

Edge-Learning AI Simplifies Electronic Inspection Challenges

Embedded smart-camera platforms with easy-to-use edge learning provide a unique solution for apps that are too difficult for conventional rule-based machine vision but don't justify the investment of a complete deep-learning solution.

For years, electronics manufacturers have relied on machine vision to process semiconductor wafers, build and inspect printed circuit boards, and assemble consumer electronics modules. Semiconductor lithography and wafer probing would be impossible without precise, visual alignment. Component presence/absence and gauging checks are critical to PCB assembly.

Meanwhile, traceability—tracking products from wafer to final modules—would be impossible without robust OCR and 2D code-reading. Where machine vision has struggled is in the more subjective task of quality inspection, as conventional rules-based tools cope poorly with unpredictable visual features such as scratches or stains.

Artificial intelligence (AI), and specifically deep learning, offers a potential solution. AI can identify inconsistencies in patterns, even if defects are hard to define with fixed rules. But the challenge of collecting and labeling the big data sets required to train deep-learning models has proven to be an obstacle to mass deployment.


Today's semiconductor fabrication plants (fabs) and electronics manufacturers not only need to keep up with the latest vision technology, but they must also contend with multiplying product variations and a dearth of skilled labor—while under pressure to reduce scrap and increase yields. The level of difficulty required to develop deep-learning solutions has presented a significant barrier to anyone but skilled data scientists.

Meanwhile, systems that perform well in a lab set-




Deep Learning

Training set



Results



Deep learning automates complicated inspection tasks by analyzing vast and highly granular image datasets, enabling swift and effective identification and classification of acceptable and unacceptable irregularities.

ting often struggle when confronted with the realities of the factory environment, where they must handle new product introductions, fast production ramps, line-to-line equipment differences, and unpredictable labor turnover. Even for experienced vision professionals, the time it takes to deploy deep learning is significant, and not every application justifies such an investment.

This article will define edge learning, compare it to deep learning, and explain how edge learning can alleviate current burdens, reduce costs, and optimize workforce utilization while tackling applications that can't be solved with traditional rules-based vision approaches. It also discusses how edge learning's ease of use makes it practical for beginners and experts alike to quickly deploy automated tasks in high-volume electronic component quality checks and final assembly verification.



Edge Learning


Pretraining



Use case training



Results



Optimizing edge-learning networks for embedded machine-vision systems requires resizing or adjusting the images so that only the important regions of interest are analyzed.

Legacy Machine Vision

Conventional machine-vision systems use rule-based algorithms to identify and measure defects that can be described mathematically based on pixel intensities. The development of such systems requires trained engineers to evaluate each vision task and apply a set of rules to solve the problem. Afterward, the system must be programmed accordingly.

To simplify this process, vendors offer low-code and no-code solutions that assist in adjusting the system's analytical pattern matching, blob, edge, caliper, and other image-processing vision tools to fit specific application requirements. However, rule-based solutions only work when defects have consistent edges, intensities, or patterns. They break down when defects are challenging to quantify, or their appearance varies unpredictably from part to part.

Maintaining and developing rule-based machine-vision algorithms is an ongoing challenge, often necessary due to changes in parts and processes. Shrinking product lifecycles or component obsolescence can cause part changes, while raw material or component variations from different suppliers, technological advances, or lighting changes in the production environment may require process changes.

Sustaining such machine-vision systems requires high-priced engineers with specialized machine-vision skills, which can be difficult to secure because they're in great demand as more and more tasks become automated.

Data-Centric Deep Learning

A decade ago, deep learning was limited to specialized professionals with ample budgets. But recent advances in theory, compute hardware like GPUs, and data availability have made it possible to use deep learning in industrial machine-vision applications.

Deep learning has proven to be particularly effective in situations that require subjective decisions, such as those that would typically involve human inspectors. It's also quite efficient in complex scenes where identifying specific features in the image are challenging due to high complexity or extreme variability.

Typical examples include identifying poorly soldered connections, spotting deformations in 3D components, or distinguishing dust from true defects on an IC. These applications are extremely challenging for conventional machine-vision tools but lend themselves to deep-learning techniques.

Similarly, conventional optical-character-reading (OCR) algorithms encounter difficulty in reading text strings that are laser marked or chemically etched on integrated circuits (ICs) due to their small size and lighting challenges. Deep learning is well-suited for scenes where highly textured surfaces and inconsistent illumination can cause characters in the image to become deformed or appear distorted.

