altamira

Revealing business value through image recognition technology

Altamira
www.altamira.ai
Email: **sales@altamira.ai**

altamira

Intro

Image recognition, a key part of computer vision, is a much-hyped field within AI. It relies on algorithms that enable machines to interpret and make sense of digital images. Virtually, image recognition trains computers to identify and analyse images, mimicking human vision but with a level of speed and accuracy that goes beyond what humans can achieve.

For businesses, this speed and accuracy in image analysis translates directly to faster decision-making, reduced operational costs, and the ability to deliver personalised customer experiences.

At the core of an image recognition system is its ability to understand and analyse a digital image. This process starts by converting the image into a format that a machine can work with, often by dividing the image into pixels and assessing these for specific patterns and features.

Machine learning algorithms, especially deep learning models like convolutional neural networks (CNNs), identify image patterns, making them essential for applications like object recognition and medical image analysis.

In recent years, image recognition has become much more common concept. Corporations worldwide have increased their investments in AI and machine learning by an average of 44% year over year, with image recognition among the top areas of focus. From improving image searches on digital platforms to advancing medical image analysis, the range of applications is endless. A good example is facial recognition, which allows systems to identify and verify people by analysing their facial features.

Overall, progress in image recognition technology, powered by machine learning and deep learning models, has paved the way for a future where machines can interpret and engage with the visual world—a task that was once exclusive to human perception.

In this whitepaper, we will:

- Explore how image recognition technology can lead to cost savings, improved operational performance, and enhanced customer engagement.
- Analyse market trends that are shaping the future of image recognition technology.
- Discover a range of use cases across various industries—including healthcare, retail, manufacturing, and security—demonstrating how organisations have successfully integrated image recognition to solve real-world challenges and achieve strategic goals.

altamira

The science behind image recognition

Image recognition is the method of identifying objects, locations, people, text, and actions in digital images. This is made possible through artificial intelligence and machine learning, which enable the system to learn from examples and make predictions or decisions without needing detailed programming for each task. As a branch of computer vision, image recognition focuses on detecting picture patterns so that software can "see" and understand the content of a digital image, much like how a human would.

The use of image recognition in healthcare for diagnostics and imaging analysis is on the rise, with an estimated 15% to 20% improvement in diagnostic accuracy reported when using AI-enhanced imaging tools.

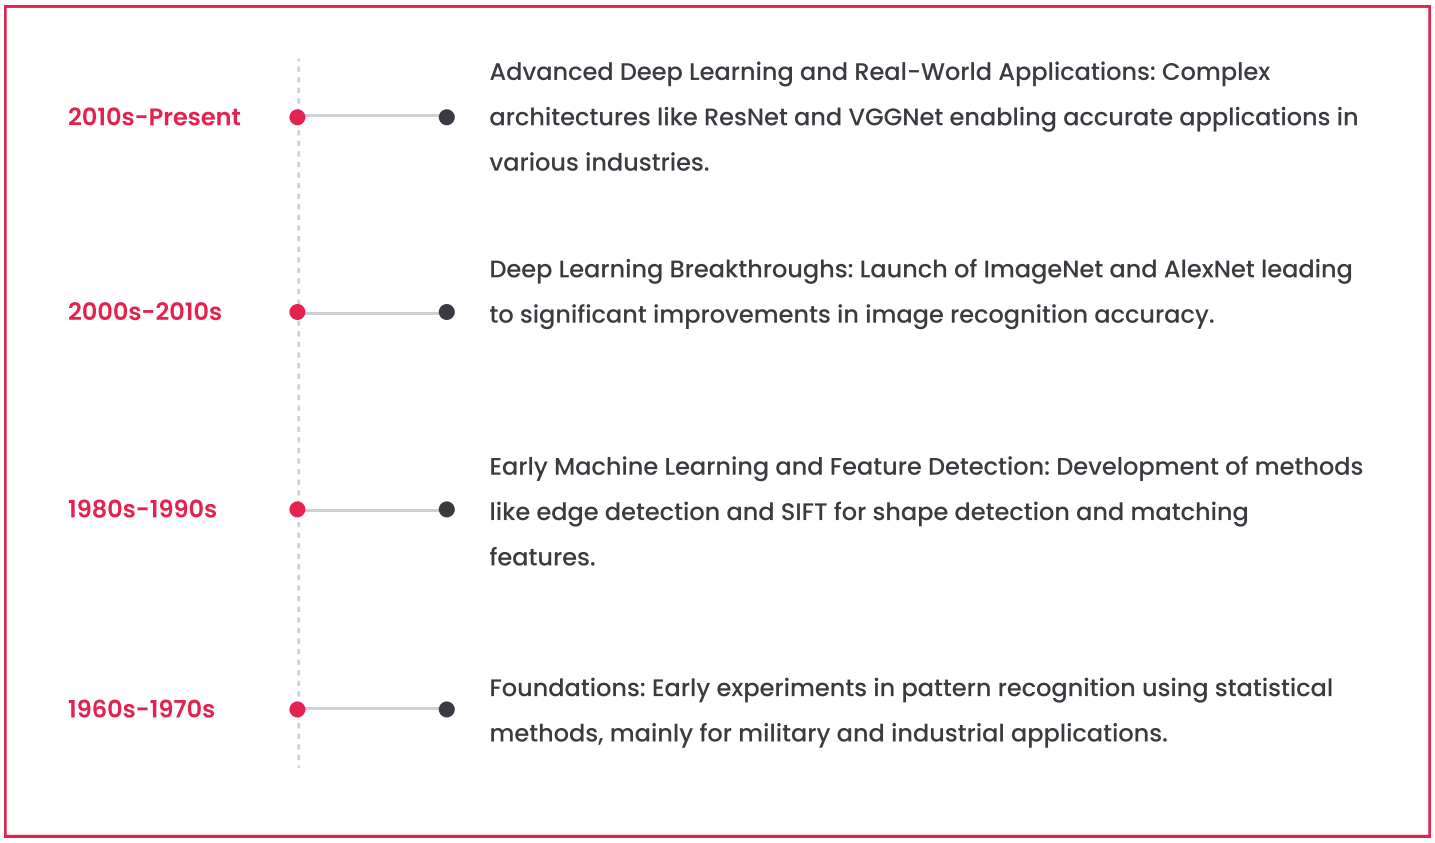

A timeline of developments in image recognition technology

altamira

Key components of image recognition process

1. Gathering data

This is the first stage of the image recognition process, where different objects, scenes, and patterns that need to be identified are collected. This data serves as the basis for building the image recognition model.

2. Preparing data

After gathering the data, it needs to be cleaned (removing noise or errors in data) and organised. This step might include fixing errors and standardising pixel values to make sure the data is prepared for training the model.

3. Labelling data

In this phase, the objects within the images are marked with labels to help the system identify them. These labelled images are then used to train the model to recognise similar objects in future images.

4. Machine learning workflow

This final step consists of assembling the labelled images, extracting important features, and developing a model that can classify these features to identify new objects.

Systems like Google's Inception v3 and Facebook's ResNet have recorded top-1 accuracy rates on benchmark datasets like ImageNet, with scores of around 78% to 85%.

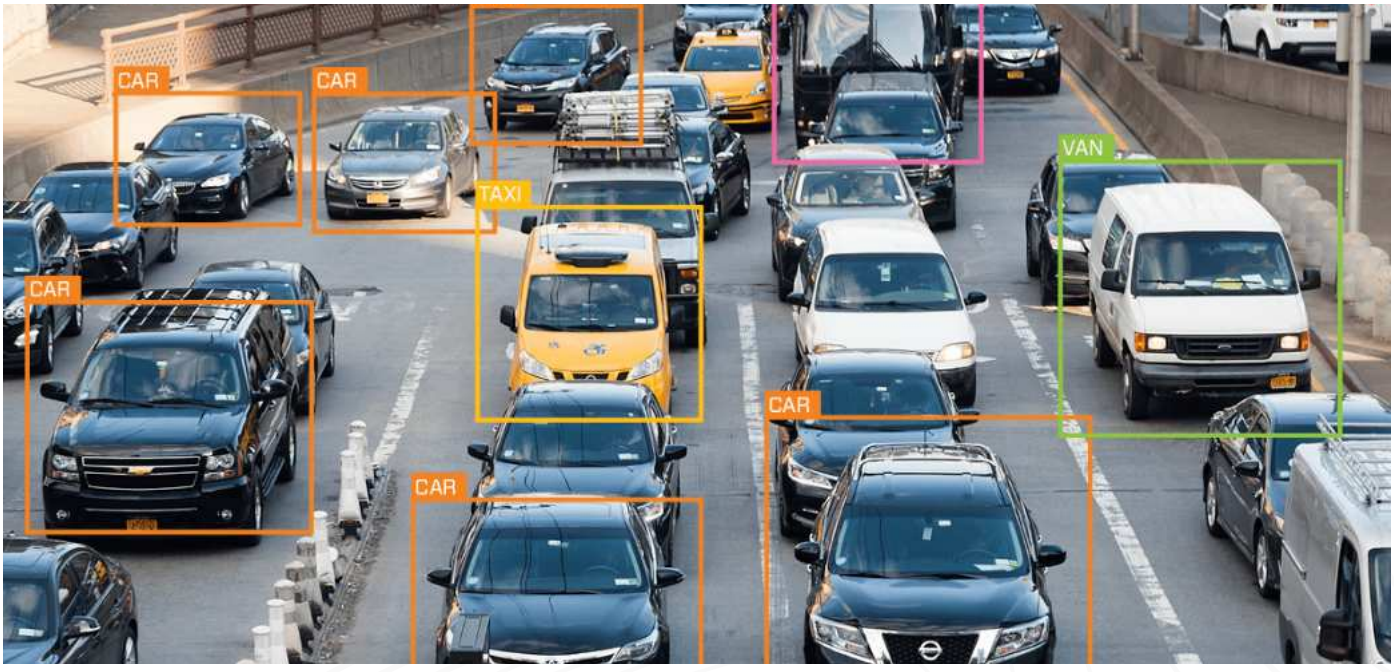

A key part of the technology is image sensors, which record various characteristics like brightness and strength. These characteristics are then processed to extract specific details from the images.

