anyvision.

AnyVision Company and Technical Overview

Don’t Let Threats Go Unrecognized

Transform passive cameras into active security systems

AnyVision takes the guesswork out of automating POI (person of interest) alerting, access control and mobile surveillance to quickly spot bad actors, identify high value customers and ensure better user experience.

Recognizing and responding to threats to safety and optimizing customer experience is what drives security and operations teams, and their success depends on accurate and fast identification of people.

The human eye alone is imperfect, and most software that automates the recognition of individuals on watchlists underperforms in real world conditions.

Powered by Vision AI, AnyVision helps protect your employees and premises with automated watchlist alerting, contactless access control and mobile surveillance -- without adversely impacting the user experience.

AnyVision’s Secret Sauce

AnyVision uses deep learning AI to eliminate a traditional systems’ shortcomings, by accurately capturing faces in real-world environments, even with low bandwidth CCTV cameras. The AI software instantly alerts security staff of unauthorized entries and individuals and dramatically reduces their false positive rates.

Highly Accurate

✓

False positives are mislabeled security alerts, indicating there is a threat when in actuality, there isn’t. These false alarms increase noise for already overworked security teams and can lead to system neglect. Most video monitoring systems suffer from rampant false positives often because of the poor quality of the video from CCTV cameras or the imprecision of their software. AnyVision is one of the most accurate facial recognition solutions, excelling in real-time and real-world environments – even in suboptimal conditions when.

Advanced Privacy Controls

✓

Streamlining the experience does not mean having to sacrifice the privacy of individuals. AnyVision’s technology was engineered to comply with privacy regulations while giving operators the tools to protect privacy and ensure compliance, including face blurring of non-targeted faces (GDPR mode), privacy mode, dynamic data retention and strict encryption standards.

Lower TCO

✓

Historically, organizations that wanted to add video analytics and intelligence to their IP security cameras or turnstiles relied on expensive on-premises servers, which were required to handle the compute loads. AnyVision is leveraging edge computing to push more of this processing to the actual cameras through embedded SDKs, enabling us to process more video streams per GPU which lessens the hardware requirements and dramatically lowers your TCO.

Our Solutions

The Vision AI Platform:
The Platform that Underpins our Technology Stack

AnyVision's Access Point AI platform leverages a number of state-of-the-art technologies including computer vision, machine learning, liveness detection and biometrics to help protect an organization's physical access points. By identifying authorized personnel and persons of interest in real-time – whether VIPs or bad actors – modern enterprises can layer in additional operational insights to streamline the customer experience. Access Point AI is being used to optimize touchless access control, video surveillance, and watchlist alerting.

OnWatch

Protect people and premises by recognizing and responding to threats

An automated watchlist alerting system that identifies persons of interest and their contact history in real-time. This video surveillance solution Identifies persons of interest (e.g., felons, shoplifters, employees and VIPs) through both face- and body-based recognition in real-time and sends automatic alerts to your security and customer service teams.

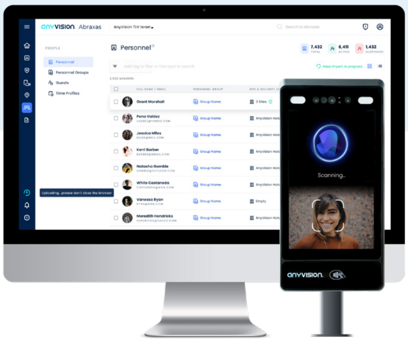

OnAccess

Improve entry experiences while enhancing safety

A touchless access control solution that uses facial recognition to open guarded points of entry for authorized people through privacy compliant and spoof-proof face recognition. By using face-based biometrics and liveness detection to verify authorization, organizations can overcome the shortcomings of traditional passcodes, card keys and fobs which can be shared, lost or stolen.

