The Complete Guide to Computer Vision for Business Leaders

Insight

Foreword

As we look to Artificial Intelligence (AI) to enable smarter, safer, more efficient ways of living and working, the ability to evaluate and respond to visual data in real time becomes increasingly essential. In a highly visual world, computer vision provides an unparalleled opportunity to assess, navigate and unlock insights from our environments.

Enabled by advancements in deep learning and near limitless access to image-based data, today's models are capable of accurately executing increasingly advanced tasks — from asset monitoring and predictive maintenance to inventory management and even disease prevention. But while this technology represents significant opportunity for growth and differentiation, relatively few organizations have realized the full potential. Achieving the momentum needed to not only initiate but also deploy these solutions at scale requires both an effective strategy and a transformative culture.

This guide is intended to serve as a roadmap for business leaders and broader advocates of computer vision to identify potential use cases, understand the technical considerations for execution, navigate potential risks and ultimately maximize the long-term value of a deployment.

Ken Seier

National Practice Lead for
Data & AI, Insight

ken.seier@insight.com

Amol Ajgaonkar

CTO, Intelligent Edge, Insight

amol.ajgaonkar@insight.com

What is **computer vision?**

Computer vision is defined as the use of advanced analytics or Machine Learning (ML) to process, evaluate and respond to digital images or video. This field of computer science leverages advancements in artificial intelligence to mimic, and in some cases surpass, the capabilities of human vision.

Computer vision works by analyzing pixel data captured by cameras as a matrix of numerical values. This data is processed at the edge or sent to the cloud where specialized software techniques, or algorithms, are used to recognize patterns associated with a target object or condition.

In order to accurately identify these targets, algorithms are rigorously and repeatedly trained using massive numbers of sample images. Historically, a highly manual process, advancements in machine learning and cloud compute now allow developers to leverage statistical learning algorithms and neural networks to automatically detect and classify image patterns. These techniques combined with access to vast image data allow many of today's computer vision models to reach high levels of accuracy.

Once trained, computer vision can be used to trigger follow-up actions or alerts in response to various inputs. This enables a wide variety of tasks to be performed with greater speed and accuracy while freeing employees to focus on more valuable, uniquely human tasks.

Section 1: The Basics of Computer Vision

Why computer vision?

With adoption of AI and the Internet of Things (IoT) proliferating across industries, there is no shortage of solutions available to extract data from a given environment. Various sensors provide critical feedback on individual data points, including temperature, proximity, vibration or pressure. Laser measurement is commonly applied to inline quality control in production. Radar, LiDAR and infrared systems are foundational to the navigation of autonomous vehicles.

While each type of sensor has benefits for specific use cases, computer vision is often able to provide more detailed information about more complex circumstances, giving systems a unique ability to identify, classify and react to varying conditions.

This includes not only evaluating visible circumstances, but also extrapolating data to infer what is unseen when objects are partially obstructed or hidden from view. When performing cycle counts in a warehouse, for example, a camera may only be able to view the first row of products. But applying 3D modeling based on product size and shelf depth enables the system to calculate the total number of items in a given space.

Computer vision can also be used in conjunction with other sensors to derive deeper insights. A system responsible for inspecting product quality may be able to classify a defect and trigger diagnostics from other embedded sensors to pinpoint malfunctions.

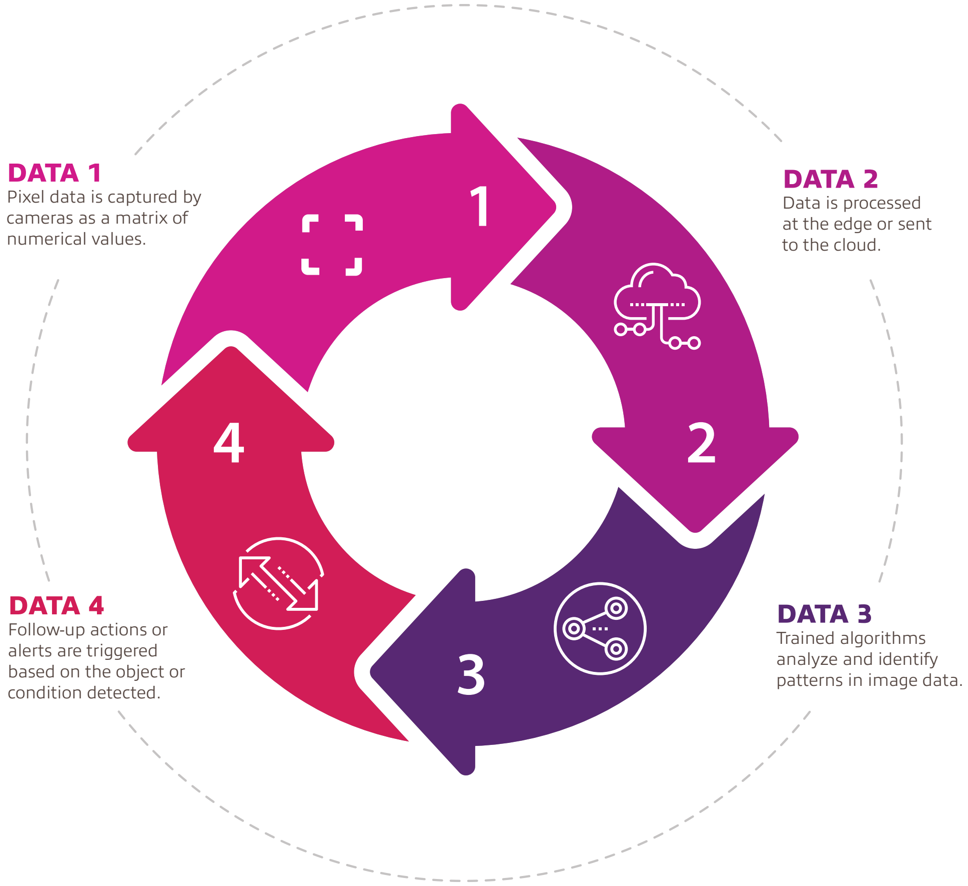

Why now?

In the last five years, the technology market has seen an influx of new products geared toward AI and computer vision. Cloud-native tools, frameworks and microservices now provide an entry point for data scientists to build and manage ML models with little prior experience. Simultaneously, advancements in edge devices have made it easier to run these models independent of the cloud. As these tools and techniques have become increasingly efficient, they have enabled more powerful, higher accuracy models to be deployed faster and at lower costs.

Until now, the availability of these technologies alone has not been enough to tip the scales from early adopters to the early majority. But in 2020, widescale disruption set a new pace for digital transformation as organizations sought new ways to optimize operations and support employee health and safety. Perceptions around the value of AI, automation and IoT shifted, and investments in these initiatives have since accelerated across industries.

75% of enterprises will shift from piloting to operationalizing artificial intelligence by the end of 2024, according to Gartner.[1]

While just **10%** of organizations are using computer vision today, **37%** have definite plans to implement and 44% report they are investigating the technology.[2]

By 2027, Grand View Research reports the value of the global computer vision market will reach **$19 billion**, up from just over $11 billion in 2020.[3]

