oosto

The Rise of Ethical Facial Recognition

FOREWORD

A Clear Need for Facial Recognition Guidance

Over more than thirty years of development, facial recognition technology has demonstrated many uniquely valuable benefits to society. At the same time, facial recognition has captured today's headlines, raising many important questions about its accuracy and use. These concerns have even led to blanket bans on use of the technology in several U.S. jurisdictions, without regard to the purpose or the benefits of specific applications.

There is a clear need for steps to build public trust without eliminating the crucial benefits the technology can provide. That's why it's important to define the elements and boundaries of ethical facial recognition. Both developers and end users have a duty to ensure that advanced technologies, particularly those enabled by artificial intelligence and machine learning, are used in a responsible manner consistent with key values and with appropriate safeguards. In 2020, members of the Security Industry Association (SIA) drafted and published SIA's Principles for the Responsible and Effective Use of Facial Recognition Technology, to help guide implementation across a range of identification applications and inform development of organizational and public policies addressing its use.

Especially for safety and security applications, customers should be confident that the technology they are using is trustworthy and effective. Suppliers must strive to use the highest-performing facial recognition technology for a given application, with accuracy validated using sound methods. While misconceptions persist about the accuracy of facial recognition, the fact is the technology is getting better over time, especially as dramatic performance improvements have driven growth in recent years. For example, the top facial recognition algorithms currently tested by NIST, the National Institute of Standards and Technology, are over 99% accurate across black male, white male, black female and white female demographics.

Facial recognition is critical to the security field because it enhances capabilities of solutions like video security, access control and identity management systems that help customers secure their facilities, employees and patrons against the threat of violence, theft or other harm. SIA strongly encourages leadership among our members in developing guidance, tools and other resources needed to support the responsible, ethical and effective use of the technology, consistent with SIA's Principles.

All the best,

Jake Parker
Senior Director
of Government Relations,
Security Industry Association

About SIA

The Security Industry Association (SIA) is a nonprofit trade association representing more than 1,100 businesses providing a broad range of security products and services in the United States and internationally. Our members include many of the leading developers of facial recognition technology; companies offering products that incorporate this technology in a variety of identity, security and public safety applications; and installers and integrators of these systems.

The Rise of Ethical Facial Recognition

Facial recognition technology offers tremendous benefits to society when used effectively and responsibly, but the industry has a duty to ensure that advanced technologies, particularly those enabled by artificial intelligence and machine learning, are used in a responsible manner consistent with our values and with appropriate safeguards.

The Many Benefits of Facial Recognition

The benefits of facial recognition are proven and growing, through a wide range of vastly different applications. For example, in the United States, the technology has been used for more than a decade to detect identity fraud that fuels other criminal activity. In other parts of the world, facial recognition has been used to identify known hooligans at stadiums and used to help find and rescue human trafficking victims, thwart potential terrorist attacks, solve hate crimes and crack cold cases.

As a means of digital identification, facial recognition can be a vital enabler for commerce by improving security, protecting identity, safeguarding our personal devices and enabling touchless access and a seamless travel experience. In the security field, facial recognition is critical, as it enhances the effectiveness of security and life safety systems to help our customers keep their facilities, employees and patrons safe.

The Growing Need for Ethical Facial Recognition

Facial recognition. It's a convenient way to unlock your phone or computer, but it's becoming more and more controversial. And the concerns have been voiced on a global scale.

Facial recognition tools are coming under intense scrutiny in Europe, with privacy watchdogs using the General Data Protection Regulation to regulate the fast-developing technology, rather than waiting for a dedicated EU law on AI to be passed. The draft Artificial Intelligence Act aims to impose strict limitations on "high-risk" applications of the technology, including market-entry authorization requirements, while applying a lighter touch to less risky uses. While there is some debate over what constitutes a high-risk application, facial recognition will certainly be included.

In fact, a 2020 draft of an EU Commission White Paper on AI caused a stir by suggesting a possible moratorium on the use of facial recognition technology for three to five years, but dropped that suggestion in the final version.