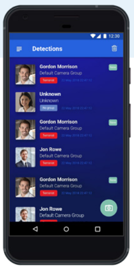

OnPatrol

Safeguard law enforcement with tactical, face-based surveillance

A tactical facial recognition application for Android smartphones that protects law enforcement and military personnel by connecting to existing body cameras and analyzing video streams in real time using on-the-go Vision AI technology to identify persons of interest such as criminals, dangerous individuals, or even missing children.

What Makes Us Unique?

Superior Recognition

AnyVision's neural nets are trained in the toughest conditions on low quality images and have been tested by the most demanding users to ensure the highest accuracy in real-world conditions. This accuracy translates to 0.1% false alarms and 0.2 ms detection speed.

Streamline Investigations

Search video for people and personal attributes through software that indexes each individual and every attribute, and get instant answers. This capability lets you search for bad actors or persons of interest within historical video footage to further investigate security matters using our newest offline file-ingestion capability.

Edge Computing

Our platform pushes more of this processing to the actual cameras (via embedded SDKs) and dramatically reduces the need for expensive, on-premise servers and expensive GPUs, fundamentally changing the historical paradigm for large-scale video security and biometric-based access control.

Advanced Privacy Controls

AnyVision automatically blurs the faces of non-targeted individuals on video playback and even offers a privacy mode which discards all detections of non-enrolled (non-watchlist) individuals.

Multisite POI Management

Centrally manage watchlists, track POIs across multiple locations, and control how POI data is managed, analyzed and distributed while receiving real-time alerts.

Crowded Environments

AnyVision is capable of simultaneously identifying multiple individuals in crowded settings with a high degree of precision. As people enter the camera's field of view, our Vision AI platform will identify persons of interest in a fraction of a second.

What Makes Us Unique?

Make Every Camera Smart

AnyVision adds visual intelligence to your existing camera network and achieves the highest stream to server ratio, which maximizes the value of your investment.

Instant Query Results

AnyVision provides the fastest video search, letting users search and instantly find people and attributes in real-time and in historical footage.

Intelligent Zone Control

Create physical and digital barriers around any space with granular control so only authorized personnel are allowed to high security areas and receive instant alerts when unauthorized people attempt to enter these zones.

Liveness Detection

AnyVision's liveness detection technology ensures that every detected face is a real person by identifying spatial inconsistency and using an array of sensors that create a 3D face map to immediately detect spoofing attempts.

Our Cutting-Edge Technology

AnyVision believes that any computer vision solution must account for real-world conditions. With that in mind, we are solving critical challenges with face and body recognition (attributes and ReID) with superior accuracy across any use case, platform, and environment.

Real world challenges include:

- Changes in appearance
- Changes in pose
- Blurriness / Compression Artifacts
- Occlusions

To Solve These Problems

Superior Data Sets of more than 50 million faces (excl. Augmentation)

Diverse, real-world data collected across challenging, security-based scenarios

State-of-the-art augmentation to prepare for the real-world problems listed above

Dedicated Algorithm Pipeline

Dedicated pipeline with the following steps:

1

Detection

Optimized for very small faces and bodies from different angles, requiring an exceptional uncertainty estimator

2

Uncertainty Estimation

To an absolute minimum, reduces the false-positive rate that would usually result from very small face & body sizes

3

Face Recognition Vector

Features in our recognition vector are well separated to support vector combination

4

Vector Aggregation

Using a tracker to capture different appearances of the same person and combining vectors to create a single vector that better represents that person

Furthermore, AnyVision believes that accuracy across face and body in real-world scenarios, which is the heart of our technology, makes us the market leader and clear choice for the western world, as demonstrated in this brief.

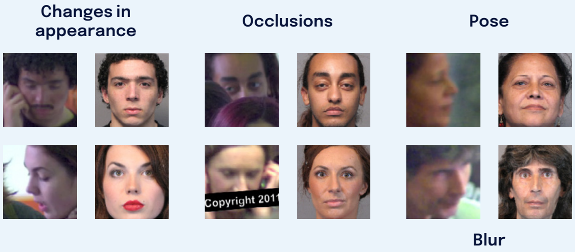

NIST's FRVT Testing

The National Institute of Standards and Technology (NIST) is a physical sciences laboratory and a non-regulatory agency of the United States Department of Commerce. As part of these efforts, NIST runs a facial recognition vendor test (FRVT) on a continuous basis.