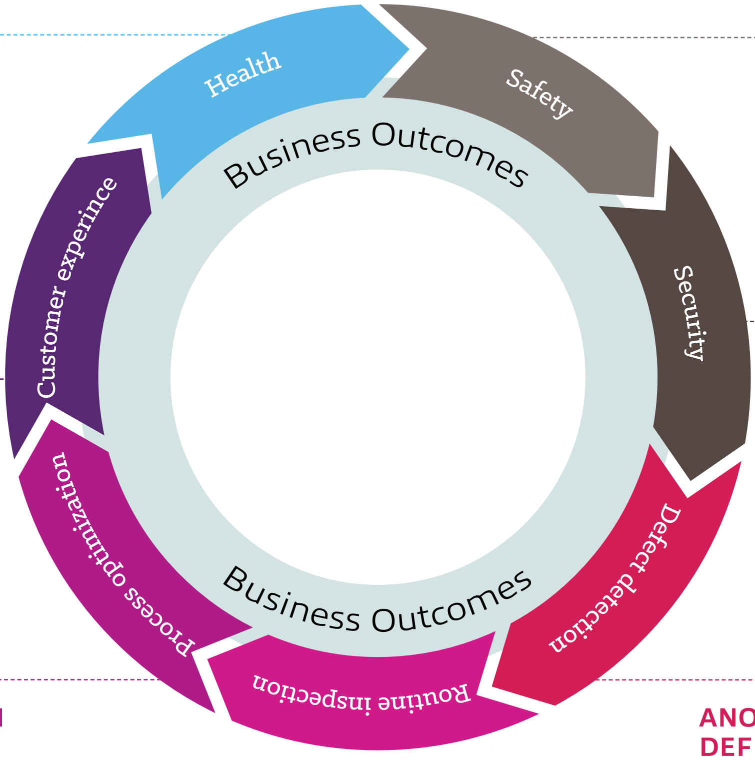

HEALTH

- Tracking food temperatures to reduce spoilage
- Identification of Personal Protective Equipment (PPE)
- Thermal detection of human temperatures
- Routine reminders for handwashing or other procedural steps

SAFETY

- Monitoring "danger zones"
- Environmental surveillance to identify hazards
- Detecting required (or missing) safety equipment
- Triggering emergency response

CUSTOMER EXPERIENCE

- Automated checkout experiences
- Immersive product testing and visualization
- Providing capacity or "long line" alerts to staff
- Supporting customer service with image-based fault detection

SECURITY

- Automated license plate recognition for restricted entry
- Surveillance to detect unauthorized objects or activity
- Object tracking and POS correlation to reduce shrinkage
- Monitoring network activity to detect cyberattacks

PROCESS OPTIMIZATION

- Automated inventory and cycle counts
- Robotic automation for assembly or navigation
- Site surveillance and selection
- Document processing with OCR

ROUTINE INSPECTION & COMPLIANCE

- Real-time asset evaluation (electrical, pipelines, machinery, etc.)
- Recognizing and responding to environmental hazards

ANOMALY & DEFECT DETECTION

- Identifying flaws in production and assembly
- Monitoring equipment to support predictive maintenance
- Detecting faults in materials or processes

Capabilities and business outcomes

Computer vision systems can be trained to execute a variety of core functions. When applied to specific use cases, these functions enable businesses to solve a wide range of common challenges.

Image classification: Models are trained to recognize a specific object.

Localization: Models locate an object using a bounding box.

Object detection: Models identify multiple objects and locate them.

Object counting: Models identify how many objects are visible.

Object tracking: Models identify and locate objects in motion over time.

Image segmentation: Models isolate object pixels using edge detection or clustering techniques.

Optical Character Recognition (OCR): Models recognize and convert typed, printed or handwritten text into digital data.

Section 2: Business Use Cases

Use cases in energy and resources

According to research from Insight and IDG, 80% of leaders in the energy industry strongly agree that computer vision has the potential to save time and money — the most out of any industry. Among those investing or planning to invest in computer vision, 88% are exploring use cases for **employee safety**.[2]

There are many opportunities to apply computer vision toward automating processes in a way that reduces human exposure to high-risk environments. Stationary cameras or remotely operated drones dramatically reduce the need for manual inspection of pipelines, electrical lines or wind turbines. Flagging anomalies that require human intervention empowers employees to shift focus from finding problems to fixing them, while lowering costs, risk and instances of human error — which accounts for up to 90% of workplace injuries in the U.S.[4]

These types of alerts can ultimately **reduce downtime** by enabling early repairs when assets begin to show signs of wear. Hydroelectric dams, which require constant surveillance to evaluate surface deterioration, were among some of the earliest potential use cases studied by researchers in the early 2000s.

Monitoring environments or equipment can also help to reveal opportunities for **process optimization**. Computer vision has valuable applications in land surveys and site selection, and can even be used to identify issues during construction or equipment installation. Mining operations are increasingly looking to computer vision to monitor and improve the efficiency of equipment like drills and ore crushers without putting human workers at risk.

In the quest to improve **energy efficiency**, computer vision is being used to analyze satellite imagery, monitor weather conditions and improve the accuracy of power requirement estimates by region. Thermal imaging can also be applied to solar facilities to assess functionality and inform projected outputs.

How energy and utilities companies plan to invest in computer vision[2]

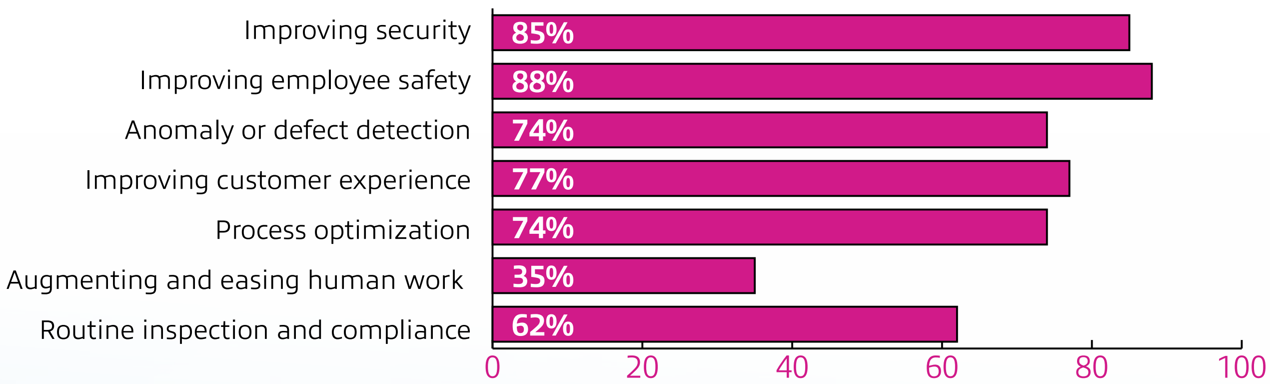

When a 100-year-old petroleum producer struggled with outdated oil field operations, Insight helped bring its processes into the 21st century by developing an automated monitoring solution for pumpjacks.

MANUFACTURING

Use cases in manufacturing

Manufacturing and production companies have traditionally been some of the earliest adopters of computer vision for **quality control** and process optimization, leveraging systems to perform inspections with greater accuracy and at higher speeds than human workers. Today, up to 78% of manufacturers have invested or plan to invest in computer vision to improve anomaly or defect detection.[2]

Beyond the production line, these systems have significant potential to augment or **automate tedious, dangerous or expensive work**, such as routine cycle counts and equipment inspections. While larger or more distributed organizations may benefit from drone-based systems, others may achieve similar gains by equipping workers with vision-based mobile apps. This approach not only reduces the potential for human error, but also provides an opportunity to remove employees from remote or high-risk environments, while enabling them to spend more time on more valuable tasks.

Visual monitoring of machinery helps organizations reduce downtime by providing early alerts when equipment is operating out of tolerance. Correlating data from cameras with other sensors can help to pinpoint the source of irregularities in production and improve the accuracy of **predictive maintenance** systems.

Beyond reducing the risk of equipment failure and exposure to hazardous environments, computer vision empowers more direct improvements to **employee safety**. Currently, workplace accidents cost American companies $62 billion per year.[5] Modern manufacturers and construction companies are using computer vision to combat this challenge by regulating active work zones, detecting required safety gear and ensuring authorization in restricted areas.

Similar models are also being applied to **improve security**, particularly in warehouse environments. By monitoring assets from shipment to shelf, organizations can better regulate activity, prevent theft and improve customer outcomes.

> When AltaSteel needed a way to prevent hazardous objects from being introduced into smelting workflows, Insight designed and implemented a custom vision solution to detect these objects — preventing costly damage to machines and helping to keep workers safer.