Unlike rule-based machine vision, which relies on experts

to develop new algorithms, deep learning depends on operators, line managers, and other subject-matter experts to label images as good or bad and classify the types of defects present in an image.

This data-centric, learning-based approach lessens the need for experienced machine-vision specialists and reduces the engineering crew required to deploy and maintain machine-vision solutions. If something changes, anyone who can visually make out what character has been marked, or is familiar with what a given defect looks like, can potentially retrain the model by capturing and labeling new images.

Various machine-vision system providers have created deep-learning software tools specifically for machine-vision applications. Examples of these tools include VisionPro Deep Learning from [Cognex](#), Aurora Vision Studio from [Zebra](#), NeuralVision from [Cyth Systems](#); the neural-net supervised learning mode offered by [Deevio](#); and [MVTec Software](#). Several other companies offer open-source toolkits for developing software tailored to machine-vision applications.

Deep-Learning Difficulties

Deep-learning toolkits make it easier for people to implement learning-based machine-vision systems, but there are still challenges to overcome. Successful deep-learning projects



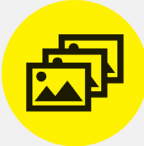


often require significant financial and time budgets, as well as expertise from vision engineers and data scientists to set up the initial system.

Some projects may not offer sufficient value to justify the operation. This can be due to several factors. The parts being inspected may be low value, whereby improvements only offer marginal savings. The defect specification may be unclear, leading to different technicians labeling parts according to different standards. If it's a new part type, the manufacturing operation itself may not be fully standardized yet, so that any deep-learning system requires frequent retraining. Or a single machine may produce so many part types that it's not practical to collect and label images of every variety.

Another challenge is lighting and image formation. Like any machine-vision application, image acquisition hardware is crucial to the success of a deep-learning solution. A well-designed imaging system is critical. The imaging techniques must be reliable and repeatable and able to clearly distinguish features or objects of interest.

Part presentation, illumination techniques, and image resolution all play a crucial role in identifying the nuances that differentiate various classifications. In a conventional rules-based vision application, changes in lighting can often be handled by adjusting a single grey-level threshold. With deep

Three Advantages of Using Edge Learning Over Deep Learning

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

Although complex and extremely accurate defect-detection tasks are better suited for traditional deep-learning approaches, edge learning can effectively handle a significant portion of applications without requiring much customization.

learning, it could potentially require collecting and labeling an entirely new set of images.

On the software side, model development can be a time-consuming process that requires tagging hundreds or thousands of images. New production lines are often unstable, leading to frequent deep-learning project “do-overs.” Conversely, stable manufacturing processes are often so well optimized that they produce few defective parts. As a result, it takes several months to even collect enough defect examples

to train a reliable deep-learning model.

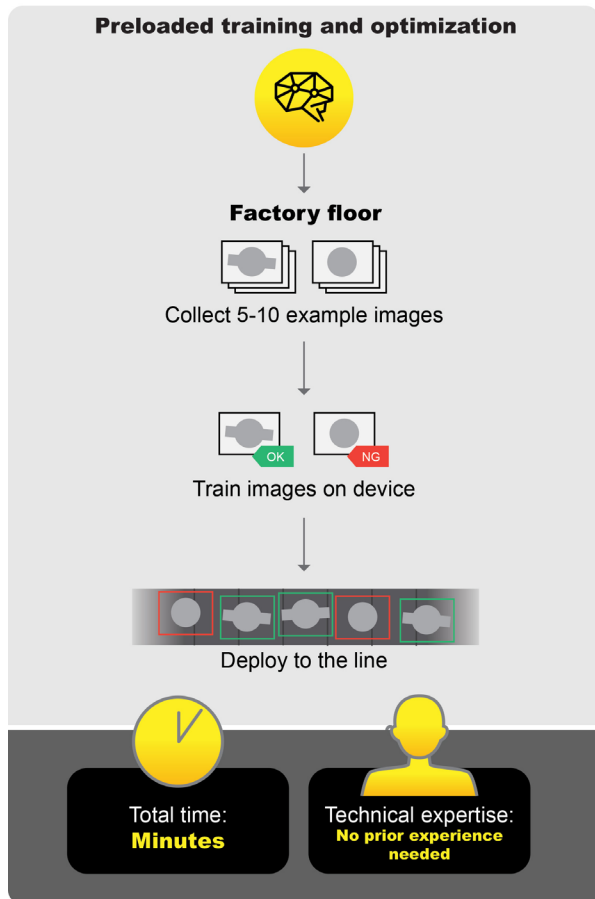
Edge Learning Fills the Gap

The challenges of many machine-vision applications are too difficult for rule-based solutions, but the ROI is insufficient to warrant the development and maintenance of a complete deep-learning solution. To address this problem, machine-vision companies have created edge learning, a type of AI that fills the gap between deep learning and conventional