The latest report, published in March, AnyVision was ranked 10th on the list. However, the top six companies are headquartered in China or Russia. This is important because many companies in the Western World prefer not to do business with China and Russia. AnyVision ranked consistently high across all categories - including two that were not part of the original ranking order.

Moreover, if you dig deeper into the report, you'll see a wide variance in accuracy across regions and ethnicities from algorithms trained on data from these regions. In contrast, AnyVision's algorithms are trained on very diverse data sets (unlike our Chinese or Russian counterparts). The result is that we are significantly more accurate when it comes to dealing with the diverse groups of people that would be encountered in real-world scenarios outside of these regions.

AnyVision's algorithms are trained on real-world footage, and perform exceptionally well in adverse conditions such as low-lighting, poor angles, occlusions (like masks), extreme poses, and on very diverse data sets.

Many of the companies participating in NIST may do well in laboratory conditions which are very controlled environments, but do not perform as quickly or accurately in more "in the wild" scenarios where subjects are less cooperative, and where one might encounter extreme poses and angles, poor image quality, low lighting or occlusions such as hats, masks, or other coverings.

What's more, many companies listed in the report, particularly those based in China and Russia, are known to train their algorithms on datasets that contain very limited diversity across gender, race and ethnicity. The result is that AnyVision is significantly more accurate when it comes to dealing with diverse sets of people that would be encountered in real-world scenarios outside of these regions.

UK Home Office

Another important test for benchmarking is provided by the **UK's Home Office**. The Home Office, also known as the Home Department, is a ministerial department of the Government of the United Kingdom, responsible for immigration, security, and law and order. The Home Office presented FRT vendors with a number of real-world facial recognition scenarios to assess their accuracy in the wild (vs. within controlled, laboratory environments). The table below provides a summary of AnyVision's performance across the Home Office tests.