The image recognition process consists of the following steps:

- Designing a Convolutional Neural Network (CNN) structure
- Training the model using sample images
- Using the knowledge gained to analyse new and unfamiliar images during recognition.

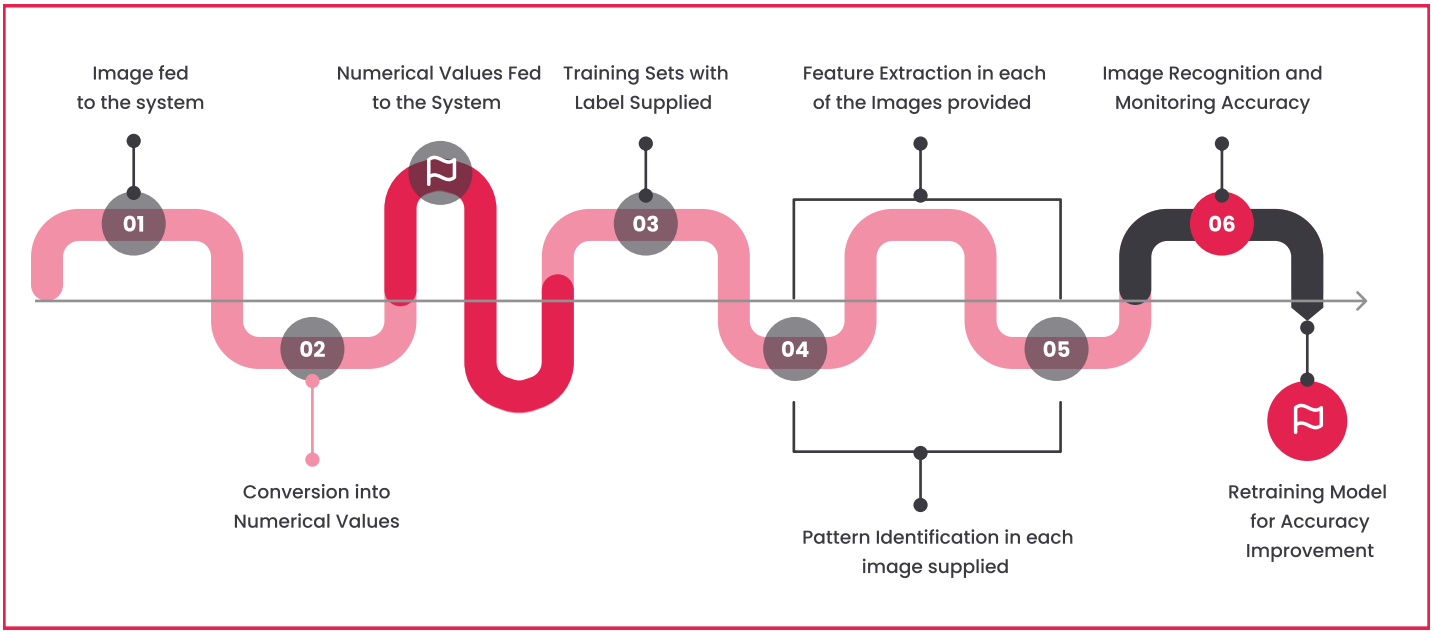

To train an image recognition system, you need a large set of labelled images. As a result, machine learning algorithms learn how to identify and categorise different features and objects in the images. Although deep learning works for image recognition, it still relies on having a wide range of data to create strong and adaptable models.

More than 60% of mobile applications now incorporate some form of image recognition technology.

Image recognition market overview

The global market for image recognition was valued at approximately USD 46.7 billion in 2024 and is expected to grow to USD 98 billion in 2029, representing a compound annual growth rate (CAGR) of 17.1%.

In 2023, North America held the largest market share at 35.23%. Specifically, the image recognition market in the US is anticipated to see significant growth, with a projected value of USD 19,193.3 million by 2032, mainly due to the increasing popularity of AI-driven technologies.

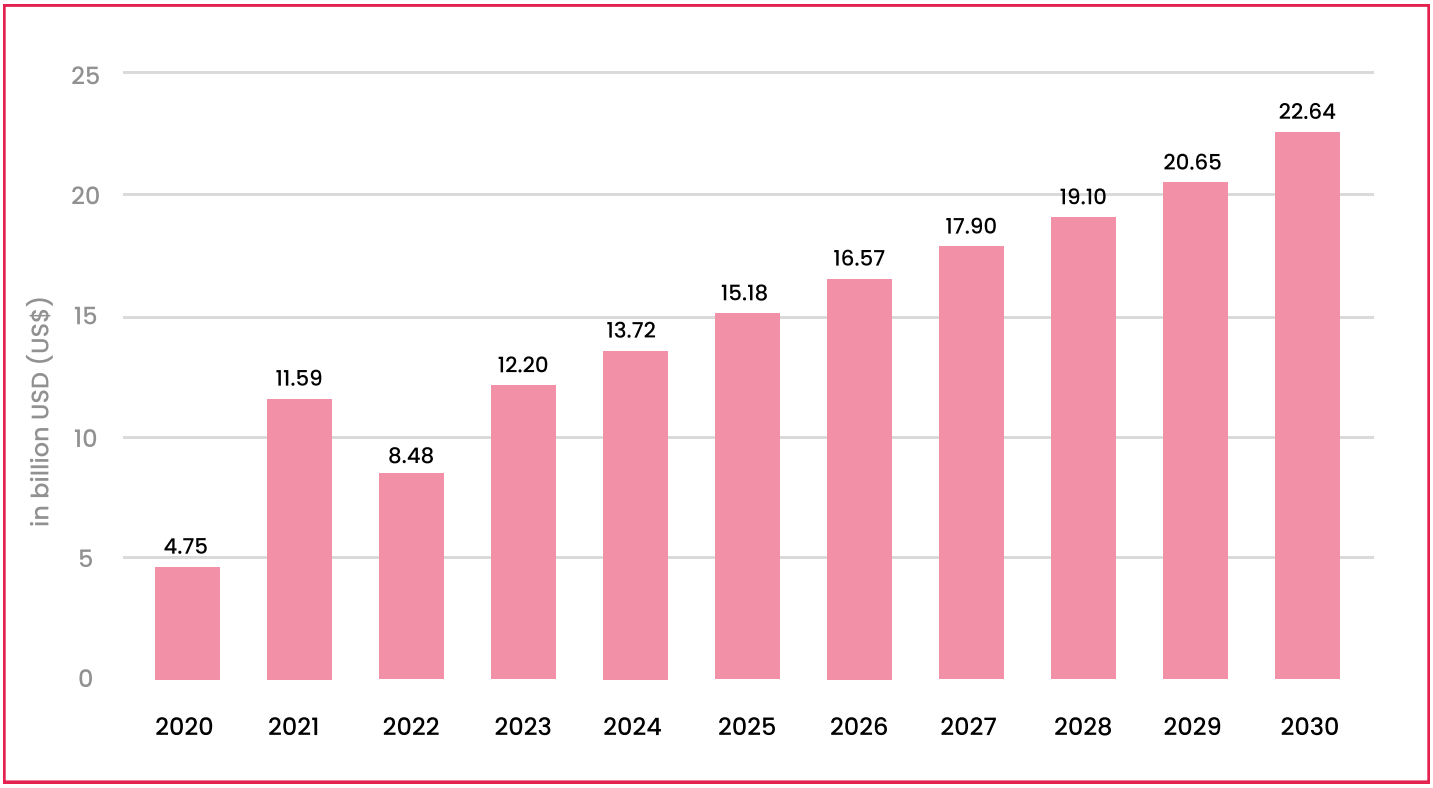

Market size

Notes: Data was converted from local currencies using average exchange rates of the respective year.
Most recent update: Mar 2024
Source: Statista Market Insights

The demand for image recognition technology is being fuelled by a growing preference for high-bandwidth data services and improved machine learning capabilities. Various sectors, including retail, media and entertainment, IT and telecommunications, banking and financial services, proactively adopt advanced technologies, contributing to market expansion.

> *As Satya Nadella, CEO of Microsoft, highlighted:* ***"Artificial intelligence is the ultimate breakthrough in human capability. It will create an intelligent cloud and intelligent edge, and image recognition is a key part of this transformation."***

Moreover, advancements in image recognition have made it possible to use images taken by smartphones to link offline materials—like brochures and magazines—with digital content, such as promotional videos, augmented reality experiences, and product details.

The market for artificial intelligence in image recognition and computer vision is currently experiencing modest growth, impacted by factors like the slow uptake of advanced technologies and a lack of consumer awareness. Still, the growing ease of using online health services is likely to help the market expand globally.

Market drivers

In the global image recognition segment of the computer vision market, there's a clear trend toward using deep learning techniques, as these models allow for enhanced accuracy in tasks such as facial recognition and object detection.

Now, we can also see a rising demand for real-time image recognition solutions, especially in retail and security, which is driving the development of new products. This trend is likely to keep on, driving market growth and creating opportunities for companies to expand their services and meet different industry needs.

Consumer preferences

As computer vision technology becomes more widespread, people are looking for shopping experiences that feel more personalised. That's why we are witnessing the increased use of image recognition for virtual try-ons and tailored product suggestions.

Also, the growth of social media and influencer marketing is driving people to look for quick ways to buy items they see online, speeding up the use of image recognition technology on these platforms.