In the United States, San Francisco, Boston and Portland, among several other U.S. cities, have actually banned the use of facial recognition by police and other agencies. These policies have stemmed from negative public perceptions about technology itself, as well as how it's used and applied especially within the law enforcement community.

In this eBook, we will explore how facial recognition works, the different types and use cases of facial recognition, its ethical challenges and how Oosto is addressing those challenges to prevent retail crime, protect schools from threats, alert security when a bad actor enters a facility, control access to sensitive areas and make air travel safer and more convenient.

How Facial Recognition Works

How does this facial recognition technology even work, and why is it sometimes seen as a controversial thing? Facial recognition technology can identify a person from a photo or video. It compares selected facial features to faces within a database for similarity, and can analyze facial textures and shapes to match them.

Historically, facial recognition systems used computer algorithms to pick out specific, distinctive details about a person's face. These details, such as distance between the eyes or shape of the chin, are then converted into a mathematical representation and compared to data on other faces collected in a face recognition database.

But how this process is actually performed varies widely by solution provider, so it's important to understand how a vendor performs facial recognition in order to assess whether it's following ethical best practices.

More recently, leading solution providers leverage neural networks to recognize specific faces based on deep learning. These networks assign a unique mathematical vector to a specific person's face. With neural network based models, vectors cannot be reverse engineered to reconstruct a person's face. This adds an additional layer of privacy and security that does not exist with older facial recognition models.

How Oosto Performs Facial Recognition

Oosto's technology is the product of years of research and is modified for a variety of use cases and implementations. Regardless of the specific product being used by Oosto business partners and customers, the core technology and the way the system detects and recognizes faces is the same.

At Oosto, facial recognition is performed using video streams in five fundamental steps:

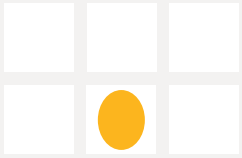

1. Face Detection

Our AI platform works by analyzing each video frame to determine if there is a face within it. If there is a face, then there is a detection. Our neural networks must also detect faces in adverse conditions such as when people are not looking directly at the camera or have a mask on. In other words, the first step answers the question, is there a face in this frame and where is the face within the frame. This is where we place the bounding boxes and crop the faces from the frames?

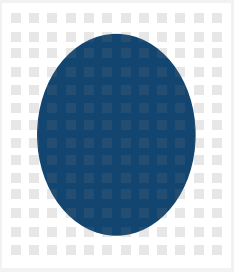

2. Landmark & Quality

Face landmark detection is a computer vision task where we want to detect the key points from a human face. Facial landmarks are used to localize and represent important regions of the face, such as: mouth, eyes, eyebrows, nose, and jawline, etc. Landmark detection also analyzes these features. For example, we can use the key points for detecting a human's head pose, position and rotation. We analyze the face to determine if the picture quality is sufficient enough to perform a high quality vector generation. Our neural networks have been trained to identify landmarks even when the face is partially occluded, captured in profile, or wearing a mask.

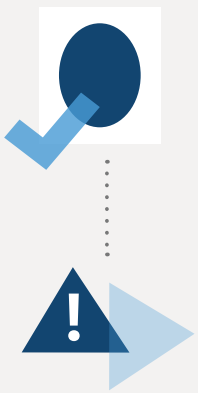

3. Vector Generation & Extraction

In this phase, the facial features of the face detected in the first step is extracted and constructed into a mathematical vector. Vector generation is a continuous process that happens within every frame of the video. If the same face appears in multiple frames, our neural nets will generate multiple mathematical vectors. Each vector will be unique, but they will not be a 100% match to each other because the pose, lighting, and other conditions change from frame to frame. It's also important to note that each vector is unique and cannot be reverse engineered — that is, a vector cannot be converted back into a picture which provides an additional layer of security.

