IMAGING
& MACHINE VISION EUROPE

For suppliers, integrators and OEMs using machine vision

February/March 2020
Issue 97

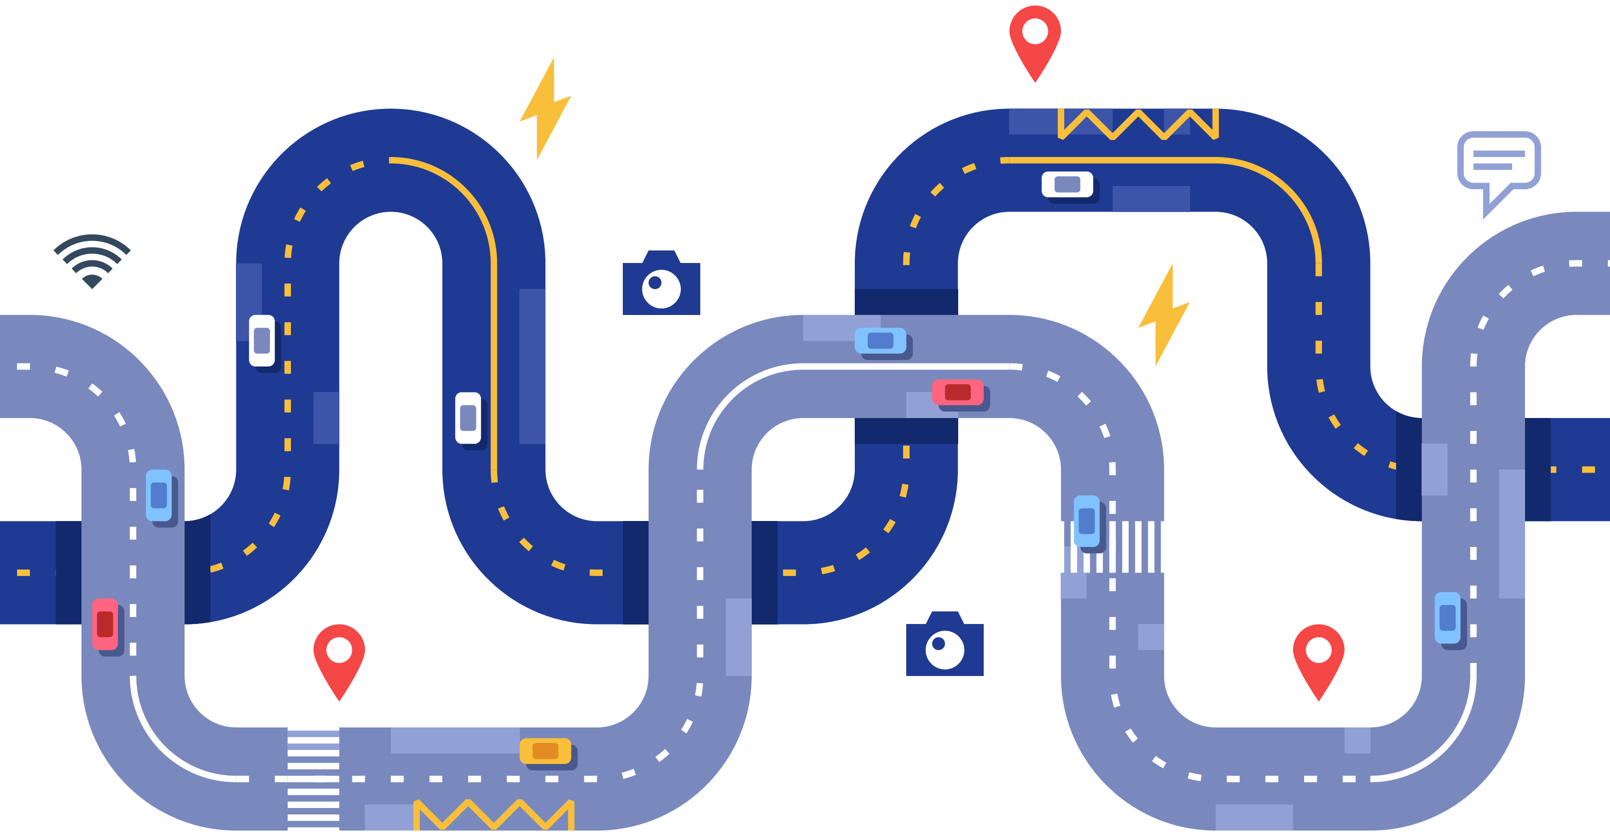

Going with the flow

Taming traffic with vision and AI

National Physical Laboratory

NPL will be going out to measure field plots in June

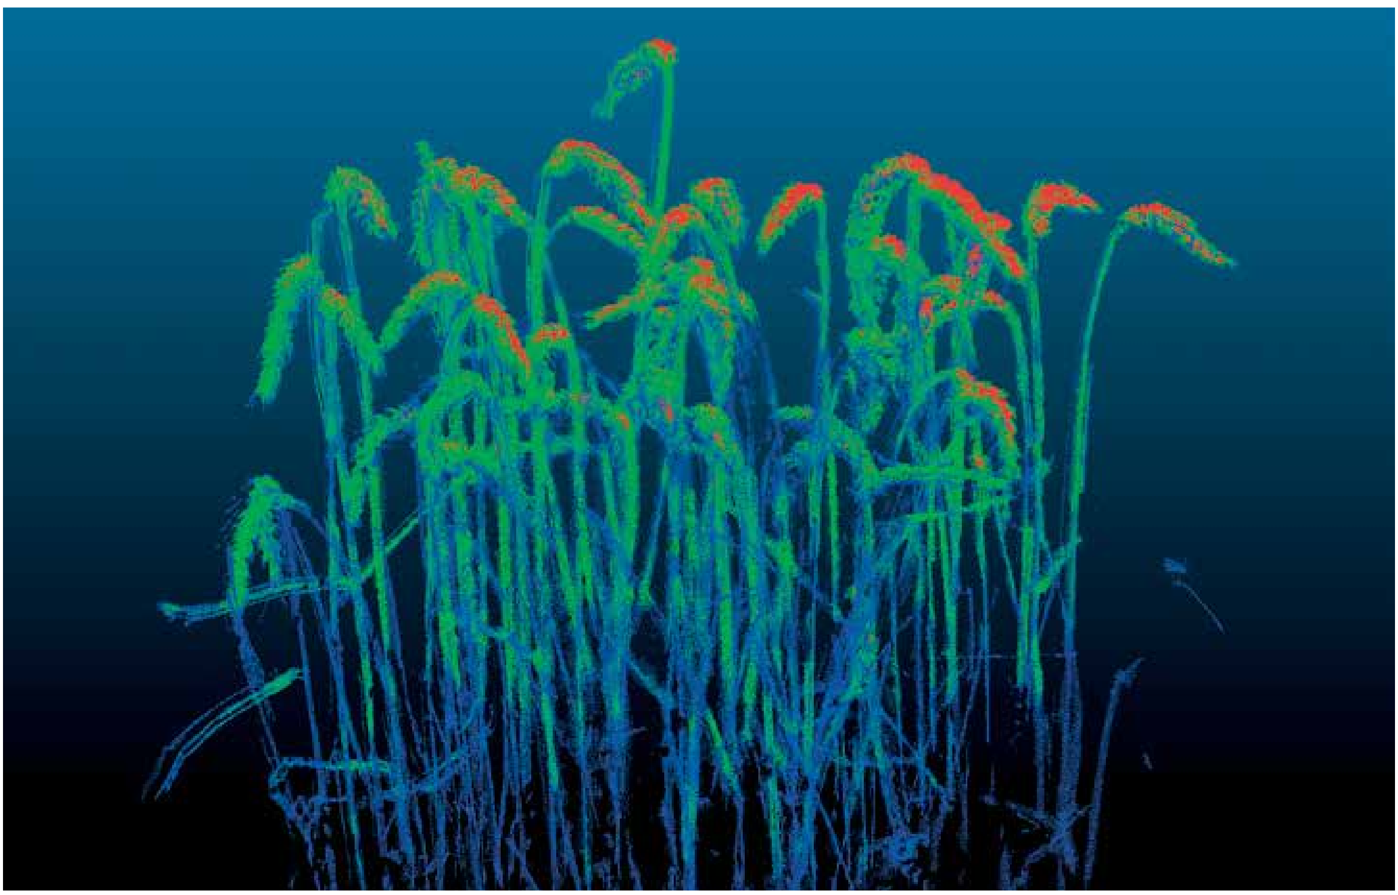

National Physical Laboratory

NPL's imaging rig can scan wheat at millimetre resolution

Wheat breeding gains from 3D data

Greg Blackman speaks to Dr Richard Dudley about the National Physical Laboratory's 3D imaging rig, which will be scanning wheat in crop breeding trials this summer

The Harvesters , an oil painting by the Dutch master Pieter Bruegel the Elder, was often a teaching aid when I was a student to illustrate how crop breeding has changed agriculture. The picture depicts labourers bringing in sheaths of wheat in 16th century Flanders. It's not the farming practices that we were asked to look at, but the height of the wheat, which reached the farm hands' shoulders in 1565 when the picture was painted – walk through a field of wheat today and it might be around your knees.

