VISION INSIGHTS

VISION SYSTEMS DESIGN®

COURTESY OF PHIL-VISION

Machine Vision Fuels Robotics

COPYRIGHT © 2023 BY ENDEAVOR BUSINESS MEDIA. ALL RIGHTS RESERVED.

How Machine Vision and Robotics Are Automating Work Processes

Machine vision systems enable robotic automation in factories, warehouses and elsewhere by helping robots "see." In fact, increasingly sophisticated machine vision systems and technologies, particularly when paired with deep learning, are helping drive the growth of robotics for tasks such as picking, sorting, and placing items. These robots are more flexible than earlier-generation robots, which focused on one specific task.

Vision Systems Design routinely covers the use of machine vision in robotics with articles discussing the underlying systems and technologies and how they are used in specific applications.

Whether you work at a components manufacturer, system integrator or end-user organization, we've provided insights here to help fuel your next project.

Linda Wilson,
EDITOR IN CHIEF
VISION SYSTEMS DESIGN

FOLLOW US

X
@VISION_SYSTEMS

FACEBOOK
VISIONSYSTEMSDESIGN

LINKEDIN
VISION SYSTEMS DESIGN

AI Offers Advantages for Control-Oriented Vision Systems for Robotics

End-to-end-trained AI-based control offers higher reliability and precision than the traditional measure-and-execute paradigm.

RONNIE VUINE, Micropsi Industries

Industrial robots excel at executing movements with submillimeter precision, notably when they have a single, well-defined target in space to navigate to. They perform path planning flawlessly—optimized for speed, distance, wear, and precision. This approach has the additional advantage of being highly accessible for a qualified human operator: A single set of 3D coordinates in a robot base coordinate system is easy for human operators to understand. Tool rotation descriptions, such as Euler angles or quaternions, are less intuitive, but, with experience, still reasonably easy for operators to read and check for plausibility.

While precise execution is a hallmark of modern robotics, precise measurement, unfortunately, is not. Cameras often capture the wrong coordinates because slight variations in lighting conditions, shapes, and colors throw them off. Measuring space in three dimensions is hard even under perfectly controlled lighting. And, a speck of dust on the lens can flip a single pixel that was needed to make an accurate measurement.

IMAGES COURTESY OF MICROPSI INDUSTRIES.

FIGURE 1: For 3D problems, feature detection needs to be performed on depth images, typically point clouds, created by dual-camera stereo systems or time-of-flight-based cameras that emit infrared light and measure the time it takes the reflection to arrive back at the camera.

FIGURE 2: To get started with the MIRAI system, designers simply connect the robot, the F/T sensor, the camera(s), and the MIRAI controller. Set up time typically takes 30-45 minutes.

Perfectly executing a path to the wrong destination is still a bad result. Therefore, many robotics engineers have developed a love-hate relationship with cameras. When cameras perform reliably, the resulting robotics applications are impressive. But, getting cameras to perform reliably is difficult and sometimes impossible. So, to mitigate project risk and keep solutions simple, engineers often avoid using cameras altogether.

Relative Movements in Real Time

Unlike robots, human movement execution is fallible, and human measurement is too. And yet, we manage to solve problems of coordinated movements that are unfathomable to robots. It's because they're doing it wrong.

Humans don't measure, and they don't move to coordinates. They make relative, real-time controlled, rough movements toward the target, and they frequently correct. Instead of solving the extremely hard engineering problems of measuring to perfection once and then executing to perfection once, they solve a much simpler problem much more often—moving roughly toward the goal.

A similar strategy was only applied to robots recently because "move roughly toward the goal" was not a directive that a robot could understand—until the arrival of deep neural networks. Now, it's possible to train a neural network to understand the goal without having to explicitly specify it. And, it's possible to perform the neural network math calculations fast enough to mimic the human ability to "visual servo" toward the goal and make many small, quick corrections during the approach. Using a camera to guide the robot also brings with it an added benefit: Manual feature engineering is no longer necessary.

Feature Engineering

Feature engineering involves carefully applying human help to computer vision algorithms. An engineer thinks about features, such as edges or characteristic points, that the algorithms should pay attention to and then configures operations on the raw data—or even the physical world—to make the features easy to find. This requires having a lot of experience, a deep understanding of the algorithms, and knowledge of the tricks of the trade. The process is complicated and time-consuming.

Naive pattern matching is the most basic approach to measuring a 2D position. In this method, the camera is calibrated to detect the exact location of every image pixel. A location is determined by calculating the difference between a predefined pattern and all overlapping pattern-size patches of the camera image. These algorithms essentially count pixels to determine the part of the image that is most similar to what is being looked for.

More sophisticated algorithms—such as BRIEF, SIFT, or SURF—use filters to emphasize features or key points of an image and then perform pattern matching in feature space instead of on the raw pixels. This approach enables the algorithm to determine positions independently of scale and rotation, and the distance calculations between two potential matches can also

FIGURE 3: This photo shows an application with ZF, a glob al technology company headquartered in Friedrichshafen, Germany. It uses the Micropsi MIRAI robot control system to automate machine tending in a high-volume milling station where gears are manufactured.

be made a lot faster than they can in pixel space. For 3D problems, feature detection requires depth images, typically point clouds, created by dual-camera stereo systems or time-of-flight-based cameras that emit infrared light and measure the time it takes the reflection to arrive back at the camera.

Even in 2D, most of these techniques work best with some form of structured light projected onto the scene—from specific colors to make certain features pop, to grids to extract surface information. Small changes in pixel brightness or color easily throw off these algorithms. It is common to deliberately reduce image complexity by ignoring color (and all the useful information it may contain) or even by projecting light in such a way that a very robust bright reflective spot becomes a reference point for a feature description algorithm to latch on to. Sometimes this is performed in a non-visible part of the spectrum to avoid interference from everyday lighting changes. But if sunlight enters the scene, results will be skewed.

The fundamental problem is that tricks are being used to mask off information that would distract brittle algorithms. A proper way of reducing the image complexity would be to have algorithms that look at a scene as humans do. Humans consider all the information that reaches the eye—colors, shapes, brightness, reflections, and refractions—and they just know which information to ignore when it is not relevant to their goal.

