

Computer Vision Projects: 7 Success Stories

Originally published at <https://www.iflexion.com/blog/computer-vision-projects>

Published: March 9, 2021 | By Martin Anderson



The global computer vision market is currently forecast to grow from US\$12,969.95 million in 2020 to US\$19,398.97 million at the close of 2025¹. From agricultural and industrial robotics through to retail innovation, the transformative power of [computer vision software](#) is clear from the wide range of sectors that have adopted image recognition as a core technology.


Let's take a look at seven computer vision projects across a range of markets, which have proved the power of AI-enabled image analysis to improve performance and profitability.

Remote Competitor Analysis with Satellite Imagery and AI

The last two decades have brought such consistent improvements in the quality of commercially available satellite imagery that it is now possible to obtain extraordinarily close detail on a target location — and thereby to use image segmentation, object classification, and custom-trained machine learning models to build up data-rich timelines on otherwise inaccessible locations.

In this regard, as one might expect, the military and state sectors are agents of change: in 2017 the US government agency IARPA (Intelligence Advanced Research Projects Activity) released a dataset² containing one million satellite images in support of its challenge³ to machine learning developers to create a 'Functional Map of the World' (fMoW).


The objective was to develop AI systems capable of predicting the functional purpose of buildings and the use of land over a period of time — a state/military objective which has crossed over fairly recently into [business intelligence development](#).



Eye in the Sky

In one case study⁴, commercial satellite intelligence company Orbital Insight conducted a years-long surveillance of construction work on the Hengyi refinery in Brunei, obtaining an extraordinary volume of information primarily from AI-driven analysis of high-quality satellite views of the


site.



Onsite activity can be gauged by identifying cars, pedestrians, the emergence of roads and buildings, and — in the case of Hengyi — the proximity of ships in the nearby harbor, with that image data correlated to AIS shipping data.

Tracking Site Progress from Space with OpenStreetMap

The proprietary database can also include output from other satellite image content sources. Here, the addition of new roads and buildings to the OpenStreetMap database images for the location are traced in polygons by the proprietary system, creating a chronology of development for the area over a two-year period:



SWIR Image Analysis

Images from short-wave infrared (SWIR) satellites can also reveal much about a remote location:

SWIR IMAGE ANALYSIS

In this example from European Space Imaging, SWIR output is overlaid on conventional satellite imagery, cutting through obscuring smoke to reveal the exact location of ground fires.




Data source: euspaceimaging.com—Short-wave Infrared Imagery (SWIR)

IFLE 

In the case of the Brunei site, SWIR image analysis revealed high-intensity heat flares. In the context of an oil refinery, these patterns match the burn-off of gases common to active refineries, revealing that the site was now either operational or actively testing production processes.

SWIR REVEALS FLARE INCIDENTS AT HENGYI REFINERY



Data source: orbitalinsight.com—Brunei Hengyi Refinery

IFLE 

Growth of Commercial Space Surveillance

It should be considered that commercially-available satellite imagery, while much improved over the last three decades, is believed⁵ to pale against the resolution and level of detail available to governments, with a maximum resolution of 50cm-per-pixel mandated by US law for non-state operatives⁶. Nonetheless, commercialized satellite surveillance, driven by computer vision, is a growth industry⁷.

Commercialized satellite surveillance, driven by computer vision, is a growth industry right now.

Livestock Monitoring Through Image Evaluation

Pig farmers need to know the current weight of an animal in order to regulate its food intake correctly, so that the animal is not starved or stalled in its progress, and expensive feed is not wasted.

The sprawl of farm environments, the generic appearance of the average pig and the sheer numbers involved make it difficult for farm staff to track a pig's weight visually, while frequent manual weighing is a drain on resources.


Wearable IoT solutions are not economically viable due to the challenging operating conditions, the cost of maintenance and replacement, and the difficulty that attached devices have in measuring weight loss or gain.

Continuous Image-Based Assessment of Animal Health

Therefore one Israeli agricultural startup turned to Iflexion to develop [a computer-vision livestock monitoring system](#), wherein machine learning workflows analyze a stream of images from pig pens and track the weight of individual animals on a daily basis by comparing their aerial profile shapes.

The technique was pioneered in the late 1980s by British researchers⁸ and was to become a subject of intense research efforts⁹ in animal husbandry over the next two decades.