The modern dwarf varieties of wheat are the result of many years of plant breeding. A stocky plant isn't flattened by high winds or heavy rain, while at the same time more of the plant's energy goes into the grain, rather than into growing tall.

Other characteristics can also be enhanced by breeding, such as yield, grain quality and disease resistance. Now the National Physical Laboratory (NPL) in the UK is offering plant breeders and agricultural researchers new measurement methods based on imaging that have the potential to accelerate plant breeding trials, and ultimately improve food security as farmers seek to feed a growing population.

The NPL has built a wheeled platform equipped with a number of 3D imaging technologies, which can be attached to a tractor or otherwise pulled across a field.

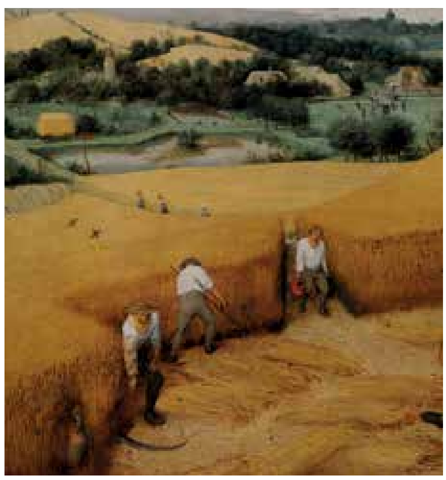

A detail from *The Havesters* (1565), by Dutch master Pieter Bruegel the Elder

One version of the imaging rig includes an 18 stereo camera system. The intention is to generate a complete 3D scan of a field trial plot to assess how varieties of wheat grow.

NPL researchers also use time-of-flight, lidar, hyperspectral imaging and other cameras, including Photoneo's 3D structured light camera supplied by Multipix Imaging, which is Photoneo's distributor in the UK.

'Currently, it takes in the region of five to ten years to develop a new variety of plant because of the throughput issues,' explained Dr Richard Dudley, science area leader for electromagnetics and precision agriculture at the National Physical Laboratory. 'With more data we feel we can reduce that – maybe down to a couple of years, a couple of seasons.'

> **'[A 3D scan] can give seed producers a three-month advantage, in time to decide what [plants] they are going to upscale'**

Data from a 3D scan can give seed manufacturers an early indication of what varieties they want to produce for next year. 'It can give them [seed producers] a three-month advantage, in time to decide what they are going to upscale production of, and that can be really valuable,' added Dudley.