Benefits for businesses

> ***"We are moving from a world where we have to understand computers to a world where computers will understand us. Image recognition is a big part of that shift."*** *- Jeff Bezos, Founder of Amazon.*

Computer vision

Computer vision is a branch of computer science dedicated to helping computers recognise and understand objects and people in images and videos. Like other AI technologies, computer vision carries out tasks that replicate human abilities, focusing on how we perceive our environment and make sense of what we see. This technology can operate either in the cloud or locally on devices.

> ***The automotive sector is a major adopter of computer vision, with around 90% of new vehicles expected to have some form of computer vision technology by 2025 for applications like autonomous driving and advanced driver assistance systems (ADAS).***

How an image is analysed with computer vision

You should also consider explainability, which relates to the quality of decision-making. The criteria for explainability include:

- A sensing device, typically a camera, but sometimes a video camera, medical imaging equipment, or another device, captures the image for analysis.
- The image is then transmitted to an interpreting device, which employs pattern recognition to analyse the image. It compares the patterns it detects with a database of known patterns to see if there are any matches. These patterns can be general, like the shape of a specific object, or more detailed, such as facial features.
- When a user requests specific information about the image, the interpreting device delivers the relevant details based on its analysis.

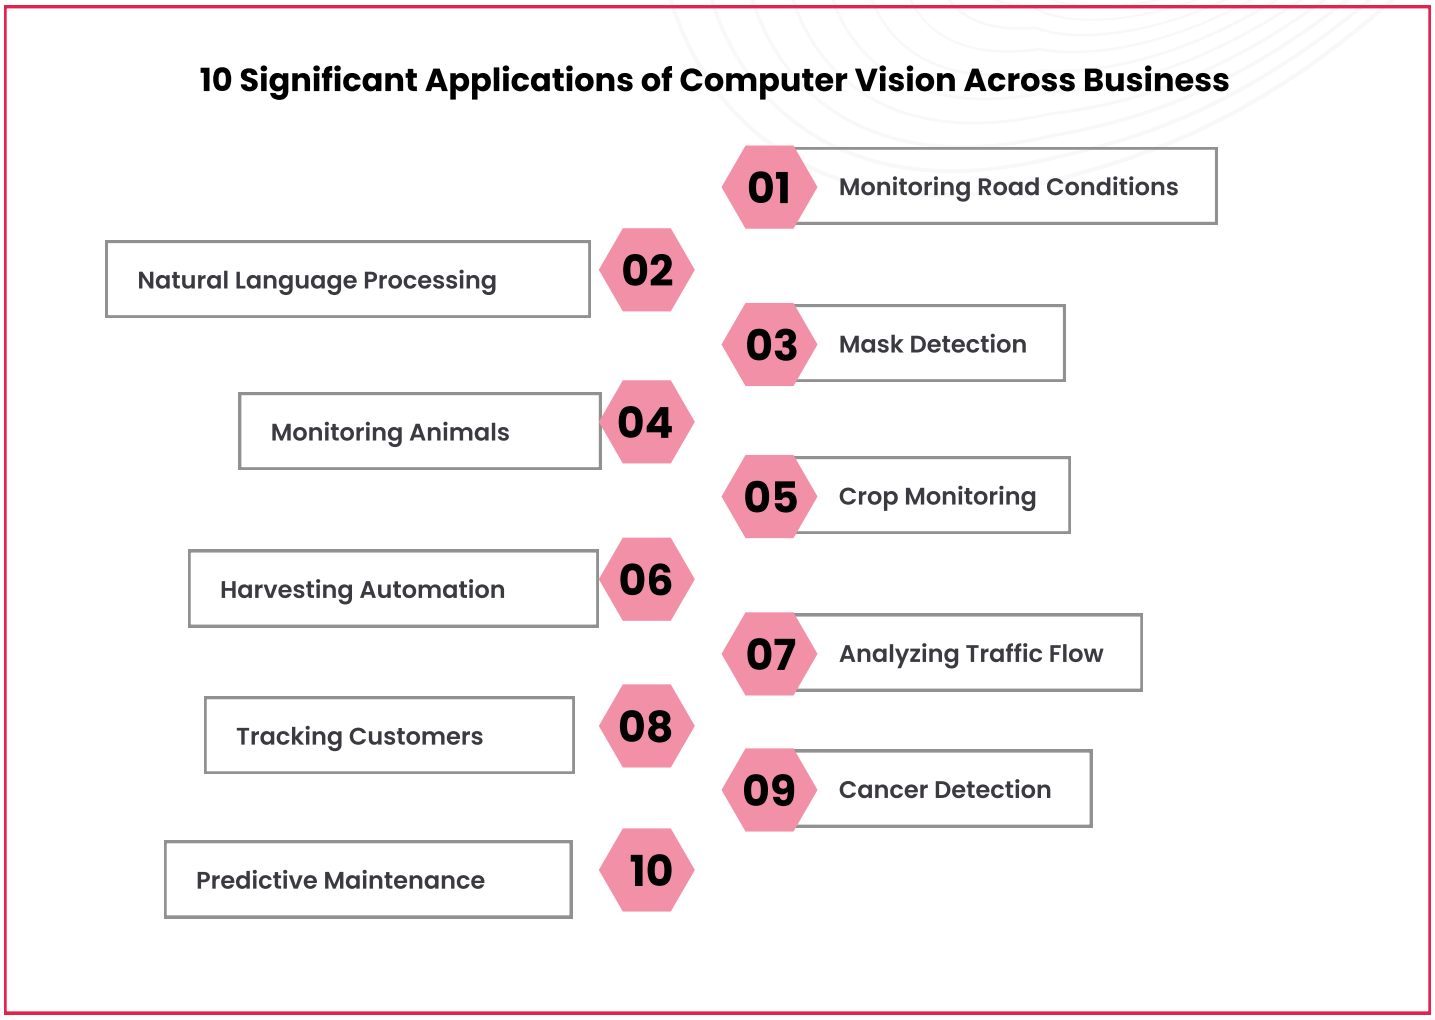

Unlike human observation, which can be error-prone, computer vision uses neural networks to analyse visual data, leading to highly accurate processing. This technology can spot flaws across multiple data sets simultaneously, which is beneficial for inventory management and speeds up object recognition, ultimately improving the quality of visual outputs.

Beyond quality control, machine vision is becoming more valuable for tasks like identification, inspection, and guidance in various industries. Here are a few examples:

altamira

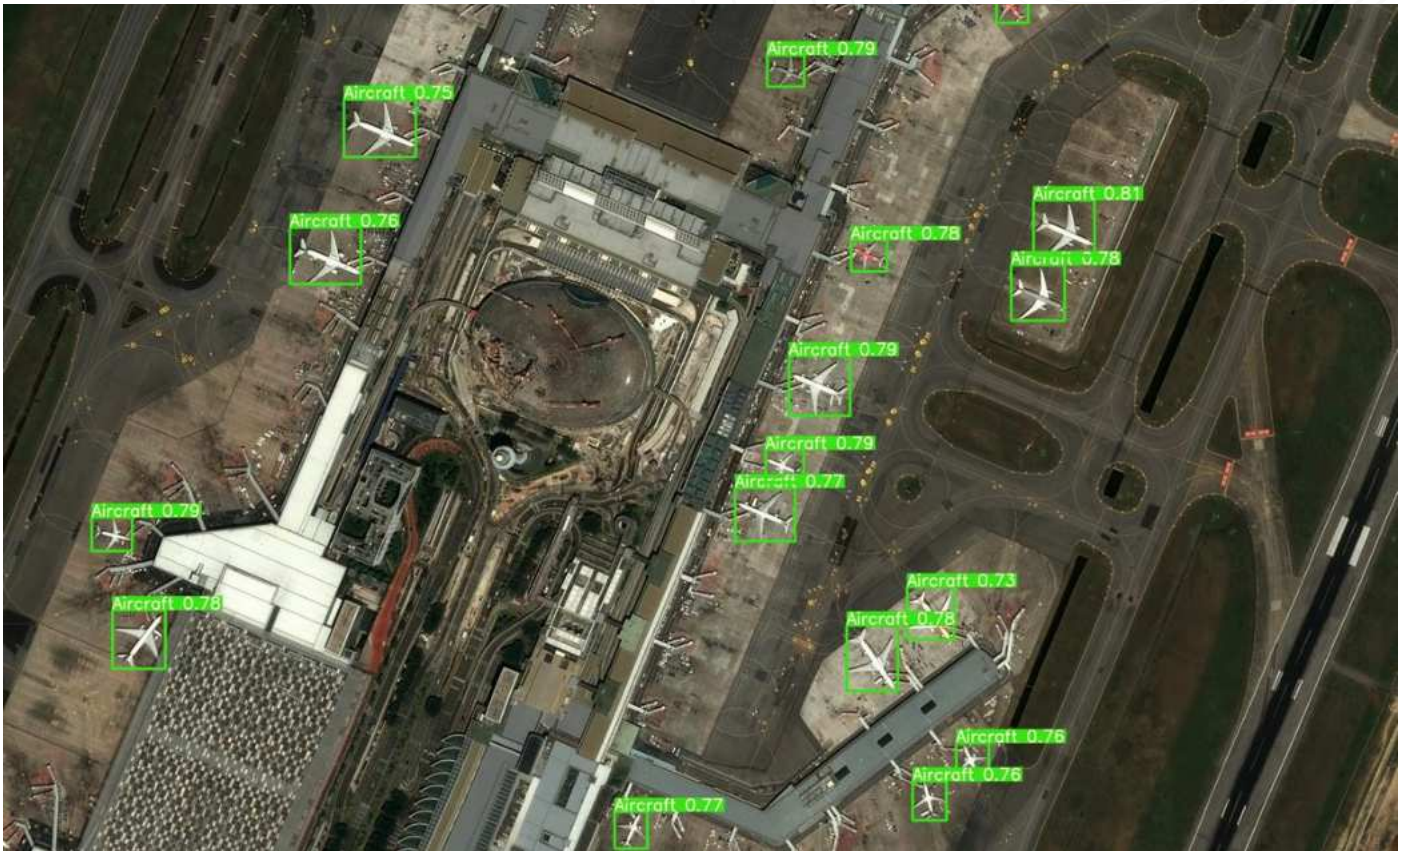

- **Production defects:** Computer vision spots defective products and pinpoints where issues arise on a production line, enabling targeted corrective measures.

