in the human brain strengthen and weaken connections to create an understanding of images. In deep learning, hundreds of layers of neural networks are exposed to a large set of images of similar objects. By slightly modifying connections within and between these layers every time it is exposed to a new image, deep learning learns to reliably identify those objects, and detect defects in them, without any explicit training.

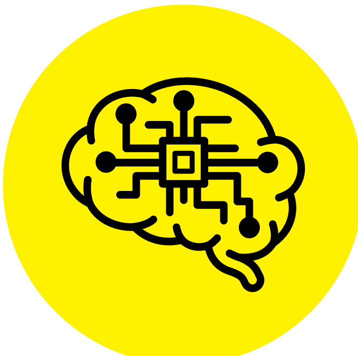

Deep learning: use cases

Traditional deep learning provides the capacity to process large and highly detailed image sets, making it ideal for complex or highly customized applications. Because such applications introduce significant variation, they demand advanced computational power and robust training capabilities. To account for this variation and capture all potential outcomes, image sets numbering in the hundreds or thousands of images must be used for training. Traditional deep learning enables users to analyze such image sets quickly and efficiently, delivering an effective solution for automating sophisticated tasks. While full-fledged deep learning products and open-source frameworks are well designed to address complex applications, the majority of factory automation applications entail far less complexity, making them better suited for edge learning.

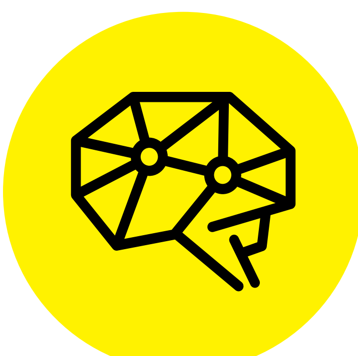

Edge learning: use cases

The power of deep learning can be applied to problems in factory automation by embedding knowledge of factory automation requirements and use cases into the neural network connections from the start. This pre-training removes a lot of the computational load, particularly when supported by the appropriate traditional machine vision tools. The result is edge learning, a light and fast set of vision tools that live on the shop floor, used daily by line engineers as part of their job.

Edge learning tools can be trained in seconds, using as few as five to ten images per class. Compare this to other deep learning solutions, which can require hours to days of training, using hundreds to thousands of images. By streamlining deployment, edge learning enables manufacturers to ramp quickly, while remaining nimble and able to adjust easily to changes.

In order to optimize the edge learning networks to run on embedded vision systems, the images are downscaled or fixtured in a way that only the specific region of interest is analyzed. If these downscaled images were to be differentiated with the line engineer's own eyes, they can be confident the edge learning tools will perform equally as well. Still, it is important to note that this optimization comes at a trade-off. It limits the use of edge learning in very advanced and high-accuracy defect detection applications, which are better solved with traditional deep learning solutions.

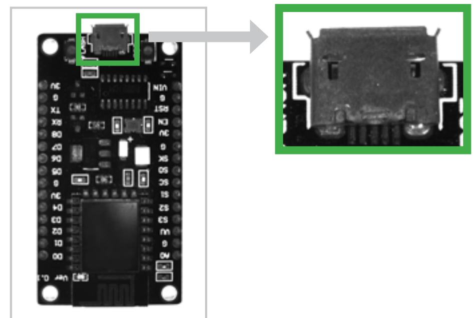

To optimize edge learning, only the specific region of interest is analyzed.

Three Advantages of Using Edge Learning over Deep Learning

Deep Learning	Edge Learning	Advantage
Hundreds to thousands of images required for training	**Five to ten** images required for training	**Fewer images** required for training
Hours to days required for processing	**Seconds to minutes** required for processing	**Faster** learning
Significant understanding of deep learning systems and programming needed	**No prior experience** needed	**Higher** ease of use

HOW DOES EDGE LEARNING WORK?

Deep learning requires sophisticated processing and vast computational resources. How does edge learning bring that power to the shop floor?