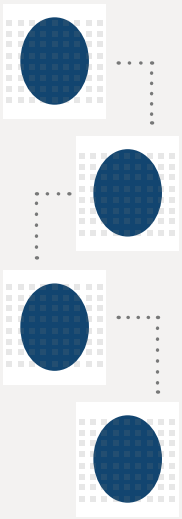

4. Tracking

A track occurs if the mathematical vector extracted from the frame matches the mathematical vector extracted from some proceeding frames. In other words, the system uses the mathematical vector to determine if the same face was in the field of view in consecutive frames. Each frame contains a unique vector for a person's face and the cumulative set of frames (i.e., we are aggregating the vectors) are used to create a single mathematical vector which is assigned to that person. This track encapsulates a number of different vantage points and head/face positions which is highly robust (since we're using a series of frames to compute this unique vector).

We can cut off a track after some specified period of time or number of frames (which is customizable) based on a particular use case. For example, with access control, we may opt for a shorter track since we need to make a faster verification decision in order to enable authorized personnel to quickly enter a building. By aggregating multiple vectors, Oosto is better able to detect people in extreme conditions, such as high contrast, profile view, low resolution, and even disguised subjects. Other facial recognition companies do not create a composite track; instead they capture a single frame to create their track. But our ability to aggregate tracks across multiple frames improves the quality and accuracy of our algorithms.

5. Facial Comparison

Once the extraction process is complete, it will automatically search for this face (more specifically, a mathematical vector) against the given company's database of enrolled images (which have also been converted into mathematical vectors). As we process each video frame, we are continuously creating new mathematical vectors from the detections and comparing them in real -time to all the other vectors in the organization's watchlist. We establish similarity thresholds to determine what we consider a match and these thresholds can be adjusted based on a particular use case.

This step answers the question, is this person in my predefined watchlist? If so, an alert is sent from the system. If not, the "non-detection" associated with the vector remains in the system, but no alert is sent. NOTE: When *Discard Detections* mode is activated, none of this information remains in the system as the data of non-watchlist people is deleted. This process happens simultaneously in multi-camera environments on all faces in the field of view.

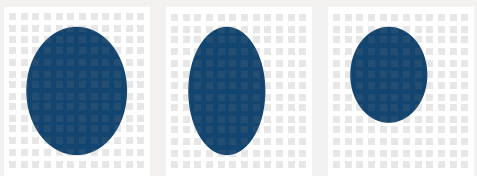

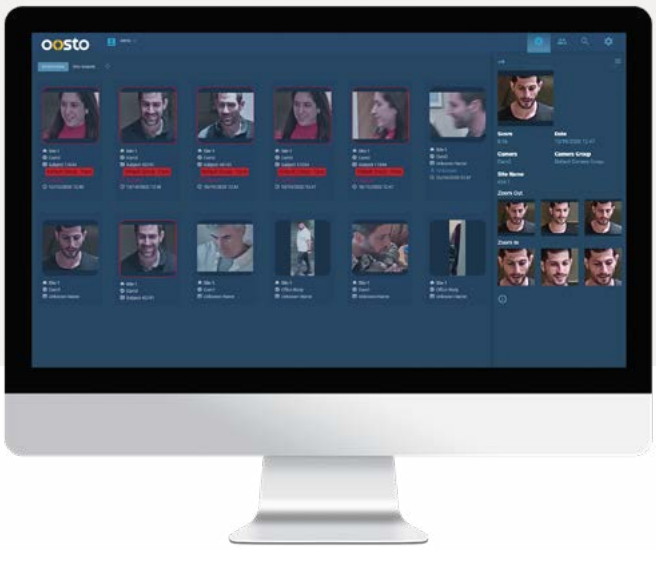

Ensuring Maximum Accuracy and Minimum Error

In order to truly determine the accuracy of facial recognition systems, they must be measured. Oosto uses the following key performance indicators (KPIs) to evaluate the accuracy of facial recognition system setups and placements. Based on the results of these measurements, stakeholders should change and adapt their system accordingly.

The four KPIs are:

True Positive

The system correctly identified a subject from the database.

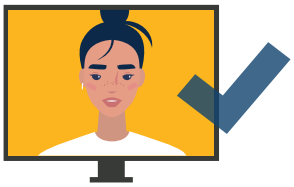

True Negative

The system, correctly, did not identify a person as a subject from the database.