- **Agriculture:** Computer vision assists harvesters by locating grapes on vines, enabling robotic systems to pick them without causing damage. It's also used in farm machinery to monitor crop health and spot plant diseases.

- **Inventory control:** Machine vision reads barcodes and labels on components and products, facilitating inventory management and ensuring the correct parts are added during assembly. Robots in warehouses also rely on machine vision for efficient bin-picking.

Retailers like Walmart are using computer vision to track stock levels and automate restocking, reducing out-of-stock items by 30%.

- **Production tracking:** In industries with strict regulations, such as pharmaceuticals, machine vision simplifies tracking ingredients, product serial numbers, and expiration dates.

- **Measurements and calibration:** Machine vision automates processes like measuring spark plug gaps to ensure they meet specifications or identifying gauges needing calibration, improving efficiency.

- **Safety:** On construction sites or in food supply monitoring, machine vision enhances safety by providing real-time monitoring and analysis.

Amazon Go stores use computer vision to allow customers to pick items and walk out without traditional checkouts, tracking purchases via cameras and sensors. This has reduced checkout times by more than 50%.

altamira

Object recognition

Object recognition is a computer vision technique that focuses on identifying different objects within images, either to generate captions or to locate specific items for analysis. To do this, computer vision models are trained to recognise specific objects.

There are two primary categories of object recognition models: zero-shot and fine-tuned models.

Zero-shot models are large, pre-trained systems capable of identifying a range of objects in images without any further training. While they tend to perform well with common categories, such as people, they may face challenges when it comes to less common categories, like product defects.

In contrast, fine-tuned models are smaller and specifically trained to recognise a limited selection of classes. Unlike zero-shot models, which target various classes, fine-tuned models focus on items. For example, a fine-tuned model might detect scratches on a specific car component or differentiate between various types of packages on a conveyor belt.

Advanced models can process images in real-time, achieving detection speeds of over 30 FPS (frames per second), making them suitable for applications in robotics and surveillance.

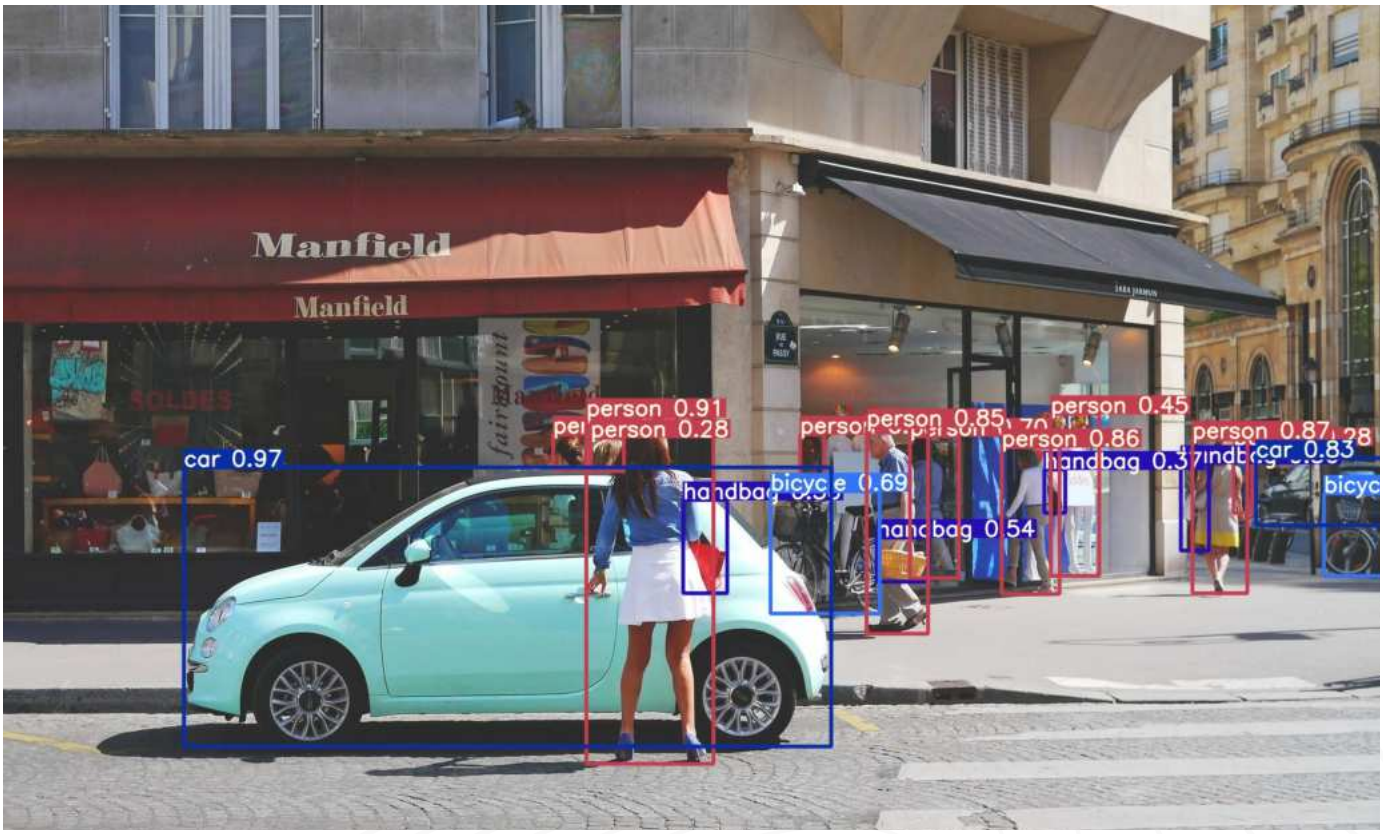

Object detection technology has many practical applications across different fields:

- **Autonomous vehicles:** Object detection is a key part of building self-driving cars, helping them spot and track pedestrians, other vehicles, and traffic signs.
- **Surveillance:** This technology helps security and surveillance systems detect and track people, vehicles, and possibly suspicious items, using camera feeds.
- **Retail:** In retail settings, object detection is used to monitor foot traffic, observe customer behaviour, and keep track of inventory levels.
- **Healthcare:** In medical imaging, object detection can identify and locate abnormalities like tumours in X-rays, CT scans, and MRIs. Studies show that AI can enhance diagnostic accuracy in radiology by up to 15%.
- **Robotics:** Object detection enables robots to recognise and handle objects in real-world environments, such as manufacturing plants and warehouses.
- **Sports:** In sports, object detection can track players' movements, identify rule violations, and gather game statistics in sports like basketball and soccer.

In a nutshell, object detection technology is scalable, letting businesses handle and analyse huge amounts of visual data. By working with thousands or even millions of images or videos, businesses can discover insights that they couldn't access before.

altamira

Feature detection

Feature detection is a key aspect of computer vision and image processing that involves identifying and extracting features or patterns from digital images or video frames. These features can be edges, corners, blobs, lines, textures, or any other unique traits used to describe and differentiate objects or areas within an image.

By recognising and analysing the unique features in an image, algorithms can make smart decisions and carry out complex tasks much like how humans see and interpret things. Feature detection algorithms generate a set of feature descriptors, which are brief representations of the identified features.

These descriptors encapsulate key information about the features, such as their location, scale, orientation, and appearance. Feature descriptors are important for the next steps in computer vision tasks, like matching features, recognising objects, and retrieving images.

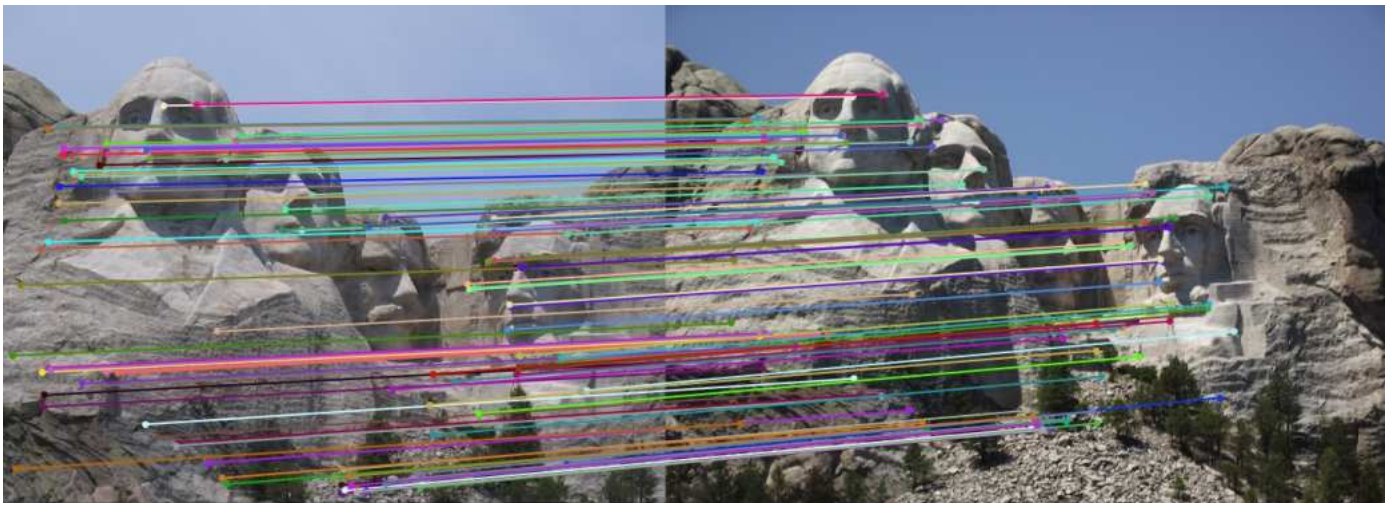

Modern feature detection algorithms can process images in real-time, achieving detection rates of up to 60 FPS (frames per second) on standard hardware.

