Machine Vision 2.0

What's next for vision systems?

ClearView Imaging

Table Of Contents

01 Introduction

To scientists and science fiction writers alike, the potential of machine vision has been obvious and exciting for over a century – even though they didn't yet know it by that name.

Machines that could 'see' like humans: what an incredible idea. What if they could move and behave like humans too? What would that mean for our lifestyles, for space exploration?

Yet, for even the simplest of robotic visions to be made real, a computer brain had to be developed that could recognise an image, process it and extract relevant and useful information from it. What those visionaries needed was an artificial eye attached to an artificial brain: computer vision.

In 1957, a breakthrough was made and not by a computer scientist or engineer, but by a psychologist. Frank Rosenblatt[i] is now known as a computer vision pioneer. His Perceptron machine, described as 'giant' and 'thickly tangled with wires', used an early artificial neuron network to sort images into very simple categories such as 'triangle' or 'square'.

By 1966, Marvin Minsky[ii], one of the founders of the AI lab at MIT, was setting a new bar for computer vision. He set his undergraduates a very specific assignment. He told them to: 'Spend the summer linking a camera to a computer. Then get the computer to describe what it sees.'

It took a little longer than a summer holiday, but by the 1970s the first applications of computer vision - Optical Character Recognition (OCR) systems - were available commercially. These primitive systems made typed, printed or handwritten words legible to computers, helping blind people to read.

Since 1982, when Time Magazine declared it the year of the computer rather than following its own tradition of naming the year after a notable human, technology has bounded ahead. For decades, Gordon Moore's 'law'[iii] remained accurate, with the number of transistors in a dense integrated circuit doubling roughly every two years whilst halving in cost. By 1975, he'd reduced this cycle to eighteen months. Technology is now advancing too quickly for Moore's law, with experts predicting that computers will reach its limits during the 2020s.

As technology continues to explode, the possibilities of machine vision are inspiring a whole new wave of pioneers. Technology that was, until very recently, most used on manufacturing assembly lines, in product and labelling inspections or in food sorting factories, is now driving the future of autonomous cars, unmanned drones, facial recognition at electronic passport control gates, biometric security systems and robot-executed micro surgery.

In turn, research and development occurring in these new sectors is driving advancements in technology that is feeding back into industry, improving reliability and accuracy in existing systems and further opening up possibilities for new applications.

In this eBook, we look ahead to the most exciting advancements in machine vision and the technologies that are making what was impossible yesterday, possible tomorrow. 3D imaging techniques, faster camera interfaces, embedded vision and deep learning: Machine Vision 2.0 is coming.

Are you ready?

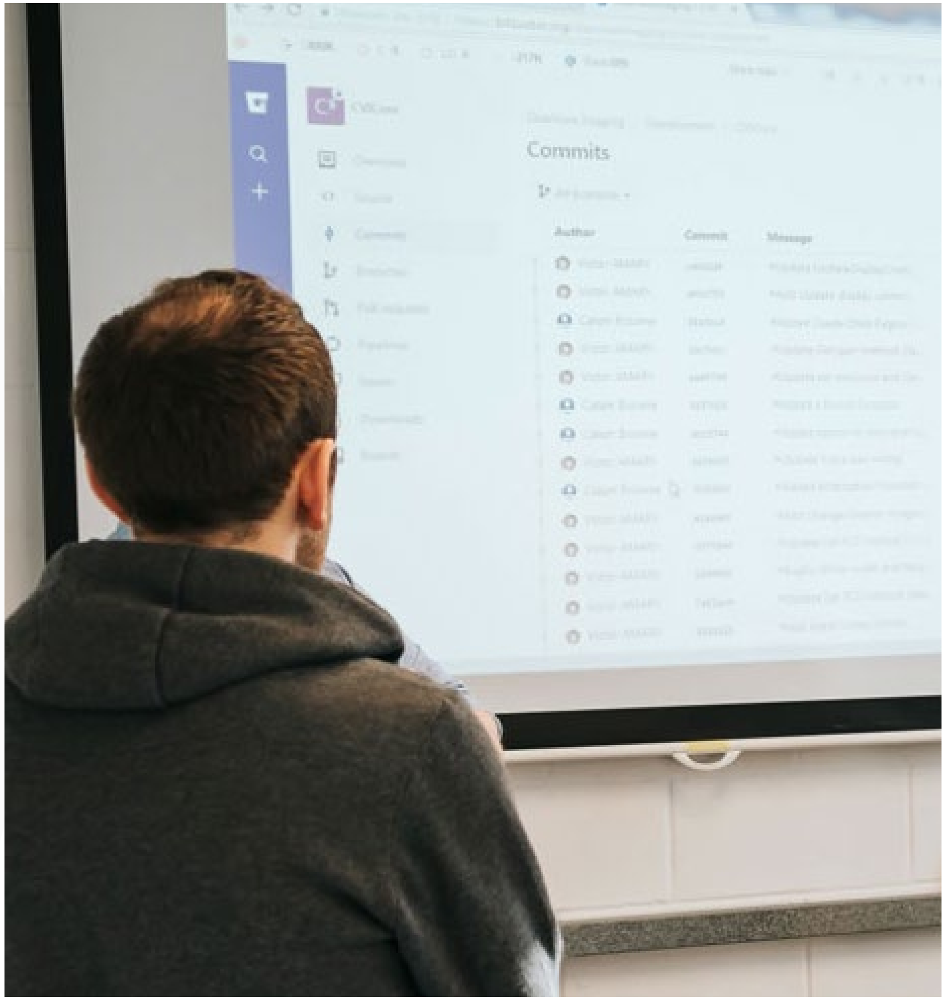

02 Machine Vision in Business, Today and Tomorrow

To understand the future potential of machine vision systems we first have to look briefly at how they are being used in business today. Automotive manufacture has used visually-guided robots on its assembly line since the 1970s, as has the food distribution and labelling sector.

Automated quality inspection is now a perpetual need because such a large number of businesses already rely on the accuracy and efficiency vision systems offer. In addition to this, governments are offering incentives to businesses who use smart factories, further increasing demand for AI and IoT integrated systems, the adoption of Industry 4.0 and the development of new connected technologies.

This increasing demand is reflected in financial analyst forecasts. According to leading analyst firms[iv], the global industrial machine vision market was worth US$ 7.91 Billion in 2017 and is expected to reach up to US$ 19.22 Billion by 2025. This represents a compounded annual growth rate (CAGR) of up to 8.9%.

But are these forecasts underestimating the true extent of our surging demand for machine vision? Have they covered all the new permutations of this technology and the way its branching into so many other industries?

In the case of artificial intelligence, for instance, our broadening need for AI is fuelling software development as deep learning increases the ability of machines to recognise objects. This will certainly further reduce the need for human intervention in the detection of product variation and defect in manufacturing. But what about its wider uses in medical diagnoses, translation software or facial recognition?

Image capture and processing equipment is also fast developing. Cameras today pick up more detailed images at higher resolutions, better networks allow faster data transfer and more powerful computers enable faster image processing. Added to that, everything is getting smaller and cheaper. This is leading to more industry sectors introducing machine vision not as a production line inspection tool, but as an integral part of their products.