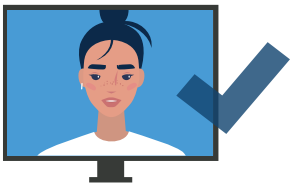

False Postive

The system incorrectly identified a person as a subject from the database.

False Negative

The system failed, incorrectly, to identify a person as a subject from the database.

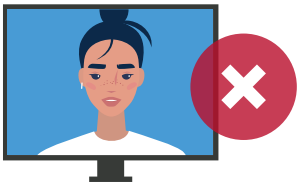

As seen above, having many true positives and true negatives are great. It means the system is working properly. However, there is always a tradeoff, where at extremely high true positive rates, false positives will also increase, and vice versa. If too many false positives and false negatives occur, the configurations of the system can be rebalanced by changing the threshold levels based on the specific use case. What's important is measuring and calibrating the number of misdetections and setting the thresholds that provide the optimal balance.

The Different Types of Facial Recognition

While face recognition has been around in one form or another since the 1960s, recent technological developments have led to a wide proliferation of this technology. While mobile phone access might be the most recognizable way face recognition is being used, it is being employed in a wide range of applications including preventing crime, protecting events and making air travel more convenient. The table below outlines two of the more popular use cases and how the underlying facial recognition technology differs.

Use Cases	Law Enforcement	Commercial
Description	The police get a picture of a suspect from a crime scene and want to find out: “Who is the person in the picture?”	Facial recognition technology is used in commercial settings to answer the question: “Is the person in the picture or video part of a pre-determined watchlist established by the user?”
How it Works	A picture is captured either statically or dynamically from a CCTV camera and then compared to a database of pictures available to law enforcement. This can range from smaller groups of arrest photos to larger database with millions or even billions of pictures.	A watchlist of people (e.g., felons, shoplifters, employees or VIPs) is created and only individuals who match their facial characteristics are identified.
Use Cases	• Crime suspect identification • Victim identification • Track down at-large criminals • Find missing people	• Touchless access control • Watchlist screening • VIP alerting • Time & attendance reporting
Reference Database	Massive	Limited (targeted)

Terminology Used in Oosto Systems

The following Oosto terms are important to understanding how we perform ethical facial recognition.

Detection	The appearance of a human face within a video frame or digital image.
Recognition	The identification of a person detected in a video frame or digital image.
Mathematical Model	A unique mathematical vector that is associated with the appearance of a face within a video frame based on a set of facial features.
Score	A value (between 0 and 1) given to each detection that represents the strength of the match between the detection image and a reference image provided from a database.
Threshold	A configurable value that specifies the minimum score required for a detection to be considered as a match.

The Ethical Challenges of Facial Recognition

Some facial recognition algorithms can be quite accurate overall, yet still show disparities in their performance across demographic groups. There are legitimate concerns that this could have an adverse impact on our daily lives. This bias could manifest itself in the form of inaccurate matching, or in the technology working well for some people but not others.

While we acknowledge that some facial recognition algorithms show unacceptable levels of racial bias, a deeper look reveals that market-leading algorithms, including ours, are far less prone to this bias, and in fact can prevent bias in security settings.