How manufacturing companies plan to invest in computer vision[2]

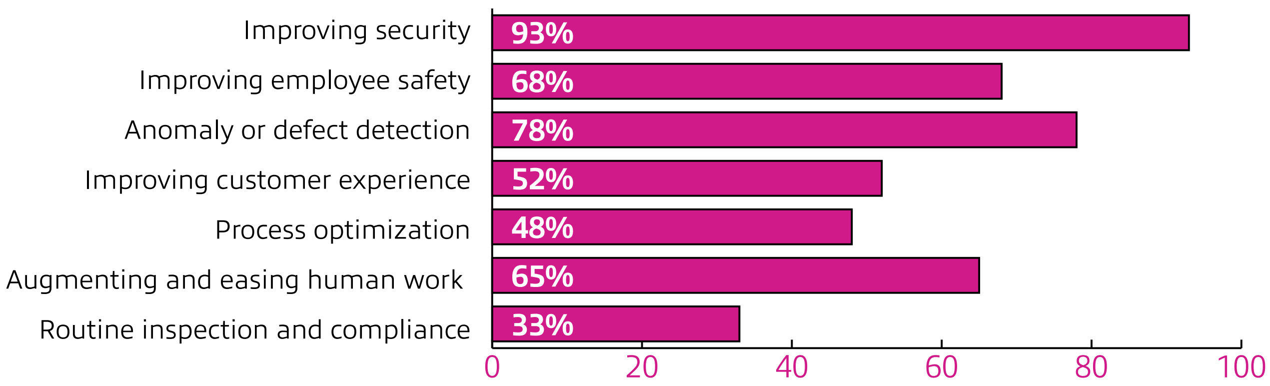

Use cases in retail

The ability to deliver the right product to the right customer in the right place and time is essential to profitability for retailers. As a result, **inventory optimization** is a top priority when it comes to investments in computer vision. 54% of retailers are currently investing or plan to invest in computer vision to optimize processes and reduce costs. [2]

Just as manufacturing and logistic companies have developed vision-based solutions in their warehouses, similar models are being piloted in stores across the U.S. to help keep shelves fully stocked. Cross referencing this inventory data with ERP systems allows organizations to quickly identify any discrepancies and analyze trends in detail, informing future purchasing decisions. This not only reduces instances of human error that can lead to over- or under-buying, but it also frees associates to spend more time serving customers directly.

Beyond supporting product availability, computer vision can **reduce shrinkage**. Today, 70% of stores report a shrink rate greater than 1%, resulting in estimated industry losses of more than $61 billion.[6] By using computer vision to identify high-value items and correlate pricing with Point-of-Sale (POS) machines, stores can reduce instances of theft via price swapping. Some retail grocers are even leveraging thermal cameras to reduce losses and improve food safety by ensuring perishables remain properly refrigerated.

Regulating store conditions in other ways can also **improve customer experiences**. Alerts about product spills, long checkout lines or curbside pickups can direct staff to areas of need — accelerating response times and enabling employees to remain focused on customer needs.

How retail organizations plan to invest in computer vision[2]

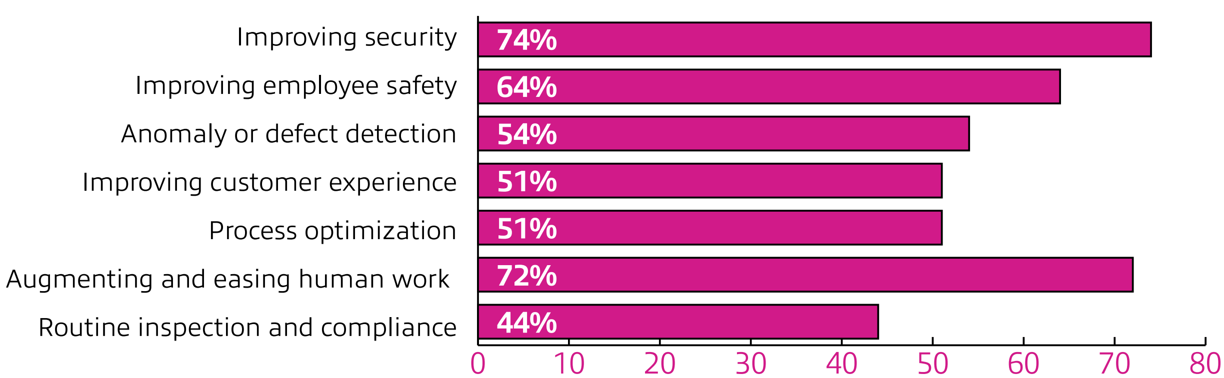

In 2020, Harris Teeter needed a smarter, safer way to protect the health of its essential workers. Insight deployed an IoT-based thermal camera solution to detect elevated temperatures and provide discreet alerts. The grocery chain is now building on this solution to monitor freezers and prevent food waste.

Use cases in healthcare

There is a wealth of compelling research on the use of computer vision to support medical diagnostics for conditions such as cancer or heart disease. But while advancements in this field will no doubt transform the future of preventive care, the potential harm caused by a misdiagnosis has far greater implications than most other use cases for this technology. Therefore, many additional safeguards are required for these models, including more rigorous training, narrower margins for error and a higher degree of human decision-making.

For this reason, many of today's healthcare providers are leveraging computer vision in other, lower-risk ways to optimize processes and improve patient care.

Optical character recognition, used to **automate document processing**, is among the most popular applications for computer vision in healthcare. Currently, administrative spending represents up to 30% of healthcare costs in the U.S., accounting for upwards of $1 trillion per year.[7] The World Health Organization estimates up to 50% of all medical errors in primary care are due to administrative or clerical issues.[8] Process automation with OCR helps to alleviate administrative burdens and reduce instances of human error while freeing care providers to spend more time with patients.

Computer vision also has valuable applications in **inventory management**, ensuring medical supplies are readily available when needed. Monitoring pharmaceuticals can even **improve security** by providing alerts of unauthorized or frequent access to high-risk medications.

In the wake of COVID-19, healthcare organizations have increasingly turned to computer vision to improve safety among patients and staff by helping to **prevent the spread of disease**. From infrared cameras used to detect elevated temperatures, to triggered reminders about routine handwashing, these solutions have significant potential to reduce healthcare-associated infections and improve quality of care.

How healthcare providers plan to invest in computer vision[2]

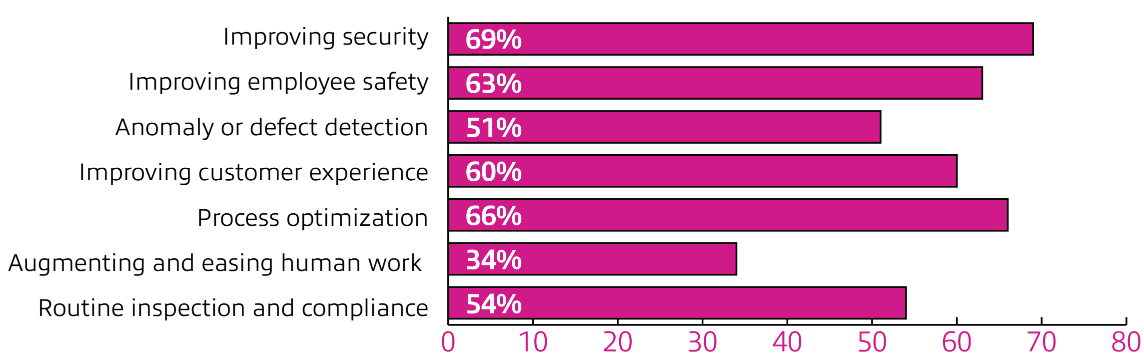

When providers at Steward Health Care wanted to optimize staffing and reduce patient length of stay at its hospitals, Insight developed a real-time analytics solution to help optimize scheduling and improve care coordination.

Selecting the right use case

When evaluating which of your current business challenges may be best suited for a computer vision solution, there are several things to consider. While computer vision can be applied to a broad range of use cases, in order to maximize the benefits while minimizing potential harm, it must have clear business value with limited complexity and scope; good, labelable data; strong executive sponsorship and responsible use of AI.