IMAGE-BASED WEIGHT EVALUATION OF PIGS IN 1989



Data source: sciencedirect.com—Evaluation of image analysis as a means of estimating the weight of pigs

IFLEXION


Real-Time Monitoring of Weight Gain

The Israeli livestock monitoring system features ceiling-mounted HD cameras that pass one aerial image of the pig every second to a cloud-based machine learning framework. Here, each pig is identified by the color of its ear-tags, and its current weight evaluated by the YOLO V3 real-time object detection system.

The pig's progress is input into a data matrix that can be accessed via a progressive web app (PWA). When the pigs are weighed periodically, the real weight is input to the system, and thus any cumulative errors in estimation are automatically corrected without the necessity for constant manual weighing.

LIVESTOCK MONITORING WITH COMPUTER VISION BY IFLEXION

Three screens in the pig-monitoring interface, revealing (respectively) averages from the pens, progress of individual pigs in a pen, and detailed weight-monitoring information about one animal.



Data source: iflexion.com—A Livestock Monitoring System for an Agricultural Startup

IFLEXION

Cloud-Based Image Recognition Pipeline


The PWA can be configured to show or send alerts for potentially ill or otherwise underweight animals. The sensors are Raspberry Pi units that are required to send images to the cloud framework even in the event of connectivity issues. If this occurs, the system will build a queue of deferred image sends and delete the oldest images as necessary. However, by tweaking the MPEG-4 compression settings on the images, the development team was able to ensure several days' worth of local storage during any downtime.

Computer Vision for Cashierless Retail

The last twenty years has seen an inexorable move away from a local shopping experience towards larger franchise outlets in the retail space¹⁰ — and longer queues.

Eliminating Queues with Computer Vision


In 2016 Amazon, a major researcher into computer vision¹¹, took a new approach to the problem by opening stores where machine systems keep tally of products that customers pick up and automatically charge them for the items they have taken when they leave the retail area.



The image recognition systems at Amazon Go stores have solved several tough challenges in [computer vision applications](#) and in the ideation of no-cashier stores, including:

- Occlusion

Keeping track of people or objects is essential to continuity. If a shopper is momentarily obscured from view by someone else, the system needs to re-acquire their identity; and if two people huddle near a shelf, and an item has been taken, the system also needs to ascertain who took the item (the store shelves have integrated scales¹², so that the removal of an item is known to the system).



- Computing Demands for Store Coverage

Though the throughput of sales effectively automates the Amazon Go inventory system and optimizes the need for in-store restocking staff (estimated to save \$90,000 p/a on a 1,000-ft sq. space¹³), the computing demands for 300 camera feeds means that occlusion detection and re-identification events alone might require nearly 4,000 GPU nodes¹⁴ on a state-of-the-art model¹⁵ — a hard prospect for smaller players than Amazon in the CV-based retail market.

- Dealing with Groups

On entering the store, the person with the payable Amazon account will need to scan their app ID for each non-shopper that is accompanying them (i.e. children, friends and family). Thereafter non-paying members are free to leave and re-enter the store without perturbing the tracking algorithm.

- **Acceptable GPU Calculation Time**

The initial launch of Amazon Go in Seattle was equipped with localized GPUs, in order to update customer app tallies in real time. However, this outfitting and maintenance expense proved unnecessary, and current stores rely on cloud-based calculations across AWS infrastructure.

Assessing Progress and Safety in Construction Sites


Construction projects employ dozens or even hundreds of workers, usually groups of subcontractors with minimal acquaintance and little understanding of each other's methods and working practices.

The unique characteristics of each site and the ad-hoc nature of the workers makes the sector prone to US\$171 billion of work injury costs per year¹⁶, while quality control failures are thought to cost the US construction industry US\$1.6 trillion per annum¹⁷.

Streaming Construction Sites into Neural Networks

Tel Aviv startup Buildots tackles the problem by outfitting the safety helmets of site workers with video cameras that stream a constant feed of imagery into a machine learning system as the worker traverses the site throughout the day.

As the data grows for each project, the machine learning system augments the accuracy and predictive power of a digital twin for the site, which acts as a Building Information Model (BIM). The use of BIM was mandated in the UK in 2016¹⁸, and has been widely adopted in the US¹⁹, despite the lack of enforcement there.



The trained machine learning model, a veteran of many other building projects, is mature enough to spot generic safety hazards, and will gradually develop specific safety parameters for the new site, incorporating what it has learned into future projects.

A Window into the Future

The centralized control system is also capable of comparing the actual geometry of the site, including the state of electrical and other utilities, with the planned finished build.



