Introduction to Machine Vision

COGNEX

Introduction to Machine Vision

Table of contents

What is machine vision?

Machine vision is an essential component of how digital systems interact with the real world. It lets automated systems see components, products, patterns, codes, or other objects and use that information to make decisions.

Since it allows manufactured parts and products to be inspected, measured, and sorted, machine vision has vastly increased the power and flexibility of industrial automation. Machine vision does all this at high speeds and high accuracy, improving product quality and reducing waste.

Machine vision also lets automated equipment locate objects, identify them, and save information about their material, condition, orientation, and other details for later analysis. That data is critical to factories looking for efficiency gains.

And every package that makes its way through today's automated logistics warehouses does so with the help of barcodes that are read by machine vision. Those codes are used to track and identify packages throughout their journeys, making sure they get to their intended destinations.

All of these activities are powered by machine vision: cameras that let computers interpret real world objects. You may also hear products in this category referred to as "computer vision," which is a broader term and sometimes describes more theoretical research systems that analyze images. Machine vision products, on the other hand, are practical, connected to industrial automation systems in factories and warehouses.

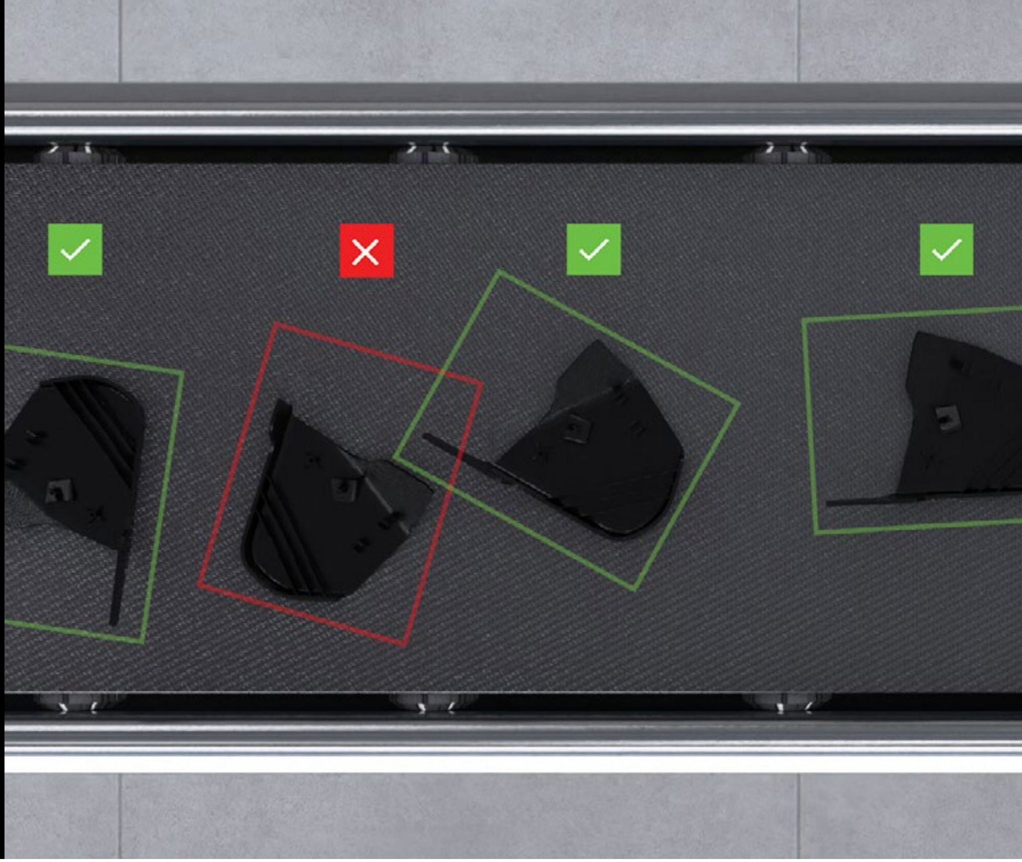

Machine vision lets automated systems see components, products, patterns, codes, or other objects and use that information to:

- **Make decisions** at high speeds and high accuracy, improving product quality and reducing waste.
- **Generate data** that is critical to factories looking for efficiency gains.

Who uses machine vision?

Machine vision is constantly evolving, and performs an ever-growing range of tasks in two main types of organizations:

Manufacturing facilities and factories
use machine vision to automate operations throughout the manufacturing process.

Logistics providers and warehouses
use machine vision to automate routing, tracking, and shipping, typically by reading barcodes or other encoding symbols.

What are the benefits of machine vision?

Because it can see with better resolution, across a wider spectrum, and process tasks faster than human vision, machine vision provides a wide range of benefits:

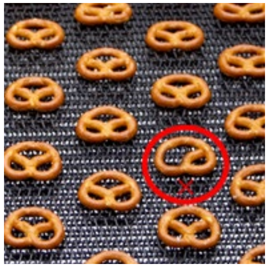

Improve product quality

Products rejected for defects are a significant source of cost, waste, and reputational damage. Automated inspection with machine vision improves speed and accuracy, catching problems of many kinds before they're packaged or shipped and allowing human inspectors to be reserved for an increasingly small number of difficult cases.

Trace parts and products

By reading codes on products and packages at every step from initial production to shipping to final sale, machine vision systems can provide critical tracking information. This lets shippers know their current location, quickly detect any delays or shipping errors, and trace any damage or other problem back to its source.

Increase productivity/overall equipment effectiveness

Machine vision systems speed up operations and decrease cycle times, and their performance doesn't deteriorate over the course of a shift. They provide the information to make the most efficient use of every piece of equipment on the floor.

Reduce waste

By catching manufacturing flaws, identifying overfill, or pinpointing the causes of defects, machine vision can reduce waste and scrap rates in multiple ways. Over time, this can help control overhead and bring down raw materials costs.

Overcome labor shortages

Manufacturing is facing significant labor shortages, and the problem is expected to get worse. The work can be tiring and difficult. By automating manufacturing with machine vision, skilled labor can be reserved for higher-value activities.

Improve processes

Machine vision instantly detects changes in product quality and keeps a visual record of every step in a product lifecycle. This form of big data reveals process bottlenecks, declining machine function, and common sources of error, making continuous process improvement possible.

Ensure compliance

The machine vision-generated data and images used to make process improvement decisions also provide the data needed to comply with reporting regulations in industries like pharmaceuticals, medical devices, automotive, food and beverage, and more.