For a project to be successful, all four of these requirements must be met. Any one element missing from this equation will lead to challenges with follow-through, delivery, results and even human outcomes. The following questions can help you identify the best use case for computer vision in your organization:

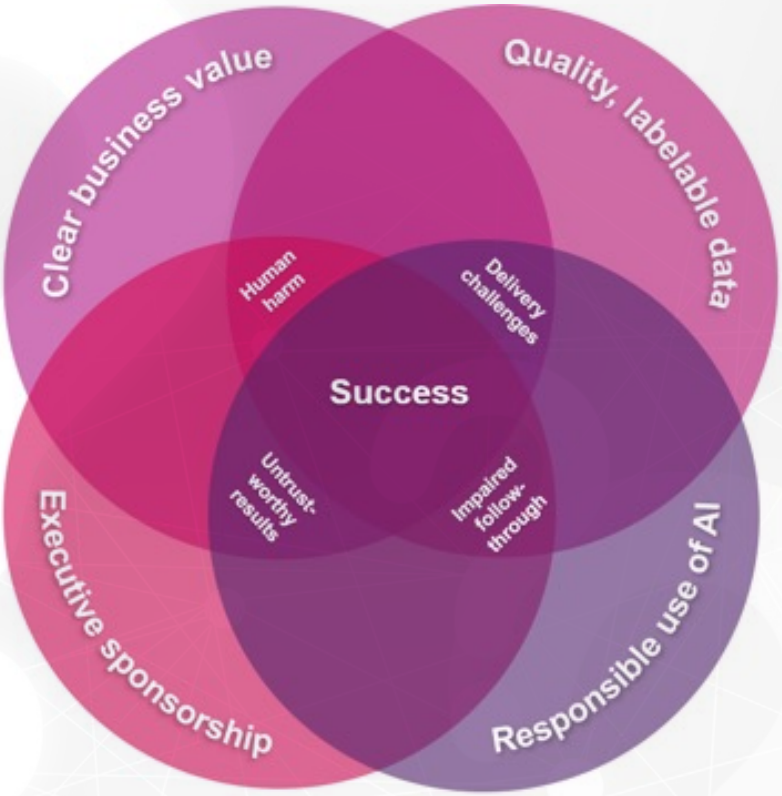

Is there value?

- Will computer vision deliver top-line or bottom-line improvements, disruption or buzz?
- Is the proposed project reasonably limited in scope and complexity?
- Could a different approach solve the problem more effectively?

Is there data?

- Do you currently have, or can you acquire, a large number of high-quality images?
- Is this data readily accessible?
- Is there a clear opportunity for labeling (using bounding boxes to identify target objects)?
- How easy or difficult will it be to train?

Is there sponsorship?

- Who will drive the initiative forward at the highest level?
- Do you have the support and momentum to overcome obstacles and corral stakeholders?

Is it responsible?

- What are the potential legal, social and ethical ramifications of the application?
- What measures can be taken to prioritize data privacy and security?
- How will you manage the solution to ensure accuracy and prevent misuse?

Choosing the right applications for computer vision will ensure a cadence of visible value that encourages future AI investments while incrementally growing expertise and expanding reusable approaches. Conversely, projects that fail to deliver value will both miss out on potential benefits related to the current use case and discourage future investment in AI.

Section 3: A Technical Overview

Once you've identified the most valuable use case for computer vision within your business, you can begin to evaluate the technical elements that will be needed to support your solution.

Hardware components

The core hardware involved in computer vision consists of edge compute devices and cameras. While cameras are used to capture imagery, edge computing is used to acquire, process and analyze the video frames and execute the vision models.

Depending on the use case, the required specifications of the edge device and cameras will vary. A fast-moving product assembly line, for example, will require a camera with a higher frame rate ("Frames Per Second" or FPS). This will also require your edge compute device to have a higher-end processor, FPGA card, graphics card, etc., in addition to larger memory and storage drives. On the other hand, there are some use cases which can be deployed using a low-power processor, 8GB RAM, 64GB storage and integrated GPU. A single camera used to monitor safety around heavy machinery, for example, might not require a high-end CPU since this type of application can often be executed by processing frames at a lower FPS rate using a smaller model.

If the cameras currently installed in your organization have good hardware specifications and support for the most common protocols, then new cameras won't be required. These existing cameras can be leveraged to implement computer vision, helping to reduce upfront procurement and installation costs. In some instances, even mobile devices currently used by staff can be repurposed as part of a computer vision solution.

Edge computing

Sensors, cameras, robots and other devices generate a great deal of data which must be processed and stored. But not all data belongs in the cloud. Processing data at the edge — closer to where the data is generated — provides faster access to data-driven insights and enables quicker decision-making.

When developing a computer vision solution, most data should be **processed at the edge**. Only the inferences or insights from that data need to be sent to the cloud to enable data visualization across multiple locations or deployments. This improves or reduces operational costs and provides low-latency responses to events.

There are many robust tools available to simplify the process of operationalizing and managing data at the edge. Depending on your current infrastructure, frameworks and services such as Google Anthos, Azure® Stack, AWS® Outposts and others can be leveraged to manage clusters and the workloads deployed on them. They can also help in scaling the solution across multiple use cases and deployments.

Of course, security is a notable concern with any technology implementation, and edge computing is no different. When it comes to securing an edge solution, you will need to consider physical security of the device, security of the data in transit and at rest, and security of the software from the ground up — including OS, frameworks and the solution itself.

Cloud considerations

While edge computing often serves as the backbone of computer vision, the **cloud remains integral** to the solution architecture.

Cloud computing complements the edge by:

- Helping to scale and manage the edge solution across multiple deployment sites
- Providing a common, scalable security framework for devices along with large scale ingestion and processing pipeline
- Making it simple to view data across multiple deployments and uncover efficiencies and inefficiencies across deployments and use cases
- Providing services to store huge amounts of structured and unstructured data that will be generated across multiple edge deployments

Working with the cloud can present some challenges as each cloud provider offers a discrete list of services that will tie your solution to the cloud provider. Selecting the right services and developing a solution with a good technical architecture is key. Security and device onboarding, for example, is not something that should be built from scratch. Leveraging existing security and device onboarding services provided by your cloud providers will mitigate some security risks.

However, all other components of the solution (including processing, storage, DevOps, MLOps, application and backend services) should be built in a cloud-agnostic manner to maximize flexibility in moving those components across cloud providers.

Custom vision models

An AI model is the mathematical result when a specific algorithm is trained on a historical data set. That model can then be used to compare new inputs to the historical data in a way that enables better decisioning or automation. A **vision-based model**, for example, can be trained on historical data to analyze products that meet the highest quality standards. New products off the line can then be analyzed by that model to identify anomalies, provide alerts and predict the need for equipment maintenance.

There are many, many algorithms that can be used to build a computer vision solution. Choosing the right algorithm is a matter of experimentation. By trying multiple algorithms data scientists can identify the one that has the best combination of accuracy, performance, maintainability and explainability.

It is important to note that no model is 100% accurate. In many cases an accuracy in the mid-80% range is considered "good" and even the best models rarely exceed the mid-90% range. It is also important to understand the impact of inaccurate results, both false positive and false negatives. Model outputs can be tuned to minimize (but not eliminate) one or the other. Organizations will need to evaluate the human, fiscal and business impacts of errors, and correctly design business processes to handle both correct and incorrect model outputs.

Additionally, because computer vision, like all AI, is based on historical data, unprecedented events — also called "black swan" events — can completely invalidate models. Even normal changes can result in AI models falling out of tune and being unable to provide accurate decision support. As a result, organizations must ensure that all AI models are automatically tested for accuracy and retrained if they fail to meet accuracy standards. Business processes must also be designed for the rare but inevitable times when models cannot be retrained to meet accuracy requirements and AI decision support needs to be removed from the process.

