IOT INTEGRATOR

SEEING IS BELIEVING WITH COMPUTER VISION

VIDEO ANALYTICS ACCELERATES OPERATIONS, ENRICHES CUSTOMER EXPERIENCES, AND IMPROVES HUMAN HEALTH AND SAFETY.

TABLE OF CONTENTS

THE COMPUTER VISION LANDSCAPE

Computer vision (CV) works with artificial intelligence (AI) to train computers to see and interpret images the way humans would. CV applications can trigger automated action based on their interpretation of visual information and generate insights that people can use to improve decisions and actions.

Valuable in multiple environments, CV systems can recognize objects and people quickly, analyze audience demographics, inspect manufactured products, help improve workflows, and more.

Demand for CV is heating up: research firm OMDIA expects nearly two-thirds (64 percent) of all network cameras shipped by 2025 to be AI-enabled. By then, the company predicts that the total worldwide market for video analytics will reach $32 billion. In addition, Gartner estimates that by 2023, 65 percent of enterprise-captured images and videos will be analyzed by machines rather than humans.

Innovative applications for CV technology are emerging in all industries. This eBook examines the CV opportunities in four of them: **smart cities/government, manufacturing, retail**, and **healthcare**. First, however, it discusses the basics of how the technology works, starting with the role edge computing plays in visual analytics applications. It also discusses some basic CV deployment considerations and includes a summary of Intel's CV solutions.