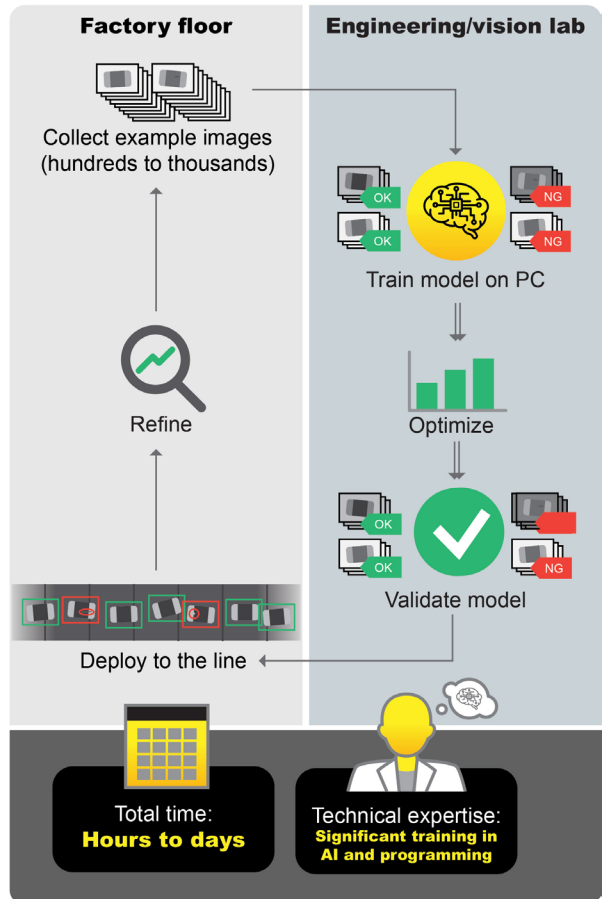
How to Deploy Edge Learning versus Deep Learning



Edge learning neural networks come pretrained and ready to deploy on the line.



Deep learning neural networks are trained by an experienced professional on a PC.



Edge-learning neural networks come pretrained by the supplier on a massive dataset optimized for industrial automation. Therefore, it takes only a few minutes for the customer to complete the second part of the training process for their particular use case.

rules-based vision.

Edge learning is a form of deep learning that's tailored for industrial automation. It involves two steps: pretraining and specific use-case training. The edge-learning supplier conducts the first step using a large dataset optimized for general industrial automation. The customer completes the second step by adding a small number of images for their specific use case. This involves far fewer images than deep learning—typically one to two orders of magnitude fewer!

Edge-Learning Platform Options

Because edge learning requires far fewer images to train, the demands on the training hardware are also lighter. Training an edge-learning model frequently doesn't even require a GPU.

For many applications, an embedded smart-camera edge-learning solution will offer quicker and easier image setup and acquisition because smart-camera platforms combine multiple elements such as sensors, optics, processors, and illumination. This minimizes hardware integration problems resulting from cabling to a PC and incorporating the inference engine, ultimately decreasing the complexity of a machine-learning system.

It also lowers costs while increasing flexibility. High-end GPUs can significantly drive up hardware expenses. They're also harder to retrofit onto existing manufacturing lines, where a vision PC without a GPU may already exist. By comparison, a smart camera with an embedded neural processing

unit (NPU), or even CPU-only, may be able to train and run edge-learning tools.

However, it's worth noting that while edge-learning embedded smart cameras offer many benefits, they're not suitable when high speed and high resolution are essential. For high-throughput applications, or in cases where the size of the product being inspected is large and the defects being identified are very small, a PC-based edge-learning system may be the most effective hardware solution.













Edge Learning Versus Deep Learning

Regardless of the platform selected, edge learning has numerous advantages over traditional rule-based and deep-learning solutions. Because it requires fewer images—especially hard-to-find defect images—users can quickly determine project feasibility from just a few example parts. Once a project is approved, the development cycle is much faster. It's also much simpler, removing the need for data science experts.

This flexibility also helps in production, where the system can be adapted quickly to changes in illumination or product type. And since no GPU is usually required, changes can be made directly on the factory floor using embedded smart-camera hardware.