Machine and computer vision are the reason autonomous cars, electronic passport control gates at airports and biometric security systems work, to mention a few examples.

Yet because the technology is developing at such a rate, machine vision is now caught at an interesting junction: the demand for more sophisticated components is outstripping user awareness of how quickly technology is changing, while many existing systems lack the necessary flexibility to accommodate the latest kit.

Enter machine vision 2.0.

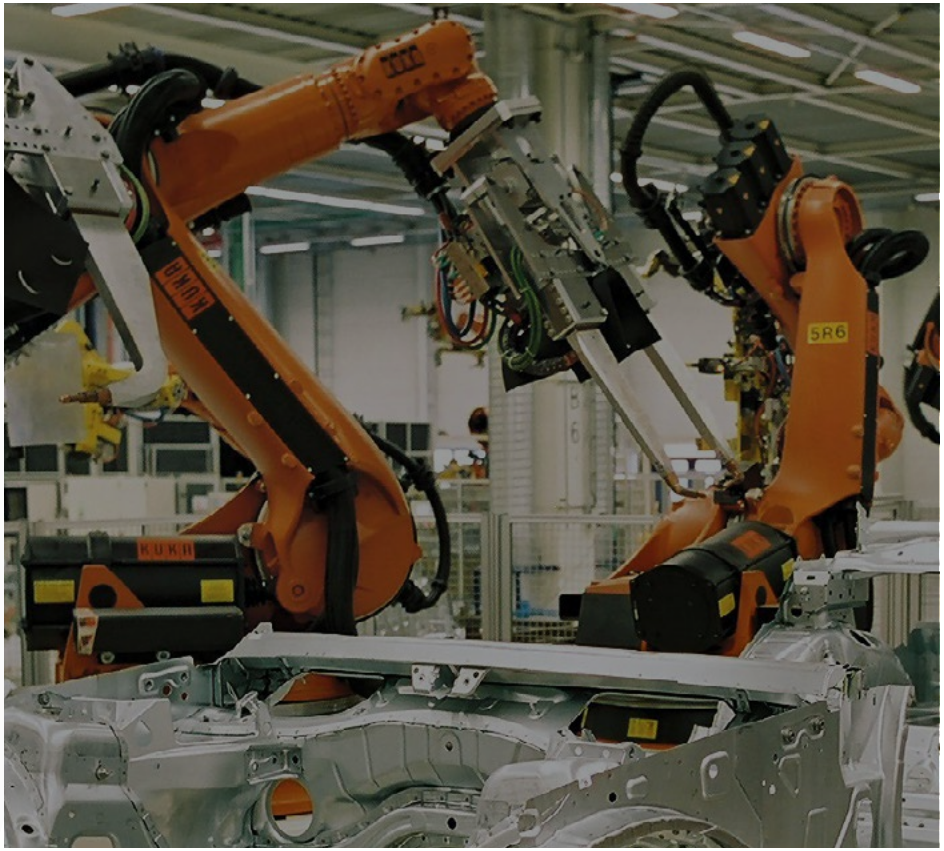

03 What is Machine Vision 2.0?

Machine vision 2.0 takes advantage of the latest technologies to introduce highly sophisticated machine vision systems into entirely new arenas. For businesses, the key to capitalising on machine vision 2.0 is to recognise the new horizons opening up to them and what that could mean for their industry as a whole.

Critical too, is the understanding that the established view of machine vision systems and their capabilities might already be outdated. As John Lewis, Editor in Chief of Vision Systems magazine wrote in a recent article[V],

"Since knowledge has an ever-decreasing lifespan, especially for those developing machine vision and image processing systems, keeping up with changes is an increasingly challenging task that is more important than ever."

If machine vision 2.0 is a combination of approach to system design, recognition of the abundance of new possibilities and awareness of the latest technological advances, it's also about maximum flexibility and building in as much forward-looking scope as possible.

Perhaps some of the greatest advances in machine vision are in perception and approach. When Roger Bannister first ran a four-minute mile in 1954, men had been attempting it for decades with many believing it to be impossible.

Within two months of Bannister's achievement it was as if the realms of possibility had been rewritten, when two runners ran sub four-minute miles at the Commonwealth Games. At least 1,400 male athletes have broken the four-minute barrier to date. It is now the standard expected of all elite male middle distance runners.

We've seen a similar speed of advancement in camera resolutions, since CMOS image sensors have taken over from CCD sensors. A CMOS device is far less expensive to manufacture and can produce extremely high-resolution images, such as the 120 megapixel Canon 120MXS sensor, which gives it approximately 60 times the resolution of a standard HD sensor.

3D vision systems, for example, that can scan products at much higher speeds, are now becoming affordable. While Hyperspectral or SWIR cameras – only recently the highest of high tech – are now easy to find.

So, machine vision 2.0 is also a response to the cycle of innovation and application. As amazing new technologies emerge, new ideas for system design appear, while new applications of machine vision come full circle to inspire the next generation of technologies.

But now that machine vision is branching out from the heavy industrial space, what are the new forces helping to drive it forward?

CMOS cameras were once considered to produce lower quality images but that's no longer the case. In fact, only a few years ago, it was considered impossible for a CMOS camera to achieve the picture quality they are now achieving as standard. CMOS sensor design is now advancing so quickly that a 20 megapixel CMOS camera costs around £500 compared with a traditional £9,000 CCD camera, which is only 16 megapixel.

Size and cost reductions of key system components, together with the ability of the latest systems to collect many layers of data at once at a far lower price point than even recent predecessors, have allowed many businesses access to machine vision that was once out of their reach.

04 What is Driving Advancements in Machine Vision?

Consumer demand

As we hurtle into the future of machine vision, it's becoming more difficult to discern the point where technology development leaves off and software development steps in.

As consumers, we are thirsty for data-led services that can be managed in the palms of our hands: we are also the most demanding consumers in history. We are spoilt for choice with incredible apps coming onto the market all the time, but the technological wizardry behind them is not enough to secure our support any more: they must also be fast, 24/7 and cheap.

The good news for us greedy consumers is that technology industry listens and responds. The great news for industry is that with every advancement that makes our smart phones smarter, research and development of new ways to use machine vision in industry also gets a boost.

Apple changed our view of mobile phones in 2007 with the iPhone, its sensitive camera and superior OS. Today, we can turn our smart phones into wine, wildlife or plant identification tools. We can use apps to identify fashions and where to buy them from a photo on social media and standard photo management apps sort our pictures using facial recognition.

In its most flippant form, computer vision allows us to add cartoon features to our faces on Snapchat. At its most significant, it helps us manage our money, park our cars, it could help foil a home burglary and even save our lives.