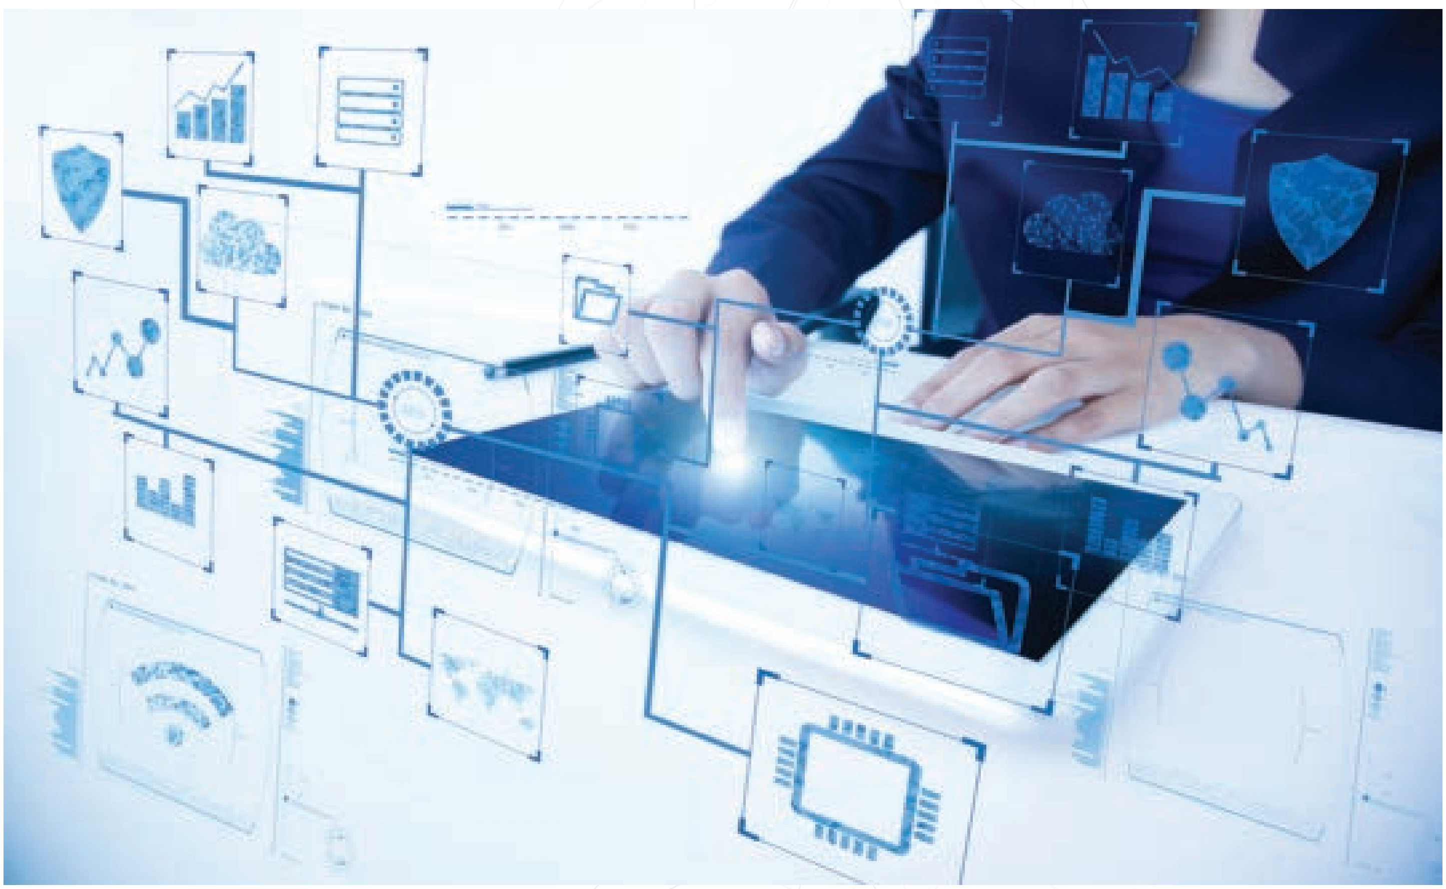

Innovative applications for CV technology are emerging in all industries.

ROLE OF EDGE COMPUTING

CV systems tend to be highly distributed. They capture, process, analyze, and store video and images at the locations where they're generated, which the computing industry typically refers to as "the edge."

The edge can be a computing or networking device, an Internet of Things (IoT) sensor, a vehicle, a TV, a camera, a wearable monitor, or any other digital entity. The visual edge requires a camera, image processor, lens, lighting source, and specialized video processors and software to gather accurate video data and process it.

Since edge computing allows visual information and other data to be processed and analyzed where it's generated, applications can access that data with little to no distance-induced delay. This enables rapid CV application response times, both in remote and urban locations. The low-latency advantages of edge computing are essential to many real-time and near-real-time CV applications, such as autonomous vehicles, public safety, transportation management, and augmented reality. These applications are also being fueled by emerging 5G networks, which in addition to slashing latency and providing significant bandwidth and data rate increases over earlier mobile networks, are highly distributed and rely on edge processing for service delivery.

CV systems capture, process, analyze, and store visual information in the edge locations where they're generated.

BASIC COMPONENTS AND OPERATION

First and foremost, CV requires cameras distributed in edge locations to capture images and video. These could be standalone IP cameras or other smart cameras; alternatively, cameras could be integrated into other devices, such as PCs, IoT devices, kiosks, medical devices, and robots. Devices at the edge should be constructed to tolerate the environmental conditions in which they're installed.

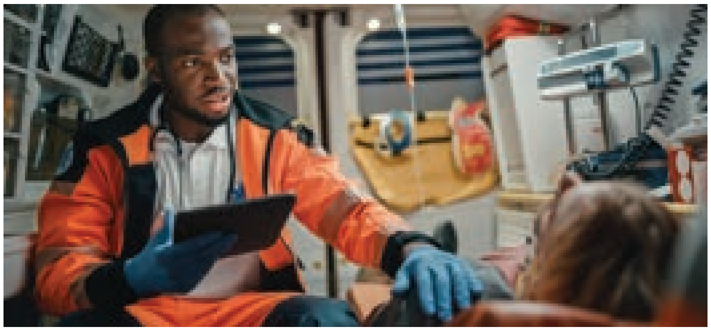

The low-latency advantages of edge computing are essential to many real-time and near-real-time CV applications.

An image or video sensor, matched with a lens, is responsible for analyzing the images captured by the camera. A processor—either embedded in the camera or, more often, in a locally connected device—runs the AI middleware and algorithms that transcode the video signals and enable the image analytics and interpretation.