End-to-End Deep Learning

To know what to ignore, an algorithm needs to know the ultimate use for the visual information. The need to separate the measuring part of the problem from the execution part is one reason computer vision used to be so hard.

When information is available on what to do—or how a robot needs to move—to a deep neural network processing an image, it can learn what to ignore. Alternatively, the network can also infer the right action. This means that combining end-to-end neural learning with real-time control provides a powerful idea.

Newer visual control systems are expanding the possibilities in robotics. Training a robot to pick up a cable and plug it into a slot in a test application requires using one of these new AI robot systems along with a hardware setup that includes a table or an automated guided vehicle (AGV) on which to place the robot, a gripper, and safety equipment. And, users will need to train the robot's deep neural networks. This process is simpler than it sounds, thanks to products that package the required math.

To perform the training, end users position the robot's tool at the point in space where the tool is supposed to go, relative to the object. In the previous example, the goal involves locating the plug at the end of a cable, which dangles in the air. A simple 2D color camera mounted on the wrist of the robot provides a stream of images to the AI controller, training the robot on where to go, allowing the camera to capture the surroundings of the target. Then, users move the cable to create another scenario of how the problem may look. Next, users place the tool in the right spot, save the location, and show the robot around again. If users repeat this process for 20 minutes, they'll collect enough data for the neural networks to learn. The AI controller will then spin up its training machinery and crunch the numbers, calculating a skill for the robot that generalizes over all examples shown to it. After a while, this skill—a neural network trained with the recorded user data—will be ready. When executed, the network will read an image from the camera stream (typically, every 50 ms) and decide how to move the robot—not where to move it in precise coordinates, just how to move it closer to where it is supposed to be.

From the examples they have been given, the neural networks learn what information to ignore, such as certain details that changed from example to example, including light brightness, background, reflections, and the precise shape of the cable. The neural networks also learn what information is relevant to the problem. For picking a plug, the relevant information may be its color and shape and the best angle to approach it. Given enough consistent data, the robot will be able to pick the dangling plug.

When users first encounter an AI system, many expect that, because it is intelligent, the system can simply guess what it is supposed to do. This is not the case. The intelligence is in its ability to learn anything that it is taught, but it still requires good teachers. Successful AI comprises both good algorithms and good data. For some AI robot controllers, the algorithms are included, but the data must come from users, and they must make sense. Showing the robot confusing examples such as conflicting movements when everything looks the same will result in a confused robot. Show the robot consistent data—for example, the target is visible, and for similar images, similar movements have been demonstrated—and the robot behavior will be significantly more capable and robust than anything implemented using the old measure-and-execute paradigm.

Using AI Controllers

Not all robot movements need to be controlled by cameras in real time. In fact, most movements do not. Camera control is typically needed only at the ends of movements—for instance, when finding an object whose rough location is known, when picking or inspecting an object, or when inserting or assembling objects. The movements in between (between two points in free space) are fairly simple and can be position-controlled and path-planned by a robot's control software. So, AI controllers for industrial robots leave that part alone. The AI controllers also never directly control the joints of the robot; they just tell the existing control stack what the controllers think should happen. The robot's own controller then decides whether the movement is safe, or whether it may need smoothing out. As a result, AI controllers do not replace robot controllers, and they shouldn't. They are add-ons, enhancing a robot's capabilities.

Industrial robots work in environments also used by humans, who depend largely on the visible portion of the electromagnetic spectrum for information intake. Robots work in factories designed to be observable to human eyes, and the robots perform work that needs to be seen by humans. Everything in a robot's environment is designed to be visually inspectable, and most available information on the state of a factory is encoded in rays of light.

If there was a way to extract that information from simple camera pictures (designed for the human eye), and decide what is relevant and what is not, cameras would be the ideal sensor. With deep learning, this way of extracting information is now possible, and with end-to-end learning, a way to tell the relevant and irrelevant apart is available too.

AI-Driven Robots Handle High-Speed Logistics Sorting and Depalletizing

Robotic system from Plus One Robotics incorporates a robot, AI-based software, and remote supervisor software.

STEVE KINNEY

Chances are that any products you've ordered online that arrive on your doorstep in poly bags, boxes, or flat padded envelopes have been touched by a robot. More than ever before, warehouses and logistics centers are relying on advanced automation technologies such as robots, machine vision systems, and artificial intelligence (AI)-based software to automate processes and help keep the flow of goods in motion.

As these facilities face a constantly changing landscape of product variability, their automation systems must be able to handle new package shapes and types while maintaining throughput rates and driving productivity improvements. Parcel-handling robotics company Plus One Robotics (San Antonio, Texas, USA; www.plusonerobotics.com) takes a novel approach to solving this issue by incorporating a complete vision system with a robot, AI-based software, and remote supervisor software that loops in a human "Crew Chief" for 24/7 support.

In a busy warehouse or facility where rapid pace is required to keep up with demands, high-speed parcel induction systems, such as a solution from Plus One Robotics, help optimize picking and placing of mixed parcels, bags, and products. Mixed packages enter a facility on a conveyor belt that moves the items to a