Sample computer vision architecture

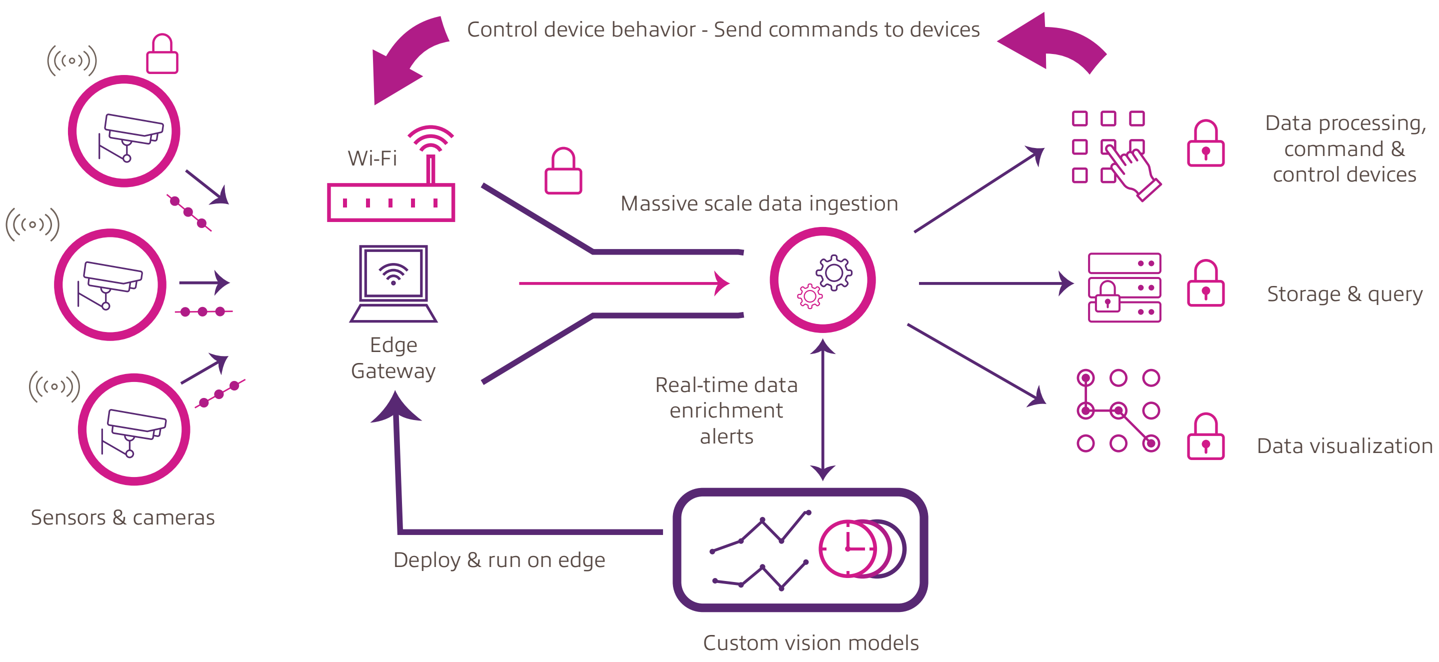

Section 4: Overcoming Potential Challenges

More than one-half (59%) of organizations report that they have high confidence in their in-house ability to implement, operationalize and manage AI-enabled technology.[2] However, overcoming the complexities involved in developing and maintaining solutions like computer vision requires more than a few skilled personnel.

Evaluating maturity

Technical and process maturity are essential to high throughput and capitalization of expertise. Unsupported by data engineers or MLOps engineers, data scientists may spend a great deal of time executing foundational tasks — and doing so poorly — resulting in significant blockers to success. Technical teams will also suffer without a ready pipeline of high-value, high-velocity opportunities from the business and a way to prioritize them. And none of this effort will bear fruit if the AI results are not properly integrated back into existing business processes and systems.

To maximize and accelerate the value of a computer vision project, begin by measuring your organization's maturity against every aspect of the AI lifecycle, as well as evaluating overarching organizational support.

OVERALL ORGANIZATION: Does every teammate have clarity on the overall vision, direction and goals of the project? Can your organization hire, grow and retain talent as needed? Is automation effectively leveraged to streamline processes and support ongoing delivery?

- **MODEL EVALUATION:** How will you confirm the model against business, compliance and technical requirements?
- **MODEL DEPLOYMENT:** How will the validated model be deployed to desired environments and endpoints?
- **MODEL MONITORING:** Are there plans or systems in place to ensure the model remains reliable and trustworthy over time? How quickly and effectively can you respond if a model becomes obsolete?
- **BUSINESS INTEGRATION:** How will you validate the final output to ensure it fulfills the intended business need? How will existing processes change and how will staff be trained on these changes? How will you capture changes to business process over time?

- **BUSINESS UNDERSTANDING:** Can you identify, evaluate and prioritize the most effective opportunities to leverage computer vision within the business?
- **DATA UNDERSTANDING:** What processes and personnel are in place for collecting, securing and understanding data for model development?
- **DATA PREPARATION:** What processes and personnel are in place for cleansing, combining and featurizing data for model development?
- **MODEL TRAINING:** Do you have the expertise and resources needed to experiment with available data to produce an effective model?

Evaluating and preparing data

Simply stated, high-quality data and labeling lead to high-quality AI, while bad data and bad data labeling lead to low-quality AI.

But there is no singular definition of "quality data." Data quality must be considered on a case-by-case basis depending on the application, accuracy requirements and risk profile for each model. Your data team will need to test and establish specific criteria for the types of images that deliver the best results.

A mature data science laboratory environment — including a data fabric, data lake and data catalog — will enable teams to effectively acquire, disseminate, collaborate on and maintain data throughout the AI lifecycle. There is no one right tactic or toolset to establish this foundation. Many resources are available through both open-source platforms and cloud providers, however, any additions to your data environment should be made strategically. Avoid the temptation of a "quick fix" approach to prevent long-term challenges associated with cost, complexity and scalability. Based on the tools and processes you choose, your organization will need to develop its own best practices to enable knowledge transfer and code reuse.

Collecting, cleansing, featurizing and labeling data is a time-consuming process. As your organization grows, it will be important to identify and reduce duplicated or misspent effort recollecting, rewrangling and reevaluating data.

Evaluating skills and partnerships

Despite confidence in in-house skills, most organizations report a high likelihood to leverage external consultants or system integrators to support their AI initiatives.[2]

90% of organizations are likely to seek external support for project strategy or planning and **87%** are likely to seek support for implementing computer vision.[2]

This is because, when it comes to AI transformation, nothing takes the place of experience.

Any organization that chooses to grow their experience organically is bound to make the same missteps that other organizations have made on their way to maturity. Hiring experience, either through full-time staff or an outside firm, can help your organization avoid common pitfalls and accelerate time to value.

When evaluating outside resources, seek out partners who will not just build a point solution, but who will also transfer real, complete organizational understanding.

Building on responsible AI

A crucial factor in the adoption of computer vision — or any AI solution — is the role of responsible AI.

Like all AI, while total accuracy is always the goal, computer vision systems are only as good as their training. Because of the risks of "black swan" anomalies and circumstantial changes, as well as the implications for privacy and security, even the most basic models must be developed and managed responsibly.

> Responsible AI provides a general set of best practices, frameworks and tools for promoting transparency, accountability and the ethical use of AI solutions.

6 keys to responsible AI

These guidelines will enable your organization to leverage computer vision to solve complex challenges without sacrificing core values such as privacy, security and human empathy.

1

Ensure computer vision is the right tool for the job.

Consider what is driving your initiative. Could the same outcome be achieved using a "lower tech" solution? Be sure to evaluate both the benefits and risks associated with your particular application.

2

Establish robust governance.

Your governance team should be made up of both technical and non-technical roles — with some representation from a statistics background. This group will need to establish a system while ensuring solutions are routinely reevaluated and performing as expected over time.

3

Commit to appropriate and secure use of data.

Best practice with any data project is to gather and store only what is needed to accomplish a given task. Processing data on the edge can help to reduce data transmission; but, regardless of your system architecture, you'll need a robust strategy outlining your data storage, encryption and auditing processes.

4

Keep humans in the loop.