For example, specialized standalone edge network video recorders (NVRs) might be used to collect and store video, either on their own or in parallel with add-on compute platforms, such as the Intel® Video AI Box, which can be installed as video volumes grow. If the standalone processing setup is used, a communications interface in the camera and a local network switch are needed to transmit the visual information to the video processing/AI software running on an external PC, NVR, or other hardware platform.

Creating CV applications that train systems on a certain set of visuals to produce relevant insights and trigger appropriate automated actions requires a specialized development toolkit that supports machine learning (ML). Finally, an Internet or other high-speed WAN connection is often required to forward processed data upstream to a centralized set of servers and storage devices that collect and aggregate analytics from multiple edge devices. These compute and storage resources might be in the cloud or a private data center.

Parifex, a partner in the Intel® IoT Solutions Marketplace, has a NANO-CAM sensing solution that combines LiDAR technology with a smart camera and incorporates AI and image processing for advanced traffic management and speed enforcement.

Beyond the hardware and software technology, CV initiatives require a human element and operational plan for applying the insights that the CV system generates. In other words, a successful implementation combines advanced software and hardware components but isn't only about technology. To optimize CV outcomes, it also relies on a sound strategy for how to act on the visual intelligence and trained personnel who know how to comply with the strategy.

INNOVATION OPPORTUNITIES

Video analytics promises to drive smart automation and enable intelligent, data-driven decision-making and actions in a nearly limitless number of situations. As with more traditional data analytics, these capabilities translate into improved operational processes, business insights, experiences, and human health and safety across multiple environments. Here are some examples of how CV technology can be applied for the betterment of business and society:

Smart cities and governments can improve public safety and security with automated, real-time transportation, parking, and crowd management. They can also improve sustainability programs with smart waste management systems that use CV to help quickly track and sort refuse.

Manufacturing floors and other industrial environments can use video analytics for real-time quality control, such as identifying defective parts before they're incorporated into production. These environments can also use CV to determine how to streamline work processes in ways that enable faster operations and greater productivity.

Healthcare providers are using CV to optimize the use of medical imaging data and improve the accuracy of radiology readings. They can also streamline operating room procedures for potentially improved surgical outcomes. The hope is that CV solutions will eventually accelerate medical evaluations, diagnoses, and treatment to enable the delivery of better patient care and, ultimately, extend the lives of patients and enhance their life quality.

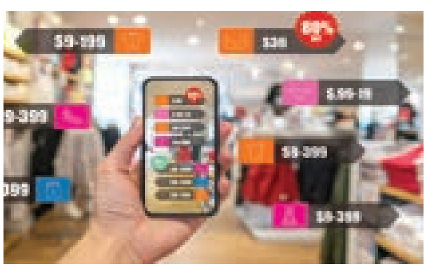

Retailers can enrich in-store experiences and improve merchandising, traffic flows, and checkout processes. For example, virtual fitting room mirrors combine CV and augmented reality to assist shoppers with product evaluations within the context of their preferences. CV can also help evaluate the effectiveness of product displays with insights as to more pleasing and efficient store layouts. Robots can wait on customers, make deliveries, and alert personnel that shelves need restocking. In the warehouse, CV identifies and resolves anomalies that could interrupt or slow operations.

Let's take a closer look at CV in each of these industries.

SMART CITIES APPLICATIONS

Many municipalities are striving to become smart cities. These initiatives involve building an urban ecosystem that uses digital technology to benefit citizens in areas such as transportation and crowd control, public safety, and waste management.

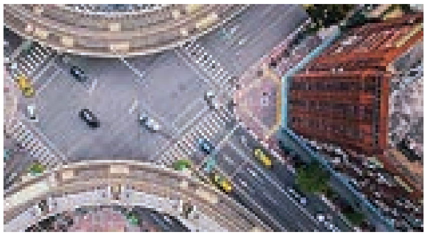

Transportation management. One CV-based smart city application is the intelligent transportation system (ITS). An ITS can control mass transit schedules and vehicle dispatching based on real-time usage requirements. It can also manage city parking by detecting available spaces and directing drivers to them.