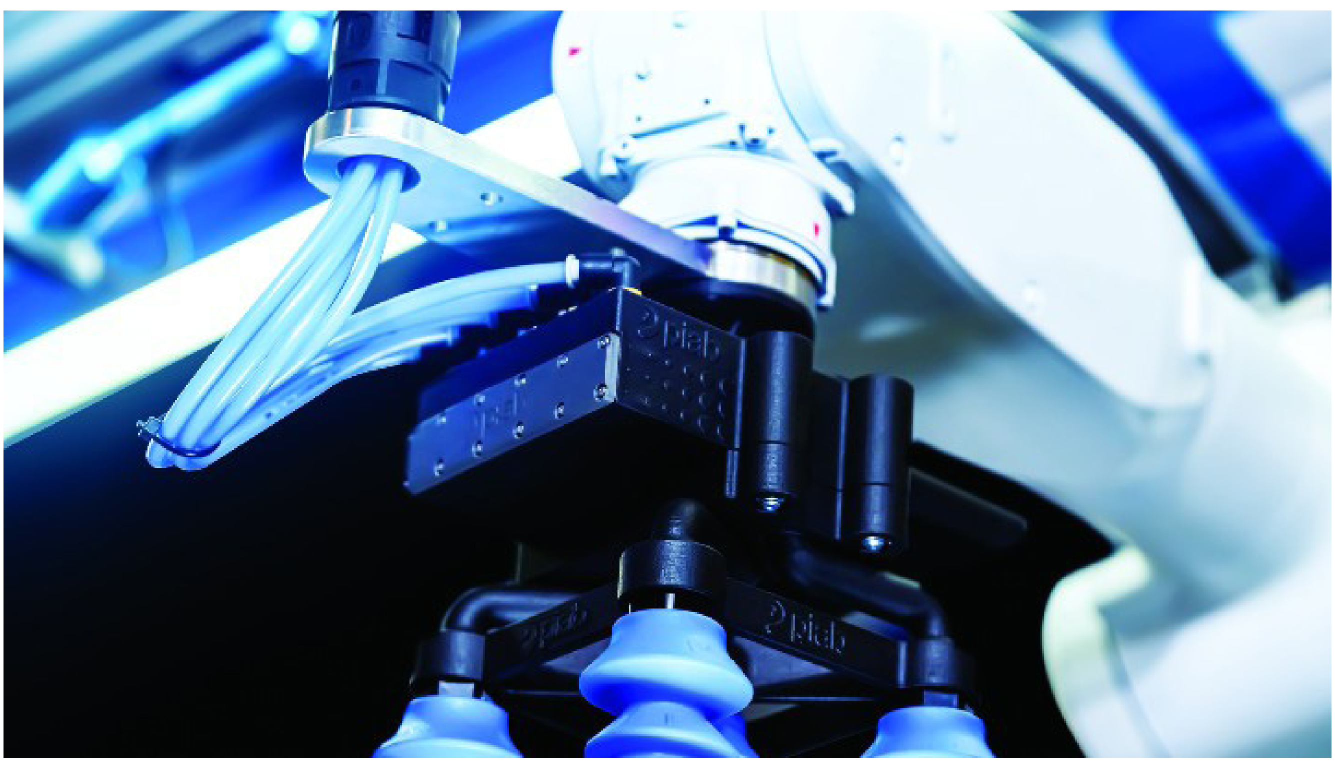

COURTESY OF SMART VISION LIGHTS

PHOTO 156480875 © PHARTISAN | DREAMSTIME.COM

FIGURE 1: By installing robots to automate sorting of small packages, FedEx kept pace with the e-commerce surge in 2020.

robot cell, where an overhead vision system identifies items and determines which products should be picked next. Once this information is acquired by the system, software commands the robot to pick and place the item onto the downstream location. In addition, a quality check is performed to ensure pick success.

While the system from Pick One is built with generic interfaces to remain hardware agnostic, it requires a 2D color image and a 3D point cloud for operation. Typically, this involves using Intel RealSense 3D RGB-D cameras (or 3D depth cameras) such as the 415 or 455 models from Intel, which are based on a structured light technique to capture depth measurement (D) while also obtaining a color image (RGB), then combining the data together, pixel-to-pixel, to create RGB-D images.

Plus One Robotics' parcel induction system typically uses a 25 lb.-payload-capacity robot from companies such as ABB, Fanuc, or Yaskawa, while the company's mixed depalletizing solution generally uses a 100 lb.-payload-capacity robot from one of the same vendors. Each of those applications also requires an industrial PC based on Intel i7 architecture augmented by an NVIDIA graphics processing unit (GPU), which can be sourced from several vendors depending on the application need. For lighting, the company uses products from Smart Vision Lights, according to Shaun Edwards, CTO and co-founder at Plus One Robotics.

"Pick success is contingent upon consistent, controlled lighting, which helps create a more robust and repeatable process, especially when considering the variability of the products our system typically deals with," he says.

Real-Time Human-Robot Collaboration

To help the system better recognize mixed objects on a conveyor, the company's software PickOne uses AI tools to extract the outlines of individual objects within that scene, according to Edwards.

"No AI would be needed here if all the items a facility was handling were the same, but these businesses are dealing with basically anything and everything that is shipped in the world today," he says. "AI helps solve this complex task by identifying individual items and outlining them for the robot to make a high-speed pick based on a sufficiently large training set."

Since the company has more than 700 million picks in production, an industry leading metric, it has amassed in-depth data and training resources. According to Edwards, that has allowed Plus One Robotics to train its AI model over time with significant amounts of data. It does not rely solely on AI, however. The robot system is also connected to the cloud and remotely monitored by a crew chief 24 hours a day, 7 days a week, who can intervene when AI can't solve a particularly difficult problem.

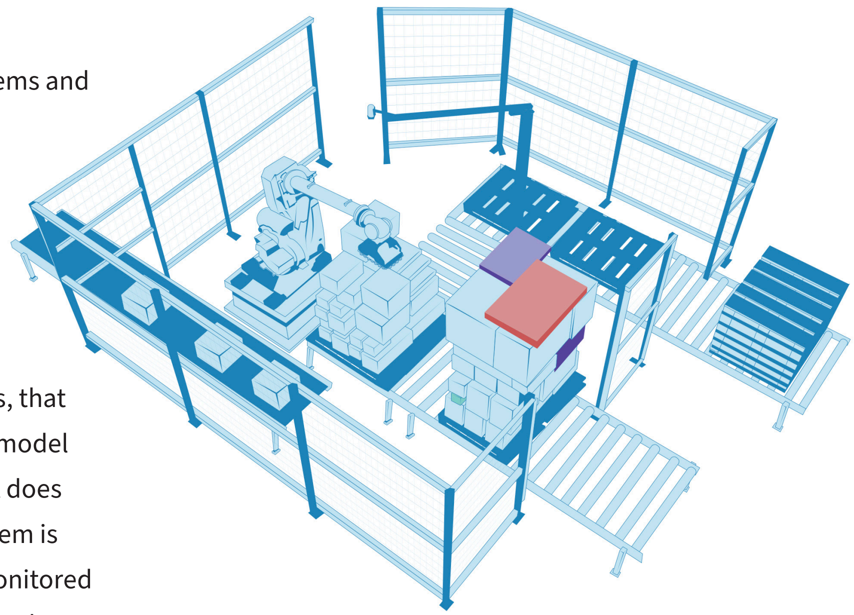

FIGURE 2: An example of a depalletizer work cell featuring a machine vision system and an AI driver.