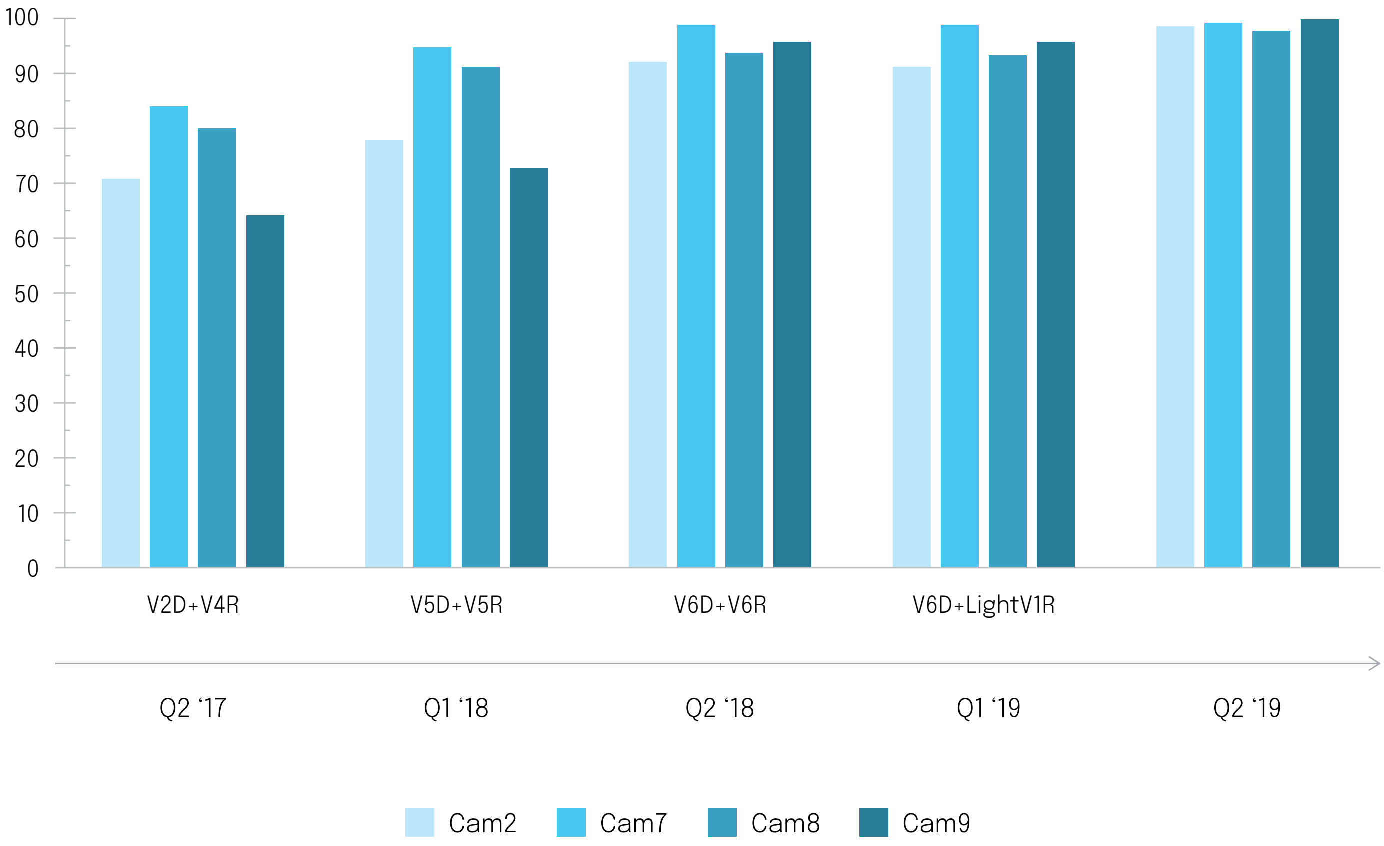

UK Home Office CCTV Benchmark Tests: AnyVision

Other Testing Methodologies

One important assesment of the technology in the wild is to test the accuracy of facial recognition solutions by looking at the minimum distance between pupils.

AnyVision is able to identify individuals with a high level of accuracy when the distance between the eyes was less than just 20 pixels and 45 pixels for the entire face or partial face (which can include only one eye).

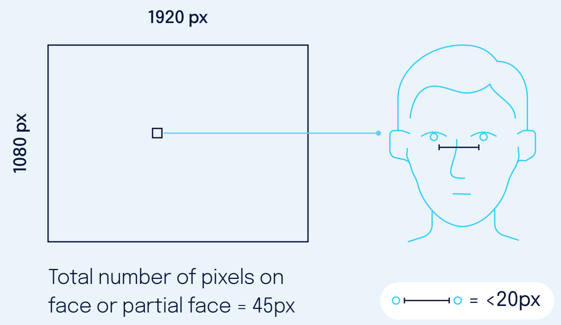

Another general test is how the facial recognition technology performs when the subject is not looking directly at the camera (which simulates real-world conditions).

This is clearly a more technically challenging problem for most facial recognition vendors. Anyvision was able to correctly identify faces with +/- 90 degree angles (horizontal orientation) and +/- 55 degrees angles (vertical orientation).

Example:
90-degree horizontal view of a face

Facial recognition vendors should also be tested when images of faces are captured in low light and different levels of blur.

AnyVision was able to recognize faces with a high degree of accuracy in a variety of lighting conditions. In low light environments, surveillance cameras produce considerable noise and due to longer exposure times, facial images are often blurred.