The ITS might be used to dynamically adjust bus and train deployments based on current congestion levels, for example. Combining CV in vehicles, stations, rails, and roadways with hundreds of other data analytics variables makes it possible to detect abnormal events that could affect scheduling—accidents, seasonality, events, or other factors—and act accordingly to ensure the smooth flow of passengers.

Allied Market Research expects the market size for smart traffic cameras to nearly quadruple from $8.36 billion in 2020 to $32.34 billion by 2030.

Case in point: The City of San Diego built a smart mass transit system to help manage train schedules and crowd congestion. Devices installed on the city's metro trains have a GPS antenna that captures location, vehicle speed, and other relevant transportation data every five seconds. The collected data is processed in an IoT software platform from Davra, an Intel partner, to create accurate estimated arrival times based on dynamically changing variables. The continually updated schedule information is forwarded to train platforms and shared on digital signage and directly on vehicles to keep passengers informed of schedule status.

Public safety. Helping keep the public safe is an important component of a smart city. Embedded cameras and sensors can detect accidents, alert first responders, and stream live video from the field to command centers so that the city is aware of events as they unfold. For a bigger picture view, such as during a natural disaster, networks of drone surveillance cameras can be trained to visually identify situations and communicate them to first responders to accelerate action that helps save lives, prevents property loss, and deters crime.

Waste management. CV allows machines to "see" and classify waste at human-level recognition or better. Once waste items are scanned, identified, and classified, waste sorting robots in recycling plants can pick and sort the materials to maximize the number of items processed. In addition, smart trash bins can improve collection efficiencies. Wireless sensors in each bin communicate their levels to monitoring and analytics platforms integrated with collection service schedules to optimize driver routes and pickup frequency.

MANUFACTURING APPLICATIONS

Video analytics has several applications in manufacturing and other industrial environments. It can be used for predictive maintenance of the machines and systems in manufacturing assembly lines, for example. If a machine should fail or underperform due to aging parts or damage, production timelines will be affected. ML-based analytics, with assistance from CV systems, can anticipate when equipment will malfunction so operators can fix it ahead of time and keep operations up and running.

Manufacturers can also record video of factory personnel and machines at work and gauge whether workflows are optimum. Such systems can identify process improvements that accelerate task completion and save time and money. Moreover, CV can identify component defects so they can be replaced before they're used in production, helping avoid quality issues that affect output and delivery schedules.

For example, world-renowned agricultural and construction equipment maker John Deere has piloted a vision AI solution with Intel to help solve the costly, age-old problem of porosity in the welding process. Porosity, created by trapped gas bubbles as the welded metal cools, weakens its strength. John Deere is testing video analytics to automatically spot these and other defects in its Gas Metal Arc Welding (GMAW) process in 52 factories around the world.

The integrated system can generate insights in real time at the edge, at levels beyond what's perceivable by the human eye. It logs defects in real time and automatically stops the welding process so John Deere can correct the issue right away.

John Deere's CV system is intended to replace a manual process requiring scarce, highly skilled technicians who are difficult to find and hire. Identifying defects early is critical, because if discovered later in the process, they often require rework or even scrapping of full assemblies, which is disruptive and expensive.

RETAIL APPLICATIONS

Retailers can capitalize on visual feedback for innovative store efficiencies and richer customer experiences. They can also use it to spot and act on new in-store traffic patterns and trends that ultimately contribute to improved sales.

For example, CV systems analyze foot traffic throughout stores to determine where shoppers tend to gather and where congestion is most likely to occur. They'll also evaluate the effectiveness of product displays and locations, suggesting changes accordingly. These insights help retailers improve store layout, merchandising, and workflows, such as the checkout process, based on data rather than relying solely on intuition.

Simpler, faster shopper experiences. Merchants interested in improving self-checkout systems, for example, might replace customer bar-code scanning with computerized CV product recognition. Having a system that automatically identifies, processes, and rings up a product eases the buying process for customers while also enhancing security.

