Eye Spy: The Basics of Robot Vision Systems

News Video posted on July 16, 2018 | 1 comment

I spoke to robotic systems integration experts for insight into robot vision basics

Listen to this story

For many manufacturing processes, automation is an ideal way to streamline throughput, improve quality and repeatability, and solve labor challenges. If you're beginning your journey in robotic automation, this video is a great place to start the benefits of industrial robotics, and this article on pitching robotics projects to management.

When looking at automation, trial and experimentation are important to discover and understand potential pitfalls. Many production managers take for granted the inherent human skill involved in even the most mundane tasks, from picking objects to inspection, motion skills and visual processing. To replicate these with an automated system can be complex.

According to Keith Voset, Product Manager for Software at Yaskawa Motoman, most customers begin their search for an automation solution with an understanding of the problem, and not much else. They know they want to get parts from A to B, but often don't know what robot to use, or how to do it.

In this article, we explain the basics of what's out there, and match common problems to solutions. For expertise, we reached out to robotic systems integration companies with robot vision capabilities, including Yaskawa Motoman, Nexis Robotics, Stäubli Systems, and Universal Robots.

The 1, 2 and 3 Dimensions of Robot Vision

Parts are conveyed and positioned in the cell

Most engineering projects are a quest for the simplest solution. In robotic vision, this raises the question: How many dimensions do you need? The answer to this question is important, and within another question how many dimensions you can get away with.

If your parts are in a line, they can be completely unconstrained in all three dimensions. Your system needs to determine the part's position and orientation in all six axes. If your parts are presented in a surface, such as a conveyor belt or a tray, the z-axis height is fixed, and a three-dimensional camera is probably overkill. If your parts are constrained in a fixture, you don't need to know the z-axis, but the camera may need to detect the part's location along one axis.

For one-dimensional detection, a laser rangefinder or optical distance sensor could be used.

According to Keith Voset, product manager for software at Yaskawa Motoman, 2D vision systems are much more common than 3D counterparts, partly due to the increased complexity and computational requirements of 3D vision.

2D robot vision is more than just a camera. Systems also include lighting, interfaces, and software, including, in many cases, plugin packages for the robot controller. The Yaskawa Motoman DX system is an example of a 2D vision system. It includes a camera and configuration software.

3D vision can include multiple cameras and laser displacement systems. Laser displacement systems work essentially like LiDAR, measuring the reflected laser light returning from a scanned surface. Multiple camera systems use triangulation algorithms to interpret information gathered from multiple perspectives on a part. Nexis Robotics's Bluevision 3D system is an example of a 3D vision system with a camera that can be configured for this type of camera. Stäubli's Robotic Vision, as well as 2D vision, Bluevision Bluebinder is an example of a 3D point cloud processing system. Software uses point cloud data to interpret the shape of objects in images.

Lighting Issues with Robot Vision Systems

Is your team ready to configure a vision system?

Many robotic systems in factories today are not case-specific perform the same task over and over, regardless of what part is in front of them. Variations in different tasks as simple as a different shape, but if you require a robotic vision system now, it's important to consider whether you have the in-house knowledge to reconfigure your vision system for a new application. According to Voset, vision systems users may require more programming know-how, but many systems are designed with user-friendliness in mind. Voset said, "Most of the time, the person who is using the camera is configuring what it's looking for, how the object is shaped, where it's looking. In the 3D vision world, it's a little different depending on the implementation. Some vendor implementations require that you actually use programming and call system provided subroutines to interpret the image. You have to provide the image, but you have to use the system's subroutines to interpret the image and then to do pick-and-place. With other systems, including one Yaskawa Motoman uses, there is a GAD model of the part and the system does the rest."

Schmidt, of Nexis Robotics, explained it this way: "Well, there's always been in the early days, it was an engineer's environment and an engineer's tool. You had to have a degree in robotics to program a robot. Nowadays, with the Industrial IoT and Industry Storage ready for artificial intelligence, manufacturers are coming up with ways to make changes in a more automated way, with a user interface. It's similar to an AR/VR the system is training the system, relatively complex, but it's easier than it used to be. Vision sort's choices. You can't open a mortgage or perform advanced banking functions without a PC, but it's easier than it used to be. Robotic vision is moving in that direction." In other words, the configuration software should be a consideration when choosing a vision system, and if you don't have the in-house expertise, making sure the configuration changes after the initial installation is feasible.

Is Robot Vision Necessary? Scenario: Robot Barista

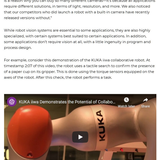

Robot barista at Café X image courtesy of Wired

Considering the advanced force and torque sensors available for industrial robots, do robots always need to "see" what they're doing?

Universal Robots recently announced the e-Series, a new line of their collaborative robots. The UR5e, the company's most popular robot, includes a force and torque sensor at the flange of the robot. I asked Esben Ostergaard, CTO of Universal Robots, why the company chose force sensing instead of vision for this upgrade. Here's what Ostergaard had to say: "In vision system was considered, and it is. A lot of our robots are using vision to do different things. If you consider adding a camera onto all robots, we did consider that, but it's very hard to know what camera to add. There is a reason why you can buy so many different cameras—it's because all applications are different. Universal Robots is a small company, and we don't have the resources that our competitors who did launch a robot with a built-in camera have recently released versions without."

While robot vision systems are essential to some applications, they are also highly specialized, with complex systems best suited to certain applications. In addition, some applications don't require vision at all, with a little ingenuity in program and process design.

