



COGNEX

DEEP LEARNING vs. MACHINE VISION

DEEP LEARNING vs. MACHINE VISION

In the last decade, the pace of technology change has been breathtaking. From mobile devices, big data, artificial intelligence (AI), and internet of things, to robotics, blockchain, 3D printing, and machine vision, industries have been thrust into a transformative era.

Strategically planning for the adoption and leveraging of some or all these technologies will be crucial in the manufacturing industry. In the United States, manufacturing accounts for \$2.17 trillion in annual economic activity, but by 2025—just half a decade away—McKinsey forecasts that “smart factories” could generate as much as \$3.7 trillion in value.

“If you’re stuck to the old way and don’t have the capacity to digitalize manufacturing processes, your costs are probably going to rise, your products are going to be late to market, and your ability to provide distinctive value-add to customers will decline,” Stephen Ezell, an expert in global innovation policy at the Information Technology and Innovation Foundation, says in a report from Intel on the future of AI in manufacturing. In other words, the companies that can quickly turn their factories into intelligent automation hubs will be the ones that win long term from those investments.

These technologies as applied in a factory or manufacturing setting are no longer nice to have, they are business critical. According to a recent research report from Forbes Insights, 93% of survey respondents from the automotive and manufacturing sectors classified AI as ‘highly important’ or ‘absolutely critical to success’. Yet, only 56% of these respondents plan to increase spending on AI—and only by 10% or less of existing budgets.



“If you’re stuck to the old way and don’t have the capacity to digitalize manufacturing processes, your costs are probably going to rise, your products are going to be late to market, and your ability to provide distinctive value-add to customers will decline.”

—Stephen Ezell, expert in global innovation policy at the Information Technology and Innovation Foundation

The disconnect between recognizing the importance of new technologies that allow for more factory automation and the willingness to spend on them will be the difference between those companies that win and those that lose in the coming years. Perhaps this reticence to invest in something like AI could be attributed to the lack of understanding of its ROI, capabilities, or real-world use cases. Industry analyst Gartner, Inc. still slots many of AI's applications into the "peak of inflated expectations" after all.


But AI, specifically deep learning-based image analysis or example-based machine vision, combined with traditional rule-based machine vision can give a factory and its teams superpowers. Take a process such as the complex assembly of a modern smartphone or other consumer electronic devices. The combination of rule-based machine vision and deep learning-based image analysis can help robotic assemblers identify the correct parts, help detect if a part was present or missing or assembled incorrectly on the product, and more quickly determine if those were problems. And they can do this at an unfathomable scale.

The combination of machine vision and deep learning are the on-ramp for companies to adopt smarter technologies that will give them the scale, precision, efficiency, and financial growth for the next generation. But understanding the nuanced differences between traditional machine vision and deep learning and how they complement each other, rather than replace, are essential to maximizing those investments.

AI, specifically deep learning or example-based machine vision, combined with traditional rule-based machine vision can give a factory and its teams superpowers.



WHAT IS DEEP LEARNING?



Simply put: deep learning allows for solving specific tasks without being explicitly programmed to do so.

Deep learning is a subset of artificial intelligence and a part of the broader family of machine learning. Instead of humans programming task-specific computer applications, deep learning uses data and then trains it via neural networks to make more accurate outputs based on that training data. Simply put: deep learning allows for solving specific tasks without being explicitly programmed to do so.

Deep learning, then, isn't just some far flung technology to help humans in the future. It's solving problems—both mundane and important—right now: facial recognition to unlock phones or identify friends on social media photos, recommendation engines on streaming video and music services or when shopping at ecommerce sites, diagnosing diseases, spam filters in email, and credit card fraud detection.



Mastering a complex strategy game

Go is an abstract and complex strategy board game for two players invented in China more than 2,500 years ago. Compared to Chess, Go has a larger board, longer play, and more potential decisions to consider with every move played. It was generally thought that no computer would ever be able to master the intricacies involved with the game of Go, unlike Chess. That belief, however, was put to rest in the spring of 2016, when Google's AlphaGo beat one of the best human Go players in the world four games to one in a best-of-five series.

How did the AlphaGo team do this? Through deep learning and neural networks. The system initially learned the game through supervised learning using human player data, but then trained with reinforcement learning by playing games against itself and using that data to further improve its gameplay.

Deep learning thrives in being able to consistently and at scale recognize anomalies and variance amongst a set of data. It's something humans do intrinsically well—spot what's different—but that until now computer systems based on rigid programming weren't good at. Computers, however, do not tire easily in their decision-making on the assembly line—unlike a human inspector.