NPL is focusing its efforts on wheat initially, because it's one of the highest value crops grown and bred in the UK.

Plant breeding requires lots of physical measurements, according to Dudley, which traditionally are made by researchers walking around the plot with a ruler and taking photos. NPL's imaging platform can capture crop trials in a lot more detail in 3D. Algorithms are then applied, to give crop breeders more data to make decisions. →

→ NPL can provide data on how wheat is distributed in the field, along with the ability to zoom into the 3D point cloud to view individual plants.

'The data is more accurate, and they've also got year-on-year comparisons, rather than ruler measurements and pictures which don't often tell them everything they need to know,' Dudley said.

Image capture

The minimum area of land the NPL is asked to measure is 3 x 2 metres. This is captured at millimetre resolution, volumetrically across about half a metre. 'It's quite a big volume and quite a high resolution, and you can't spend too long because things like wind and movement in the crop get in the way. So you need to capture it relatively quickly,' explained Dudley.

Time-of-flight imaging is reasonably fast, offering 24fps or more, according to Dudley, but these cameras often don't have the resolution required. The Photoneo camera, an area-scan device operating by what the Slovakian firm calls parallel structured light, has higher resolution and can capture movement in 3D – the Phoxi XL camera captures up to 3.2 million 3D points for each scan, at 16 million points per second throughput – although it still doesn't meet all the demands of NPL.

To cover a 3 x 2-metre plot, NPL would use up to 18 cameras, possibly including three Photoneo scanners. The point clouds from these cameras are then combined to give an image from different angles. Which imaging technique is used comes down to the area of land covered, the speed of capture, and the resolution required.

There are difficulties in overlapping the images and getting the complete view, Dudley said, and errors can be produced when combining point clouds.

The measurements are made early in the day or late in the evening when the sun is low in the sky, because practically every camera technology used struggles with direct sunlight, according to Dudley. 'I still think there's a long way to go on most scanners out there in terms of [operating in] daylight,' he said. 'Midday, bright sunny day, most systems still fail.'

He added that lidar is the only technology that can just about deal with bright sunlight, but that lidar has shortcomings for the NPL team in other ways.

Making sense of the data

NPL is developing algorithms that fit to individual ears of wheat. The algorithms first identify where the ears are, then calculate the volume of each ear along with other aspects like length and number of grains.

'It's not trivial; it's pretty tough stuff,' Dudley said. 'Fitting shapes into complex point clouds, although it's been done a bit for autonomous vehicles, it's still pretty challenging.'

The researchers are taking two approaches to analysing the data. The first is clustering the 3D points using a combination of MVTec's Halcon image processing software, also supplied by Multipix Imaging, and Matlab vision algorithms. These look for densities of points and try to separate noise and stems, which are generally thinner than the ears. Shapes that resemble an ear of wheat are then fitted to find the ears in the point cloud.

The other approach is based on machine learning. The algorithm learns what an ear of wheat looks like from classified ears, then hunts them in the point cloud. But these shapes aren't individual objects sitting with nothing around them; often the ears are overlapping or touching, which creates strange double-headed shapes.

'Fitting shapes into complex point clouds, although it's been done a bit for automomous vehicles, it's still pretty challenging'

Dr Richard Dudley, science area leader for electromagnetics and precision agriculture at NPL

Dudley said that the team is starting to get up to 90 per cent identification in an image of around 100 to 200 ears of wheat. The data produced after this pre-processing is then much higher quality.

At the moment, the NPL team captures the data in the field and then post-processes it in the lab. That can be done within a day. However, the team is also starting to work with groups that have 5G capability out in the field. It then becomes possible to do some cloud computing on these large datasets.

'This is much more [relating to] commercial farming, where you might want to make decisions on the vehicle that's out in the field,' Dudley said. 'That's the future, trying to do data processing on-the-fly and have feedback as a vehicle is progressing through the field. The 5G network, with its capability of sending high data rates, is certainly something we're looking at using.'