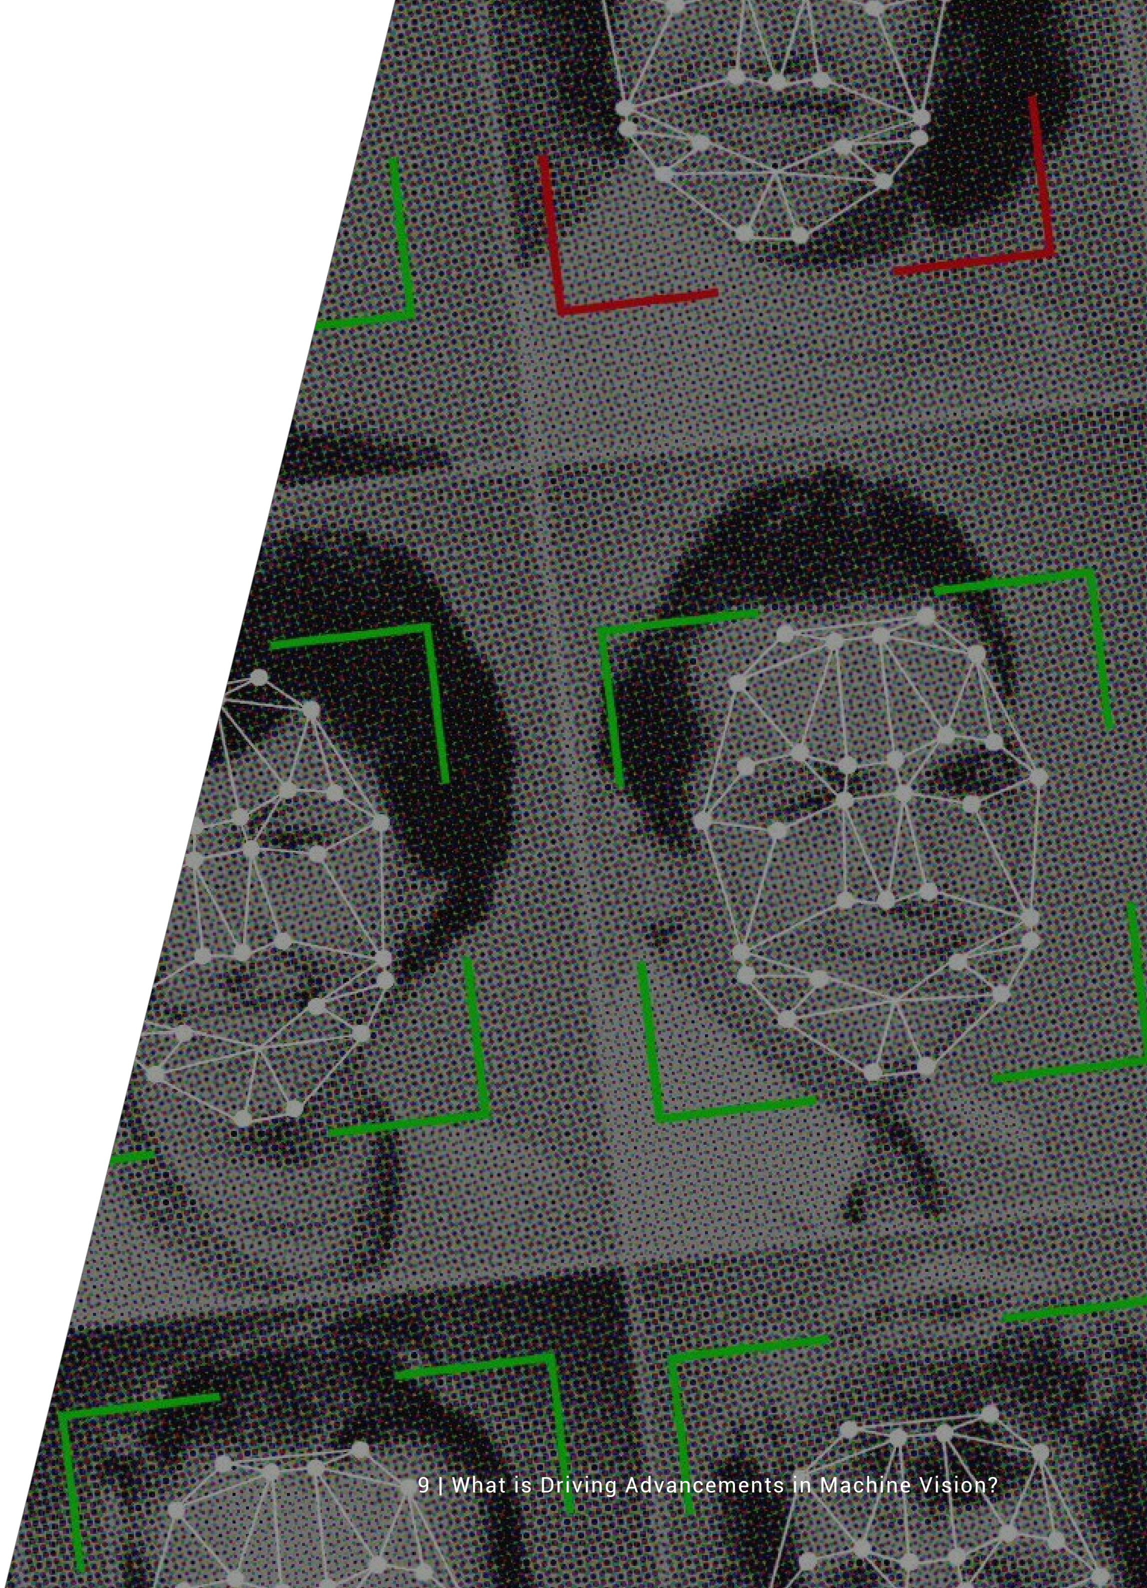

Public health and safety and law enforcement

Machine vision is being used extensively to help people in their daily lives and to keep communities safe. What began with the earliest OCR software is now taken for granted as part of the everyday, helping visually impaired people to 'see' the images on their social media feeds.

Meanwhile, Microsoft's InnerEye[vi] is a computer vision project that uses a variety of visual detection methods to offer faster, more accurate medical diagnoses and recommend treatments for certain cancers. Cancer Research London and the University of Edinburgh are also developing a new technique known as Revolver[vii] (Repeated Evolution of Cancer). AI picks out patterns in DNA mutation within cancers and uses the information to forecast future genetic changes.

Banks too are using computer vision to help customers securely deposit cheques into their online accounts using their smart phones. While home security systems use inexpensive cameras, smartphones and computer vision software to distinguish homeowners and their guests from intruders.

Police forces all over the world use similar facial recognition software to identify known criminals and they also use automatic number plate recognition (ANPR); OCR software that was originally developed and introduced in the UK in the 1980s. In 2005 ANPR first helped solve a crime; the murder during an armed robbery of PC Sharon Beshenivsky[viii]. Now cameras and image processing software are routinely used by the police to spot mobile device infringements, people not wearing seatbelts and other dangerous driving behaviours.

If a serious accident were to happen, injuries might be treated by a robotic surgeon[ix]. Currently, robotic surgical arms carry out keyhole surgery guided by human doctors. However, there are already plans to automate simple procedures such as suturing.