For example, consider this demonstration of the KUKA new collaborative robot. At timestamp 2:07 in this video, the robot uses a barcode scanner to confirm the presence of a paper cup on its gripper. This is done using the torque sensors equipped on the axes of the robot. After this check, the robot performs a task.

So, what is it about this demonstration that makes vision unnecessary? Well, as you can see in the video, the cup is placed in the gripper by a human operator. As a human operator, the robot could probably be part of the robot's program to locate and pick up the object by itself.

Robots without vision capabilities can locate objects if they are in a predictable location, usually fixtured in place. This demonstration of the Kawasaki Duiko, a short-arm collaborative SCARA robot, shows the fixturing used for tests and parts. However, custom fixturing and tooling can get expensive.

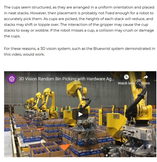

Cup dispenser (courtesy of Ian Jamieson)

Then again, a dispenser like this would require a human operator to refill the dispenser. It would be better if the robot could pick cups out of a case, so that human would be limited to delivering and preparing cups.

Cups on tray sopper in a case

Could a robot do this task?

Now, the task of picking a cup is similar to what's known as bin picking, in which the robot locates and picks objects from a bin. In these tasks, neither the location in X, Y or Z is controlled, nor the orientation in A, B or C. Our cups look something like this:

The cups seem structured, as they are arranged in a uniform orientation and placed in neat stacks. However, their placement is probably not fixed enough for a robot to accurately pick them. As cups are picked, the heights of each stack will reduce, and stacks may shift or topple over. The interaction of the gripper may cause the cup stacks to sway or wobble. If the robot knocks a cup, a collision may crush or damage the cups.

For these reasons, a 3D vision system, such as the Bluevision system demonstrated in this video, would work.

While Bluevision's software is hardware agnostic, the setup in the demonstration uses a Mitsubishi robot sensor to create a 3D point cloud of the target area. Next, the software determines which object to pick without colliding with other objects or with the bin. This solution could be ideal for our case of cups, as it could correctly identify the height of the stack of cups and position the gripper accurately.

However, 3D vision has much higher computational demands than 2D vision processing. Voset explained, "In 3D, there is much more image processing involved in comparing a 2D image of the part to point cloud or 3D image data, and that's why computer intensive. It's not as time sensitive as 2D vision. If you're really looking for fast cycle time, 2D is still the king of the hill."

In the scenario of the robot barista, cycle time is not really a factor. The robot needs a cup for every new order, leaving plenty of time for the X-axis robot to move parts that worse but still takes a long time, in an application where the robot is trying to capture the image and do the processing, maybe after the robot has picked up the part and gets out of the way, the vision system is starting to try to get the next image ready. Voset said, "Sometimes customers will say, 'I want to be able to start processing time within seconds, rather than waiting for it to be ready." Ostergaard suggested Schmidt. With this sort of different programming, our barista robot would have to wait for the cup, delay before placing a new cup, as the vision processing would be complete by the time it finished making the previous order.

However, 2D vision may be a more cost-effective option. Schmidt explained that, "Sometimes, you can take a 2D camera and configure it so that it's looking at the part and the so-called masks and use that information to interpret the image. For example, if a cup's bottom surface is parallel to the floor, it should appear in the image as a circle. If it appears as an oval, the software determines the angle of the cups. To determine the height of the cup, cameras on the end-effector or a second camera at a different perspective could be used.

Besides choosing the differences between 2D and 3D vision, this scenario also illustrates one very important factor: that the design of your process is critically important to the design of your vision system. In fact, design changes to your process, or to your part, could save you money, as well as installation lead time for the robotic system integration, and even cycle time when it's all up and running.

Bottom line: whatever your problem, a robotic system integrator can probably get you set up with a vision system that gets the job done. But without examining the bigger picture, you won't get the most efficient solution.

Redesigning your Part and Process for Better Automation

If your factory uses industrial robotics, it's unlikely that they were just dropped into place one day without a full study of the costs, benefits and workflow changes. Similarly, if you want to assume a vision system would be different, the same thing applies. You're the problem is getting your parts from A to B, the problem could be the X and the Y.

Schmidt gave this example: "If you are trying to pick up a pencil, and you need to orient the eraser and a certain way, if you can have the pencils coming down so that your eraser is always within 90 degrees and pointing in a certain direction, that's a much simpler problem than if the pencils are coming down in a random orientation. Anything you can do to trim those variables tends to less complexity or greater success."

Beyond the process, even the part itself could be redesigned or changed to improve automation performance. Voset said, "If a customer is considering using a 3D camera, could the surface finish be changed? If a part is challenging for a robot to grip, could a feature be added? While it may seem ridiculous to change your part designs to suit the automation, doing so could save you big time in costs and headaches down the road.

For more coverage of automation and robotics, follow me on Twitter at @AmandaAutomotive!

Subscribe
Sign up for news
Submit

Manufacturing 101:
Lean Manufacturing in Healthcare
Quality Control Laboratory Conditions
Optimizing Machining and Moldmaking for Metal Additive Manufacturing
The Seven Basic Tools of Quality
How Traditional Machine Tool Alignment Processes Compare to Laser Calibration
Understanding the Cases of Faults in the Laser Particle Sizer
A Brief History of Interferometry and Dimensional Measurement in Manufacturing
3D-Organizing the Laser Workspace
Resources
Whitepaper – Switches and Sensors for Harsh and Hygienic Applications
Research Report: The Connection Between Operator Experience and Quality
Webinar: How Metal Additive Manufacturing Delivers New Efficiencies in Lightweight Parts
Research Report: 2019 Industrial Process Instrumentation
Webinar: How 3 Medical Device Manufacturers Increase Productivity with Modern Design Tools