Improve safety

With industrial automation made possible by machine vision, workers can minimize their interactions with noisy and potentially dangerous machinery. If there is an equipment problem, they can often diagnose and adjust it remotely, through the information machine vision provides.

Types of machine vision systems and what they can do

While machine vision systems offer a wide range of features and options, there are three main categories to consider.

Line scan

Line-scan cameras are used in continuous inspection applications such as web manufacturing. They take a wide but very thin image, typically as material moves past the scan area, and use software to reconstruct the image line by line.

They are significantly faster in these applications than standard 2D cameras. Examples include inspecting fabric, paper, and other soft goods.

Area scan

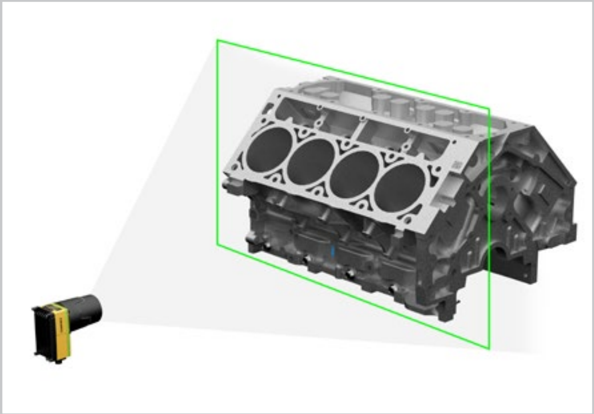

Most machine vision uses 2D cameras, also called area scan systems. These can be simple, single-purpose sensors or more full-featured systems. Sensors are lower cost, smaller, easier to deploy, and often more rugged than cameras with a wider range of features. Generally, full-featured systems are more configurable, perform more complicated tasks, and can take larger and higher resolution images.

The decision between sensors and more complex machine vision systems comes down to the task being performed, the format of the data output, cost, and ease of use.

3D vision

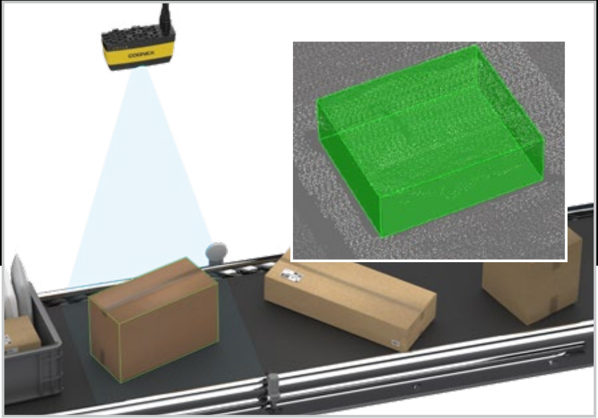

3D vision systems add depth to their images, sometimes using lasers to measure distances and compute depth. The addition of depth can increase complexity and cost to some degree, but that additional information is critical to some applications.

For example, 3D vision is essential in guiding a robot arm to precisely reach out to grab an object in the correct orientation, no matter where it is in space. 3D vision is also used to automate difficult cutting and welding processes.

What does machine vision do?

Machine vision providers sometimes classify what their systems can do with the acronym GIGI, for Guidance, Identification, Gauging, and Inspection. This guide breaks down the types of applications into slightly more detail:

- **Defect detection**
- **Object detection and counting**
- **Measuring/gauging**
- **Locating/guiding/positioning**
- **Barcode reading**
- **OCR/OCV**

Defect detection

Defects can occur in any part of a manufacturing process, from problems with the quality of raw materials or parts through final inspection. Inspections to catch defects have traditionally been done manually, by trained workers.

Machine vision is a significant improvement over human inspection: it operates at production line speeds, doesn't get tired, can detect even small and unexpected defects, and stores information for continuous operational improvement.

EXAMPLES:

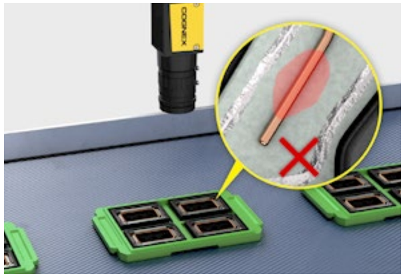

Looking for loose connectors, poorly soldered wires, bad seams, or improperly crimped tubes

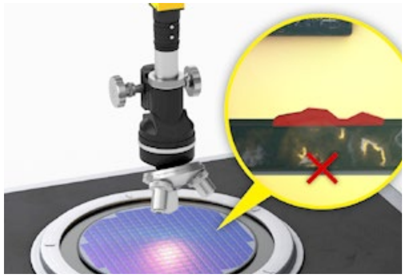

Detecting flawed photovoltaic cells in solar panels or defects in semiconductor wafers or EV battery assemblies

Finding contaminants or other problems in food products

Object detection and counting

Determining the presence or absence and counting objects is a widely used function in inventory management, on production lines, and before releasing or accepting shipments. Both manual inspection and mechanical counting are slow and prone to error, compared to high-speed and consistently accurate machine vision systems.

EXAMPLES:

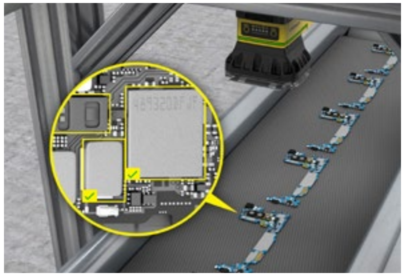

Confirming the presence of all electronic components on a printed circuit board

Verifying the presence of components such as clips, screws, springs, labels, seals, manuals, inserts, or accessories

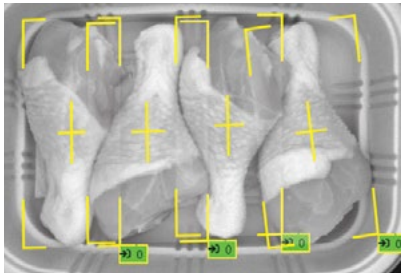

Counting products in a package or on a pallet

Measuring/gauging

Precise manufacturing requires accurate measurement of distances, areas, diameters, and more. Manual measurement with gauges, calipers, or inspection jigs is slow and introduces inconsistencies.

Machine vision consistently and accurately measures down to the micron level, while parts are on the line, at high speeds. Each image and its associated data can be stored in case of warranty or compliance issues.