"AI is not going to solve every problem out there, so when the software can't solve a problem, it loops in a human by sending a real-time request that typically results in a response in just seconds, which will help the robot determine its next action," Edwards says.

Powering up to 1,000 Parcels Per Hour

Deploying advanced automation technologies such as the system from Plus One Robotics enables higher throughput at 24/7 fulfillment operations; it also decreases operational costs associated with picking by up to 70%, while allowing human employees to add value in other areas that are less repetitive and safer, according to Plus One Robotics.

FedEx Corporation experienced that first-hand after it installed four robots to automate its small package sorting facility at its Memphis, Tennessee headquarters in March 2020, at the beginning of the COVID-19 pandemic. In addition to integrating the system into FedEx's existing process, Yaskawa Motoman supplied the robots, arms, and grippers, while Plus One Robotics supplied the 3D cameras, software, and industrial computer.

By installing the robots, FedEx was able to keep pace with the e-commerce surge by maintaining warehouse throughput and mitigating labor shortage issues. The robots continue to improve efficiency while also freeing former package-sorting personnel to focus on less physically demanding and higher value-added tasks.

Total Robotic Units Ordered in North America

Year	Units Ordered
2022	44,196
2021	39,708
2020	31,044
2019	29,988
2018	35,880

Source: Association for Advancing Automation (A3)

TABLE 1: Sales of Robotic Units in North America have been growing over time.

In another case, after a large apparel company found itself wasting time and decreasing throughput when handling certain package exceptions, it sought an automated solution that could increase picks-per-hour and effectively eliminate lost time. Prior to installing a new system, the retailer had people shutting down its robots to manually correct package

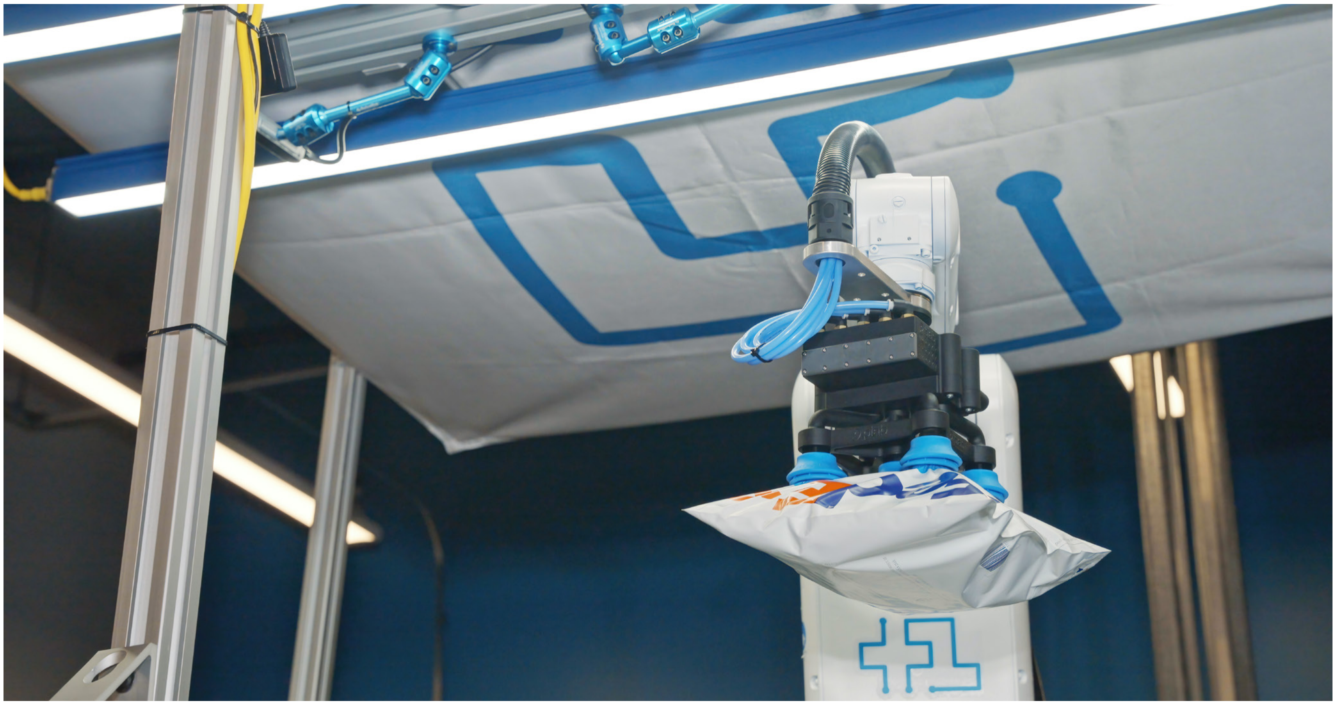

COURTESY OF SMART VISION LIGHTS.

FIGURE 3: The robotic system from Plus One Robotics typically includes a lighting source, 3D camera, robotic arm, and industrial PC. (Picture does not show actual installation at a customer's site.)

mishaps, resulting in increased lost time due to repeated starts and stops. That's when the company enlisted Plus One Robotics to install its system to improve the parcel induction process.

In addition to the PickOne system, the retailer leveraged Plus One Robotics' Yonder, software that facilitates access to remote crew chiefs, to quickly resolve any exceptions and minimize downtime. After the new technology was installed, the company achieved its goal of increasing pick rates to 1,000 parcels/hour. In addition, since the system's AI software learns continuously and the robot becomes more efficient over time, it was able to handle exceptions even faster within the subsequent six months. In yet another benefit, Yonder delivered valuable performance metrics indicating needed maintenance or training opportunities that further optimized throughput and operations over time.