The explosion of deep learning technologies this decade owes itself in no small part to the explosion in popularity to modern video games. According to MIT, “the complex imagery and rapid pace of today's video games require hardware that can keep up, and the result has

been the graphics processing unit (GPU), which packs thousands of relatively simple processing cores on a single chip. It didn't take long for researchers to realize that the architecture of a GPU is remarkably like that of a neural net.”


These modern, high-performance GPUs enabled the 50-layer neural networks of today. And the new, low-cost GPU hardware has made it practical to deploy biology inspired, multi-layered “deep” neural networks that mimic the human brain. Starting from a core logic developed during initial training, deep neural networks can continuously refine their performance as they are presented with new images, speech, and text.



What is a neural network?


At this point, you might be wondering “what is a neural network?” Essentially, it's a computer system modeled on the connections of the human brain.

When a neural network is being trained, training data is fed to the input layer and it passes through succeeding computational layers, or connections, getting multiplied and added together in complex ways, until it finally arrives, radically transformed, at the output layer. During training, the weights and thresholds are continually adjusted until training data with the same labels consistently yield similar outputs.



HOW DEEP LEARNING SYSTEMS LEARN


According to O'Reilly Media, there are five broad categories for machine learning algorithms:



Supervised learning

consists of mapping input data to known labels, which humans have provided. The


recommendation engines of streaming music and movie services use supervised learning techniques.



Unsupervised learning

is where the input data is unlabeled and the system tries to learn structure from that data automatically,


without any human guidance. Anomaly detection, such as flagging unusual credit card transactions to prevent fraud, is an example of unsupervised learning.



Semi-supervised learning


is often a combination of the first two approaches. That is, the system trains on partially labeled input data—usually

a lot of unlabeled data and a little bit of labeled data. Facial recognition in photo services from Facebook and Google are real-world applications of this approach.



Reinforcement learning is mostly a research area, but industry use cases are starting to emerge. Reinforcement learning occurs when a computer system receives data in a specific environment and then learns how to maximize its outcomes. Google's

DeepMind AlphaGo computer, which successfully learned to master the game Go, is a recent example of this technique.



Transfer learning involves reusing a model that was trained while solving one problem and applying it to a different but related problem. An example of transfer learning is where a deep learning model was trained on millions of images of cats, then

“fine-tuned” to detect melanoma in medical imaging.

Deep learning technology is being used to predict patterns, detect variance and anomalies, and make critical business decisions. This same technology is now migrating into advanced manufacturing practices for quality inspection and other judgment-based use cases.


If implemented for the right types of factory applications, in conjunction with machine vision, deep learning can benefit factory and manufacturing companies by scaling manual processes; especially when compared against investments in other emerging technologies that might take years to payoff.

HOW DOES DEEP LEARNING DIFFER FROM MACHINE VISION?

At a fundamental level, machine vision systems rely on digital sensors protected inside industrial cameras with specialized optics to acquire images. Those images are then fed to a PC so specialized software can process, analyze, and measure various characteristics for decision making.

Traditional machine vision systems perform reliably with consistent, well-manufactured parts. They operate via step-by-step filtering and rule-based algorithms that are more cost-effective than human inspection at scale. They can be executed at extremely fast speeds and with great accuracy. On a production line, a rule-based machine vision system can inspect hundreds, or even thousands, of parts per minute. The output of that visual data is based on a programmatic, rule-based approach to solving inspection problems.

In a factory setting traditional, rule-based machine vision is ideal for:



Guidance

Locate the position, orientation, and key feature of a part to leverage other vision inspection tools downstream.



Identification

Read barcodes, data matrix codes, direct part marks, and characters printed on parts, labels, and packages.



Gauging

Calculate the distances between two or more points or geometrical locations on an object and determines whether these measurements meet specifications.



Inspection


Find flaws or other irregularities such as the absence of safety seals, broken parts, etc.

Rule-based machine vision is great with a known set of variables: Is a part present or absent? Exactly how far apart is this object from that one? Where does this robot need to pick up this part? These jobs are easy to deploy on the assembly line in a controlled environment. But what happens when things aren't so clear cut?

Enter deep learning, which combines the scalability and reliability of traditional machine vision with humans' innate ability to detect variance.

Deep learning uses an example-based approach instead of a rule-based approach to solve for certain factory automation challenges. By leveraging neural networks to teach a computer what a good image is based on a set of labeled examples, deep learning will be able to analyze defects, locate and classify objects, and read printed markings, for example.