EXAMPLES:

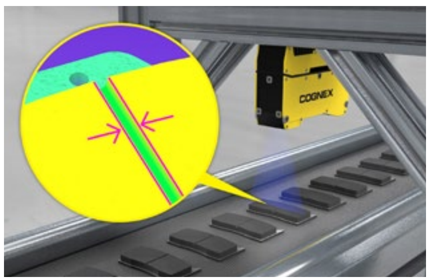

Capturing the dimensions of cast and injection molded parts

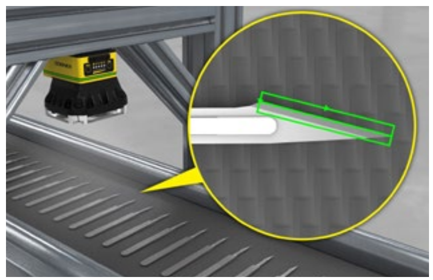
Measuring the roundness and angle of tips on parts

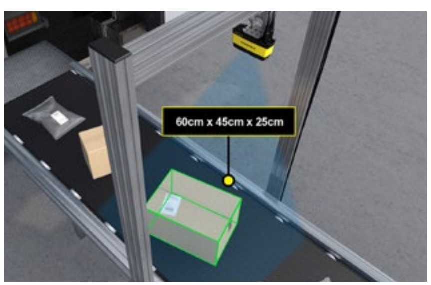

Determining label positions or package sizes

Locating/guiding/positioning

Functions such as assembly, pick-and-place, and inspection depend on machine vision's ability to locate a part, whether on a conveyor or in a bin. Machine vision guides robots in picking and placing parts and ensures accurate micron-level positioning in precision assembly.

If a part is out of place, machine vision measures the difference between desired and actual location and orientation and communicates that to a robot or programmable logic controller (PLC) to realign the part.

EXAMPLES:

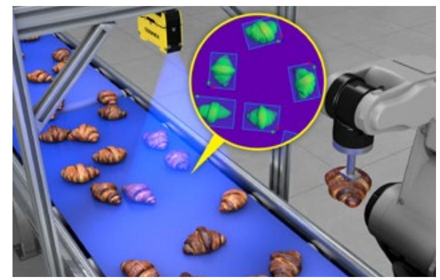
Locating parts on a conveyor for inspection

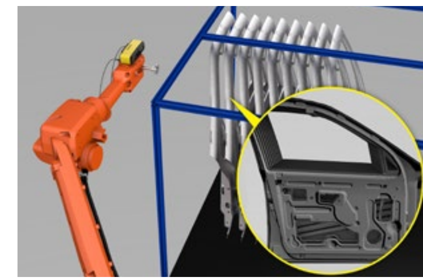
Guiding robots on automated automotive assembly lines


Assembling microchips that require micron-level precision

Barcode reading

Barcodes are used to track and identify raw materials and finished goods throughout the supply chain. They include 1D and 2D codes, such as UPC codes on retail products and data matrix codes on packages.

Barcode readers, process information at high speeds, and correctly read codes that are partially torn, obscured, smeared, or distorted. Image-based barcode scanners collect images of the barcodes they read, allowing for analysis of no-reads or misreads to diagnose problems, such as a clogged print head, damaged codes, or inadequate lighting. Multiple camera systems can scan from multiple sides of a part or package at once to increase read rates when not all codes are oriented the same way.

EXAMPLES:

Tracking packages as they travel through a logistics warehouse

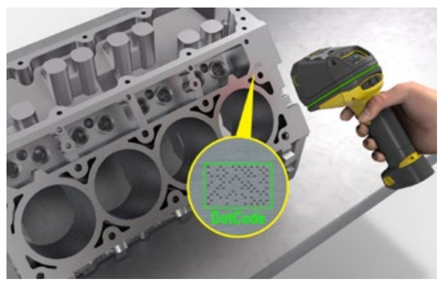

Ensuring the correct components are assembled

Providing accurate traceability of medical supplies

Reading text: OCR/OCV

Barcodes are everywhere in modern industry, but human-readable printed text is still essential, particularly for retail, manufacturing, pharmaceutical, and food and beverage supply chains, to identify sell-by dates, lot numbers, and other important information. For this text to be useful in modern high-speed processing, it must be reliably readable by machines. Machine vision systems can read these codes in milliseconds with 99.99% accuracy.

Optical character recognition (OCR) and optical character verification (OCV) both identify and interpret text in images, with one key difference. OCR reads text to trigger another action, while OCV is used to verify the quality of text against a known standard.

EXAMPLES:

Classifying a part based on the characters printed on it (OCR)

Checking if a sell-by date or lot number is printed correctly (OCV)

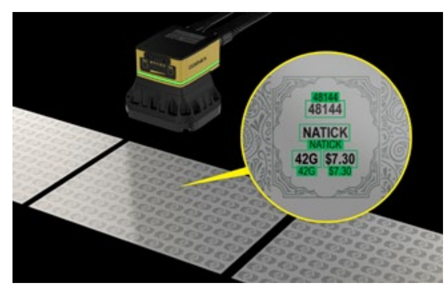

Distinguishing authentic from counterfeit products (OCR)

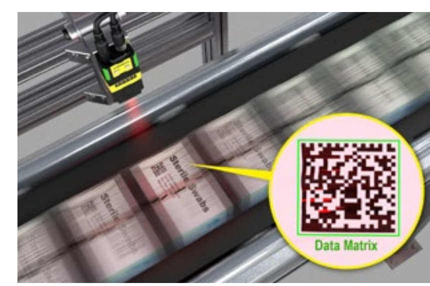

Establishing traceability to meet regulatory demands (OCR)

Rule-based vs. AI-powered machine vision

There are two ways machine vision can be used to make decisions such as counting, classifying, or approving and rejecting items. Rule-based systems follow user-programmed, step-by-step instructions to interpret images and make decisions. In contrast, artificial intelligence or AI-powered systems use a database of reference images to “learn” how to make decisions.

While rule-based machine vision is still the prevalent technology, various types of AI-powered machine learning have become capable and flexible enough to take over in many applications.

Often, a combination of rule-based and AI-powered machine learning can provide the most efficient solution.

Rule-based systems

Traditional rule-based machine vision uses specific if-then rules to make decisions about an image. “If the bottle is filled to at least this level, let it pass.” The rules are programmed by vision engineers with deep knowledge of the best way to get the desired output.