So why not replace deep learning with edge learning everywhere? It depends on the complexity and accuracy of the application. Deep learning is ideal for subtle or very challenging inspections needing high precision and repeatability. For ex-

Recommended uses

 Rule-based vision	 Edge learning	 Deep learning
 Measurement and gauging	 Classification	 Complex defect detection
 Location and fixturing	 Assembly verification	 Tasks with significant variation
 Robotic guidance	 Character reading	 Highly customized jobs

Edge learning bridges the gap between conventional rule-based machine vision and complete deep-learning solutions, thereby addressing the needs of intermediate applications.

ample, some applications can tolerate small defects, but reject parts with defects larger than a specified size. For these cases, the vision system must be able to not only locate and evaluate defects, but precisely measure their extent. Deep learning is essential for these cases.

But given the general advantages of edge learning, it's often best to start projects with edge learning, and only upgrade to deep learning if it proves necessary. Many projects begin with only a handful of example parts. Edge learning can jumpstart the evaluation, giving reasonable performance quickly. In many cases, that may prove to be sufficient. And if not, they can migrate their application to a full deep-learning approach.

Edge Learning for All

Edge learning is an excellent fit for electronics and semiconductor makers. These industries must balance high-volume manufacturing, tight quality standards, a wide range of product variations, and strong price pressure.

Edge learning offers an efficient approach to check component quality, as well as module and PCB assembly. It can detect defects such as scratches, stains, chips, and dents. It's able to classify part types as well as their quality. And edge learning OCR is robust enough for the challenges of high-magnification etched and scribed character strings.

Edge learning benefits different players in the semiconductor and electronics manufacturing space. Original equipment manufacturers (OEMs) can use edge learning to address challenging machine-vision problems more easily and empower their customers. In addition, edge learning's ease of retraining enables end users to quickly resolve issues and add new products, reducing the need for after-sales support and service costs.

Moreover, edge learning helps OEMs deliver machines that can achieve better yields and reduced scrap rates, while allowing smaller OEMs to compete with larger establishments by providing complex solutions with fewer resources.

Systems integrators can also benefit from edge learning as it allows them to perform more feasibility studies in less time, reducing the burden of image-acquisition setup and machine-vision tools selection. Thus, they can determine the feasibility of an application quickly and win more business faster, all while taking on more projects.

End-users can leverage edge learning to automate many manual optical inspections or automation tasks that don't warrant the development costs of a sophisticated machine-vision or deep-learning system. With edge learning, manufacturers can easily deal with changes in parts and processes and iterate without developing new algorithms for each new generation of product.

On top of that, edge learning can simplify existing rule-based machine-vision applications and reduce costs associated with expensive image-acquisition components, such as

telecentric optics, illumination, or part-handling systems. This simplification and elimination of costly components helps achieve a lower-cost setup, resulting in savings on image formation, fixturing, or complex image-processing requirements.

Conclusion

Embedded smart-camera platforms that feature easy-to-use edge learning provide a unique solution for applications that are too difficult for conventional rule-based machine vision but do not justify the investment of a complete deep-learning solution. If more processing power or higher resolution is required, edge learning can be deployed on a PC-based vision solution that has the compute power to handle higher line speeds and can use industrial cameras with the resolution to handle microscopic defect detection.

Compared to traditional machine-vision analytical tools, edge learning has proven to be more capable, especially in situations where human inspectors need to make challenging subjective decisions, such as identifying specific features in a highly complex or extremely variable image.

Edge learning is also more user-friendly and cost-effective than traditional deep-learning solutions, making it economical for many applications. It takes just a few images per class and a few minutes to train edge-learning tools, whereas traditional deep-learning solutions may require hours to days of training and hundreds to thousands of images.

Edge learning also streamlines deployment and enables fast ramp-up for manufacturers, allowing for quick and easy part and process changes with just a few images needed for on-the-fly model updates. Ultimately, edge learning is an additional tool that can improve workforce utilization for OEMs, machine builders, systems integrators, and end-users serving semiconductor fabs, electronics, and other industries.

John Petry is the marketing director for Cognex's Advanced Vision Technology business unit. He works closely with the R&D team to identify opportunities and test AI solutions across a range of critical applications. These include EV battery manufacturing, cell-phone assembly, and high-speed logistics automation across the U.S., Europe, and Asia.



John has worked in machine vision for over 25 years as a software developer, an engineering manager, and a marketing director. He has been based at Cognex's headquarters in Natick, Mass., and at Cognex's Tokyo office. He holds a Bachelor of Science degree from M.I.T.