Key advantages include the following:

- **Quality control:** Feature detection identifies specific characteristics or defects in products during manufacturing, ensuring consistency and quality. This automation reduces human error and speeds up quality checks.
- **Inventory management:** In warehouses, feature detection technology helps accurately track inventory by detecting item characteristics, labels, and barcodes. As a result, businesses have real-time visibility into their processes, which prevents stock issues and minimises losses.
- **Customer insights in retail:** Feature detection can analyse shopper behaviour, such as dwell time in specific areas and reactions to product displays. As a result, retailers optimise store layout, stock popular products, and improve customer experiences.
- **Improved decision-making with real-time data:** Businesses can use feature detection technology to gather real-time data from visual feeds, enabling faster, data-driven decisions, for example, monitoring factory equipment to detect issues early or observing retail foot traffic for immediate adjustments.
- **Product tracking:** Recognising labels, batch numbers, or other specific features make it easier to ensure that products meet safety standards and regulatory requirements.
- **Marketing and personalisation:** For eCommerce and digital platforms, feature detection helps recognise customer preferences by identifying products viewed or interacted with. This data can refine marketing strategies, recommend relevant products, and create a personalised shopping experience.
- **Waste and resource efficiency:** Feature detection can monitor production processes to reduce material waste by spotting defects early, thus avoiding the use of faulty components.

Feature detection is used in AR applications to recognise and track objects or surfaces. For example, applications like Pokémon GO use feature detection to anchor virtual objects to real-world locations based on identified features in the environment.

Feature detection technology helps businesses operate more efficiently, adapt quickly, and make better-informed decisions. This ultimately leads to increased customer satisfaction, cost savings, and a stronger competitive edge in the market.

Deep learning and neural networks

Deep learning is a branch of machine learning that uses artificial neural networks to analyse large datasets for image recognition—inspired by how the human brain functions. These neural networks play a critical role in this process, allowing the system to learn features directly from the data.

During training, an artificial neural network functions as a filter, matching input images to their corresponding output labels. Deep neural networks, which consist of multiple layers, improve the model's ability to predict by effectively learning various mapping functions.

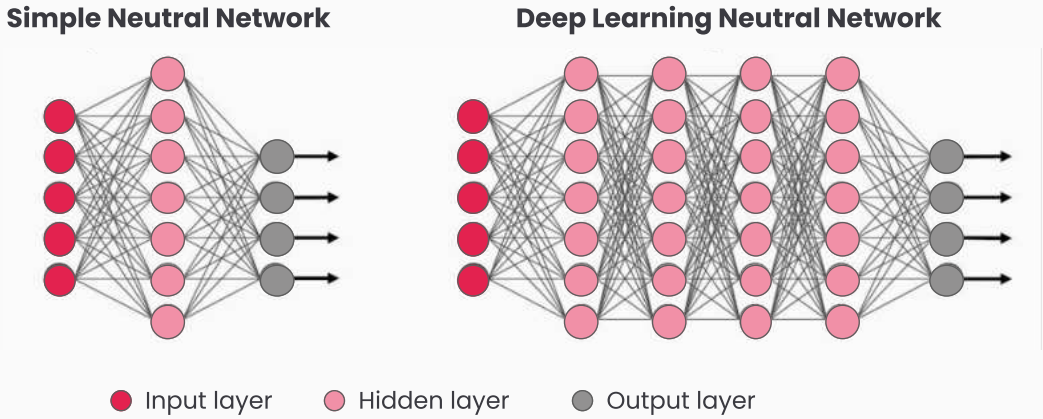

There are three main approaches to learning in image recognition systems: supervised, unsupervised, and self-supervised. Each of these methods plays a different role in training image recognition models.

In supervised learning, models are developed using labelled datasets, where each image has a specific category attached to it. This leads to high accuracy in recognising the visual traits associated with each category. In contrast, unsupervised learning does not depend on labelled data. Instead, it employs algorithms to uncover underlying patterns within the dataset, often using clustering methods.

Self-supervised learning, a branch of unsupervised learning, trains models using a secondary task that automatically generates labels. This approach helps the system identify complex image features, allowing for more effective learning.

Deep learning models, particularly convolutional neural networks (CNNs) for image recognition, have achieved accuracy rates exceeding 95% on benchmark datasets like ImageNet and CIFAR-10.

Google Translate's use of deep learning has increased translation accuracy by around 60% compared to previous statistical methods.

Deep learning and neural networks offer businesses powerful tools to automate complex tasks, gain deeper insights from large datasets, and improve decision-making across various sectors:

Customer experience and personalisation

Neural networks analyse large volumes of customer data to tailor recommendations and improve engagement. For example, **Netflix uses deep learning algorithms to recommend shows based on individual viewing habits, which has led to a 75% increase in viewership** from recommendations alone.

> *Personalisation can increase revenue by 10-15% for retail and entertainment companies that adopt and implement AI effectively.*

Medical diagnostics

Deep learning can analyse medical images and identify conditions faster and often more accurately than human experts. For example, **DeepMind's algorithms diagnose eye diseases with an accuracy comparable to that of ophthalmologists. It has also been used to detect breast cancer from mammograms with an accuracy of 94.5%.**

> *AI diagnostics could save up to $150 billion annually in healthcare costs by 2026 through faster and more accurate diagnostics, leading to earlier treatments and better patient outcomes.*

Financial analysis and fraud detection

Neural networks analyse transaction patterns to detect fraudulent behaviour and automate financial forecasting.

> *Financial institutions using AI for fraud detection reduce false positives by 80% and decrease investigation times by 70%, helping save billions annually in fraud losses.*

Predictability in manufacturing

Deep learning models predict when equipment will likely fail, enabling proactive maintenance and reducing downtime.

> *Siemens uses neural networks to monitor the health of wind turbines, predicting malfunctions before they occur. Predictive maintenance can increase equipment uptime by 20% and reduce maintenance costs by 25%.*

Sales forecasting

Amazon uses deep learning to forecast product demand, optimise inventory levels, and reduce the risk of overstocking or stockouts.

> *Companies using AI in supply chain management see improvements in logistics costs by 15%, inventory levels by 35%, and service levels by 65%.*

Automation in customer service and virtual assistants

Neural networks power chatbots and virtual assistants that provide real-time responses to customer inquiries.

> *H&M uses a deep learning-powered chatbot to assist customers with style recommendations and product searches, improving customer engagement.*

Businesses using AI-driven chatbots see a 30% reduction in customer service costs and a 60% improvement in response times.

Quality control in manufacturing

Machine vision-based quality inspection can increase defect detection rates by up to 90%, enhancing overall product quality.

> *For example, General Electric uses deep learning to analyse defects in aircraft engines and turbine blades, achieving a 15% reduction in production costs.*

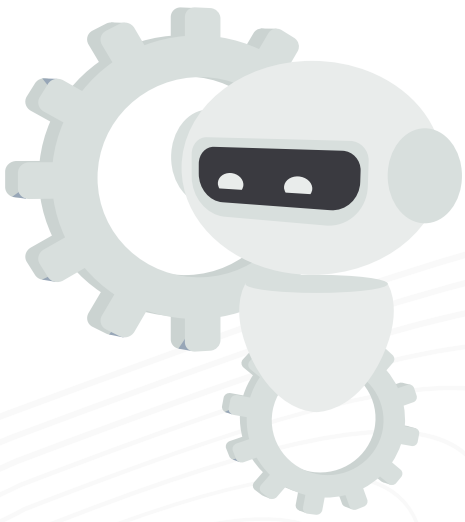

Transfer learning

Transfer learning is a technique in machine learning where knowledge acquired from one task or dataset is applied to streamline the performance of a model on another related task or dataset. Basically, it uses insights gained in one context to improve outcomes in a different context. This approach is often applied for tackling regression problems in data science and training deep learning models.

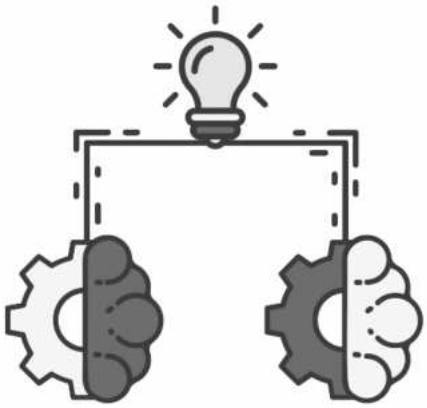

Traditional learning methods, as a rule, create a new model for each task using the available labelled data. This approach assumes that the training and testing data come from the same feature space. Therefore, if the data distribution changes or a trained model is used on a new dataset, users must retrain a model from the beginning, even when dealing with a similar task, such as classifying sentiment in movie reviews versus song reviews.

In contrast, transfer learning uses models or networks already trained as a foundation. It uses the knowledge acquired from an initial task, such as classifying movie reviews, and applies that understanding to a new, related task, like classifying song reviews.

Here are some of the key benefits of transfer learning:

Reduced training time and costs

Transfer learning cuts down on training time since businesses can use pre-trained models instead of starting from scratch. So, you can expect lower computational costs and quicker time-to-market.

> *Transfer learning can reduce training time by up to 90% compared to traditional training methods, especially in fields like natural language processing (NLP) and computer vision.*

Improved model performance

Businesses can achieve better performance with smaller datasets by starting with a model that has already learned relevant features from a large dataset. As per recent research, transfer learning can boost model accuracy by 5-10% in classification tasks compared to models trained from scratch, mainly when the target dataset is limited.