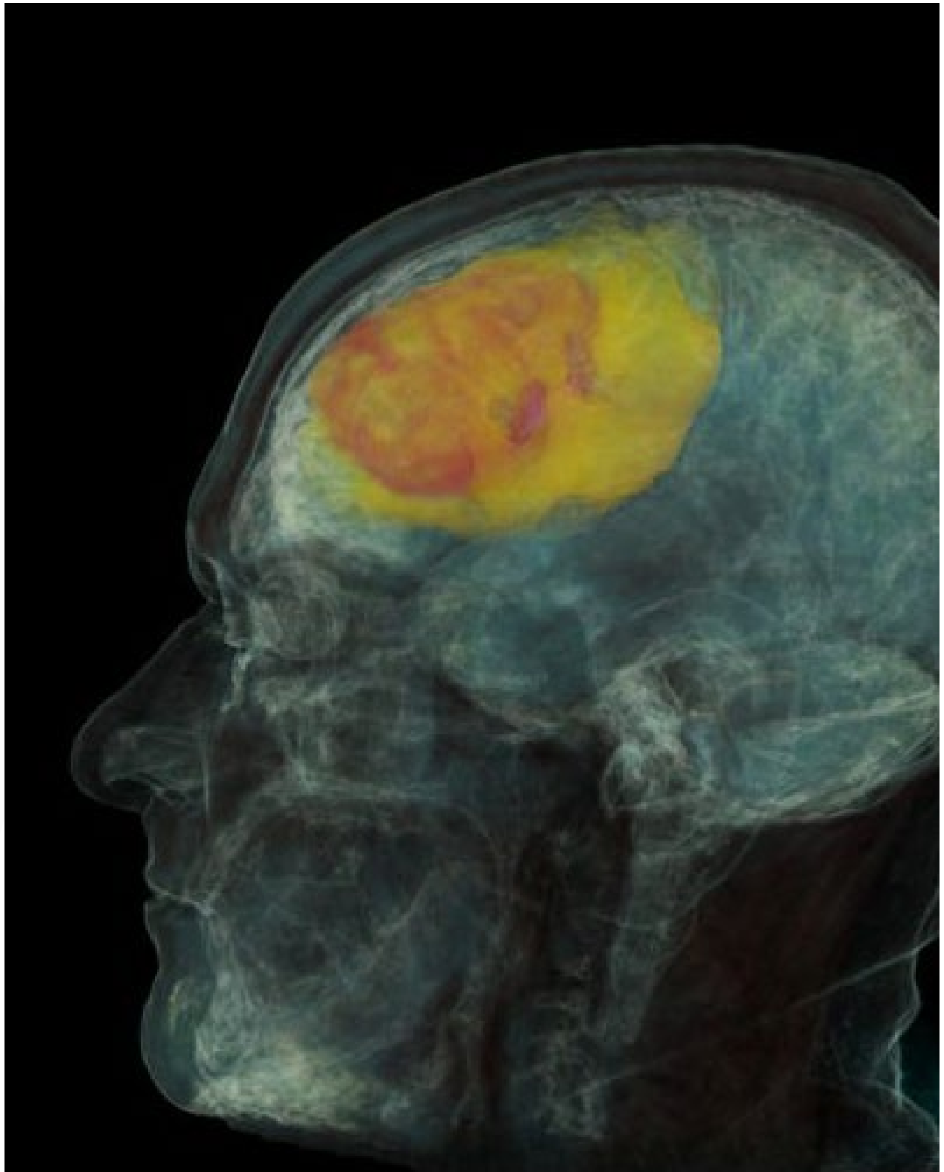

Microsoft's InnerEye

05 Key Technologies That Make Machine Vision 2.0 Possible

3D Imaging Techniques

Ensuring a high-quality image is captured is critical for any vision system. Currently most commonly used in industrial vision systems for precision, multiple angle inspections, measurement of complex free-formed structures and pick and place systems, new applications for 3D imaging are being explored all the time. Most obvious is its use in entertainment media such as gaming and VR.

The more closely 3D imaging systems mimic the human eye, the better the results for vision systems. We see in 3D thanks to receiving almost-the-same images through two eyes set slightly apart. The brain combines these two slightly different images into one; a process called stereopsis.

3D machine vision uses a variety of techniques to create point cloud (3D) data. Popular methods include Laser Triangulation (profiling), Pattern Projection (using structured-light), Time-of-Flight and Stereo. Each of these approaches have their pros and cons, so understanding depth resolution, speed and implementation cost is key when choosing which technology to use.

The main drawback to 3D imaging systems in the past was the time it took to process point cloud (3D) data. Today, with multi-core processors available plus powerful GPU's, 3D imaging systems can function in real-time.

As with every vision system component today, massive advances are being made across the board in 3D imaging too: laser triangulation, pattern projection, time-of-flight, stereo and also in the software used to process the 3D data.

Time is still a factor, because 3D image processing is so computationally intensive. However, as Moore's law fades into obscurity – and consumer urge for increasingly realistic gaming graphics strengthens – processors are becoming smaller and more powerful very quickly.

As more powerful, portable hardware becomes ubiquitous, the options for smaller businesses to use 3D imaging techniques open up. To give you an idea of how broad the appeal of 3D imaging is becoming, in 2017, Facebook launched its 360 cameras. Comprising of 17 FLIR USB3 cameras, combined with Facebook's proprietary architecture, the user is able to shoot in six degrees of freedom (6DoF).

This is great news for VR gaming fans and also for advertisers. Airlines can show off their first-class cabins in 360, for example. And perhaps most interestingly, the same cameras (with different 3D technology) are used in 3D vision systems industrially to help robots place objects more accurately.

3D Scanning: the future's part in rebuilding the past

Not only are 3D scanners available commercially, some are already small enough to hold in your hand. As Sony's Xperia XZ1 shows, what was once extremely heavy, slow processing of 3D images can now be done on a phone processor. This is a massive leap forward. If 3D scanning can be done on such a tiny scale, imagine the implications for future applications.

Computer-aided design (CAD) has been around since the 1960s. As 3D scanners develop and the prices reduce, they are being used increasingly within CAD processes, increasing accuracy by up to 50% as standard. This is very important when precision, measurement and quality control is paramount to your business.

Already integral to motion capture, industrial design, orthotics and prosthetics, reverse engineering and prototyping, quality control and inspection, 3D scanning is also key to the future of 3D printers. As 3D image capture becomes more precise and processors even more powerful, any process that requires perfect replication will benefit.

Model makers are already using 3D scanners to help them create perfect miniatures of famous landmarks from 2D images. But what about real architects using the same technology to recreate actual historic buildings using modern materials as a way of preserving the best bits of the past?

New York architecture and engineering firm, EDG, are pioneering a new cost-effective restoration technique using 3D printing technology[x]. Director, Richard Unterthiner, said:

"This is where modern technology and traditional craft meet: 'modern ornamental' is a new form of architectural digital sculpture that transforms renderings into printable objects."

Laser scanning software allows the company to recreate virtually anything, from colonnades and cornices to a whole building. Their digital catalogue of parts means architects anywhere could reprint the same mould.

The same technology is being used in a Beijing funeral home to reconstruct the faces of any deceased person who suffered injuries in an accident. And once again, to recreate treasures from ancient cities such as Palmyra in Syria, hugely damaged by Isis in 2015.