The rules are used to create programs that tackle specific jobs, such as detecting the edge between light and dark areas or measuring the distance between two points.

These rules can then be used one after another on an image, in what is called “tool chaining,” to perform a wide range of sophisticated tasks. They’re also very good at high-speed, high-accuracy inspections on products or parts that are predictable and consistent.

Rule-based systems work on consistent parts

Technical expertise: Vision expert with programming knowledge

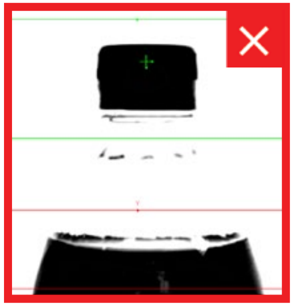

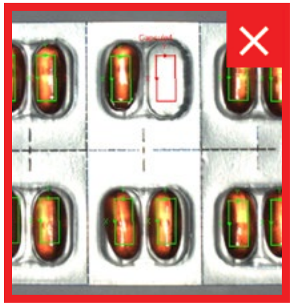

Often, a combination of rule-based and AI-powered machine learning can provide the most efficient solution.

Deep learning

Deep learning uses AI to power machine vision systems by using examples to train the software with labeled images until it can make distinctions on its own.

This training mimics how humans learn, helping the system learn how to make accurate decisions without being thrown off by irrelevant variation.

Operators train deep learning systems with hundreds of labeled images, such as a variety of possible defects in a manufactured part, or several different assemblies that need to be classified.

One significant improvement over rule-based systems is that deep learning systems learn to distinguish between a wide range of real defects and merely cosmetic variations without needing to be shown every possible result — something time-consuming or even impossible to accomplish using rule-based programming.

In contrast to rule-based systems, deep learning systems don't require experienced machine vision programmers or expertise to set up. They do need someone with good knowledge of the products being inspected to provide hundreds of images of both defective and good components.

Deep learning excels at assembly verification, defect detection, and classification of complex parts that vary unpredictably. It's also good at providing accurate results even when parts are reflective or images are distorted.

Deep learning is designed for complex applications

Example images required: hundreds to thousands
Total time: Hours to days
Technical expertise: Training in AI and programming

Use case training: hundreds of example images

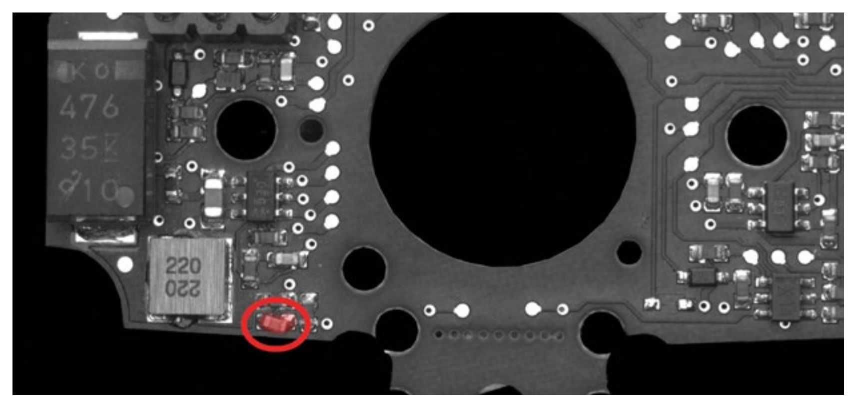

Results: complex applications with significant variation

Edge learning

Edge learning is another type of AI, optimized to meet the needs of industrial automation. It comes pre-trained to solve the types of problems typically faced by industrial automation.

As a result of the pre-training, an edge learning system can be trained on a specific industrial inspection problem with as few as five images, making it even easier and faster to deploy than a deep learning system.

Edge learning enables line engineers to quickly implement an optimized classification or defect detection function on their line, without specialized training. Since it doesn't require a sophisticated processor, it can be quickly deployed where it is needed.

The advantages of edge learning lie in its speed, low demand on computing resources, and flexibility, as well as its user-friendliness.

Edge learning is designed for ease of use

Example images required: 5-10
Total time: Minutes
Technical expertise: No prior experience needed

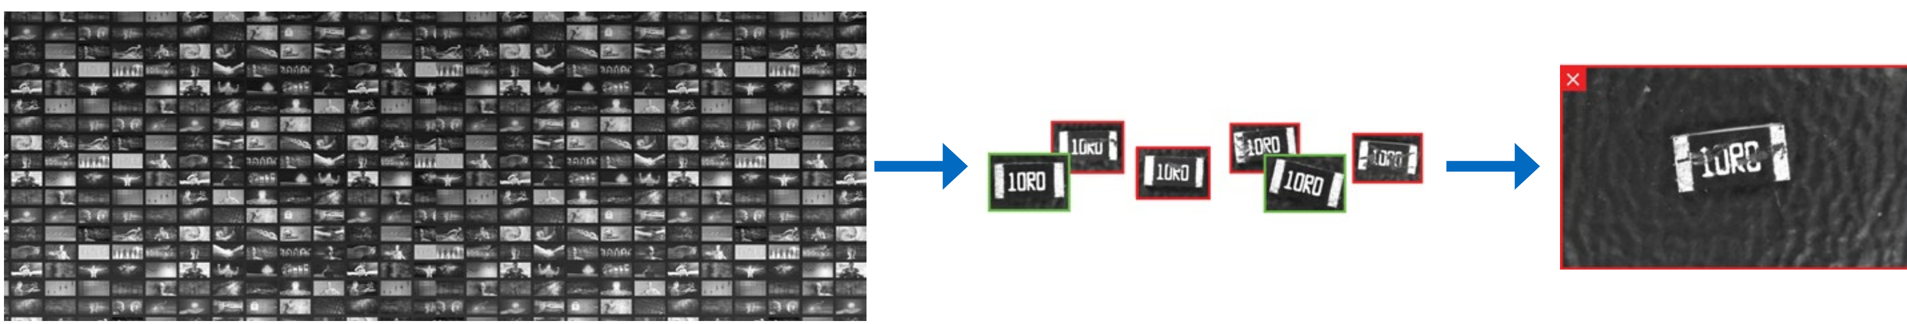

Pre-loaded training and optimization **Use case training: 5-10 examples** **Results: easy, fast**

Comparing rule-based and AI-powered systems

The decision between rule-based, deep learning, and edge learning systems mainly comes down to the type of application you're implementing.