In the real world, that means a company might be trying to inspect electronic device screens looking for scratches, chips, or other defects. Those defects will all be different in size, scope, location, or across screens with different backgrounds. With deep learning it's possible to tell the difference between a good part and a defective one, considering those expected variations. Plus, training the network on a new target, like a different kind of screen, is as easy as taking a new set of reference pictures.



That makes deep learning particularly adept to:

- Solve vision applications too difficult to program with rule-based algorithms
- Handle confusing backgrounds and variations in part appearance
- Maintain applications and re-train with new image data on the factory floor
- Adapt to new examples without re-programming core networks

Deep Learning Compared to Other Inspection Methods



Compared to Human Visual Inspection, Deep Learning is:

More consistent

Reduces inconsistencies between different human inspectors.

More reliable

Operate more reliably even when scaled or reproduced to other lines.

Faster

Identifies defects in milliseconds, supporting high-speed applications and improving throughput.



Compared to Traditional Machine Vision, Deep Learning is:

Designed for hard-to-solve applications

Solves complex inspection and classification applications impossible or difficult with classic rule-based algorithms.


Easier to configure

Applications can be set up quickly, speeding up proof of concept and development.

Tolerates variations

Handles defect variations for applications that require an appreciation of acceptable deviations from the control.

Complex surface textures and variations in part appearance introduce serious inspection challenges. Rule-based machine vision systems struggle to appreciate variability and deviation between very visually similar parts. “Functional” anomalies, which affect a part’s utility, are almost always cause for rejection, while cosmetic anomalies may not be, depending upon the manufacturer’s needs and preference. Most problematically, these defects are difficult for a traditional machine vision system to distinguish between.





A piston’s welding seams are highly variable, making abnormalities difficult to distinguish. Certain welding anomalies, like missing or underfilled welds, cracks, or porosity, are unwanted. Other anomalies, like overlapping seams, are desirable and required for safety reasons. Dark image areas introduce additional complications. Given the many possible flaws and lighting challenges, deep learning-based analysis offers a simple and robust alternative to traditional machine vision inspection.

Certain traditional machine vision inspections, such as defect detection, are notoriously difficult to program due to multiple variables that can be hard for a machine to isolate such as: lighting, changes in color, curvature, or field of view. This is not a problem, in and of itself, but it is problematic when companies attempt to solve applications with traditional machine vision when there are now other appropriate tools available to them.

While traditional machine vision systems perform reliably with consistent, well-manufactured parts, the applications become challenging to program as exceptions and defect libraries grow. In other words, at a certain point some applications needed for factory automation will not be best served by relying on rule-based machine vision.

Understanding those differences will be vital for any company embarking on a factory automation journey. Because those differences are key to determining when it makes sense to leverage one or the other in a factory automation application.


While deep learning is transforming factory automation as we know it, it's still just another tool that operators can employ to get the job done. Traditional rule-based machine vision is an effective tool for specific job types. And for those complex situations that need human-like vision with the speed and reliability of a computer, deep learning will prove to be a truly game-changing option.




The high-level differences between traditional machine vision and deep learning include:

- Development process (tool-by-tool rule-based programming vs. example-based training)
- Hardware investments (deep learning requires more processing and storage, for example)
- Factory automation use cases appropriate for each tool

Deep Learning solves applications that traditional vision cannot



Deep Learning makes solving tough applications easier



DEEP LEARNING'S BENEFITS FOR INDUSTRIAL MANUFACTURING

Rule-based machine vision and deep learning-based image analysis are a complement to each other instead of an either/or choice when adopting next generation factory automation tools. In some applications, like measurement, rule-based machine vision will still be the preferred and cost-effective choice.

For complex inspections involving wide deviation and unpredictable defects—too numerous and complicated to program and maintain within a traditional machine vision system—deep learning-based tools offer an excellent alternative.

To learn more about Cognex deep learning solutions, visit
cognex.com/ViDi-deep-learning





www.isotron.eu
info@isotron.eu

Isotron Systems B.V.
Afrikalaan 21-23
5232 BD, 's-Hertogenbosch
Nederland
Tel: +31(0)73 639 16 39

Isotron Systems BVBA
Antwerpsesteenweg 45
B-2830, Willebroek
België (luxemburg)
Tel: +32(0)3 450 70 45

COGNEX

© Copyright 2019, 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. DLvsMV-EN-07-2019