It is almost impossible to predict what other applications will benefit from 3D imaging techniques a year from now.

3D imaging and scanning the key to robotics

In many ways, of course, industrial robots are already among us. Automotive manufacturing has relied on them for years and there are businesses already in existence who hire out robots to factories on contract. A quick search online brings up articles such as: Expert Advice for Buying Entry-Level Robots – making it easy to feel that the automation revolution has already happened.

Today, one of the most popular uses for 3D robotic vision is in pick and place applications. 3D vision allows the robotic arm to pinpoint the orientation of the part that requires handling, even if the location and position of the parts vary. In comparison, older 2D systems required the parts to be carefully placed prior to inspection to avoid error.

Now, a robot with 3D vision has gone even further. It can recognise objects. This means it can make intelligent, real-time decisions on behalf of the automation system, based on captured visual data. These latest robots are flexible, cost-saving, can work even in low-light conditions and can easily adapt to recognise new products as they are introduced.

Cameras with better resolution and better interfaces – such as the latest models by Zivid – are coming onto the market with increasing frequency. So too is highly advanced image processing software – such as Matrox Imaging Library. As these two core components of any robotic system grow more sensitive and powerful, robots are able to recognise shapes, textures and 3D objects faster, more accurately and process where they should go.

In industry, these advances mean fewer human overrides, vastly improved productivity and fewer product recalls. They are narrowing the need for any human intervention at all, as they broaden the potential for robotics applications.

We are still some way off seeing humanoids manning reception desks and serving us during flights. However, if MIT Technology Review predictions prove to be as on point as those around the saturation point of Moore's law, we will see this as part of our reality within the next 40 to 75 years. A critical part of this humanoid evolution lies in better use of deep machine learning: something that is currently occupying some of the brightest minds in the world.

06 Making Machines More Like Man With Deep Learning

As with computer vision, deep learning as a concept was inspired by the desire to create artificial neural networks (ANNs). These systems progressively improve their ability to accomplish tasks by analysing and comparing examples, allowing them to make decisions based on their findings.

For example, in image processing, an ANN will use rule-based programming to determine which images contain dogs and which do not. They will first learn to recognise images containing dogs that have been manually labelled as 'dog' or 'no dog'. They will then use these analytical paradigms to identify dogs in any following images.

ANNs are used in computer vision, speech recognition, machine translation, social network filtering, playing board and video games and in medical diagnosis.

Deep neural networks (DNNs) contain multiple layers between input and output layers. They are more complex and sensitive than ANNs, able for instance to recognise the style of painting in an image.

Deep learning has now become 'superhuman' in image recognition, beating human competitors in tests. It helps Google Translate support over one hundred languages, helps businesses manage customer relationships and shows very specific ads to an ideal mobile audience.

Somewhere in between the two lie convolutional neural networks (CNNs). These use a variation of multilayer perceptrons designed to require minimal processing and are most commonly applied to analysing visual imagery.

According to Forbes, in manufacturing, deep machine learning is increasing defect detection rates up to 90%, improving semiconductor manufacturing yields by up to 30% and reducing scrap rates. It's also reducing supply chain forecasting errors by 50% and lost sales by 65% with better product availability.[xi]

Deep learning is also helping cars learn how to drive themselves.

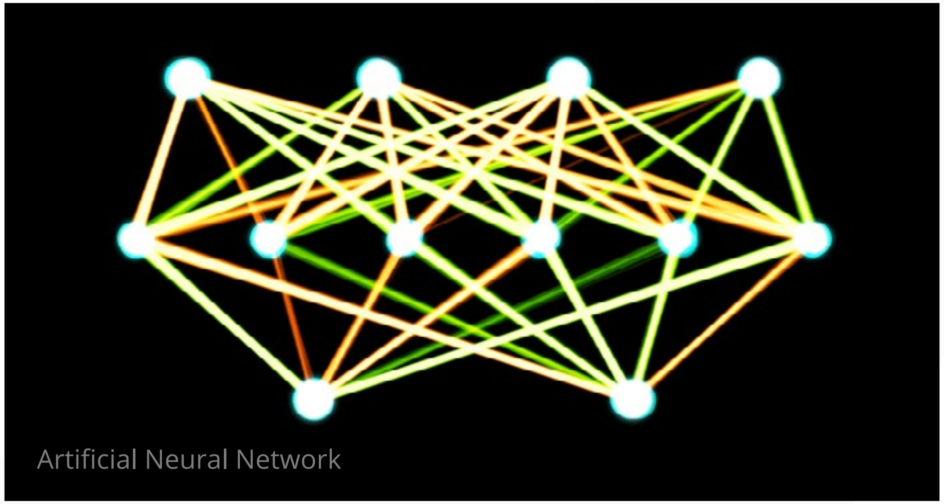

Artificial Neural Network

Self-driving vehicles

For autonomous vehicles to become a reality, developments in engineering, vision systems, computing and processing power have all been necessary. Staggering then, to realise in the past five years, autonomous driving has gone from 'maybe possible' to 'definitely possible' to 'inevitable' to 'how did anyone ever think this wasn't inevitable?'[xii]

Deep learning advances are making it possible for autonomous cars to learn as they drive. Whilst they will come with tens of thousands of scenarios built in, to help them do things like detect lane lines and identify pedestrians crossing the road, they will also be able to learn from experience and also to navigate on their own.

Core components of any self-driving car are:

- **Computer vision system –** cameras and sensors, real-time data processing
- **Sensor fusion –** to integrate sensor and visual data to determine environment
- **Localisation accurate to within 2CM –** high performance GPS plus 3D representation of direct locale
- **Path planning software –** AI and machine learning updated via the cloud constantly
- **Control –** all data is processed and used to drive the car safely

This means a single car could carry multiple megapixel cameras, ultrasonic sensors, radar, lidar, GPS, maps, cloud-based data processing and management, smart technologies, connectivity, and of course a powerful computer to coordinate it all. The chips needed to process all the data in these cars make them 'super computers'. It wasn't so long ago that such machines took up an entire room.

Intel, Qualcomm and Nvidia are all vying for top place in the price-and-size-lowering, while processing-more-for-less power battle of the chips. In yet another example of the march of technology, according to Waymo, creator of Google's Alphabet car, driverless vehicles will be everywhere by 2020.

Just how much data these computers will process daily is a point of debate. Estimates have it at anything from 380 TB to 5,100 TB per year.[xiii] Storing so much data has been one of the issues faced by developers of these cars, together with ensuring there is enough battery power to feed all the systems and 3D mapping. The results? Massive companies such as Tesla, TomTom and Here are devoting some of their biggest and brightest brains to these problems.[xiv]

With such attention, not only are autonomous cars even more likely to be everywhere by 2020, but the horizon for applications that can be developed or improved using advanced deep learning and the technologies necessary to support is about to get much broader.