More Robot Sales

Robot adoption continues to grow worldwide as companies protect themselves against labor shortages and strive to increase overall throughput, efficiency, and revenue. In fact, Association for Advancing Automation (A3) figures show that 2022 saw record robot sales in North America – to the tune of 44,196 robots valued at $2.38 billion – which represents increases of 11% and 18%, respectively, over the previous record highs in 2021. And remember: Robots aren't taking away jobs; instead, they are freeing up humans and opening the door to better careers, according to Edwards.

"It isn't necessarily the fact that people don't want to do these jobs anymore, it's the fact that many people didn't want to do these jobs in the first place," he said. "Robots work and people rule, we like to say. Managing a team of robots or operating one on the plant floor is a much better job than sorting packages non-stop."

Amazon Testing Robotic Arm that Identifies Individual Products

The robot can manipulate 65% of Amazon's sortable inventory of more than 100 million products, not including large products such as appliances.

LINDA WILSON

Amazon (Seattle, WA, USA; http//:www.amazon) is using computer vision and artificial intelligence (AI) to automate its logistics operation with robotics, and it's touting its endeavors publicly.

It has unveiled several experimental robots in 2022 and showcased a production robot that sorts packages.

In November 2022, Amazon introduced Sparrow, which the company is testing in the field. The core of the robotic system is an arm (M-20*t*D25) from FANUC that Amazon customized.

Sparrow uses suction cups to grip and then move individual products, such as vitamins or board games, from one tote, or bin, to another.

Leveraging computer vision and AI, "Sparrow is the first robotic system in our warehouses that can detect, select, and handle individual products in our inventory," Amazon says in the November news release.

In a video Amazon created in its robotics lab near Boston, Massachusetts, Sparrow sorts products of

COURTESY OF AMAZON

FIGURE 1: Sparrow sorts packages at Amazon's fulfillment centers.

different shapes and sizes, moving them from a yellow tote to one of four gray totes.

Machine Vision and AI

While not referring specifically to Sparrow, Amazon says it uses cameras positioned at different angles combined with machine learning to help its robots visualize individual objects within a crowded scene and determine how to pick them up, according to in a 2022 article posted on Amazon Science.[1]

In the Amazon-created video of Sparrow, which we have posted here with the company's permission, three cameras are visible—one mounted on the ceiling, one on the robotic arm, and one on the wall. In the video.

David Dechow, Machine Vision Expert and *Vision Systems Design* Contributing Editor, reviewed the video of Sparrow in action. He has designed similar applications and uses that knowledge to speculate on how Sparrow works.

Here's how he breaks it down:

- A 3D camera mounted on the ceiling takes an image of the items in the yellow tote after the robotic arm clears that tote with an object in its gripper. The image provides information on the basic shapes of objects in the yellow bin, allowing the vision application to determine how to pick up the next object.
- There are eight vacuum tubes in the suction-type gripper. By retracting some of the tubes, the robotic arm can vary how many of the tubes it uses to pick up a given item. The number of tubes selected would depend on the available surface area and shape of an item. "They've done a really nice job with this gripper because it gives them the flexibility to pick all of those really different shapes and different sizes," Dechow says.
- Two 2D cameras (one mounted on the robot and a second on the wall) take an image of an object after it is in Sparrow's grip. The arm pauses at this point, and Dechow speculates that this allows the vision application time to determine what the item is and which of the four gray totes it should be placed in. "If they are using AI for the identification, that's interesting," Dechow says.
- The illumination for the application appears to be LED. Dechow notes that the lighting devices have multiple segments, giving Amazon the flexibility to determine which segment to flash in each situation.

Challenges with Robotic Bin Picking at Scale

Tom Brennan, President of Artemis Vision says the operation depicted in photos and videos accompanying the news release—picking an item from one tote and moving it to another tote—is a challenging machine vision application.

It is difficult enough to develop an application in which a robotic arm takes a part out of a tote and puts it on a production line, but that is only one part, Brennan notes. It is much more challenging to develop a vision system that can recognize and pick as many arbitrary items as would be necessary in Amazon's environment, he says.

"I've seen tons of robotic arms 'doing' this at tradeshows in constrained environments but not really at distribution centers," Brennan says.

However, it would be important to know what error rate Sparrow has logged during testing to fully evaluate its potential, he adds. Nonetheless, "if it works on a very broad set of products, that's pretty amazing," Brennan concludes.

Dechow says that other organizations are experimenting with this type of robotic process, but two stumbling blocks often prevent these applications from moving successfully to a production environment:

- Accurately recognizing all objects, particularly those with few unique features or identifying information, such as printed words or logos on a label.