Outdoor

Sunny Day (Direct Sunlight)	Sunny Day (Shade)	Overcast Day (Shade)	Dawn	Dusk
10000-1000000 Lux	2000-4000 Lux	100-1000 Lux	5-25 Lux	5-25 Lux
Street Lighting	**Full Moon**	**Quarter Moon**	**No Moon**	**Overcast Night**
0.5-10 Lux	0.1 Lux	0.01 Lux	0.001 Lux	0.001 Lux

AnyVision & Edge Computing

Anyvision is making big strides to optimize the algorithms towards edge platforms. The focus is on the following platforms:

Ambarella™
CVflow

NVIDIA
TensorRT

intel.
OpenVINO™

Qualcomm

Qualcomm
snapdragon

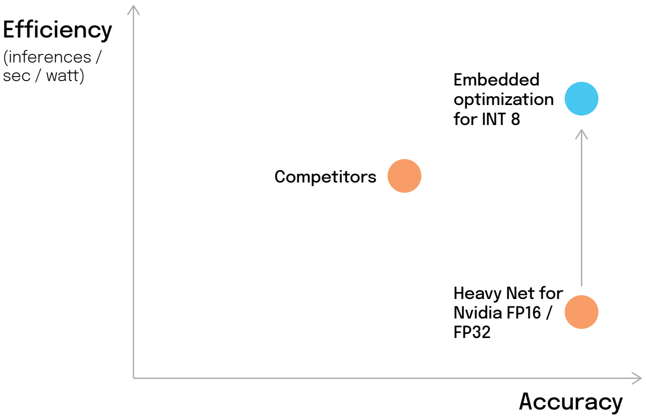

Jetson Xaviar NX

Release Dates

September 2021
Jetson only (no OnWatch) with onboard enrollment ·

December 2021
Final release with OnWatch and OnAccess

The Jetson Xavier NX can run up to 6 channels of HD RTSP connected cameras. The functionality of this is identical to the SDK on the camera just running on a Xavier NX. We are fully compatible with Jetson Xavier NX running up to six cameras.

Performance is about 8-10 FPS in HD resolution including liveness detection:

Face detection
Uncertainty
Vector generation
Tracking
3D liveness (depth map)

Honeywell S70 Camera

Release Dates

November 2021
Integration of AnyVision's SDK in the S70 camera. Data will be processed on the camera, vectors shared with OnWatch. Metadate available via Onvif. Camera will support optimal settings for facial recognition.

3D Facial Access Reader

Release Dates

September 2021
Demo using the reader to enroll a face and check liveness and facial recognition on the device. No OnWatch needed.

December 2021
OnWatch/OnAccess integration

Edge Access Control Reader is a professional design of Ambarella Janus board. The functionality will include face recognition, 2D and 3D liveness detection in addition to card and mobile reader.

Access Control Terminal

Fully edge capable with no processing needed at server and incorporates modern smartphone like design.

High level specs:

CPU Qualcomm 6125

RAM 4GB DDR, 64GB eMMC

Screen 5.5-inch, 800x1280

Camera 3D Depth Sensing Camera-A1

Communication Ethernet/ WiFi/Bluetooth

Other Contactless Access

Reader NFC

Interfaces Relay, Wiegand 26/34, RJ45, DC 12V, USB-C, HDMI, POE, RS485

Audio Digital Audio Speaker, Microphone

Environmental IP56 IK 5 (outdoor usage)

OS Android 10

Carnegie-Mellon University Research Sponsorship

In October, AnyVision will announce a research partnership with Carnegie Mellon University's (CMU) CyLab Biometric Research Center to develop advanced object classification and behavior recognition algorithms. As part of the partnership, Professor Marios Savvides, the founder and director of the Biometrics Center at CMU, will join AnyVision as the Chief AI Scientist to expand AnyVision's AI team led by CTO, Dieter Joecker.

AnyVision pioneered Vision AI to automate watchlist alerting, identifying security risks as well as valuable customers in real-time to personalize customer experiences and enhance an organization's physical security. This collaboration will help AnyVision address a broad range of safety-related use cases, including object detection (e.g., weapons, unattended luggage, etc.) and behavioral analysis (e.g., when someone falls down).