Less data requirement

Transfer learning allows businesses to train models effectively even with limited data available. This is particularly plays in favour of those with limited or expensive data. For example, in medical imaging, **transfer learning has been used to improve the detection of diseases with small datasets, allowing hospitals to develop diagnostic tools without the need for extensive labelled data.**

> *Stanford researchers developed a model that outperformed human radiologists in identifying pneumonia from chest X-rays. The model achieved an accuracy of 94.6% compared to 88% for human radiologists.*

Faster prototyping

Businesses can quickly test and deploy models, facilitating rapid iteration and experimentation.

> *Companies like Facebook and Google have used transfer learning to prototype various NLP applications quickly, reducing the time required to roll out new features.*

Image classification

Computer vision covers several key issues, including image classification, localisation, image segmentation, and object detection. Among these, image classification is seen as the foundational aspect, laying the groundwork for addressing the other problems in the field.

Image classification is widely used in areas such as medical imaging, identifying objects in satellite photographs, traffic management systems, detecting brake lights, and machine vision.

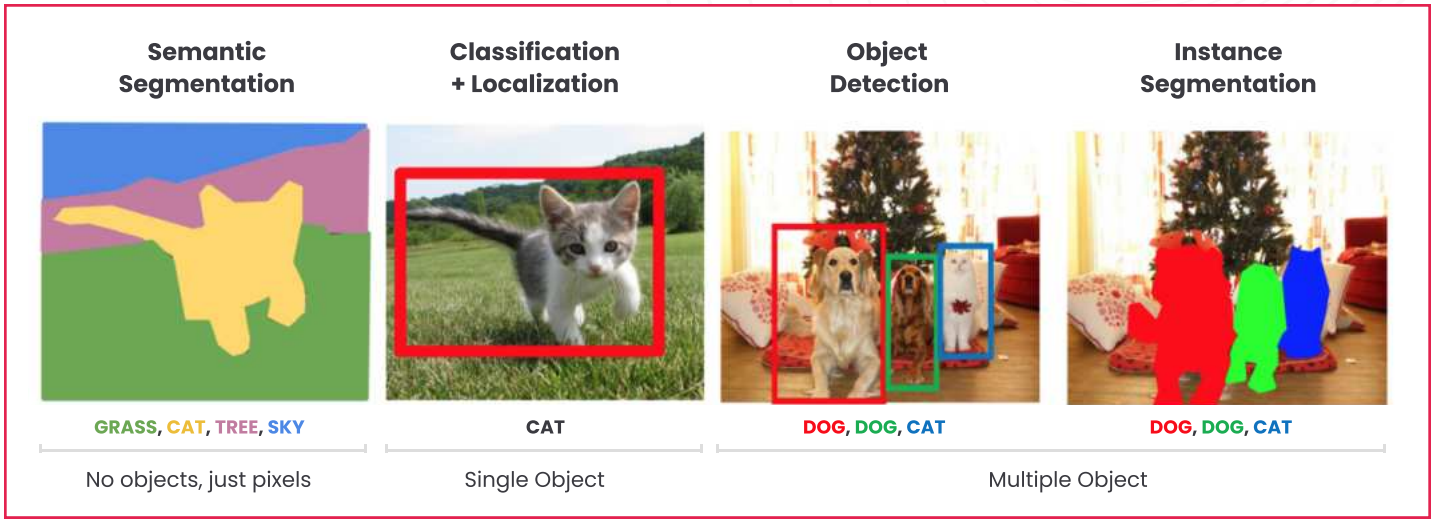

Image classification involves sorting an image and assigning class labels to specific groups of pixels or vectors based on set criteria. This process can use one or several spectral or textural features to make decisions.

In the past, image classification mainly focused on security applications, but today, image classification plays an important role in healthcare, industrial manufacturing, smart cities, insurance, and space exploration.

This dynamic adoption trend can be explained with the growing volume of visual data and rapid developments in computing technology. Visual data can become a valuable asset when used effectively, as its management costs are outweighed by the benefits it brings to various business operations.

> ***For example, Pinterest's Lens feature allows users to take pictures of items and find visually similar products, using image classification technology to enhance search results.***

Some specific applications include:

- **Skin cancer classification using AI:** Dermatologists often examine numerous skin conditions to find malignant cells. Image classification can automate this labour-intensive process.
- **Facial recognition for security purposes:** In settings like airports, image classification can analyse security footage to identify individuals quickly.
- **Traffic monitoring and congestion detection:** This technology can automatically count vehicles on roadways and identify traffic jams.
- **Customer segmentation in retail:** Image classification can categorise retail customers into different groups based on their buying behaviours, helping businesses target their marketing efforts.
- **Land use mapping:** This technique can also automatically classify land use, identifying areas like forests or agricultural land. Additionally, it can monitor environmental changes, such as deforestation or urban development, and assist in estimating agrarian yields.

In the spotlight: current trends
Edge computing

Edge computing is a network model that processes data at the edge or close to where the data is generated. This approach can remove the need for sending image data to a central server or cloud for processing. For example, edge computing in road toll collection systems can enable license plate recognition to be performed on a low-power single-board computer near the camera. This means only the license plate numbers are sent, rather than full images of vehicles or the roadway. The availability of affordable and efficient single-board computers makes this feasible today.

Processing image data at the edge reduces system latency and jitter by minimising the number of switches and hosts that data packets must pass through between their source and destination. Each network node the data travels through adds to the delay from image capture to action. This latency is further reduced by using edge computing because there is no need to upload image data, which also speeds up response times. Additionally, edge nodes can anonymise the data before sending it to the cloud for further analysis, helping to enhance system security and address privacy concerns.

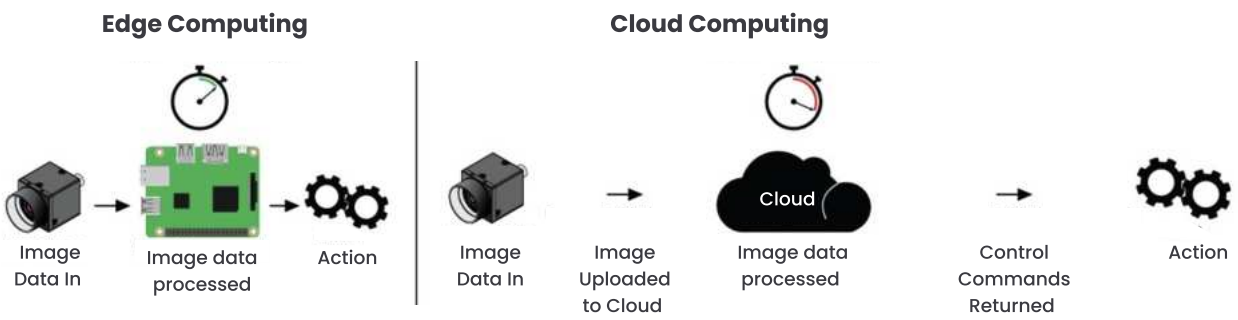

Figure 1: Edge computing processes image data close to the source for low system latency

Figure 2: Cloud computing results in a long signal path for image data, which increases system latency

Reduce Bandwidth	Handling your data at the source removes the need for sending images to a central server. By only transmitting relevant information, there is a reduction in the amount of bandwidth needed. Businesses can see a reduction in bandwidth costs by up to 30% when implementing edge computing strategies.
Reduce Latency	By decreasing the volume of data transmitted from the edge, the system operates more quickly and reduces the delays between image capture and the receipt of information. Edge computing can reduce latency to as low as 1 millisecond, compared to typical cloud latency of 20-100 milliseconds.
Improve privacy and security	Confidential data, such as license plates and facial images, is not sent to the cloud. Companies implementing edge computing have reported a 20% decrease in security incidents related to data transmission.

Examples and use cases

Retail

Retailers use edge computing for personalised customer experiences, using local analytics to tailor promotions and product placements based on real-time customer behaviour.

> *Retailers implementing edge solutions report an average sales increase of 5-10% through improved customer engagement.*

Transportation and logistics

Fleet management systems use edge computing to process vehicle data on-site, optimising routes and improving fuel efficiency based on real-time traffic conditions.

> *Companies using edge computing for logistics can achieve up to a 15% reduction in fuel costs through optimised routing.*

Healthcare

Edge devices in hospitals allow for real-time patient monitoring, ensuring that critical data is processed instantly for better patient care and emergency response.

> *Hospitals employing edge computing technologies report a 30% improvement in patient response times and overall care quality.*

altamira

Real-time performance in image recognition

Real-time object detection is catching on and showing up in all kinds of applications. OpenAI has contributed to this field, improving the precision of object detection methods.

In analysing stock market data, real-time object detection plays an important role in sophisticated trading AI systems. It helps with investment decisions by showing stock market information in easy-to-read formats like charts, graphs, and tables.