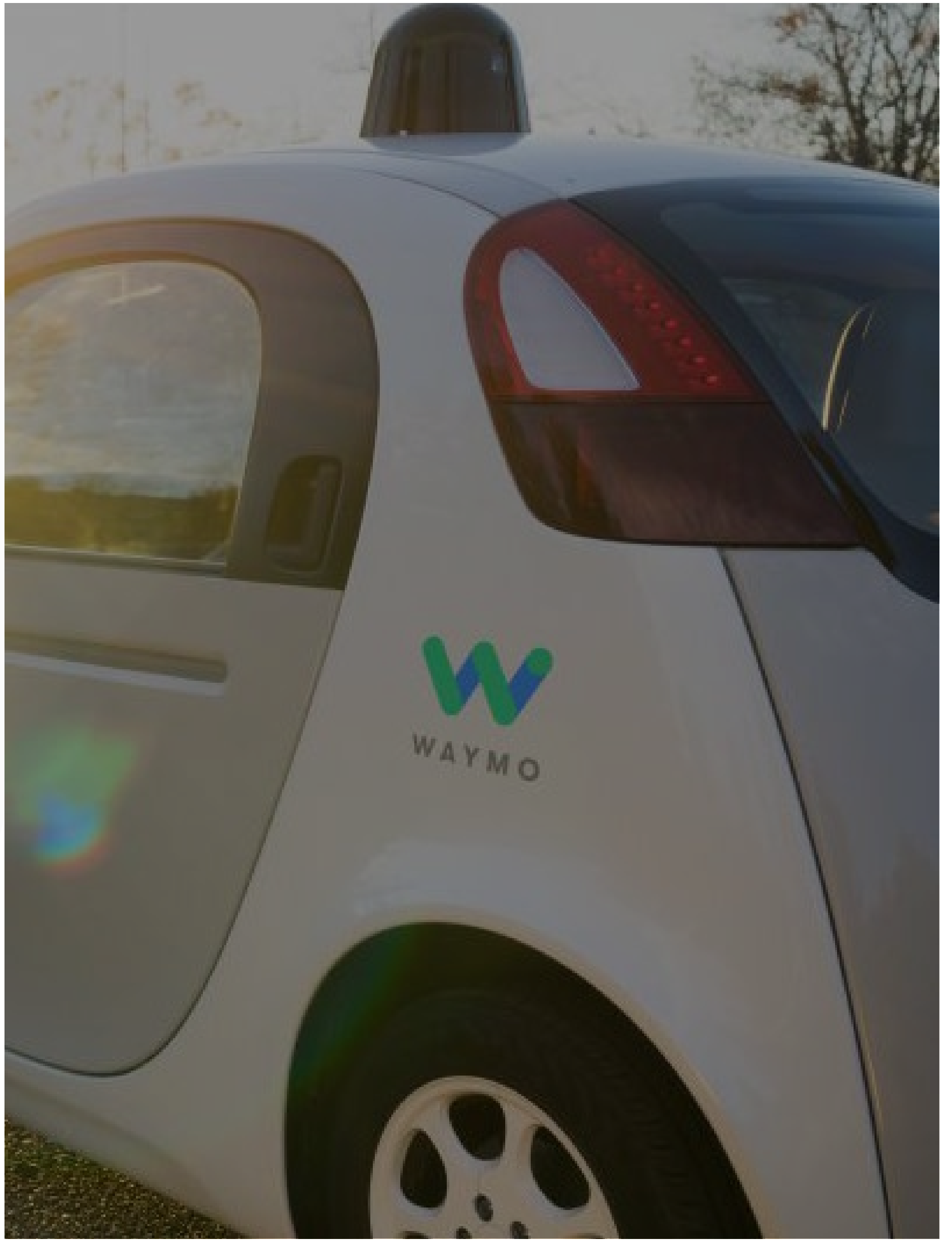

07 Embedded Vision: Cameras and AI All-in-One

What could be better than a camera linked to a computer?

A computer with a camera built in, plus all the processing power needed to classify the images as well. This is embedded vision.

Heterogenous computing, where systems use more than one kind of processor, made embedded vision commercially viable. Now, with heterogenous systems the norm, it's about making them smaller, more lightweight, with lower energy consumption and more cost-efficient.

Let's go back to the iPhone. The iPhone X, released ten years after the first model, came with facial recognition software built in. This allows the phone owner to use their face to unlock the phone. It is powered by a dual-core neural engine built into the new A11 Bionic Chip, with data stored in its Secure Enclave and all processing done on-device.

Embedded vision is still a relatively new area of machine vision, thanks to recent advances in the miniaturisation of cameras and processors. However, its enormous potential is already being explored.

FLIR is an obvious leader in this field. In 2016, they released a camera with a Boson Thermal Core to enable intelligence into a thermal imaging camera – think 'Predator vision', from the 1980s Arnold Schwarzenegger movie. Powering this incredible piece of technology is Movidius's Myriad 2 chip. So tiny is this 12-core, low power processor that The Boson camera is half the size, one tenth the volume, one seventh the weight and twice as power efficient as the model that came before.