- Rule-based vision tools are effective at a variety of specialized tasks with very consistent targets, such as location, measurement, and orientation.
- Edge learning excels at repeatable tasks with moderately consistent targets.
- Deep learning shines when looking for complex defects or analyzing images with significant variation.

Recommended uses for each type of system

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

Machine vision hardware components

Choosing the right hardware for a machine vision system

While most vision systems include a similar collection of hardware, there are specific choices to be made for each component and they all need to work seamlessly together.

The exact configuration of any machine vision system is dictated by the job requirements: by starting with the application, you can work with a vendor to determine the precise setup you need.

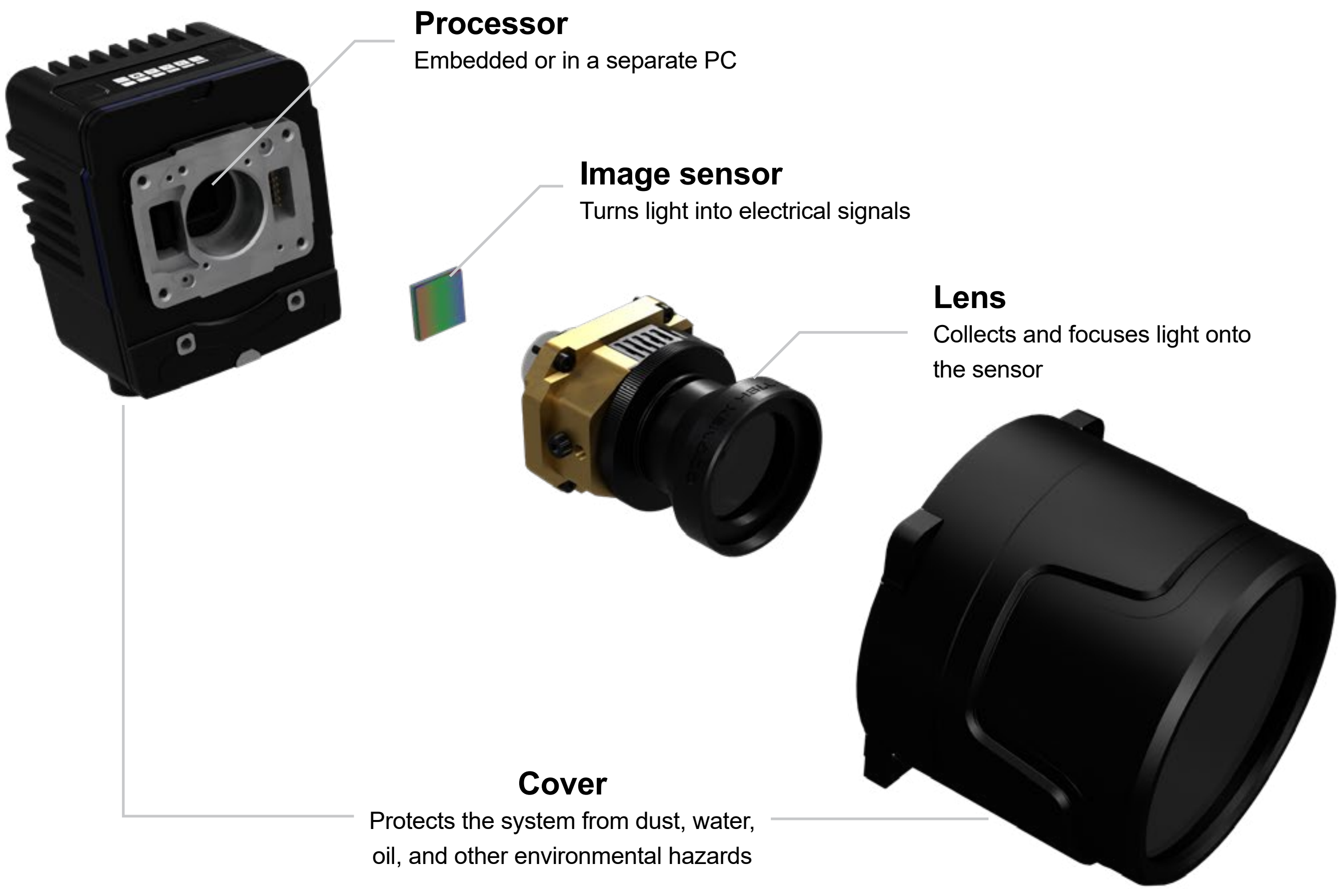

Image sensors

At the center of any vision system, an image sensor converts light energy into electrical signals that can be analyzed by software.

Image sensors are solid-state semiconductor chips comprising millions of photodetectors called pixels, which convert light into electrical signals. The two main technologies used are charge-coupled device (CCD) or complementary metal oxide semiconductor (CMOS). CMOS technology is newer and taking an increasing share of the market, thanks to lower costs and faster processing speeds, but CCD sensors are still used in applications that require extremely high-quality images.

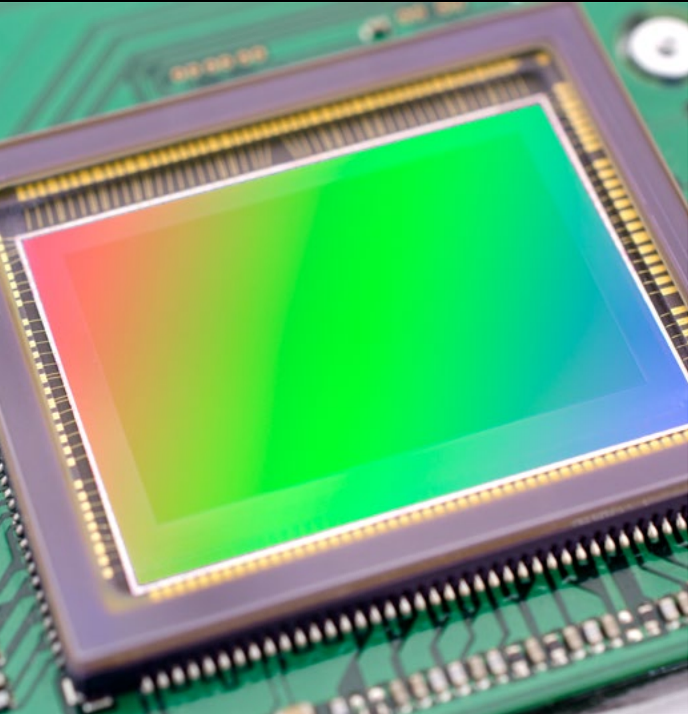

Mono vs. color sensors

Machine vision image sensors are monochromatic, meaning each pixel only detects the intensity of the light falling on it, not what color it is. They can capture color images using red, green, and blue filters, but in the vast majority of industrial applications, monochromatic imagery is a more effective choice. Since color filters reduce light intensity, they're only used when color is absolutely required for a specific task.