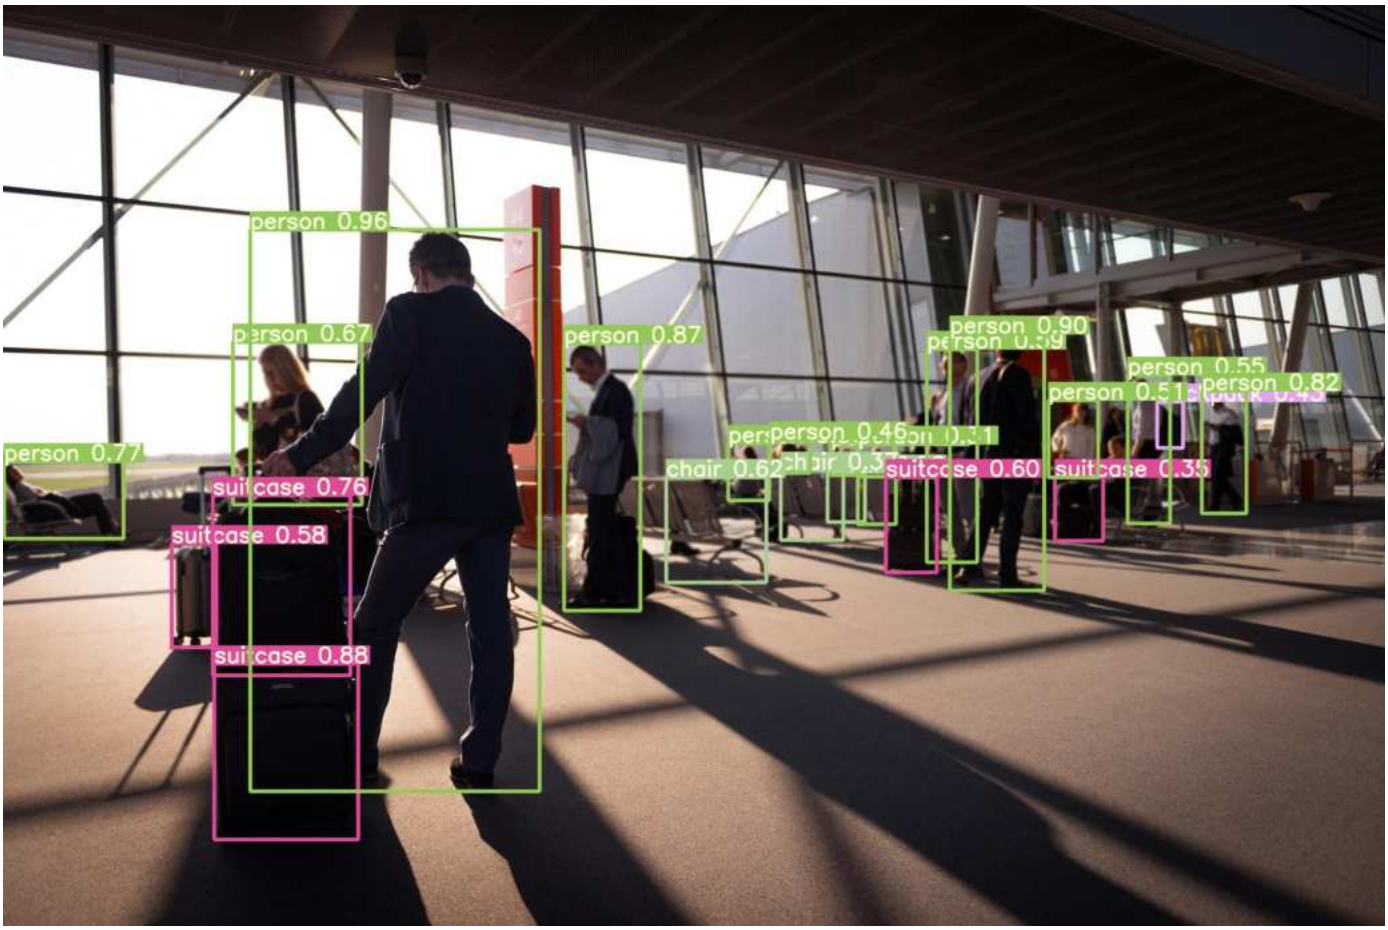

By automating the analysis process, traders can save time that would otherwise be spent on examining and forecasting stock market trends.

To make the most of object detection, traders must become familiar with key tools and libraries, including point cloud annotation tools for labelling 3D data and the PYOD collection for detecting anomalies.

Real-time image recognition extends beyond merely capturing images; it provides immediate insights when an image is taken. Here are some benefits of this technology:

Enhanced user experience

Real-time image recognition allows applications to respond to visual data instantly. For example, users can try virtual outfits in an online store or receive product suggestions based on their surroundings in augmented reality applications.

> *AR applications that use real-time image recognition can increase user engagement by 20-30% compared to traditional applications.*

Improved security and safety

Facial recognition technology can identify individuals with clearance, while algorithms designed to detect unusual behaviour can highlight suspicious activities immediately, facilitating quicker interventions.

> *Airports using real-time facial recognition have reported a 15-20% reduction in security incident response times.*

Streamlined operations and automation

Real-time image recognition is often used to streamline manufacturing quality control or expedite logistics sorting and shipping. **Companies implementing real-time quality control solutions can achieve a 30% reduction in production defects.**

Data-driven decisions

Businesses can gain insights into customer behaviours, product efficiency, and market dynamics by analysing visual data in real-time. This information helps them make informed decisions and refine their strategies. For example, **farmers using real-time monitoring systems can increase crop yields by up to 20% by addressing issues promptly.**

From concept to reality: use cases
Image recognition applications in healthcare

Market growth

The global healthcare AI market, which includes image recognition technologies, is expected to grow at a compound annual growth rate (CAGR) of around 37.5% from 2022 to 2030.

Diagnostic accuracy

AI algorithms can match or even exceed the diagnostic accuracy of human radiologists. For example, a study published in Nature found that AI outperformed human radiologists in breast cancer detection by 94.6% compared to 88.0% for human experts.

A few years back, Facebook made big progress in helping blind and visually impaired users interact with the platform by using image recognition technology. While it may seem simple on the surface, the ability to scroll through Facebook and stay updated on friends' activities can take hours for someone who cannot see.

To address this, Facebook integrated face recognition, image classification, and automatic alternative text features to provide accurate descriptions of photos, identifying not just the objects in the images but also the people present, even if they are not tagged.

This initiative was led by an accessibility team that included Matt King, Facebook's first blind engineer, who has been legally blind since his early twenties.

The technology behind this feature is similar to the system that suggests friends to tag in photos. It uses machine learning to analyse facial pixels, creating a "template" that helps identify individuals in future images.

Another area benefiting from this technology is radiology. According to IBM's estimates from last year, images make up at least 90% of all patient data. This huge volume has put a lot of pressure on radiologists, who have to analyse more and more images every day.

The massive number of medical images can be overwhelming for human radiologists, but it creates an opportunity for deep learning algorithms, which are key to many image recognition technologies. These algorithms thrive on data; the more they have to work with, the more effective they become. We're now seeing many cases where deep learning algorithms and image classification tools are outperforming human radiologists and are being incorporated into healthcare practices.

Another notable example is Enlitic, an Australian company founded by Jeremy Howard, the former president of Kaggle. Enlitic specialises in identifying tumours in lung CT scans to enable earlier diagnoses. In internal testing, their software demonstrated a 50% increase in accuracy for diagnosing malignant lung tumours compared to a group of radiologists.

altamira

Image recognition in the security industry

Market growth	Surveillance systems	Facial recognition accuracy
The global facial recognition market size is projected to grow at a CAGR of around 15.4%, reaching about $8.5 billion by 2027.	By 2026, the global video surveillance market, which proactively uses image recognition technology, is expected to reach $62.6 billion, growing at a CAGR of 10.4% by 2026.	Modern facial recognition systems achieve accuracy rates of over 95%.

The ability to identify and recognise faces is a valuable tool for the security sector, particularly in safeguarding properties from unauthorised access.

Home security systems are not a new concept. Many homeowners choose systems with motion detectors connected to a 24/7 monitoring service. However, these systems can be pretty basic. They often lead to false alarms since they depend primarily on motion or heat detection, which cannot distinguish between a homeowner who has forgotten their code, a pet wandering around, or an intruder.

Home security systems have improved a lot with the addition of image recognition technology and QR code generators. These systems can now recognise and remember household members, regardless of lighting conditions or angles, and can tell the difference between people and pets.

For example, the Netatmo Welcome camera only begins recording when it detects an unfamiliar face. Similarly, Ulo, a cute owl-shaped personal security device, takes this further by streaming live video to a designated device when it identifies unknown individuals.

Law enforcement agencies are also adopting image recognition technology. In the UK, the South Wales Police use facial recognition to monitor large events and crowds while looking for suspects. The system complements officers' efforts rather than replacing them. If the technology identifies someone with at least a 59% resemblance to a suspect, the information is sent to an officer for verification before further action is taken. As a result, the use of image recognition has notably reduced costs and enhanced the overall efficiency of police operations.

Image recognition in the automobile industry

Market growth

The global automotive AI market, which includes image recognition technologies, is projected to reach $14.6 billion by 2030, growing at a CAGR of 33.8%.

Surveillance systems

As of 2021, about 85% of new vehicles sold worldwide were equipped with at least one form of ADAS. Many of these systems rely heavily on image recognition technologies for features like lane departure warning, adaptive cruise control, and automatic emergency braking.

Facial recognition accuracy

According to Markets and Markets, the autonomous vehicle market is expected to grow from $54 billion in 2022 to $556 billion by 2026, with image recognition being a critical component for navigation, obstacle detection, and decision-making in self-driving cars.

Although fully autonomous cars aren't common yet, they are making impressive progress in their development. Image recognition technology is a major factor contributing to their ability to navigate without human drivers. Alongside lidar and radar sensors, multiple cameras monitor traffic lights, read road signs, keep track of surrounding vehicles, and detect pedestrians and other obstacles.

Driverless vehicles can potentially lower accident rates, enhance compliance with emission standards, and reduce traffic congestion. This is mainly because machines can adhere to rules more effectively and respond to sudden changes more quickly than people can.

Google's self-driving initiative, Waymo, has been testing and refining autonomous vehicles for nearly a decade. They have even constructed a small town in the Arizona desert to evaluate their algorithms in various real-life scenarios.

AI car damage recognition

Tech stack:
React, Python, OpenCV, Pillow, TensorFlow, EfficientNet, ResNet, Mask R-CNN, ChatGPT, Django, Azure

Altamira's AI-powered car damage detection solution uses photos captured with a smartphone to instantly detect dents, scratches, and minor body defects in vehicles. This revolutionary technology makes it possible to activate insurance policies on the spot without any need for expert assessment.