Nourish & Bloom, a family-owned grocer in Georgia, is using a CV system powered by Intel RealSense cameras, processors, and edge technologies to create a frictionless shopping experience. Its setup involves contactless checkout, allowing customers to use an app to check into the store as they arrive and pay without having to wait in a queue.

The grocer also uses CV-driven robots that take hot and cold orders and deliver them to customers located up to three miles away. When the robot arrives at its destination, the customer uses natural voice language to tell the robot the order number, which activates the opening of the robot's belly, where the order is contained.

Automated retail systems like these could also support autonomous stores, kiosks, and pop-ups, where appropriate.

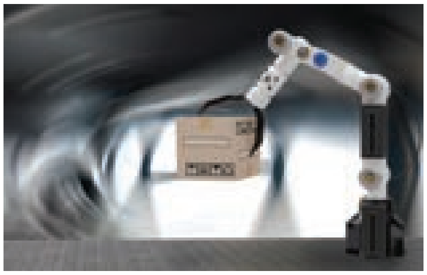

Inventory management. Another retail application for visual analytics is smart shelves. These use visual intelligence to discover and alert staff about places where product is out of stock or running low so they can be replenished. Similar capabilities communicate that smart vending machines need refills. CV can also identify damaged packaging and inaccurate pricing to keep transactions running smoothly.

Customer personalization. Virtual mirrors or fitting rooms use CV and augmented reality to help personalize and enrich customer experiences in the store. For example, the virtual mirror might make “complete the look” fashion recommendations based on the clothes the shopper is currently trying on, taking into account color, cut, size, style, fabric, and other variables.

Retailers can also capitalize on video analytics for security to reduce shoplifting and other criminal activity. And in warehouses, CV can identify and resolve irregularities in receiving, stocking, storage, picking, packing, and shipping that might slow down operations or result in dissatisfied customers.

HEALTHCARE APPLICATIONS

CV applications are being created to help the medical community improve the quality of care with faster, more accurate diagnoses and improved processes that also free up caregiver time to spend with patients.

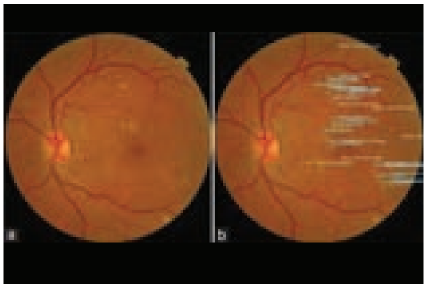

Improved imaging interpretation. CV applications that read and convert scanned images into interactive 3D models can potentially help medical professionals better understand a patient's health conditions and accelerate diagnosis and treatment. Using the technology to analyze CT and MRI scans, for example, helps doctors detect tumors, internal bleeding, clogged blood vessels, and other life-threatening conditions, because machines can now identify details invisible to the human eye.

> **"At least one study has determined that visual AI systems are more accurate than human radiologists when analyzing mammograms for signs of breast cancer."**

In fact, at least one study determined that visual AI systems were more accurate than human radiologists when analyzing mammograms for signs of breast cancer, reducing both false positives and false negatives.

In another case, GE Healthcare turned to Intel to help develop its Critical Care Suite, which detects a life-threatening lung condition called pneumothorax in chest X-rays. GE Healthcare optimized its algorithms in part using the Intel® Distribution of OpenVINO™ toolkit. By improving the algorithm performance, Intel helped GE Healthcare more than triple its pneumothorax detection.

Operating room efficiencies. CV is also contributing to the creation of smart operating rooms (ORs) that could help reduce errors and improve patient safety, according to the Healthcare Information and Management Systems Society (HIMSS). To comply with electronic health record (EHR) mandates, for example, surgical nurses typically make up to 100 clicks for EHR documentation during an OR procedure, HIMSS reports. CV systems observe and document the data automatically to reduce or eliminate the human input, freeing up healthcare providers' time to focus on patient care quality.

Smart OR technology can also help reduce errors, such as the 1,500 incidents each year in which the National Institutes of Health National Library of Medicine estimates that a foreign object is left inside a patient post-op. By keeping track of surgical supplies and tools used throughout the procedure, CV has the potential to help minimize these events and improve surgical outcomes.