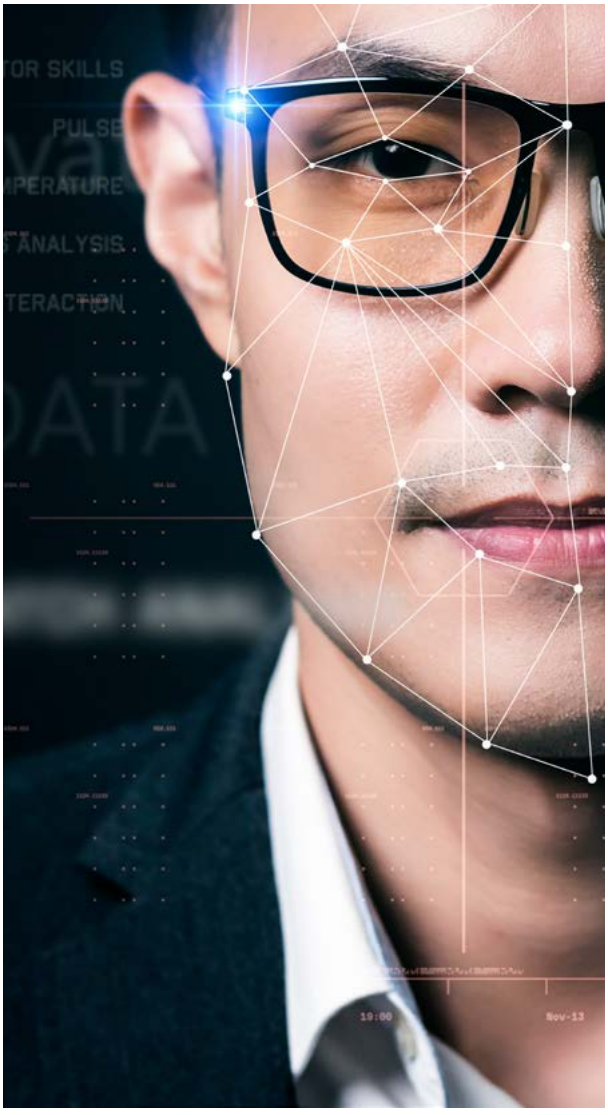

Consider these stats:

- Recent studies performed by NIST, the National Institute of Standards and Technology, looked at the top facial recognition algorithms around the world and found "no good evidence for a difference in the face detection or failure-to-enroll rate between the African-American and Caucasian cohorts."
- Facial recognition accuracy, as a whole, has significantly improved over time. As of April 2020, the top facial identification algorithms had an error rate of just 0.08%, compared to 4.1% in 2014. This improvement was uniform across all races.
- Oosto held the Fair Face Recognition Workshop and Challenge in late 2020, which evaluated the accuracy and bias of facial recognition algorithms with regards to gender and race on 1:1 face verification. The challenge found that the top-10 facial recognition teams exceeded 99.9% accuracy and were "able to minimize bias to the point where it was almost negligible."

Facial recognition, by its very nature, is intended to minimize racial bias by taking much of the guesswork and "gut feeling" out of security operations. Without the aid of facial recognition, security teams may spend their time stalking and investigating individuals who appear to be "suspicious." Unfortunately, unconscious bias may guide security guards to label someone as "suspicious" whereas facial recognition technology should be significantly more objective (if the training models were based on representative populations spanning different genders, races, ethnicities, ages, and facial poses).

Insufficient Training Data

AI training data is the information used to train a machine learning model. Machine learning models use the training dataset to learn how to recognize patterns and apply technologies such as neural networks, so that the models can make accurate predictions when later presented with new data in real world applications. When it comes to AI, size matters. The larger and more representative the training data set, the better its ability to perform consistently and minimize demographic bias.

Source of Data

When companies don't have enough of their own data to build robust models, they often turn to third-party data sources to backfill this gap, and these purchased datasets or data sets from publicly scraped websites (e.g., Facebook, LinkedIn, Instagram) can introduce unintentional bias. For example, a dataset of facial images captured under perfect lighting conditions with high-resolution cameras is not representative of the facial images that are captured in the real world. Not surprisingly, AI models built on unrealistic models will struggle with faces that contain blur or glare or were captured in dim lighting. Algorithms that were built with real-world production data, on the other hand, will contain real-world imperfections. As a result, these AI models are more robust and less susceptible to demographic bias.

Image Quality

The performance of facial recognition relies on a variety of environmental factors. This includes the quality of the cameras, how they're positioned, and the surrounding lighting. These factors can have a significant impact on the image quality and the facial recognition software's ability to reliably detect faces within the video frames and compare those faces to individuals on the watchlist.