You can overcome the following challenges with our AI car damage detection solution:

- **Personalised customer experience:** Generate personally adjusted prices based on the car owner's driving patterns, vehicle's condition, speed patterns, traffic environment, and the number of miles driven.
- **Better risk assessments:** Make more informed and accurate decisions about the risks of insuring different vehicles.
- **Faster incident investigation:** Assess the damages remotely, allowing quicker processing of insurance claims.
- **Exact vehicle damage detection using AI:** Exactly evaluate the vehicle condition based on gathered and processed data generated by internal car sensors and get accurate calculations.
- **Detection of fraudulent cases:** By providing objective, consistent, and accurate information about the condition of a vehicle, this technology can help insurers identify potentially fraudulent cases.

Recognise a car

The system automatically assesses the car in the photo, its colour, model, and angle while detecting any signs of car damage. It also recognises whether the image was modified to prevent fraud.

Know the severity of the damage

Damage recognition contributes to rapid assessment of the car damage level and the location of the damaged components. It also identifies whether the car is damaged due to a car incident or merely surface-level dirt.

Get a repair estimation cost

By analysing data from a large number of repairs, a car damage detection solution powered by machine learning identifies patterns and trends in the cost of different types of repairs. So that the repair costs are estimated accurately, and insurance policies are priced reasonably.

altamira

Image recognition in the retail industry

Market growth	In-store experience:	Customer insights
The global image recognition market in retail is expected to grow at a CAGR of about 14.5%, reaching around $6.78 billion by 2027.	According to a survey by Deloitte, around 59% of retailers are expected to invest in AI technologies, including image recognition, to enhance in-store experiences by 2025.	Retailers using image recognition to track shopper movements and behaviours increase sales by up to 15% by better understanding customer needs and preferences.

Thanks to advancements in image recognition technology, trying on clothes before buying may soon be a thing of the past. Several well-known brands, including Topshop and Timberland, have adopted a device known as the visual mirror, allowing customers to virtually try on various items from their collections. This visual mirror can be placed inside or outside a store to draw in shoppers.

Essentially, the mirror is a large screen equipped with multiple cameras that track different body parts of the person standing before it. It automatically selects the correct size, allowing users to turn and see how the clothing fits from every angle. Customers can also browse for different colours and styles, making the shopping experience more pleasant.

Some versions of the visual mirror enable users to capture images of the outfits they create, send them to their phones, and compile a complete inventory of items available in the store.

While visual mirrors aim to enhance the shopping experience, a Japanese company has developed a security system called AI Guardian, designed to combat shoplifting. This technology scans entire bodies rather than just faces and identifies suspicious behaviour based on a dataset that outlines characteristics typically associated with shoplifters.

A test conducted in a Japanese store revealed a 40% reduction in shoplifting incidents after implementing this system. While this technology isn't widely used yet, the developers of AI Guardian and similar security cameras believe it's just a matter of time before it becomes more common and its accuracy improves.

Recently, Altamira delivered a ready-to-scale smart data capture and retail shelf analytics solution

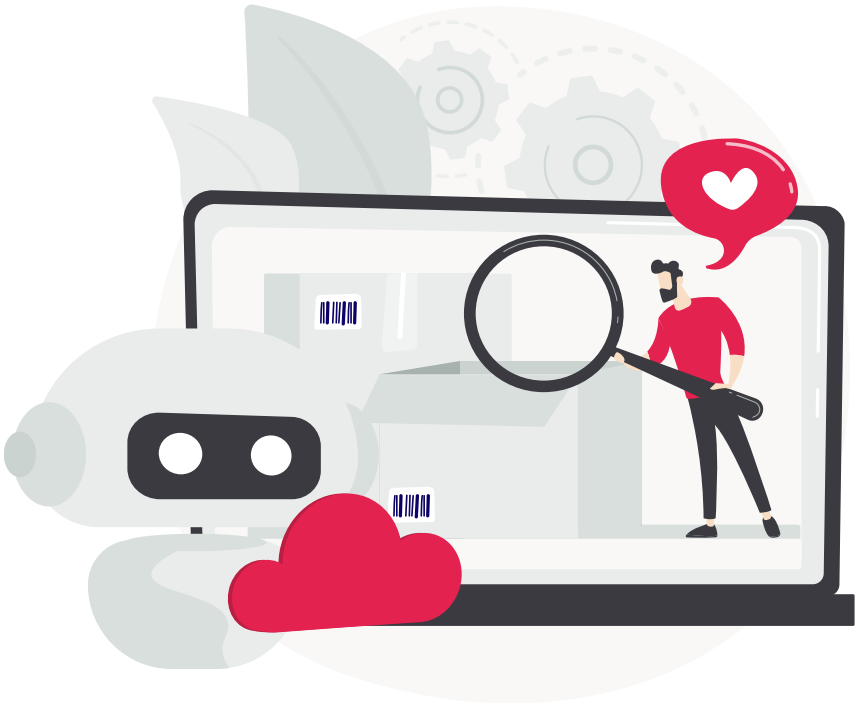

The product detection solution for retail company is powered with a product shape detection technology which employs sophisticated algorithms to identify and delineate the bounding boxes of products in the images captured by the mobile app. This detection process is important to accurately locate and identify different products on crowded shelves.

Feature matching entails calculating the unique “features” of the recognised product bounding boxes. The calculated features are then compared against the features of products stored in the product management platform.

The AI system is designed to perform this highly precise comparison, enabling it to match the product captured in the image with its corresponding entry in the database.

Combining these AI-driven processes within the PoC development, our solution efficiently identifies front-facing products. The integration of the mobile app with the product management platform ensures accurate product recognition, greatly improving inventory management and operational efficiency.

Conclusion

The future of image recognition looks bright and is set to bring about significant changes. We can expect improvements in advanced image recognition algorithms due to:

- Deep learning techniques that enhance object detection and classification
- Neuromorphic chips that enable faster processing with less energy consumption
- Quantum computing for managing large datasets
- Generative adversarial networks that produce improved synthetic training datasets

> *However, as Ofer Razon from superwise.ai points out,* ***"Image recognition is an AI-based technology, which means it is only as good as the data it was trained on. As such it is naturally inclined to develop biases. To ensure ethical and fair use of AI, leaders should make use of best practices and technologies to mitigate the risks, including wide data sets, proactive bias detection, end-to-end visibility and transparency regarding the health of these models."***

altamira

How does Altamira help with your computer vision?

Altamira provides a variety of artificial intelligence and machine learning consulting and development services for Clients of all levels of experience. Altamira has a diverse portfolio of computing, networking, and storage solutions designed to meet various performance and budget requirements.

For businesses looking to implement a standardised image recognition solution, we simplify the process of preparing data, building, training, and deploying machine learning models.

For clients who may not have machine learning expertise, seek a quicker time-to-market, or want to enhance an existing application with artificial intelligence, we offer a selection of machine learning-based services.

Challenges we solve

- Security issues addressed through facial recognition, anomaly detection, and real-time surveillance, alerting unusual activities without human intervention
- Poor quality control processes improved by identifying defects without the need for manual inspection, thus reducing costs and improving efficiency
- Inventory tracking, warehouse management, and logistics streamlined by automating the process of identifying, counting, and managing stock levels and tracking assets across different locations
- Time-consuming monitoring of crop health, pest detection, and yield prediction optimised through real-time monitoring and analysis
- Traffic and urban planning enhanced by analysing traffic flow in real-time from camera feeds, identifying congestion, accidents, and other anomalies

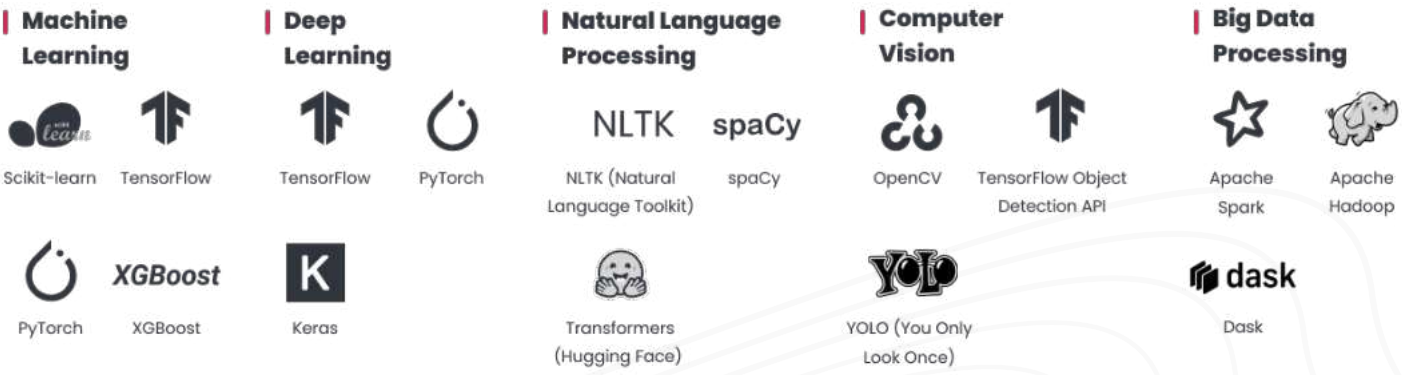

altamira

About Altamira

We provide solutions that make a tangible impact to our clients' growth and productivity. With domain knowledge across product and technology development, we aim to provide cost-efficient solutions without compromising quality. We are driven to deliver the best, every single time.

We specialise in providing AI/ML software development services that can spread the benefits of AI throughout your entire organisation, from top to bottom.

- Understanding the business goals and assessing the possibility of achieving them with AI/ML solutions
- Identification of possible use cases of AI/ML models and their implementations within your business environment
- Feasibility validation and understanding the quality of your data and their usage by the proper AI/ML models
- Identification of the right pre-built AI/ML models to be implemented within your projects

Let's turn your AI concepts into tangible realities through Artificial Intelligence and Machine Learning.

Have a question? Let's talk.
We love to help our clients be successful!

Email: **sales@altamira.ai**
www.altamira.ai