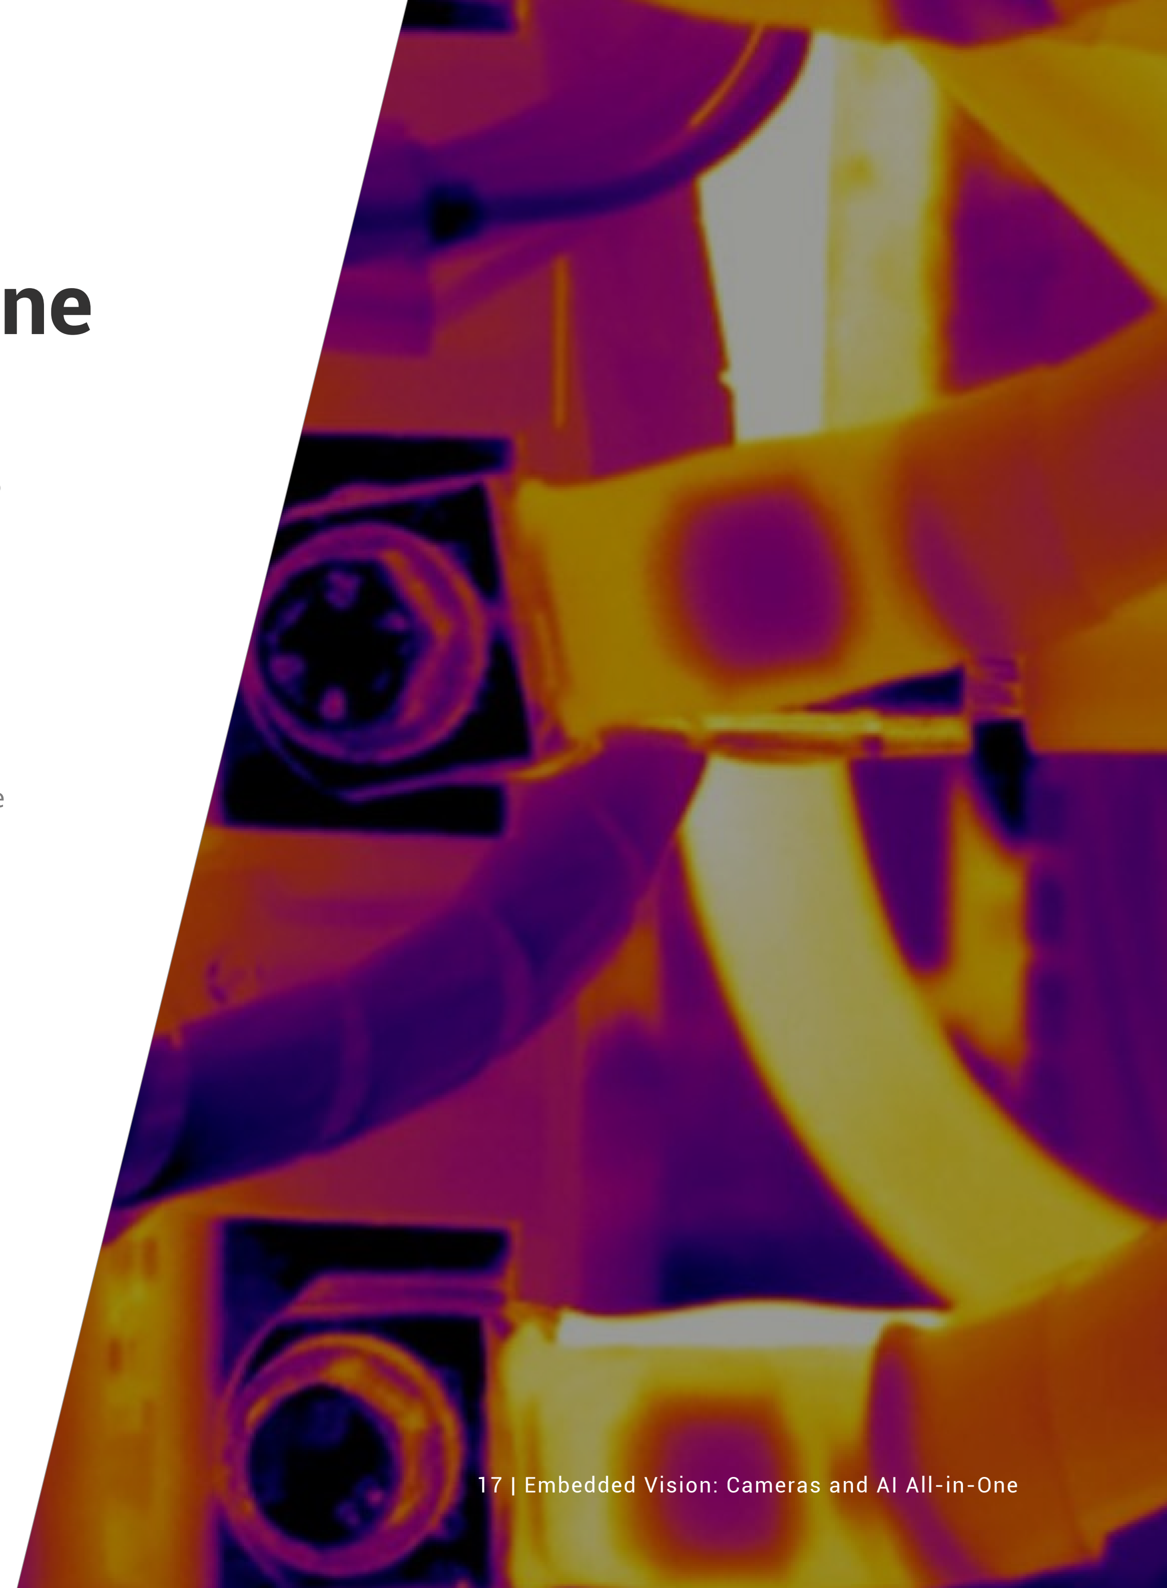

The next FLIR camera with a Myriad 2 chip is going to be released shortly (by FLIR IIS, previously Point Grey). This visible light camera will be so powerful that it could, theoretically, be used by medical organisations to screen microscopic human cells for cancer without pathologist intervention.

The only thing that is stopping this advancement at the moment is a lack of software to enable the processing of the images. If the right deep learning training set existed, it could be loaded onto the Myriad 2 chip and the world could have a miniature cancer screening embedded vision machine.

Robotics and autonomous cars are both set to benefit from embedded vision systems and, as we've seen across the board with these technologies, when they do there will be another drive for advancement. This in turn will open up new possibilities and so the cycle continues.

Already it's likely that the next generation of embedded vision systems will incorporate 3D imaging and deep learning, requiring more powerful processors without a dip in energy efficiency. Our humanoid robot friends are feeling closer by the day. In the meantime, it's the retail sector that is perhaps benefitting most from embedded vision, thanks largely to consumer reliance on hand-held, dual processor, connected devices.

Embedded vision in retail

Online shopping has massively disrupted the high street but now, with embedded vision, the high street is redefining itself. This technology is helping retailers cut operational costs, communicate with their customers and offer entirely new experiences.[XV]