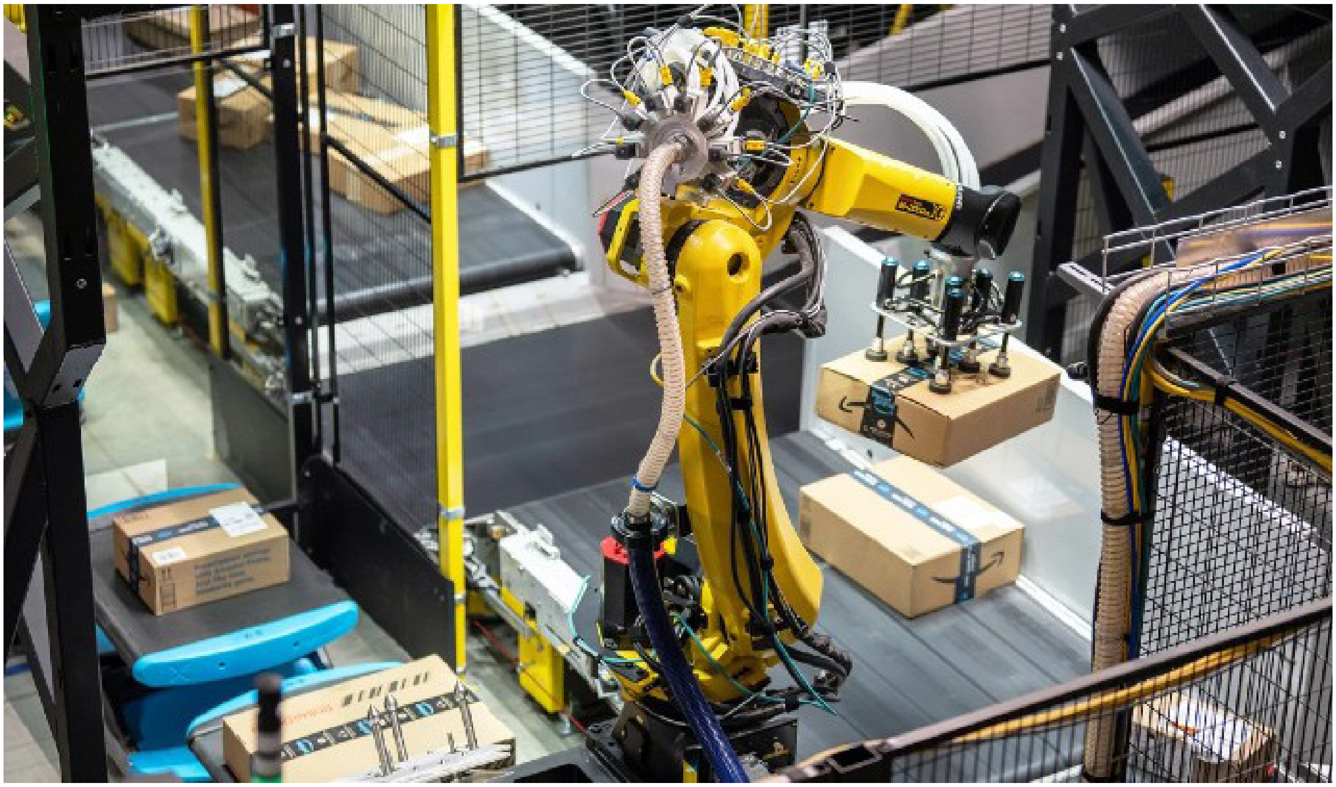

COURTESY OF AMAZON

FIGURE 2: Robin sorts packages at Amazon's fulfillment centers.

- Picking and placing fast enough to keep up with production rates.

Field Testing

Amazon is field testing Sparrow at a site in Texas. In those tests, "Sparrow is working on a process known as inventory consolidation within our fulfillment process," explains Amazon Spokesperson Xavier Van Chau. "It is handling items from one tote and consolidating them in another tote, densely packing that tote to ensure it is completely full and helping maximize inventory we can hold within an operating site. This means Sparrow has to deal with a great deal of clutter and have the ability to pick items in a crowded tote while also densely packing them in another bin," he says.

Van Chau adds that Sparrow can manipulate 65% of Amazon's sortable inventory of more than 100 million products, not including large products like appliances.

In 2021, Amazon "picked, stowed, or packed approximately 5 billion packages" worldwide, according to the news release.

Summing up the testing with Sparrow, Van Chau says, "We're excited by the progress we are making, but it's too soon to share our plans for broader deployment."

However, Dechow says that there are at least two other points in a logistics operation where Sparrow could be deployed: picking customers' orders or sorting stock for storage after it arrives at a warehouse from a manufacturer's operation.

Many Robots at Amazon

At Amazon, numerous robots are in various stages of development or production use.

Robin is another robotic arm from FANUC (M-20*t*D25) that Amazon has customized. Robin, which also has suction-type grippers, sorts parcels and mailers as they move down a conveyor belt. To do so, Robin grabs each package, rotates it, and scans the label to read the zip code. Robin also removes packages with rips, tears, or illegible addresses from the conveyor, so that employees can fix those issues before the packages are shipped to customers, according to articles posted on Amazon Science.[2, 3]

Currently, "We have deployed over 1,000 Robin robotic systems across our operations network," Van Chau says.

For safety reasons, Robin operates in restricted areas of Amazon's fulfillment centers.

In an article in Amazon Science from 2021 describing Robin, Amazon says Robin "brings many core technologies to new levels and acts as a glimpse into the possibilities of combining vision, package manipulation, and machine learning." [2]

Training Robots

Training robots in visual perception tasks requires a dataset of annotated images to teach robots how to distinguish between types of packages (a mailer vs. a box, for example) or types or products (such as different shampoos and conditioners).

Since Amazon has found that publicly available image datasets aren't good enough for training robots on the products shipped from the company's fulfillment centers, it's developing an in-house trove that the company believes will cut the "setup time required to develop vision-based machine learning solutions from between six to 12 months to just one or two," Amazon explains in a November 10, 2022, blog post on Amazon Science.[4]

Next Steps

While Sparrow and Robin use suction to grip items, Amazon also is experimenting with pinch grasping, which more closely mimics the actions of human hands, the company explains in the 2022 article posted on Amazon Science.[1]

Sparrow and Robin aren't the only robots Amazon has discussed publicly this year. The company also says it's experimenting with an autonomous robot, Proteus, which uses perception and navigation technology developed in-house to move in and around Amazon's employees, meaning it isn't restricted to robot-only areas of the fulfillment centers.

Proteus will initially move GoCarts, which are tall, wheeled shelving systems, around the outbound area of its fulfillment centers. However, the goal is for the robot to bring the shelving units directly to employees situated at workstations in the fulfillment centers. Currently, employees move the shelving units around manually, Amazon says.

Overall, the company says it has more than 185 fulfillment centers, including 50 of them that incorporate more than 520,000 individual robotic drive units.

Amazon says the impetus for embracing computer vision and machine learning is not only to improve operational efficiency but also to reduce employee's risk of injury and provide them with jobs that are more satisfying.

"The design and deployment of robotics and technology across our operations have created over 700 new categories of jobs that now exist within the company—all because of the technology we've introduced into our operations," the company says in the November press release introducing Sparrow.