Over the past 10 years, more than 400 startups linked to CMU have raised more than $7 billion in follow-on funding. Those investment numbers are especially high because of the sheer size of Pittsburgh's growing autonomous vehicles cluster – including Uber, Aurora, and Argo AI –all of which are there because of their strong ties to CMU.

> **"Under the leadership of Prof. Savvides, CMU's CyLab Biometric Research Center has an impressive track record of successfully transferring AI research out of a lab environment and into reliable and scalable solutions," said Avi Golan, CEO of AnyVision.**
>
> **"Visual intelligence is in its infancy and there is so much more work yet to be done. With this partnership, we now have an elite U.S.-based AI research center that will work in concert with our existing AI teams to accelerate the development of advanced deep learning algorithms and exploration of new use cases, markets, and industries, including medical, payments, and smart cities."**

CMU has a long history in **artificial intelligence** including the creation of the first AI computer program in 1956 and pioneering work in self-driving cars, facial recognition, and natural language processing. Professor Marios Savvides is the founder and director of the CyLab Biometrics Center at Carnegie Mellon University and was named one of the "2020 Outstanding Contributors to AI" awards from the former U.S. Secretary of the Army.

His research has been focused on developing core AI and machine-learning algorithms that were successfully applied for robust face detection, face recognition, iris biometrics, and most recently, general object detection and scene understanding. The CyLab Biometrics Center has generated over 35 patents and patent publications, and over 50 unpublished patent applications to date.

AnyVision's Marquee Customers

AnyVision has helped Fortune 500 companies around the globe, across a myriad of industries, protect their employees, partners, and customers with state-of-the-art facial recognition. We have transformed how financial service, retail, stadiums, casinos, critical infrastructure, healthcare organizations and government agencies provide a safe and secure environment while adhering to the highest standards for ethical AI.

Consulting

Deloitte.
OnAccess

EY
OnWatch

Banking/Financial Services

Morgan Stanley
OnAccess

American Financial Group, Inc.
OnWatch

Fidelity Investments
OnAccess

Freddie Mac
OnWatch

BlackRock.
OnAccess

Auto

Ford
OnWatch

Mercedes-Benz
OnWatch

Sports/Stadiums

CHASE CENTER
OnAccess

SAINTS
OnAccess

OnAccess

Retail

MENARDS
OnWatch

macy's
OnWatch

Staples
OnWatch

Saks Fifth Avenue
OnWatch

Health

Cedars Sinai
OnWatch

CDC Centers for Disease Control and Prevention
OnWatch

SHEBA Tel HaShomer City of Health
OnWatch

Telco

verizon
OnAccess

vodafone
OnWatch

Casinos/Gaming

CAESARS ENTERTAINMENT
OnWatch

SOLAIRE RESORT & CASINO
OnWatch

MGM RESORTS INTERNATIONAL
OnWatch

Sands LAS VEGAS SANDS CORP.
OnWatch

Energy

bp
OnWatch

The Global Leader in Vision AI

We are a core visual AI platform company with unmatched expertise in AI and deep learning, and a proven track record of designing, developing and deploying software to solve real-world problems.

$355M+ in Funding, Backed By A High-Quality Growth & Strategic Investor Syndicate

SoftBank
Lightspeed
FIFTH WALL
Qualcomm
DFJ GROWTH
IQT IN-Q-TEL
Schneider Electric

Built On The Most Cutting Edge Research, Deployed with Tier-1 Companies Around the Globe

World-leading algorithms fueled by our proprietary neural networks, designed, developed, tested and commercially deployed with leading organizations across the globe to ensure the highest accuracy in challenging, real-world conditions.

500+

Customers, including many Fortune 500

100+

Experts across four major geographies

20+

PHDs with deep knowledge and real-world expertise

Appendix A: Building Ethics and Security into our Algorithms

Privacy protections are becoming increasingly important against the backdrop of an estimated 700 million video cameras in operation throughout the world, recording and storing images of people going about their daily lives.