Hardware

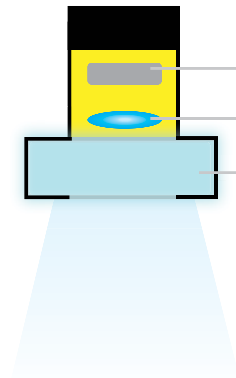

Edge learning packs a lot of sophisticated hardware into a small form factor. It runs entirely in a smart camera with integrated lighting, an autofocus lens, and a powerful sensor.

Lighting is key for a high-quality initial image as it is needed to maximize contrast, minimize dark areas, and bring out the necessary detail.

A high-speed autofocus lens ensures that the object of interest is always in focus, even as distance changes. The lens can instantly adjust its focus as the region of interest (ROI) changes. Liquid autofocus lenses are smaller and lighter than equivalent mechanical lenses, reducing the size and weight of the camera while making it resistant to the shock and vibration of a production line.

A large and capable sensor offers greater resolution and a wide field of view (FOV).

All of these hardware features help make edge learning possible.

Machine vision tools

Rule-based vision tools are enormously effective at a variety of specialized tasks, such as location, measurement, and orientation. For the purposes of edge learning, they are combined in ways specific to the demands of factory automation, eliminating the need to chain vision tools or devise complex logic sequences.

These tools provide fast preprocessing of any image, extracting density, edge, and other feature information shown by experience to be relevant to detecting and analyzing manufacturing defects. By identifying and clarifying the relevant parts of the image, these tools reduce the computational load of deep learning.

Deep learning capabilities

In the past few years, deep learning has transformed factory automation. Instead of using rules created by human programmers, it learns by example, building a neural network and devising effective pass/fail thresholds from labeled examples of acceptable and unacceptable parts. Effectively, it mimics the way humans learn.

These deep learning capabilities have large computational requirements. Edge learning, on the other hand, takes advantage of the fact that factory automation images have specific structural contents, and so pre-trains its algorithms with that domain knowledge. Not starting from scratch results in a less computationally intensive application.

USING EDGE LEARNING

Deploying

Edge learning works entirely on-camera, so it does not require pulling data to be processed on a PC, which consumes space, incurs delays, and requires IT intervention, resulting in an overall more expensive solution. Because of its small form factor, it fits easily, even on lines with a lot of other equipment, and provides its own sophisticated lighting, which can be calibrated for the needs of the location.

Training

Training edge learning is similar to training a new employee on the line.

What the user of edge learning needs to know is not how vision systems or deep learning works, but what classification problem they need to solve. If it is straightforward, for instance classifying acceptable and unacceptable parts as OK/NG, the user needs to know which parts are acceptable and which are not. This can include knowledge not readily apparent, derived from testing down the line, that reveals defects hard for a human to detect. Edge learning is particularly effective at determining which variations in the part are significant, and which variations are purely cosmetic and do not affect functionality.

Edge learning is not limited to binary classification into OK/NG but can classify into any number of categories. If parts need to be sorted into three or four different categories, depending on components or configurations, that can be setup just as easily. Edge learning is also capable of directing attention to multiple regions of interest (ROI) in the image. And, of course, both multiple ROIs and multiple categories can be handled together, as in the frozen meal tray example on page 5.

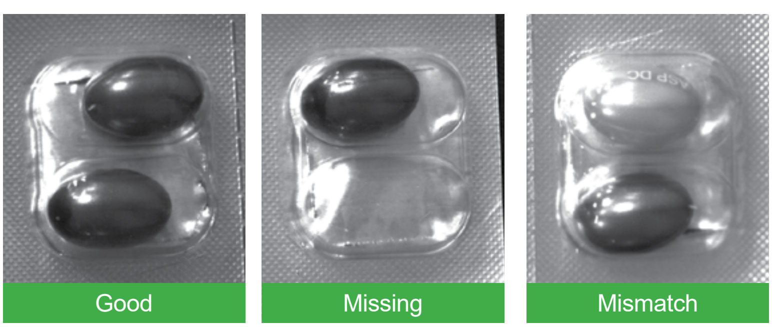

Multi-Class Classification: *Edge learning tools perform multi-class classification on pills in a blister pack, classifying each section as either a good pair, missing a pill, or having a mismatched pill.*