RESOURCES:

1. "Pinch-Grasping Robot Handles Items with Precision," Amazon Science. https://www.amazon.science/latest-news/pinch-grasping-robot-handles-items-with-precision
2. "Amazon's Robot Arms Break Ground in Safety and Technology," Amazon Science. https://www.amazon.science/latest-news/amazon-robotics-see-robin-robot-arms-in-action
3. "Robin Deals with a World Where Things Are Changing All Around It," Amazon Science. https://www.amazon.science/latest-news/robin-deals-with-a-world-where-things-are-changing-all-around-it
4. "How a Universal Model is Helping One Generation of Amazon Robots Train the Next," Amazon Science. https://www.amazon.science/latest-news/how-a-universal-model-is-helping-one-generation-of-amazon-robots-train-the-next

Automating the Production of Prefabricated Houses with Robotics and Machine Vision

FingerHaus uses a robotic cell with a gripper system in combination with a 3D vision system to automate the production of prefabricated houses.

JIM ROMEO

FingerHaus, a manufacturer of prefabricated houses, has been producing energy-efficient and sustainable housing for many years. However, with an uptick in the prefabricated housing market, and a shortage of skilled employees, the company has been looking for ways to automate its production processes while still producing high-quality products.

"Fingerhaus had the vision to increase output and relieve the workload of employees by integrating robot automation to their wall manufacturing line," says Christopher Köster, an executive in general management and sales for BETH Sondermaschinen GmbH (Medebach, Germany; www.beth-germany.com), which was involved in the project. "The prefabricated housing market has been growing for the past 20 years. The rising demand for prefabricated houses can't be satisfied based on the existing production technologies and the increasing shortage of skilled workers.

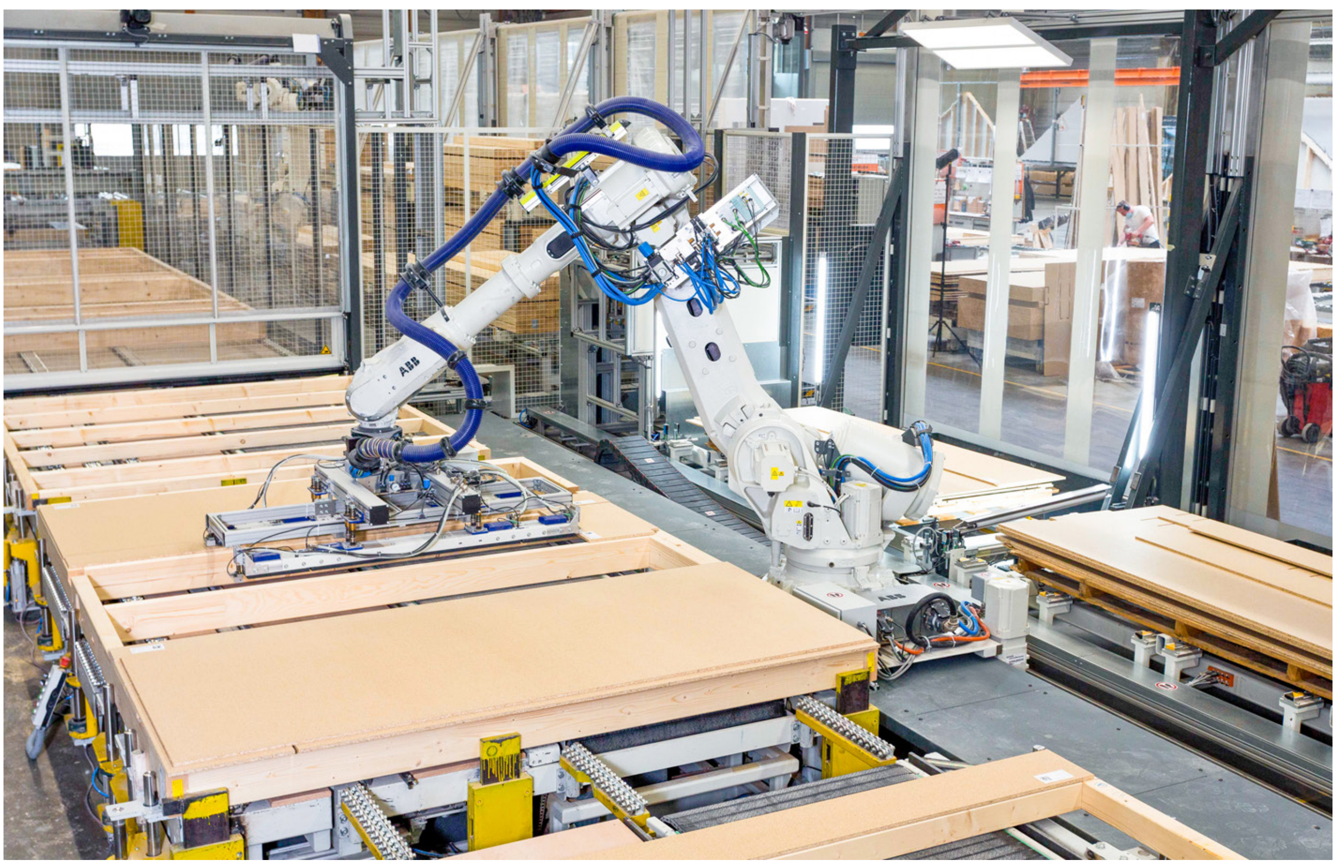

COURTESY OF PHIL-VISION

FIGURE 1: A robotic cell with a gripper system was used in combination with a 3D camera.

COURTESY OF PHIL-VISION

FIGURE 2: Wall elements have various dimensions in length, width and thickness.

Automation can help to increase the productivity of the manufacturers by shifting from craftmanship to modern production technologies."

Building a Solution for a Complex Process

Automating the paneling (or assembly) of wall elements for prefabricated houses is a complex process. Each prefabricated house is unique. They are customized to a customer's wishes. Thus, they may have different geometries created from an assortment of panel sizes.

Wall elements can have various dimensions, such as a length of 800 (about 2.62 ft) -12000 mm (about 39.37 ft) with a tolerance of +/- 10 mm (about 0.39 in), a height of 1600-3000 mm (about 9.84 ft) with a tolerance of +/- 5mm (about 0.2 in), and a thickness of 80-300 mm (about 11.81 in) with a tolerance of +/-2 mm (about 0.08 in). The dimensional stability of the materials can also vary, with a range of +/- 3 mm (about 0.12 in).