Computer vision is at its best when it is driving human decision-making — not replacing it. The more important the decision, the more involved a human should be. Models should also be structured to prevent over-reliance that occurs when humans interface with highly accurate systems.

5

Safeguard against bias.

Fundamentally biased data leads to fundamentally biased results. Decision-makers and developers must use the best of their human intelligence to safeguard against the inequitable design or application of any given model.

6

Prioritize explainablity.

In order to have value as a decision-making tool, users need to be able to trust the system and its outputs. Explainability and interpretability are key to ensuring stakeholders within and outside your organization are empowered to understand how and why your model has come to a given conclusion.

Section 5: Calculating Costs, Timelines and ROI

Making the business case

Securing the executive sponsorship required to drive a computer vision implementation begins with building a compelling business case.

This will inevitably come down to dollars or a near proxy for dollars. How much money will this earn or save? How much valuable employee time will this save? How much will this reduce purchasing, materials, scrap or waste.

Often these values are hard to estimate or measure. Even models that might have high organizational impact might have that impact spread too thin to be measurable — or might be in a risk-prevention case where success is measured by the absence of infrequent costs. For initial projects, focus on cases with clear and measurable value to justify investment. Once your organization is confident in the project value then you can begin to expand use cases to include models with harder-to-demonstrate Return on Investment (ROI).

Establishing proof of value

Understanding the initial costs, time to value and potential return on investment is critical when building the business case for computer vision. Although the timelines and ultimate value of a computer vision solution will vary greatly by application, company and scale, there are relative benchmarks which can be used to begin planning and advocating for your project.

72% of organizations currently investigating or implementing computer vision expect to see a return on investment within two to three years. However, those that have already implemented computer vision have a higher propensity to estimate ROI can be achieved within just one year.[2]

While ROI may be the long-term measure of success, proof of value can often be established much faster with a Minimum Viable Product (MVP), serving as the first point of validation in the AI lifecycle. With the right strategies in place, project value can be demonstrated in just a few months.

1. Begin with clear strategic and technical direction. Lay down initial projects using repeatable tools and techniques that can be leveraged in later projects.

2. Limit the scope and complexity of your MVP — don't try to boil the ocean. Instead, identify the minimum amount of work required to start driving value.

3. Expand from initial success. Each novel technology or pattern will dramatically increase the risk of cost and time overruns. So, prioritize features and MVPs that reuse as many existing components as possible, then extend capabilities incrementally to drive a cadence of visible value.

Estimating ROI

When evaluating the potential ROI of your computer vision project, consider the following factors.

TOP-LINE IMPROVEMENTS:

- Increased customers
- Increased purchases and product attach
- Improved pricing

BOTTOM-LINE IMPROVEMENTS:

- Reduced labor
- Reduced materials purchases
- Reduced scrap/waste
- Better materials purchasing costs
- Reduced downtime

Hardware costs

The upfront costs of implementing an edge-based computer vision solution comes down to the hardware components, including cameras, edge compute and networking infrastructure. Calculating the cost of this hardware depends on the use case, area covered, number of cameras and number of frames that need to be analyzed. As you tweak these variables, the cost of the hardware in the solution will vary.

When evaluating hardware options, it is worth noting that a larger edge device (with more compute) is not always necessary. There are frameworks available to leverage existing CPU and integrated GPU to run computer vision models, reducing the overall cost of the solution.

Development costs

The cost of hiring internal or external experts to develop, deploy and manage the computer vision model will represent the bulk of your investment. While this may serve as a stumbling block for some project stakeholders, it is important to recognize this is also the most critical factor in the success of your project — and therefore worth the investment. Additionally, when the solution is architected with scale in mind, the repeatability can help to justify these costs. Providing a foundation for other use cases and locations will help to spread the financial load and increase the return on investment.

Network requirements are often overlooked when considering developmental costs but should be factored into initial estimates. While most organizations typically assume their current setup is sufficient to support computer vision, the additional demands of cameras and edge processing often require some optimization.

Recurring costs

The recurring costs of a computer vision solution will depend on the architecture. Sometimes, third-party computer vision models are sold as a subscription business model. These costs can add up quickly based on the number of cameras and the number of models being used.

However, if the model is custom built and owned by your organization, then the only potential recurring costs will come from cloud service consumption and any managed endpoint monitoring services. If the solution is architected correctly, these costs should not be high.

Accelerating timelines

It is essential that computer vision initiatives focus on delivering clear, visible value to the organization as quickly as possible. Initial projects must deliver real ROI quickly to dispel doubt and justify ongoing and future investments in AI technology.

Beginning with a sufficient quantity of good data is essential for timely execution. Data collection can add months of additional time and effort to a project, so it is important to ensure this data is readily available. Data labeling can also be a time-consuming step, depending on the complexity of your solution and the level of experience on your team. Only after these elements are in place can your data science and operational teams build, deploy and integrate a computer vision model.

Working through these processes often takes significant collaboration across groups which is greatly facilitated by a strong executive sponsorship. Without this sponsorship, organizational friction, communication challenges and competing priorities will often stall or even kill a project.

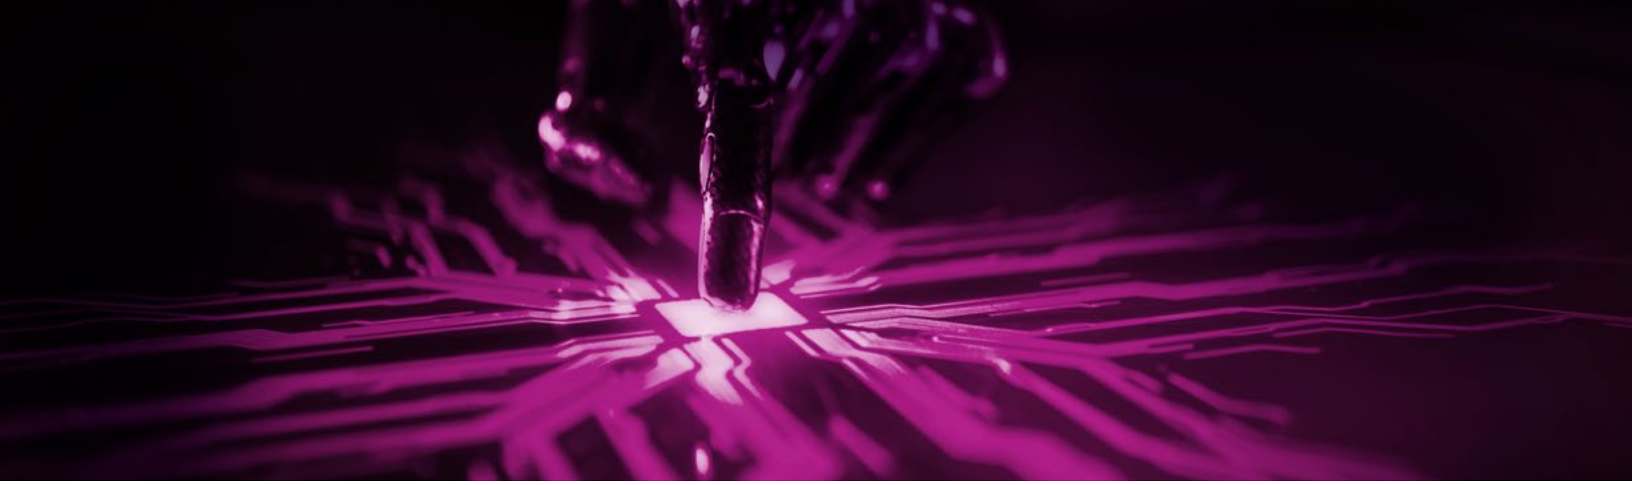

Section 6: Maximizing Long-Term Value

Successfully building and deploying a computer vision model is only the first phase of project implementation. There are several key considerations to maximize long-term value as you move from proof of concept to full production.

Documentation and sponsorship

Documenting the value of a solution before it is built is critical for a project to progress through the AI lifecycle. Without a clear business driver to secure sponsorship or propel operationalization, projects are doomed to fail. Beginning with a well-defined, well-documented business outcome is the first step to long-term success.