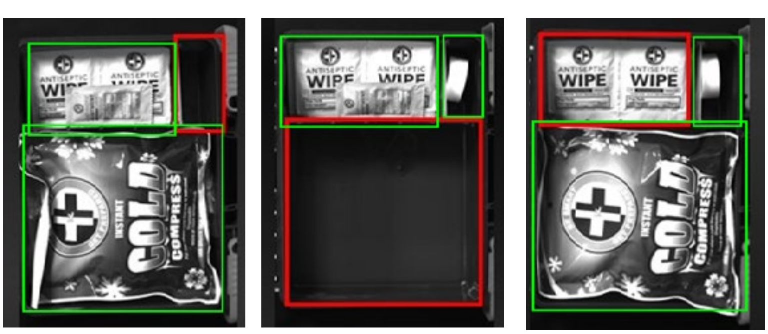

Multi-ROI Assembly Verification: *Edge learning tools check three different regions to ensure all necessary components are present in a first aid kit.*

APPLICATIONS

Edge learning is useful in a wide range of industries, simplifying factory automation and solving tasks of varying complexity.

Electronics

Many printed circuit boards (PCBs) include LED indicator lights to show status. The goal of this application was to identify which indicators showed a power on (PWR) condition, a transmit (TX) condition, or an off (OFF) condition.

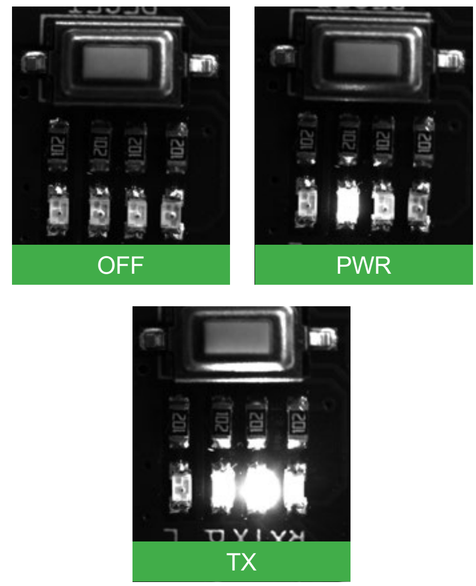

The typical way to make these determinations using traditional machine vision is with a pixel count tool. This involves setting thresholds for brightness at specific locations for each condition, a complex process that requires machine vision programming experience.

Edge learning was trained on a small set of labeled images of the OFF, PWR, and TX conditions. After this brief training, edge learning reliably sorted the PCBs into the three different states.

Medical/Pharmaceutical

Glass vials are automatically filled with medication to a predetermined level. Before they are capped, the level must be confirmed to be within proper tolerances. The transparent and reflective nature of both the glass vial and its contents makes it difficult for traditional machine vision to consistently detect the level.

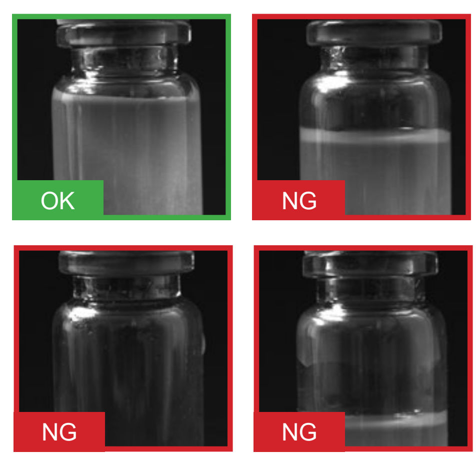

Edge learning was trained to identify fill level, without getting confused by reflections, refraction, or other confusing variable parts of the image. Fills that were too high or too low were classified as NG, while only those within the proper tolerances were classified as OK.