DEPLOYMENT CONSIDERATIONS

Numerous architectural considerations come up when planning a video analytics implementation. For example, in your edge sites, do you process visual intelligence in the camera, an NVR, another device, or in the cloud? Are there privacy compliance rules that necessitate keeping all the visual information stored and computed in private edge equipment, rather than in the cloud? What, if any, visual data do you forward upstream to the cloud (or data center) for aggregation with other edge data for more holistic analytics?

All options have tradeoffs, and decisions depend on the industry and the use case at hand. For example, an organization wishing to provide foot traffic statistics and analytics to multiple stores in a shopping mall location might choose to process and analyze video streams in the cloud for aggregated and more meaningful results, rather than at individual locations.

Platform components. Visual data creates huge files, requiring a platform that can store, process, communicate, and secure them. Here are some considerations for each:

- ✔ **Store –** Storage can be provided in the camera, a connected edge device, or the cloud. Again, it's recommended to store the bulk of the data at the edge so that it can be acted on quickly and reviewed locally without latency-prone, long-haul network dependencies. Once stored and processed, it's most economical to send only the most meaningful data to the data center or cloud, as this approach reduces network traffic loads and associated communications costs, as well as the amount of data stored in costly cloud services or on-premises data center servers.

- ✔ **Process –** Specialized processors analyze visual information and accelerate CV applications. Intel IoT processors, for example, are common in edge appliances, such as NVRs and the Intel® Video AI Box, and high-end Intel visual processors are used in the cloud for visual data aggregation and analytics.

- ✔ **Communicate –** A resilient network is necessary for sending visual data to points where it can be analyzed and distributing the resulting insights to key personnel quickly. High-speed local switching networks share data for edge processing, and the Internet or other WAN service forwards relevant data from edge locations to sites where it's necessary for decision making or actions or to a data center or cloud for aggregated analysis.

- ✔ **Secure –**CV data is only meaningful if it can be trusted and protected. With the massive amount of data being generated by networked cameras and sensors, an ever-increasing number of threat vectors need to be taken into consideration. Executing a strong cyber-security implementation will protect the data and the network from threats and bad actors.

Partner for expertise. Enterprises and Intel partners can benefit greatly by working with domain expert ecosystem partners to assist in architecting a solution for a given organization's custom requirements. Intel's rich solution partner network can provide the hardware, software, and services required to create and enhance systems that enable desired capabilities, fill technology gaps, and help maintain solutions over their lifecycle.

For example, the Intel AI Inference Software and Solutions Catalogue compiles hundreds of Intel solution partners that have developed visual AI offerings using the Intel® Distribution of OpenVino™ Toolkit to accelerate CV deployments at the edge.

INTEL COMPUTER VISION PORTFOLIO AT-A-GLANCE

Intel has a rich portfolio of hardware, software, and solutions to enable and accelerate visual AI. The following matrix summarizes Intel computer vision offerings and how they can be used.

Hardware Processing Platforms	Video Recording/ Storage	Development Tools	Management Tools	Video-Optimized Silicon
Intel® Video AI Box Integrates video decoding on multiple streams with video analytics **Intel Video Processing Platform** **Edge Video Servers** Local storage, processing, analytics, I/O **Intel® Hyper Converged Edge Server (HCE)**	**NVR/AIO** All-in-one servers record video in a digital format and store it directly on internal disk drives or direct-attached storage **Edge Video Servers** Local storage, processing, analytics, I/O	**Intel Edge Software Hub** Preset software packages allow testing and experimentation with computer vision and deep learning applications on Intel® architecture **Intel DevCloud** A development sandbox in an Intel private cloud for learning about pro-gramming cross-architecture applications **Intel® Distribution of OpenVINO™ Toolkit** Open-source, write-once-deploy-anywhere toolkit for optimizing and deploying AI inference **Intel® oneAPI Video Processing Library** Video processing and compression/ decompression of video analytics and video streaming applications **Smart Video Evaluation Toolkit** Enables prototype implementation for custom use cases	**Intel Smart Video and AI Workload Reference Implementation** Set up and adjust video analysis workloads for optimum performance	**IoT (Edge) and Embedded CPUs** **Intel® Iris® Xe Integrated Graphics (GPU)** **Intel® Movidius™ Vision Processing Units (VPUs)**