GDPR & Privacy Modes

AnyVision's technology includes a GDPR mode which effectively blurs all faces of people not explicitly listed on an organization's watchlist. When this feature is activated, only individuals identified on the watchlist are visible -- all other people in the field of view of the camera are blurred. Privacy Mode goes a step further as it discards all face detections of non-enrolled individuals. That is, when Privacy Mode is activated, we do not retain any data of non-watchlist individuals. This means that organizations cannot capture any metadata from non-watchlist detections which further protects the identities of bystanders. These advanced features are designed to help organizations capture and collect data on individuals that is strictly necessary for the purposes of the processing (which conforms to the GDPR principle of data minimization).

Continuous Learning

The industry around facial recognition technology is rapidly maturing due to advances in AI, ML, and deep learning technologies. Facial recognition employs machine learning algorithms which find, capture, store and analyze facial features to match them with images of individuals in a pre-existing database. But, unlike other "static" software programs that are deployed on premise and get updated periodically, leading facial recognition providers deliver continuous improvements, even after it's deployed in production environments.

As neural networks collect more real-time data and identify more potential watchlist matches (i.e., detections), the algorithms naturally improve and yield better, more accurate results. In fact, a U.S. government study found that facial recognition technology is getting better at identifying people wearing masks. This continuous learning means that customers should expect that their underlying facial recognition technology should be better tomorrow than it is today.

Skin Color

At AnyVision, we do not use the Fitzpatrick Skin Type scale in production. Instead, our AI-based technology gives us much higher resolution in the skin color as we literally capture millions of skin tone categories (vs. the six categories from the Fitzpatrick scale) that we're able to extract from video streams.

We use RGB color codes which reflect a combination of the red, green, and blue colors. The red, green, and blue use 8 bits each, which have integer values from 0 to 255. This makes 256*256*256=16,777,216 possible color values (though clearly many of these combinations would not be legitimate skin colors). For Greyscale we use values from 0 = black to 1024 = white. This type of granular color coding ensures a higher facial recognition accuracy and dramatically improves our ability to reduce demographic bias.

Data Capture & Storage

AnyVision does not collect or share user data. In fact, we don't even capture images. The reference data that our facial recognition algorithms are used to search against is created and uploaded from our commercial customers. We do not provide any such data, including scraped images from third-party sources (e.g., Google images, Instagram or LinkedIn).

The data that we capture is rendered using mathematical vectors that act as secure cryptography, preventing identity hacking even if data is stolen. AnyVision allows only the collection of data that is strictly necessary for the purposes of data processing (data minimization), and we enable our customers to automatically and/or selectively clear data from an active database at the end of a certain period, or to easily delete unused data.

Large Data Sets

AnyVision has acquired millions of images and state of the art augmentations (one video contains hundreds or thousands of still images) from a variety of sources to create a massive data set which is used for building a powerful neural network. This includes images of people from most ethnicities, skin colors, races, and genders from around the world. With 50 million images, this gives us a significant capability to address bias since our training data includes large representative populations of different demographics.

User Roles and Permissions

AnyVision software functionality provides our commercial customers with different user roles and permissions which can be applied based on the operator's role within the company. This means that operators can only see the data that is relevant to their role and need for watchlist access. Admin users can give control to specific operators who have access to certain cameras (e.g., cameras that monitor specific employee access points) or grant access to only specific groups on the watchlist (e.g., VIP customers).

Watchlist Creation & Privacy

Unlike other popular forms of facial recognition, AnyVision does not start with a pre-enrolled database of photos. AnyVision does not, for example, provide access to billions of pictures from popular websites such as LinkedIn, Facebook, and Google, or endorse this practice. Instead, watchlists are created from scratch by our commercial customers based entirely on their needs, which vary widely. Such lists may consist of known bad actors, authorized employees or even VIPs.