Because of the interdisciplinary nature of the AI lifecycle, initial projects often require high-level support. Ideally your executive sponsor should be in a position to have all the required business, data and operations stakeholders within their management structure. C-suite sponsorship is ideal, but VP-level sponsorship is also often very effective. Initiatives can work with director-level support but can become challenging if other director stakeholders are not invested in overall success.

Consider who within your organization has the propensity to provide the most value and momentum to your project. Work to establish sponsorship by leading with the benefits for your business, supported by as many relevant metrics as possible. Use the top- and bottom-line improvements along with other relevant details in Section 5 of this guide to make the case for ROI and prioritize quick proof of value.

The following steps can help to ensure the long-term success of your project:

1. PINPOINT VALUE.

Define a clear business outcome and the potential impact on the business. Use this to secure sponsorship and drive your project forward.

2. TAKE A PHASED APPROACH.

Your implementation plan should include:

a. Pilot deployment site/location
b. Initial roll out sites/locations
c. Post deployment device and solution monitoring
d. DevOps and ML Ops strategy
e. Device and solution support

3. WORK ITERATIVELY.

A computer vision solution is an iterative process and with each iteration, the system/solution will get better. Expect iterations and tweaks as part of the process.

Management and maintenance

Once your computer vison model has been deployed, you will need to take steps to maintain performance and ensure the system continues to deliver value to your business over time.

For this reason, device lifecycle management and monitoring are core requirements of any edge solution. To ensure visibility into the performance of the edge device as well the computer vision model, the right device and management solution must be selected.

Choosing the right edge device is important to enable the following functionality:

- Out-of-band management capabilities enable solutions to monitor and take actions on the device remotely.
- A Trusted Platform Module (TPM) can store cryptographic keys.
- Depending on the use case, secure code enclaves may be needed to protect code execution.

The value of MLOps

Building and operationalizing machine learning or AI models is an iterative process.

Static models may provide value for a finite period, but over time the data points will vary too significantly to maintain relevancy. As new data is collected and the model retrained, newer versions must continuously be deployed and tested. This process can quickly become complicated and will not scale if not automated.

Using existing frameworks and creating automation scripts, an MLOps pipeline can be created to simplify the process of training, testing, deploying and validating the model.

MLOps reduces challenges associated with manual processes, complexity and siloes between teams by:

- Enabling rapid development and experimentation with new models through continuous training
- Supporting automated testing and modularized pipeline components through continuous integration
- Improving collaboration and alignment between the development and production environments to drive continuous delivery
- Accelerating the deployment of newly trained models with automated triggering
- Providing ongoing feedback on performance based on live data, allowing models to be optimized and retrained over time

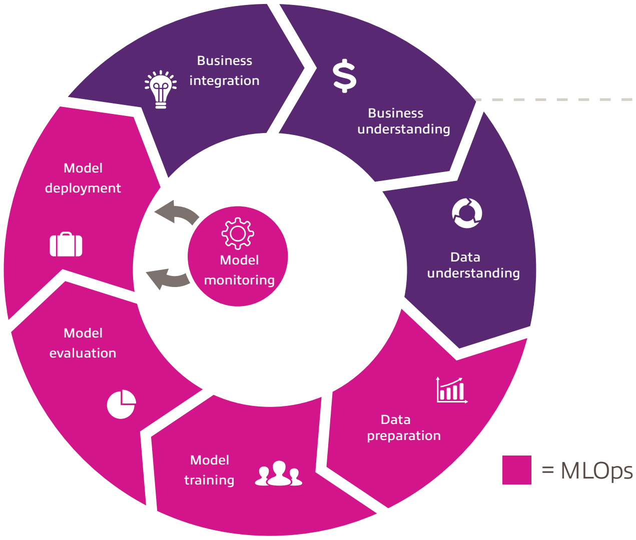

Discussions around MLOps should be included in the earliest stages of your computer vision project. Start by gathering all impacted teams — including data scientists and engineers, infrastructure and DevOps teams, software developers, business analysts, architects and IT leaders — to begin researching, developing and documenting a comprehensive MLOps strategy.

The value of change management

Computer vision only drives value if it is used to make better decisions faster and more effectively. Organizational Change Management (OCM) provides the framework to ensure that a computer vision solution is adopted and that the resulting decision support is successfully integrated into business processes. Since the implementation of computer vision is likely to disrupt existing processes, OCM is also critical for facilitating a smooth transition to new operating models.

The **five levers of OCM** will help to drive adoption, build resilience, support meaningful transformation and establish long-term agility.

Building an AI Center of Excellence

While establishing a Center of Excellence (CoE) is not strictly necessary, it is often the fastest way to grow AI competency and success within your organization.

Implementing AI successfully requires interdisciplinary expertise across business understanding, data systems, subject matter expertise, data science, DevOps, Agile, OCM and more. Bringing these skills together under a CoE fosters collaboration and reduces normal organizational friction, allowing the team to focus on growing the specific skills and practices needed to succeed with AI. Initially your CoE should be responsible for the end-to-end success of pilot initiatives. Later, as the broader organization becomes more AI-ready, the CoE should act as a training ground to transfer proven knowledge and expertise to teams across the organization.

Section 7: Conclusion

Today, more than 90% of organizations agree computer vision has the potential to drive value for their business.2 While there are many factors to consider when evaluating and planning to invest in these types of solutions, the benefits of successful implementation far outweigh the initial costs.

By following best practices to develop and operationalize solutions, organizations across industries have an opportunity to leverage computer vision to optimize processes, increase revenue, and improve both employee and customer experiences.

As more organizations begin to embrace computer vision as part of their strategic business model, those that get ahead of the adoption curve will be equipped to uncover new insights and gain a key competitive advantage as this technology continues to redefine the role of AI in the workplace.

To learn more, visit: insight.com/computer-vision.