Juice Flair/Shutterstock.com

Food security

Dudley made the point that, currently, global production of wheat is struggling. 'It's difficult to produce as much grain from the area of land we've got, so it's really important to develop these new varieties of crop that yield more, so we've got food for the future. Wheat is one of the main staples in our diet; it accounts for a lot of calories and protein,' he continued. 'That's the driver, both economic and sustainability for the future.'

NPL is trying to secure its first customers for next season, starting in June. It will be going out to customers and measuring their field plots.

Dudley said that he is still looking for new imaging technologies to capture these 3D plots, both faster and in higher resolution, and with lower sensitivity to sunlight.

Fulfilling promise of pick and place

Keely Portway on the importance of vision for robots working in warehouses

The market for online shopping is continuing to grow at a rapid rate, which, in turn, has led to a rise in automated warehouses. To use the obvious example, Amazon now has nearly 200 operating fulfilment centres globally, spanning more than 150 million square feet, where employees work alongside the retail giant's robots to pick, pack and ship customer orders to the tune of millions of items per year.

This would not be possible without machine vision, and it is not just the hyperscale organisations benefiting from this kind of technology in their warehouses. German electrical supply wholesaler Obeta is now running a robotic station to pick and sort components at its facility in Berlin. The wholesaler partnered with AI robotics company, Covariant, and logistics technology firm, Knapp, to bring the system into production.

Managing expectations

Michael Pultke, head of logistics at Obeta, explained: 'Customer expectations for fast, affordable package delivery have never been higher. To stay competitive, we need to modernise our operations and keep order processing and delivery running quickly and smoothly. The Covariant-powered robot is an integral part of our live operations, exceeding our performance requirements and adapting quickly to change. AI robotics is a foundational part of our future strategy.'

Covariant was founded in 2017 by Rocky Duan, Tianhao Zhang, Pieter Abbeel and Peter Chen. Abbeel is a professor at the Department of Electrical Engineering and Computer Sciences at Berkeley Engineering, California, the engineering school from which the other three founders all graduated.

Abbeel's lab at UC Berkeley has enjoyed a number of recent breakthroughs in robot learning, including a robot that organises laundry, robots that learn (simulated) locomotion, and robots that learn vision-based manipulation from trial and error and human VR teleoperation.

The firm is backed by lead investor Amplify Partners, plus some big names in AI, such as Jeff Dean, Geoffrey Hinton, Yann LeCun and Raquel Urtasun. To date, it has raised $27m in funds, and the last two-and-a-half years have been spent researching, developing, testing and deploying its AI technology at Obeta, as well as facilities across North America and Europe.

Photoneo

Photoneo's MotionCam-3D camera can inspect objects moving at 40m/s

Covariant

Robots learning general abilities such as 3D vision at AI robotics firm Covariant


zhu difeng/Shutterstock.com

Brainwave

Based on the team's background and research in deep imitation learning, deep reinforcement learning and meta-learning, Covariant's vision is the Covariant Brain – a universal AI system for robots designed to be applied to any use case or customer environment. The robots learn general abilities, such as robust 3D vision, physical properties of objects, few-shot learning and real-time motion planning. This enables them to learn to manipulate objects without being told what to do.

As with all new technology, it is reaching the commercial stage for continuous production that often proves to be the challenge. This is because there are an almost unlimited number of scenarios that a robot could encounter, so it's impossible to program each one. In order to operate and add value to customer environments, a robot must be able to see, understand, make decisions, learn from mistakes and adapt to change, on its own.

The Covariant Brain is designed to enable the robots to do all of these things continuously in real-world environments.

> **'Developing safe and reliable robotic technologies to address challenges in the real world is extremely hard'**

Abbeel said: 'Even though we are just getting started, the systems we have deployed in Europe and North America are already learning from one another and improving every day.'