AnyVision encourages our commercial customers to enroll more than just one image for anyone on their watchlist. The more images of a person used for comparison, including pictures of the person captured at different angles or light levels, the better the system will detect those persons of interest in operation. We also encourage clients to add high quality reference images where the person's face and facial characteristics are used to better train the system. The AnyVision system provides feedback to our commercial customers when they attempt to upload images of poor quality to their watchlist to help increase the accuracy of our matching algorithms. If a customer does not wish to expose the personal information of anyone enrolled in a watchlist, such as their name, they can decide to use numbers as unique identifiers instead. For example, a subject can be named 123 or "subject 1" instead of their actual name "John Smith."

Real-World Conditions

AnyVision's technology can overcome the most challenging conditions – from large crowds to low light environments, extreme angles, and obscured faces. Many facial recognition systems struggle to correctly identify people under these conditions which is often the norm. Our data augmentation algorithms help improve the quality of our facial recognition in low light conditions, when the person is looking away from the camera, and in low light conditions, when the person is looking away from the camera, and in low bandwidth settings which result in compression artefacts (e.g., flickering, blurring, and speckling).

Encryption & Data Storage

Data from the camera to our servers is encrypted in transit with AES-256 bit.
When our solution runs on-premise, no data is passed over the Internet from our commercial customers to AnyVision, which effectively makes it a closed network.
There are two types of data stored by our commercial customers: the images of persons of interest uploaded which form the watchlist and all detections from the live video streams. If you recall, our technology does not store any images of faces or bodies, just mathematical vectors. The watchlist photos are uploaded and managed by our commercial customers based on their unique security needs. The detection data whether they be watchlist individuals or non-watchlist detections are managed and stored based on the retention policies of our commercial customers.

Object Recognition

AnyVision's Watchlist Alerting system can be used in "body" mode to find specific individuals based on their body or clothing worn (e.g., specific blue shirt, red backpack) instead of their faces. When a video stream is set to body mode, no facial features are collected. This technology does not include any personal identifiable information and helps commercial enterprises track individuals across multi-camera environments and facilitates after-the-fact investigations.

Setting the Right Thresholds Based on the Use Case

It may be that more false positives or false negatives occur due to the setup configurations of the system, based on the established scoring and threshold feature. The threshold is one of the most important configurations to take into consideration when setting up a facial recognition system as it is used to evaluate whether a score is high or low enough to be considered a recognition.

Depending on the size of the database, as well as whether the use case is one-to-one, one-to-many, or many-to-many, the threshold can have very different impacts. A high threshold should be used when dealing with massively large databases as our interest is to lower the false positive rate. Additionally, a high threshold should be used when dealing with one-to-one as the main purpose of this use case is to ensure that the right person is present in the video frame.

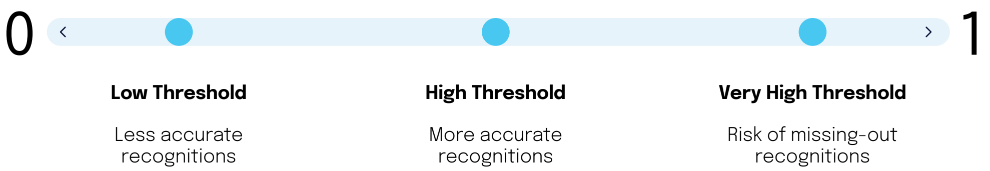
0
1
Low Threshold
Less accurate recognitions
High Threshold
More accurate recognitions
Very High Threshold
Risk of missing-out recognitions

The figure above shows the impact of low, high, and very high thresholds. When the consequences of letting a bad actor on premise is high, you will want a low threshold to ensure there are no missed watchlist detections, but this will increase the probability of possible false positives. In lower risk use cases, commercial customers can set a higher threshold level to optimize your acceptance rates.

Data Retention

Data retention of watchlist detections can be configured based on your customer's retention policy. If the organization wants to retain their detection data for X amount of time (e.g., 30 days), it will be automatically deleted from the system once that period expires.