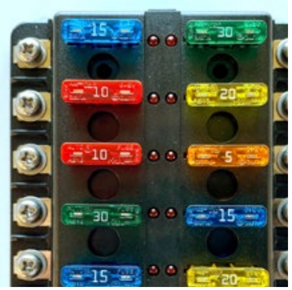

Rolling vs. global shutters

Machine vision cameras use electronic shutters to determine how the pixels are exposed to light. Despite the use of the word "shutter," as in an old mechanical camera, the shutter is part of the way a sensor processes the light falling on it, not something separate.

A rolling shutter image sensor exposes lines of pixels sequentially. This kind of sensor is lower cost, but if the target being imaged is moving rapidly, the image can be distorted.

A global shutter image sensor can expose all of its pixels at once. They can be more complex to manufacture, but prices have come down and they are widely used in industrial applications.

A rolling shutter exposes a single line of pixels at a time.

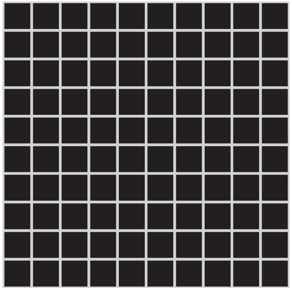

A global shutter exposes all of its pixels at once.

Additional sensor features

Sensors range widely in a number of characteristics, including size, resolution, pixel size, frame rate, sensitivity, and dynamic range. Each choice can affect other characteristics, as well as cost. Sensors can also have a wider spectral range than the human eye and detect light into the infrared (IR), useful in some industrial applications.

There's no single answer to "which sensor should I use?" Sensors should be chosen based on their specific purpose, whether for a fast-moving assembly line, where a high frame rate might be essential, or detecting tiny details in a complex product, where resolution and dynamic range might be more important.

Lenses

Lenses collect and focus light onto image sensors. They are matched with the sensor so they completely illuminate the photosensitive area with a resolution equal to or greater than the sensor's pixel size.

Manual lenses need to be refocused if the objects being observed change, or if the camera moves.

But this can be time-consuming if done frequently, or if there are a large number of cameras that need to be adjusted. In these situations, autofocus lenses are optimal.

Autofocus lenses adjust their focus automatically as needed, saving time if frequent changes are expected.

Liquid lenses are a specialized subset of autofocus lenses that adjust their shape on the fly to maintain sharp focus. They're also smaller and more resistant to shock and vibration than mechanical autofocus lenses of similar spec.

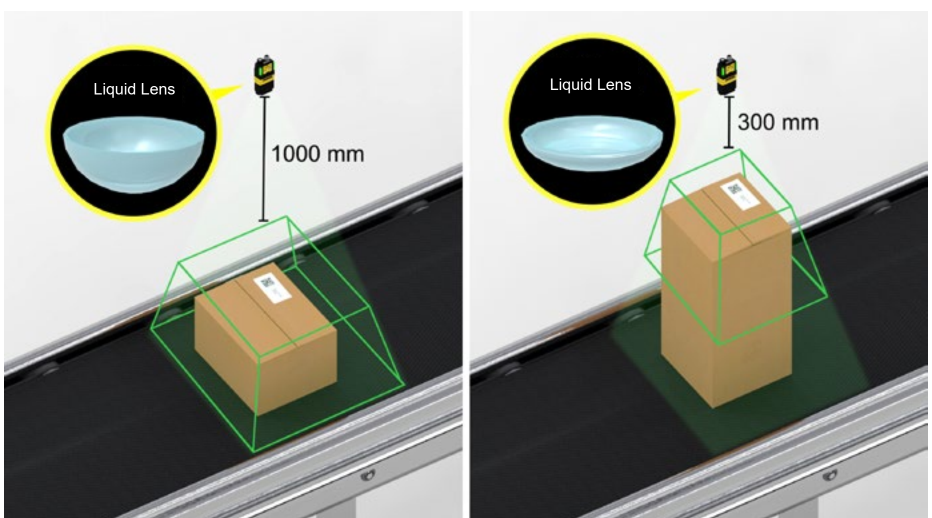

Covers

A camera cover, also called a housing or enclosure, protects your expensive machine vision camera from the dust, oil, water, food particles, impacts, and other hazards that come with operating in an industrial environment.

Covers can also dissipate heat, protect cable connections, and prevent tampering.

Covers are rated by how well they protect the camera, with IP67 being the usual rating for a dustproof camera in an industrial environment. An IP67 rating also means the enclosure is waterproof, so the operating area can be washed down safely.

Hardware requirements

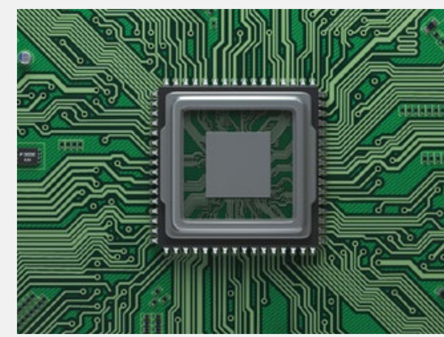

Another factor in choosing the type of processing is the required hardware. Because of the amount of parallel processing it requires, deep learning generally needs to run sophisticated parallel processors called GPUs. Edge learning requires less processing, and rules-based algorithms the least.

Some industrial cameras have onboard or embedded processing. Such cameras can typically do both rule-based and edge learning processing entirely on their own. Deep learning may require a more powerful processor located in a separate PC.

The importance of lighting in machine vision applications

Poor lighting is the most common cause of poor machine vision performance. Even sophisticated cameras and software can't make up for inadequate lighting.

Good lighting for machine vision maximizes contrast on the features of interest and minimizes it everywhere else. It also needs to be consistent; the lighting setup should make sure normal variations in parts or their arrangement don't affect the contrast in the image.

Depending on the part's texture (such as matte, reflective, or refractive) and its shape, different lighting choices will change which portions of the image will be brighter than others. Poorly chosen lighting might create glare, when light reflects off a smooth part directly into the camera. A different position might occlude an important spot or leave it in a shadow.