Raquel Urtasun, chief scientist at Uber Advanced Technologies Group, and a Covariant investor, added: 'Developing safe and reliable robotic technologies to address challenges in the real world is extremely hard. Whether you're talking about self-driving cars or warehouse operations, robots encounter an endless number of unexpected scenarios. Covariant has demonstrated exceptional progress on enabling robots to fill orders in warehouses, which could unlock many other robotic manipulation tasks in other industries.'

Take your pick

One of the biggest challenges when it comes to automation in warehouses, said Mark Williamson, group marketing director at Stemmer Imaging, is bin picking. Programming a robot to pick up random objects from boxes relieves people of this task. The issue in logistics, says Williamson, is that robots are asked to handle a far wider range of items than is the case in a factory, which makes the imaging and gripping aspects a lot more complicated.

While a number of companies produce general bin picking tools, Williamson feels they all typically target one kind of imaging technique. 'They might use a low-cost stereo camera; they might use a laser profiling system, or a fringe projection kind of system,' he explained. Each technique has advantages and disadvantages, but while one 3D method might work well in a certain scenario and with certain equipment, it might not work for a different application.

Stemmer Imaging's solution to this dilemma came following the acquisition of Infaimon, a provider of software and hardware for machine vision and robotics. →

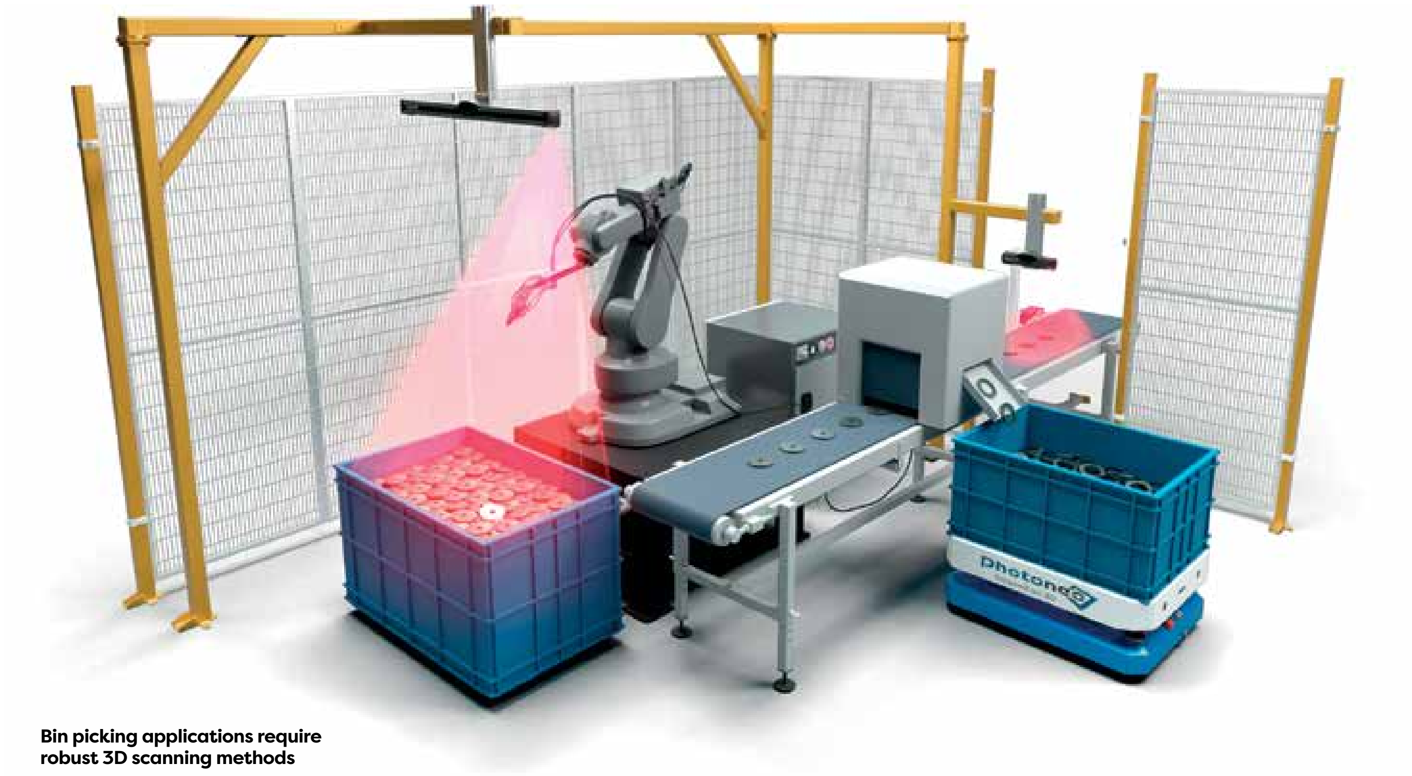

Photoneo

Bin picking applications require robust 3D scanning methods

→ The Infaimon InPicker 3D bin picking system uses a variety of 3D imaging techniques, such as passive stereo, active stereo or laser triangulation to recognise and determine the position of objects randomly piled up in a container. It also interfaces directly with a variety of robots.

'The unique thing about the InPicker,' said Williamson, 'is that it is robot agnostic, so it doesn't matter what make of robot it is. It is also camera agnostic.' An application using InPicker, he explained, becomes bespoke without the customer having to develop everything from scratch, which would traditionally be the case when building a bin picking system.

Get a grip

InPicker requires a CAD model of the gripper and a CAD model of the item that the robot will handle. There is also the functionality to upload the CAD model of the box the objects are housed in.

This is important, said Williamson, because if the robot does not know the shape of the box, it will not know how to approach the box to pick up products from it – if the box has a high sidewall, for example, the arm could try to pick an object from the wrong side.