WHERE TO FIND MORE INFORMATION

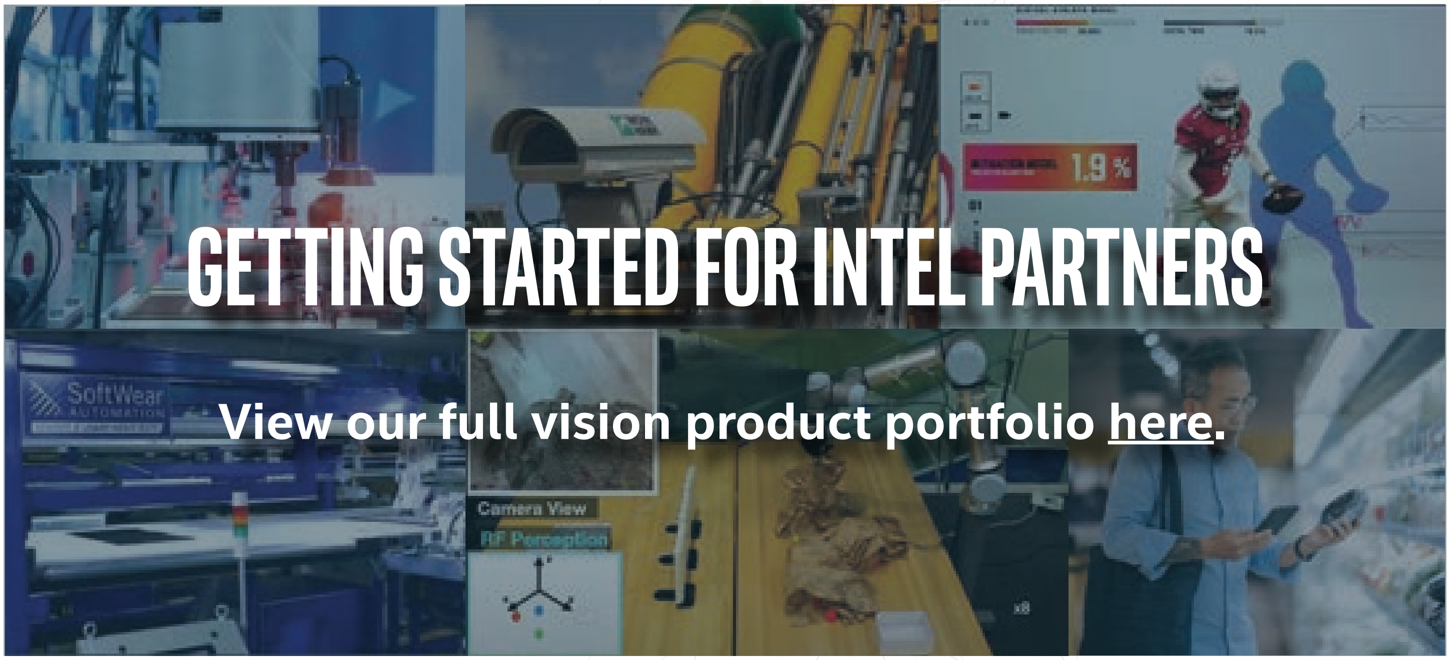

Stay in the Know: Visit IoT Solution Integrator

Get expert advice on IoT technology, find new partners, discover use cases, and learn more about IoT and edge computing innovation and solutions at http://www.theiotintegrator.com

Listen to our IoT Podcast

Hear engaging stories and candid interviews from integrators and IoT experts rolling out the latest Internet of Things and edge computing innovations in our IoT Integrator Wire podcast. You can subscribe on Spotify, Apple Podcasts, and Amazon Music.

THECHANNELCO.