Some cameras come with **integrated lights**, which provide uniform illumination that's important when working with matte objects, among other uses. They're easy to deploy and work well in many situations.

Other conditions can require a separate lighting setup, which can involve experimentation with different lighting angles and types to get the most effective image. There's no one best lighting setup for all machine vision applications. Instead, lighting optimization depends on the application, specific requirements, and the desired results.

Good lighting for machine vision maximizes contrast on the features of interest and minimizes it everywhere else.

Types of machine vision lighting

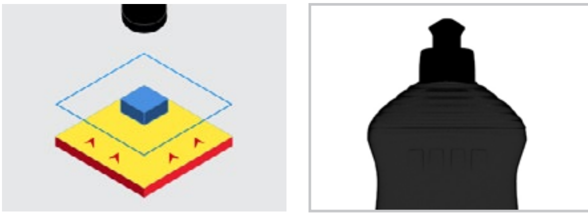

Backlighting

A flat, diffuse light behind the target forms a black silhouette with maximum contrast and clear edges.

Full/partial bright field lighting

Lighting an object from above creates strong contrast so flat parts reflect into the camera.

Dark field lighting

Shining a light from the side at a low angle (like a rising or setting sun) helps show edges, height variations, texture, and contours. Background areas reflect light away from the camera, and appear dark, while the camera receives light from the features of interest.

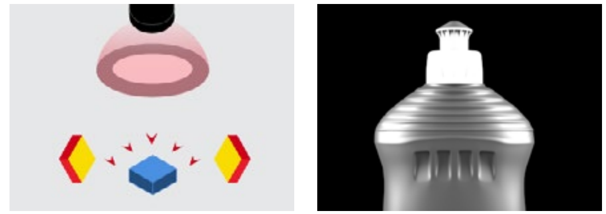

Diffuse lighting

Even lighting from multiple directions minimizes confusing reflections from a curved or shiny surface. Also called cloudy-day lighting, diffuse lights can be domed, axial, or flat, and help create an even background, making it easier to detect features of interest.

Multicolor/multispectral lighting

Using infrared (IR) lighting can make it easier to find small surface defects within visually busy images, see differences in material, and work in environments with challenging visual conditions.

Filters

Filters can keep unnecessary ambient light from entering a camera or increase the contrast of a part or other object.

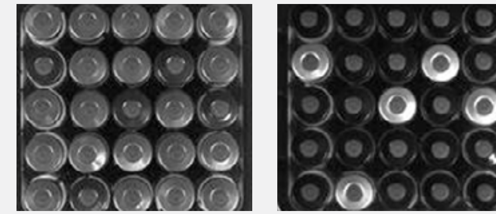

Color filters only let red, green, or blue light through. Depending on the objects being examined, one color may provide higher contrast.

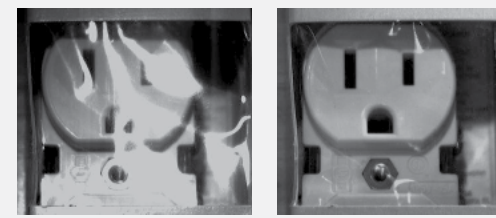

Polarizing filters can eliminate glare from light being reflected into a camera.

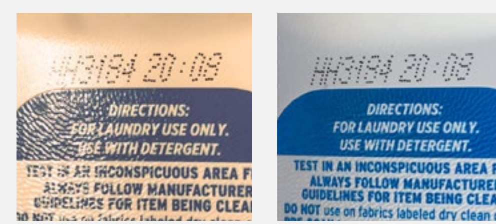

IR and UV filters only let through light outside the visible spectrum and can often show contrast differences that don't show up in visible light.

Image analysis software

With the right combination of sensor, lens, lights, and other components in place, a machine vision system is ready to capture images. But that's only the first step.

Once it acquires an image, the system needs to process and analyze the result to make a decision: to read a barcode, detect a defect, confirm that a kit contains all of its items, or measure a part.

Historically, the only way to get enough processing power to do image analysis was to send the image to an external PC. Depending on the size and complexity of the image, that could take time, potentially slowing down the line, requiring a separate PC dedicated to the task.

Today, highly capable vision systems have embedded processing, which speeds computation and reduces possible points of failure.

Types of image analysis tools

Depending on the task, a vision system can leverage one or several specific image analysis tools.

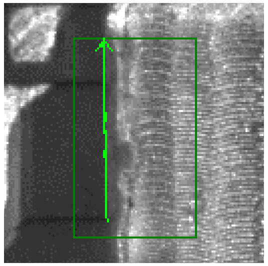

Edge detection/ extraction

In an image, an edge marks a distinct change in intensity. Such changes mark discontinuities in depth, orientation, and material, and allow for finding the boundaries of objects.

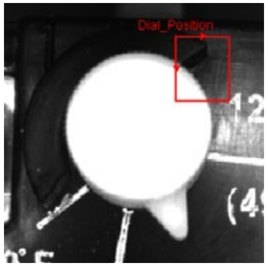

Pixel counting

A straightforward algorithm can count the number of pixels at each grayscale level in a region of interest or the overall image, providing information used by other functions.

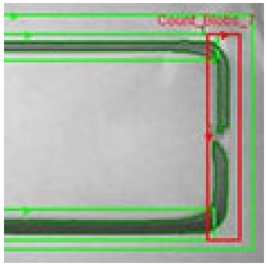

Blob detection

Different regions can be identified, providing information used by other functions.

Pattern/template matching

Guiding and positioning functions use pattern matching or template matching on parts and products to determine object location and orientation.

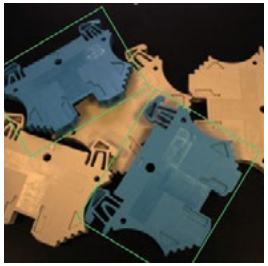

Color analysis

Where object color is an important feature, there are several color tools, including color extraction and segmentation, and color matching.

Putting image information to use

The information derived from that image analysis isn't much good if it stays in the camera. It needs to be shared with other systems that can put it to use.