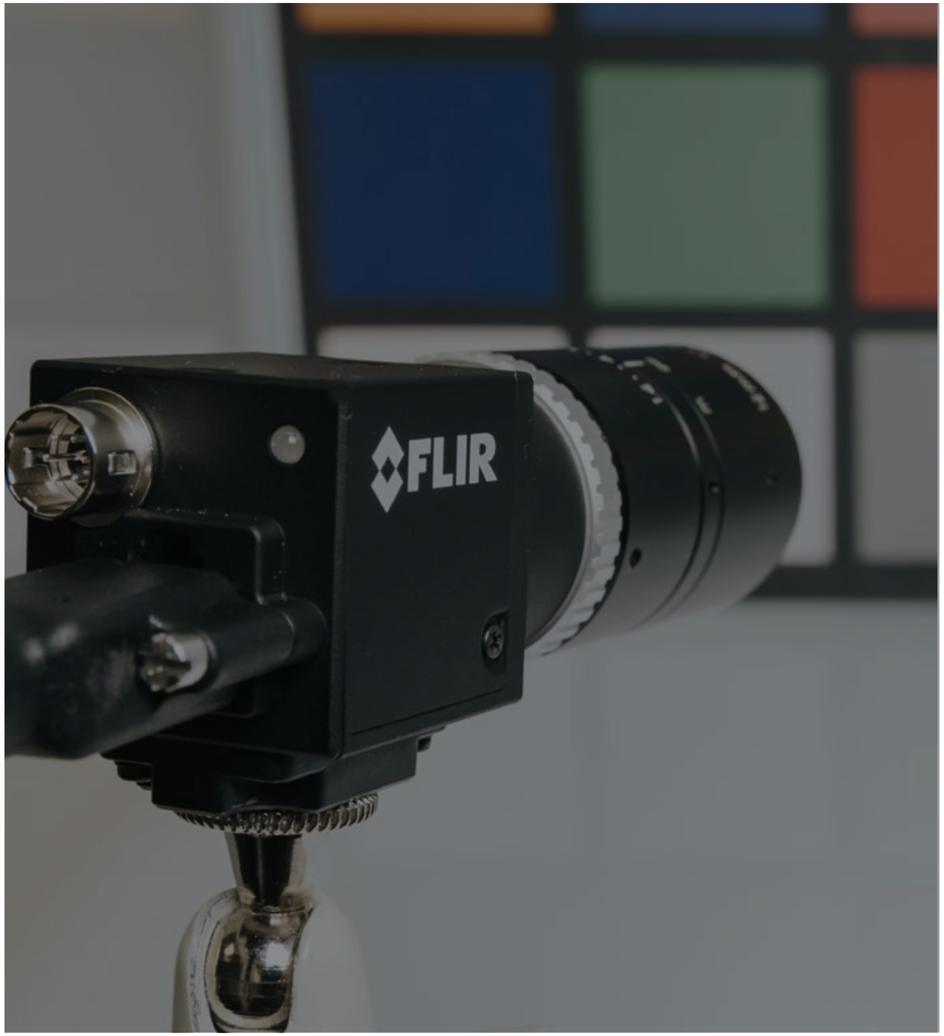

In store:

Touchless checkouts:

Think Amazon's new high street store experience. Customers get an app on their smartphone as they enter the store. From there, they can simply pick up items and leave. Advanced vision systems identify the products and charge the customer accordingly.

Customer service:

Robots with embedded vision are able to guide customers through a store and even recognise facial expressions.

Augmented reality:

'Smart mirrors' use embedded vision to allow shoppers to virtually try on clothes without actually having to change.

Store monitoring and marketing intelligence:

Embedded vision can help with store security monitoring, and also to monitor customer behaviour and shopping patterns. The data captured can inform promotional activities and other marketing strategies.

Behind the scenes:

Warehouse automation:

Autonomous mobile robots (AMRs) use embedded vision systems, along with other sensors, to move goods around the warehouse.

Inventory management:

Robots can also be used to roam stores, monitor stock levels and inform management of most popular products and items that need replenishing.

08 Machine Vision Interfaces

The final part of any machine vision system is how all the components are linked together. The best image in the world is useless if it can't be processed and used in a meaningful way, data integrity and transfer speed is critical to the efficacy of any system.

Interfaces are being developed and updated constantly as they strive to keep up with speed of technology.
Here are the most widely used:

USB 3.2

USB 3.2, launched in September 2017, has doubled the maximum transfer rate of the previous USB 3.1 SuperSpeed interface. The increase is a result of multi-lane operation over existing wires that were intended to flip-flop capabilities of the Type-C connector. This double side connector enables two USB 3.1 Gen 2 lanes in parallel. Subsequently, since each USB 3.1 Gen 2 has a maximum transfer rate of 10 Gbit/s, the combined new transfer rate of USB3.2 is 20 Gbit/s.

Thunderbolt 3

Thunderbolt 3, which uses the USB-C connector, is the first generation that supports USB. It boasts a pretty impressive bandwidth of 40 Gbit/s, halves power consumption, and simultaneously drives two external 4K displays at 60 Hz (or a single external 4K display at 120 Hz, or a 5K display at 60 Hz when using Apple's implementation for the late-2016 MacBook Pros) instead of just the single display previous controllers can drive. Furthermore, its USB power deliver specifications enable source or sinking up to 100W. The downside of Thunderbolt 3 is the lack of availability of affordable optical extenders, which will slow down its uptake.

CoaXpress 2.0

CoaXPress (CXP) 2.0 double single lane output is 12.5 Gbit/sec. However, it still necessitates a frame grabber. CXP can support multi-destination transmission, through which multiple host PCs can process frames in parallel. Furthermore, doubled uplink speed allows for 500kHz+ triggering rates. Another superb characteristic is the ability to use the interface over a 40m standard coax cable, with cable cost reductions taking place when shifting over to bespoke multi-core cable assemblies.

10 Gigabit Ethernet

This computer networking technology has been designed to transmit Ethernet frames at a rate of 10 gigabits per second and is the latest iteration of Ethernet. The interface caters for cable length of up to 30m over CAT5e, with even longer lengths for CAT6A. It requires no frame grabbers, supports IEEE 1588 PTP, and its large install base has led to setting up of favourable consumer environments.

CameraLink HS

CameraLink is a serial communication protocol standard designed for camera interface applications based on the national semiconductor interface Channel-link. It requires a host side frame grabber, as well as offering a trade-off for speed and length via multiple cabling options. This interface is best suited for high-speed applications (e.g. line scan cameras).

09 What's Next for Machine Vision?

It might be easier to look at industries that are not, in some way, linked to machine vision today. It's extraordinary to think how much this technology has advanced since Rosenblatt and his Perceptron machine in the 1950s.

Ultimately, the future of machine vision applications lies in the developments of its four key components: 3D imaging techniques, embedded vision, deep learning and the interfaces that enable real-time processing of the highest resolution images.

One thing is certain: the forum is well and truly open for ideas, innovations and new applications for machine vision. The key is also to keep an open mind. Limits to what this technology can achieve are being blown away every day.

Machine Vision 2.0 is about thinking outside the box. If you have an idea that involves a machine being able to process an image today, it's highly likely that the technology already exists to turn your idea into reality.

And it if it doesn't exist today – have another look tomorrow.

10 References

i https://en.wikipedia.org/wiki/Frank_Rosenblatt

ii https://www.britannica.com/biography/Marvin-Lee-Minsky

iii https://www.investopedia.com/terms/m/mooreslaw.asp

iv Research and Markets

Grand View Research

Stratistics Market Research Consultancy

v https://www.vision-systems.com/articles/print/volume-23/issue-7/departments/my-view/keep-your-knowledge-fresh.html

vi https://www.microsoft.com/en-us/research/project/medical-image-analysis/

vii https://www.bbc.co.uk/news/uk-scotland-45381947

viii https://en.wikipedia.org/wiki/Murder_of_Sharon_Beshenivsky

ix https://www.bbc.co.uk/news/health-45370642

x https://www.independent.co.uk/life-style/design/3d-printing-buildings-architecture-design-futuristic-restoration-a8297911.html

xi https://www.forbes.com/sites/louiscolumbus/2018/03/11/10-ways-machine-learning-is-revolutionizing-manufacturing-in-2018/#588755e023ac

xii https://www.wired.com/story/guide-self-driving-cars/

xiii https://www.tuxera.com/blog/autonomous-cars-300-tb-of-data-per-year/

xiv https://www.ft.com/content/2a8941a4-1625-11e8-9e9c-25c814761640

xv https://www.retaildive.com/news/smart-dressing-rooms-and-touchless-checkouts-4-promising-retail-tech-innov/353049/

Contact Details

United Kingdom:

+44 01844 217270
info@clearviewimaging.co.uk

Office Address

Unit 4, Thame Park Business Centre,
Wenman Road, Thame,

Oxfordshire OX9 3XA

Spain:

+34 680204123
info@clearviewimaging.es

Office Address

Centro Tecnológico Habitec.
Calle Marie Curie, 22.
Parque Tecnológico de Andalucía,
29590,
Málaga

ClearView Imaging