Additionally, the sheets for one wall element are combined in a stack with up to 70 layers, and several sheets can be in one stack layer. All the sheets for one wall element form one coherent stack.

To further complicate things, these sheets are not in any fixed order, making it impossible to process them side by side from left to right on the wall element. Instead, each individual panel part must be placed in its specific position on the wall element. A gap of 3 mm (about 0.12 in) is provided between the individual sheets of a panel system, which can be used to compensate for tolerances.

Exact positioning of the individual sheets is essential because faulty paneling leads to a stand-still of the complete production process.

The Solution: A Robotic Cell with an Accurate and Effective Machine Vision System

To solve these challenges, executives from FingerHaus

COURTESY OF PHIL-VISION

FIGURE 3: A robotic cell and vision system coordinates the placement and positioning of sheets used in wall fabrication.

(Frankenberg, Germany; www.fingerhaus.de) called on two German firms, phil-vision GmbH and BETH Sondermaschinen GmbH, to improve the manufacturing process.

Phil-vision (Puchheim, Germany; www.phil-vision.com) was chosen as the supplier for the vision system. BETH Sondermaschinen GmbH developed and implemented an automated plant concept that includes both the semi-automated production of wooden frameworks and the automated paneling of wall elements.

The new system uses a robotic cell with a gripper system in combination with a 3D camera system. To determine the exact position of the sheets for the gripper system, the sheets are measured with the 3D camera system, and the subsequent data provides the absolute position to the coordinate system of the robot cell.

"The material is automatically loaded to trolleys based on the best stacking stability. The orientation of the material is not precise and must be located by a vision system. As the tolerance for placing the planks and boards to the framework is quite low, it is necessary to detect the position very precisely," explains Köster.

The robot, using a vacuum gripping system, removes the sheets and positions them precisely on the timber frame. The integrated 3D camera system identifies the sheets, and then they are automatically fastened together using a clamping device that the robot activates.

The 3D camera system recognizes the respective position of the wood-based or plasterboard sheets, while an intelligent light system supports position recognition. The exact position data is immediately processed by software developed by BETH and transferred to the robot, which then takes over the precise paneling.

The solution, which is mounted above the paneling station, is based on a system consisting of eight 20 MPixel monochrome cameras with GigE vision interface, corresponding compact lenses and two pattern projectors. The cameras are combined into four stereo systems, which generate a complete point cloud with just one shot. Subsequent processing is performed on a high-performance industrial PC. A “PK Construct” software solution developed by BETH ensures fast and secure communication between the camera system and the robot.

The interface between the robotic cell and the vision system used highly flexible communication between panels, as each wall part is unique. A highly precise robotic system was used with an accuracy of 100 µm over 30 m of coverage. This robotic system was combined with a highly precise vision system and gripper to pick different sizes of walls.

The walls on the trolley are 3200 x 1250 x 1000 mm³. In order to place the walls correctly with the robot arm, an accuracy of .5 mm was required.

The new system is different from other systems that had been tested, and it has a combination of specific unique but effective parameters. It has a short cycle time of 8 seconds, and a large measuring volume of 4.1m³ with an accuracy of well below 1 mm (about 0.04 in). The image is acquired while the robot is still placing the previously measured sheet, meaning that very fast cycle times are possible.

Automating the Paneling of Wall Elements

In addition to the vision system developed by phil-vision, BETH developed the robot cell. The handling and clamping of the panels on the framing unit is carried out by an industrial robot from ABB, type IRB 6700, on a traversing axis that is over 20 m long. The robot is equipped with a SCHUNK changing system to be able to switch automatically between the suction gripper for picking up the panels and the clamping device. After the sheets have been placed in the correct position on the framing unit produced in the previous cell, they are fastened with clamps according to the customer’s design specifications. BETH’s “PK Construct” software is responsible for calculating the optimum travel paths.

The overall automation solution includes a 50-meter-long (about 164 ft) line with a width of approximately 10 m, which offers maximum flexibility for paneling with standard dimensions and special geometries, including feeding and transfer to the next production step.

When asked about the main challenges facing the vision system, Patrick Gailer, Managing Director of phil-vision GmbH, says it was one of handling the paneling of wall elements with adequate stability, without knocking the robotic and vision system out of calibration.

“To keep the large system mechanically stable enough for the large measurement volume is a major challenge for the calibration,” explains Gailer. “Only a few recalibrations per year are required to correct for temperature. In addition, it is difficult to give a high-precision position, while at the same time reacting flexibly to components.” Gailer says there are large variations from actual production when compared to simulated models. Plus, he says cost restraints came into play as they strived to develop a system that is within budget and less expensive than other high-accuracy 3D scanners currently available on the market.

“The initial precision of the system was working OK, but we got quite some setbacks when the system stopped working or got inaccurate without knowing why,” says Gailer. “With this precision and these sizes, it was difficult to know when [and if] the camera and robot were decalibrated. “

Gailer said that in subsequent comprehensive infrastructure testing, they combined different calibration methods to catch any inaccuracies early and resulted in greater overall precision.

"We learned a lot about mechanical and thermal decalibration," says Gailer. "The development took a good 6 months longer than anticipated. To reach effectiveness much more (than anticipated) communication interaction between robot and vision system was needed. [Resultingly] the communication interface was broadened a lot."

Fast and Efficient Production Throughput with Minimal Waste

Developing a solution of a robotic cell with an accurate gripping system, combined with sophisticated machine vision system using a 3D camera was impactful and effective for FingerHaus' prefabricated housing production.

Not only was the line speed improved, but the placement of the panels and their assembly was more accurate than previously. Production workers were spared some of the monotonous and time-consuming tasks and able to focus on other areas within the production process that could best use their attention. In addition to the faster speed, improved accuracy and efficiency, and better utilization of human capital resources, the new solution using the robotic cell and vision system reduced overall waste from production than the previous process.

"FingerHaus can now produce prefabricated houses with higher speed, cost efficiency and accuracy," says Gailer. "The workplace is more attractive for skilled workers because the monotonous and physically demanding tasks are done by robots, and there is less waste because of maximized material utilization."