The ability to superimpose unfinished elements over the real geometry of the site helps to automate scheduling, with relatively little human oversight.

The progress of the site is monitored by AI-trained algorithms capable of prioritizing tasks as necessary, accessible through a single integrated dashboard.

Panopticon View

The system can monitor the location of any helmet camera to a tolerance of a few centimeters and is therefore able to pinpoint the location of installation details²⁰. It monitors 150,000 objects multiple times each week and assigns one of four possible statuses to each object.

The control center collates information from a number of contributing machine learning models assigned to specific tasks including security clearance, feature-based interior navigation, image stabilization, status classification, and image quality assessment²¹.

Computer vision is entering construction sites to assess progress and ensure safety.

Weeding Agricultural Fields with AI

Weeds in agricultural crops costs the farming industry US\$4.8 billion per year in the US alone²², and have been traditionally managed by chemicals, crop spacing, and manual intervention. Another approach has been to fertilize only crop plants rather than adjacent weeds.


However, considering the soaring prices of pesticides, it would be better if really aggressive weed killer chemicals could be targeted against interspersed weeds without being applied to the crop itself. New research in the area of [machine learning in agriculture](#) and plant datasets for computer vision are making this a reality.

California-based agricultural AI startup Blue River Technology leverages image recognition to identify weeds in harvest fields and target them with rapid and effective herbicide jets²³. Dedicated equipment can cover a field at a rate of 2-8mph — incomparably faster than a manual approach.

To achieve this, the company created a convolutional neural network (CNN) over [PyTorch](#) and trained a model on the target species profile shapes and other indicators²⁴.

The cameras run at a high frame rate, taking 5,000 images per minute. Since the visual inference and response must occur in milliseconds, the latency requirements preclude a cloud-based workflow. Therefore, the agricultural robots are powered by a mobile edge compute unit running on an NVIDIA Jetson AGX Xavier System.

The system, originally titled LettuceBot, has been adapted to soy and cotton, with other crops anticipated. It can differentiate between a weed and a harvest plant in 0.02 seconds²⁵ and will also target harvest plants that are either undeveloped and unlikely to become saleable, or else growing too close together to mature optimally. In internal tests, the AI had a higher recognition rate than the company's agronomist:




The system is claimed to reduce chemical usage by 50-90%, and currently operates on 10% of the entire lettuce harvest in the US.

Object Identification for Vision-Based Fashion Retail Search

Singapore-based ViSenze is one of a growing number of startups in image-based search, where pictures instead of words form the basis for a search query. In 2018, the company's own research suggested that 62% of millennials had a high interest in image-based search over other possible technologies²⁶.

FASHION IMAGE SEARCH BY VISENZE



ViSenze's computer vision search platform links²⁷ over 400 million products across 800 high-profile merchants around the world, and has partnered with many smartphone manufacturers, including regional partnerships with Samsung. The company claims that its system powers over 350 million searches a month²⁸, showing 30% higher conversions than text-based search, and a 160% engagement increase²⁹.

The machine learning system behind ViSenze is trained on a Pascal TitanX and 1080Ti, and a NVidia DGX-1, with centralized training management on the cloud. The visual embedding model produces four scales of match: exact match (same item); product variations; same category; and similar category.


Image recognition-based search can drive 30% more conversions in fashion retail, ViSenze states.

Computer Vision for Facial ID Verification

In 2019 it was estimated that the global facial recognition market would grow at a CAGR of 17.2% from US\$3.8 billion in 2020 to US\$8.5 billion by 2025, with North America as the largest market³⁰.

Though COVID restrictions have changed the nature of [intelligent video analytics](#) to an extent, the consequent growth in online activity has increased the scope for device-based facial authentication³¹.

California-based Identiv is one of the market leaders in facial recognition ID technologies. Among a broad portfolio of face ID services, the company's facial database can be used to identify fraudsters through the cameras embedded in ATM machines:



Among various implementations, in Singapore Identiv has developed a video analytics system to monitor class sizes in a higher learning institution, including analysis of students' interests³².

Though the company acknowledges the growing concern around privacy in [facial recognition software](#), it does not believe that the movement will affect on-premises security solutions (such as using facial ID to gain access to private buildings and sensitive public areas, such as airports). Identiv has also stated that growth in the facial recognition market will return after the COVID crisis has abated³³, and that the demands of the pandemic has actually accelerated³⁴ contactless face technologies in the mobile space.