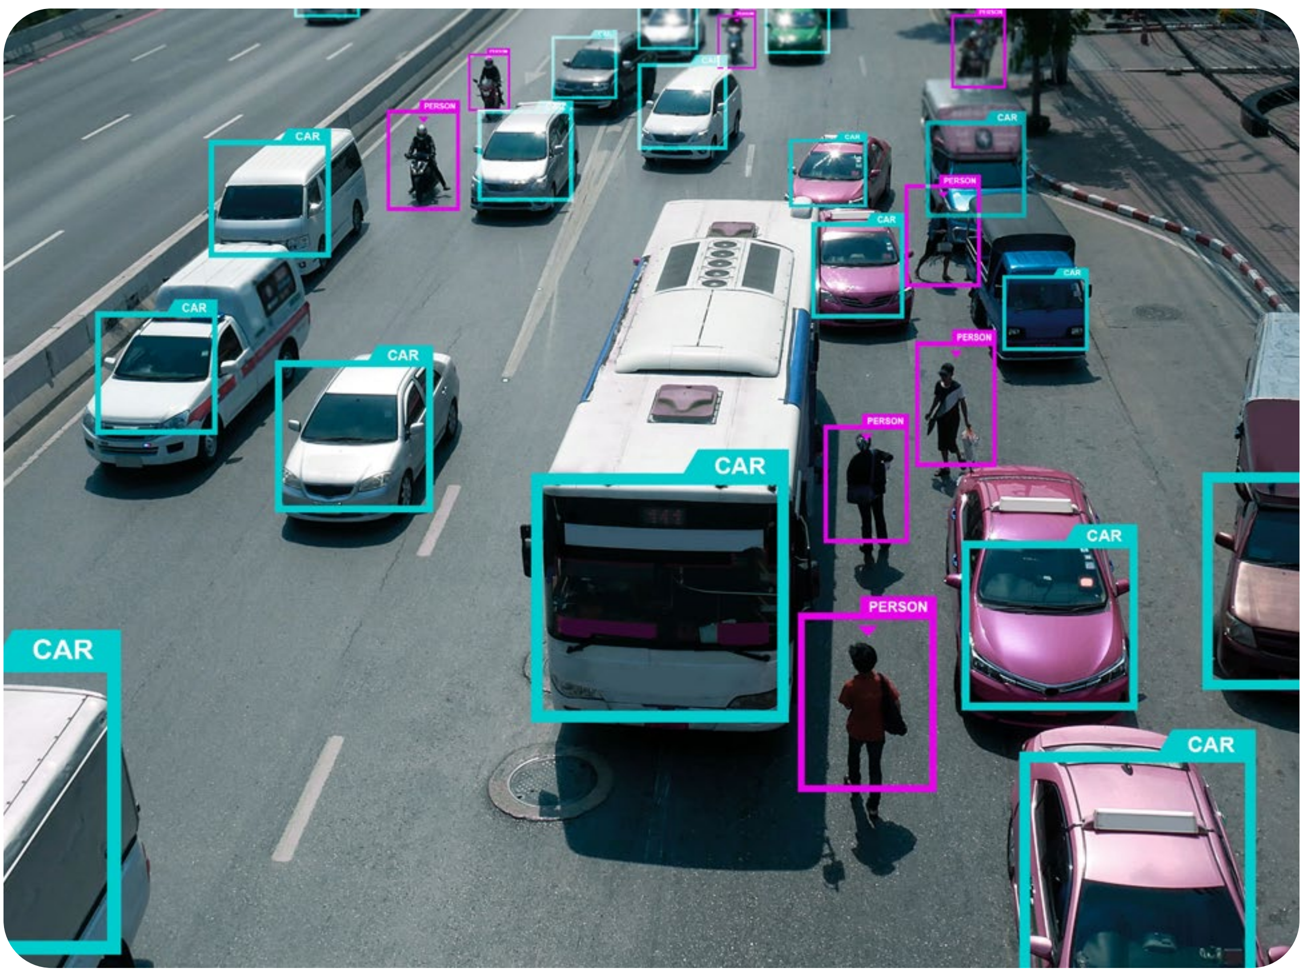

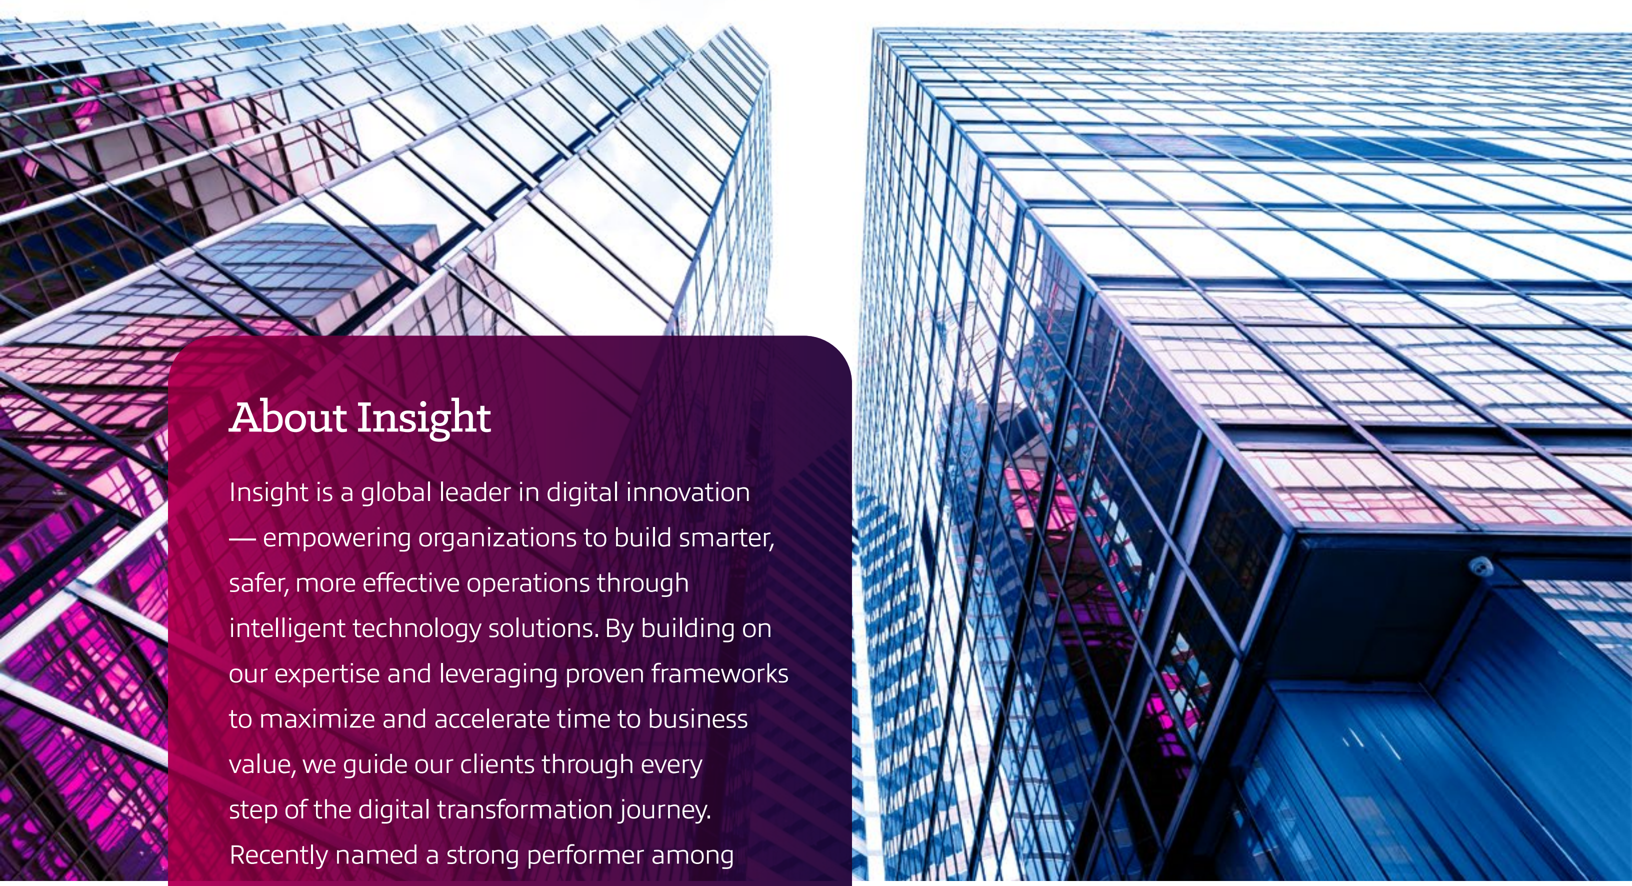

About Insight

Insight is a global leader in digital innovation — empowering organizations to build smarter, safer, more effective operations through intelligent technology solutions. By building on our expertise and leveraging proven frameworks to maximize and accelerate time to business value, we guide our clients through every step of the digital transformation journey. Recently named a strong performer among top computer vision consultancies by Forrester, Insight delivers the end-to-end expertise and support to help clients successfully adopt, operationalize and manage AI at scale.

[To learn more, visit: insight.com/computer-vision.](https://insight.com/computer-vision)

Sources

1 Gartner. (2020, May 11). Top Trends in Data and Analytics for 2020.
2 MarketPulse Research by IDG Research Services. (May 2021). Computer Vision: Adoption and Application. Commissioned by Insight.
3 Grand View Research. (September 2020). Computer Vision Market Size, Share & Trends Analysis Report.
4 Koen, S. (2015, Oct. 25). Safety Leadership: Neuroscience and Human Error Reduction. Safety and Health Magazine.
5 Liberty Mutual Group. (2018). Liberty Mutual Workplace Safety Index.
6 National Retail Federation. (2020, July 13). National Retail Security Survey.
7 Keating, E., Tollen, L., Weil, A. (2020, Feb. 20). How Administrative Spending Contributes to Excess U.S. Health Spending. Health Affairs.
8 World Health Organization. (2016). Administrative Errors: Technical Series on Safer Primary Care.