'That's a key thing,' said Williamson, 'when you start to do proper bin picking, and certainly for things like warehousing, where they have these plastic boxes with a pile of bits in them.'

Warehouse robots also might use multiple grippers, and potentially different vision techniques. In this scenario, said Williamson, the first thing is to calibrate the space of the robot into the space of the InPicker. The InPicker tells the robot its co-ordinates and identifies what gripper is needed to pick up certain objects.

Williamson explained: 'Using the 3D matching technology, it looks at all the shapes of the products in the box and says: "I can pick up 30 of these and I see this particular one has the minimum inclusion of other products on top of it".'

> 'The market for imaging to support automated warehousing will only continue to grow, particularly when it comes to bin picking'

Picking at speed

3D imaging firm Photoneo is close to launching its MotionCam-3D camera, which won the 2018 Vision Award at the Vision trade fair in Stuttgart, and which can be used for bin picking.

MotionCam-3D is based on the company's parallel structured light technology, implemented on a custom CMOS image sensor. It is able to inspect objects moving as fast as 40m/s, and can be used for warehouse applications such as autonomous delivery systems, object sorting, as well as bin picking applications.

The precision of the camera is designed to allow robots to handle smaller and sensitive objects in palletising, de-palletising, machine tending applications, quality control and metrology.

Photoneo has been selling engineering samples of the camera, but plans to release an official version before the summer, according to the firm's CTO, Tomas Kovacovsky.

Last summer, Photoneo developed an autonomous mobile robot, which was designed for transporting materials in factories and other industrial facilities. It works thanks to two 270° laser scanners, two ultrasonic sensors, and its vision is courtesy of two 3D cameras. It can carry up to 100kg direct weight, pull up to 300kg, supports a dual-way operation with an interchangeable back and front, and has a maximum speed of 1.125m/s in both directions.

This year, there will be a new version of the robot launched. Kovacovsky explained: 'The new one we are going to release this year will be called Phollower 200. It will have new features and is designed for transport of materials in larger factory environments.'

Looking to the future, Williamson believes that the market for imaging technology to support automated warehousing will only continue to grow, particularly when it comes to bin picking. 'Obviously that is very much an exploding industry,' he explained. 'InPicker started around three to four years ago, and we are working on second generation products, so it's evolving all the time.'