Packaging

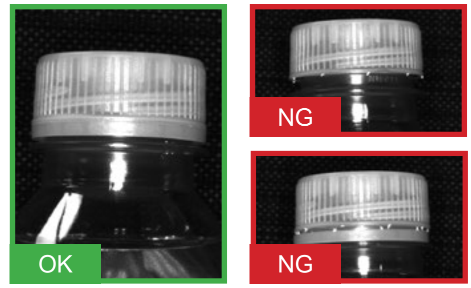

Bottles of soft drinks and juices are filled and sealed with a screw cap. If the rotary capper mistreads the cap, or it gets damaged during the capping process, this can leave a gap that allows for contamination or leakage. Both the speed and the wide range of ways in which a cap may be almost, but not quite sealed, make this a challenging application for traditional machine vision.

Edge learning was given a set of images labeled as good and a set of images that showed caps with slight gaps that were almost imperceptible to the human eye. Only fully sealed caps were categorized as OK and all other caps were classified as NG, at line speeds.

EDGE LEARNING BRINGS THE BENEFITS OF DEEP LEARNING TO ALL

Edge learning is a game-changing technology that is more capable than traditional machine vision while being extremely easy to use. Its powerful capabilities are easily deployed and used by line engineers in the course of their daily work, without requiring advanced machine vision or deep learning training.

Automation engineers with more knowledge of traditional machine vision tools can leverage their existing knowledge to use the power of edge learning in developing sophisticated and robust factory automation processes.

BUILD YOUR VISION

2D VISION SYSTEMS

Cognex machine vision systems are unmatched in their ability to inspect, identify and guide parts. They are easy to deploy and provide reliable, repeatable performance for the most challenging applications.

www.cognex.com/machine-vision

3D VISION SYSTEMS

Cognex laser profilers and area scan 3D vision systems provide ultimate ease of use, power and flexibility to achieve reliable and accurate measurement results for the most challenging 3D applications.

www.cognex.com/3D-vision-systems

VISION SOFTWARE

Cognex vision software provides industry leading vision technologies, from traditional machine vision to deep learning-based image analysis, to meet any development needs.

www.cognex.com/vision-software

BARCODE READERS

Cognex industrial barcode readers and mobile terminals with patented algorithms provide the highest read rates for 1D, 2D and DPM codes regardless of the barcode symbology, size, quality, printing method or surface.

www.cognex.com/barcodereaders

COGNEX

Companies around the world rely on Cognex vision and barcode reading solutions to optimize quality, drive down costs and control traceability.

Corporate Headquarters **One Vision Drive Natick, MA 01760 USA**

Regional Sales Offices

Americas

North America	+1 844-999-2469
Brazil	+55 11 4210 3919
Mexico	+800 733 4116

Europe

Austria	+49 721 958 8052
Belgium	+32 289 370 75
France	+33 1 7654 9318
Germany	+49 721 958 8052
Hungary	+36 800 80291
Ireland	+44 121 29 65 163
Italy	+39 02 3057 8196
Netherlands	+31 207 941 398
Poland	+48 717 121 086
Spain	+34 93 299 28 14
Sweden	+46 21 14 55 88
Switzerland	+41 445 788 877
Turkey	+90 216 900 1696
United Kingdom	+44 121 29 65 163

Asia

China	+86 21 6208 1133
India	+9120 4014 7840
Japan	+81 3 5977 5400
Korea	+82 2 530 9047
Malaysia	+6019 916 5532
Singapore	+65 632 55 700
Taiwan	+886 3 578 0060
Thailand	+66 88 7978924
Vietnam	+84 2444 583358

© Copyright 2022, Cognex Corporation. All information in this document is subject to change without notice. All Rights Reserved. Cognex is a registered trademark of Cognex Corporation. All other trademarks are property of their respective owners.
Lit. No. WPEL-EN-03-2022

www.cognex.com