To do this, machine vision systems use standard industrial formats or protocols to share data with other devices, usually via wired connections. That data can then be used in many ways:

It can tell a nearby programmable logic controller (PLC) to fire a piston to **bump defective parts off the line.**

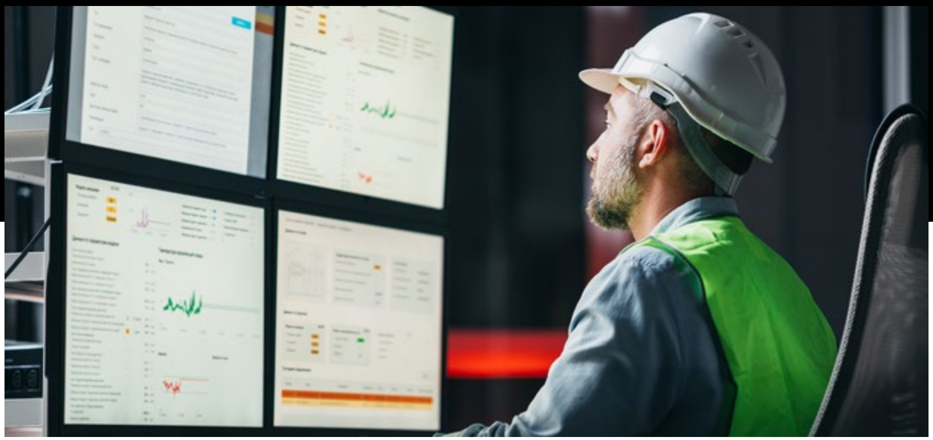

It can be used with the factory's process control system or a manufacturing execution system (MES), so the **system can make adjustments to optimize the manufacturing process.**

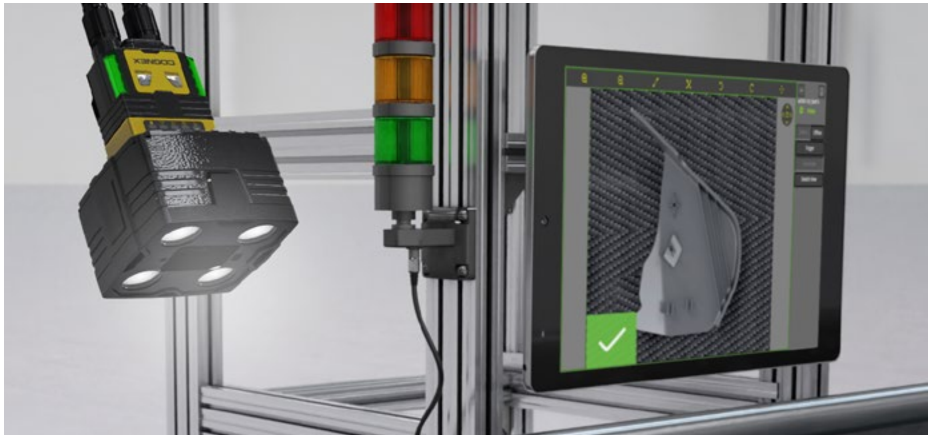

It can pop up on a human-machine interface (HMI) display so an **operator can check up on a production process and make adjustments.**

It can trigger sorting mechanisms in an automated sortation system to **get a package on to the right truck.**

Image-based data drives larger decisions, too

In addition to individual actions on single parts, products, or packages, the data generated through machine vision applications can also be aggregated and used to tackle larger business issues.

For example, businesses can use the data to find the root causes behind delays, defects, and missed shipments. Is one supplier using inferior raw materials for their part? Or is the problem caused by how parts are stored or shipped? Are workers being trained inadequately? Machine vision data can be analyzed to identify the source of the problem.

Other uses of machine vision data can include:

Learning that a certain type of part causes rejection of the assemblies it is installed in.

Locating and identifying parts, constantly updating inventory, and guaranteeing accurate order fulfillment.

Analyzing large data sets with an enterprise resource planning (ERP) system to make long-term operational decisions.

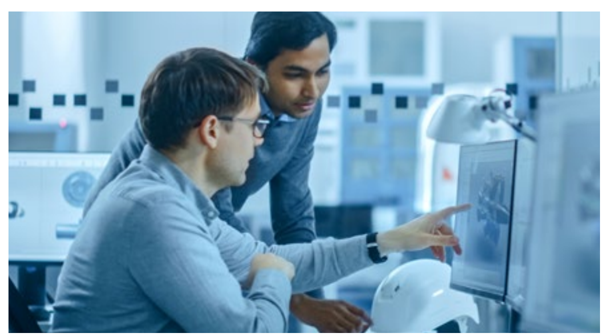

Reviewing failure rates and types to pin down where productions errors are being introduced.

Types of communication protocols

There is a wide range of protocols used in automated manufacturing, from simple on/off signals sent by a direct connection, to sophisticated industrial protocols such as Ethernet/IP, Profibus, and DeviceNet, among many others.

While interoperability is common, meaning you can buy the best device for a specific job and be confident it can communicate with all your current equipment, it remains true that various manufacturers favor one protocol or another, and there are a lot of proprietary standards.

Machine vision systems can be configured to be interoperable with virtually any set of industrial protocols, providing their information in a form that is easily used to improve automated industrial operations.

EtherNet/IP™

PROFI BUS®

DeviceNet™

Conclusion

Industrial automation enabled by machine vision is increasingly essential for competing in the market and meeting compliance requirements. Machine vision reduces costs, generates data for process improvement, increases OEE, enables new types of automation, and improves safety.

Machine vision constantly improves, both through technical innovations such as higher-resolution sensors, improved lenses, more capable processors, and smaller, more rugged casings, as well as software advances such as new AI algorithms that make machine vision more accurate, faster, and easier to use.

As these advances continue, machine vision systems will be used in more ways, by more types of businesses, with greater impact on overall results.

www.cognex.com

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	+43 800 28 16 32
Belgium	+32 289 370 75
Czechia	+420 800 023 519
France	+33 1 76 54 93 18
Germany	+49 721 958 8052
Hungary	+36 800 80291
Ireland	+353 21 421 7500
Italy	+39 02 3057 8196
Netherlands	+31 207 941 398
Poland	+48 717 121 086
Romania	+40 741 041 272
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-Pacific

Australia	+61 2 7202 6910
China	+86 21 5875 1133
India	+91 7305 040397
Indonesia	+62 21 3076 1792
Japan	+81 3 5977 5400
Korea	+82 2 539 9047
Malaysia	+6019 916 5532
New Zealand	+64 9 802 0555
Phillipines	+63 2 8539 3990
Singapore	+65 3158 3322
Taiwan	+886 02 7703 2848
Thailand	+66 6 3230 9998
Vietnam	+84 98 2405167