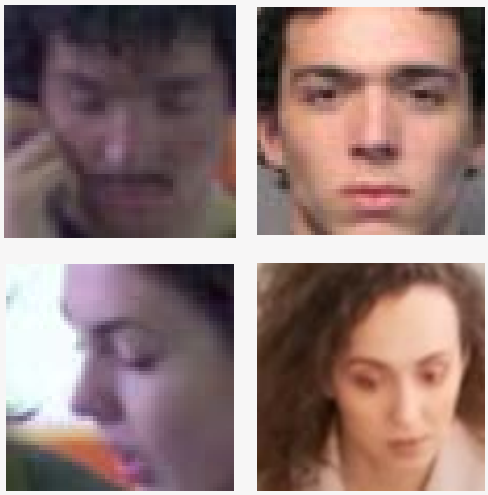

User Consent

For facial-recognition algorithms to work well, they must be trained and tested on large data sets of images, ideally captured many times under different lighting conditions and at different angles. In the 1990s and 2000s, scientists generally got volunteers to pose for these photos — but many facial recognition solution providers often collect facial images from publicly available data sources.

Skin Tones

Some facial recognition technologies rely on Fitzpatrick Skin Type — a six-color scale which dermatologists have used since the 1970s. Tech companies now rely on it to categorize people and measure whether products such as facial recognition systems perform equally well across skin tones. At Oosto, when we train our neural networks, we do not explicitly consider skin tones. We leverage very large datasets that contain a wide variety of skin tones — that go well beyond a six-color scale — which helps minimize demographic bias.

Data Labels

In most AI projects, classifying and labeling facial data sets takes a fair amount of time and subject matter expertise. Initially, humans are used to label the data (e.g., creating the bounding boxes around the face) to establish the ground truth. Then, neural networks leverage that ground truth to create the algorithms moving forward. If the wrong labels are used when tagging individual facial data, the AI models will bake that information into the algorithms which will make the models less accurate and more subject to failure over time. When facial recognition companies purchase third-party data sets, some of this data has already been labeled and these pre-populated labels can introduce bias into the models. At Oosto, we train and label our own data and do not train our models based on open source data sets.

Quality of the Data Scientists & Researchers

Reducing bias is also about the people who are developing the AI algorithms and tagging the datasets. When the environmental conditions are challenging, the greater the need for "explainable AI" which requires an AI team that is well versed in this nuanced area of artificial intelligence. In other words, it's fair to ask about the experience and composition of the AI team:

- Experience in the field
- Experience with difficult or very specific scenarios based on real world challenges (e.g., bad lighting, non-direct facial poses, etc.)
- Experience work with state-of-the-art technologies, applications and best practices

Lack of Operational Due Diligence

Police say facial recognition technology has been instrumental in helping crack some tough cases, but in the last year, there have been several claims of wrongful arrests. In many of these cases, the wrongful arrests were a result of poor process vs. shortcomings within the facial recognition software. In one notable case, Robert Williams was wrongfully arrested in Detroit.

Detroit's Chief of Police, James Craig, explained the process: "Once we insert a photograph, a probe photo — into the software — the computer may generate 100 probables. And then they rank these photographs in order of what the computer suggests, or the software suggests is the most likely suspect. It's then up to an analyst to compare each of those possible matches to the suspect and decide whether any of them should be investigated further." According to Craig, "This old drivers' license photo of Robert Williams popped up. We learned it was ranked 9th among 243 possible matches. An analyst then sent it to Detroit police as an investigative lead only, not probable cause to arrest." This lack of due diligence obviously reflects poorly on facial recognition, but it highlights the importance of human review and investigation when applying this powerful technology. Craig added: "But it wasn't facial recognition that failed. What failed was a horrible investigation." (source: Police departments adopting facial recognition tech amid allegations of wrongful arrests).

Incorporating Ethics When Using Facial Recognition

Here are four best practices for ensuring ethical facial recognition:

Training
Train commercial customers to properly use Oosto's technology and system.

Tools
Provide privacy and safety features that are customizable for everyone, at every level.

Verification
Employ human oversight when conducting facial recognition systems, especially in highly populated environments.

Permissions
Limit system access to a specific number of certified personnel through various roles and levels of permission.

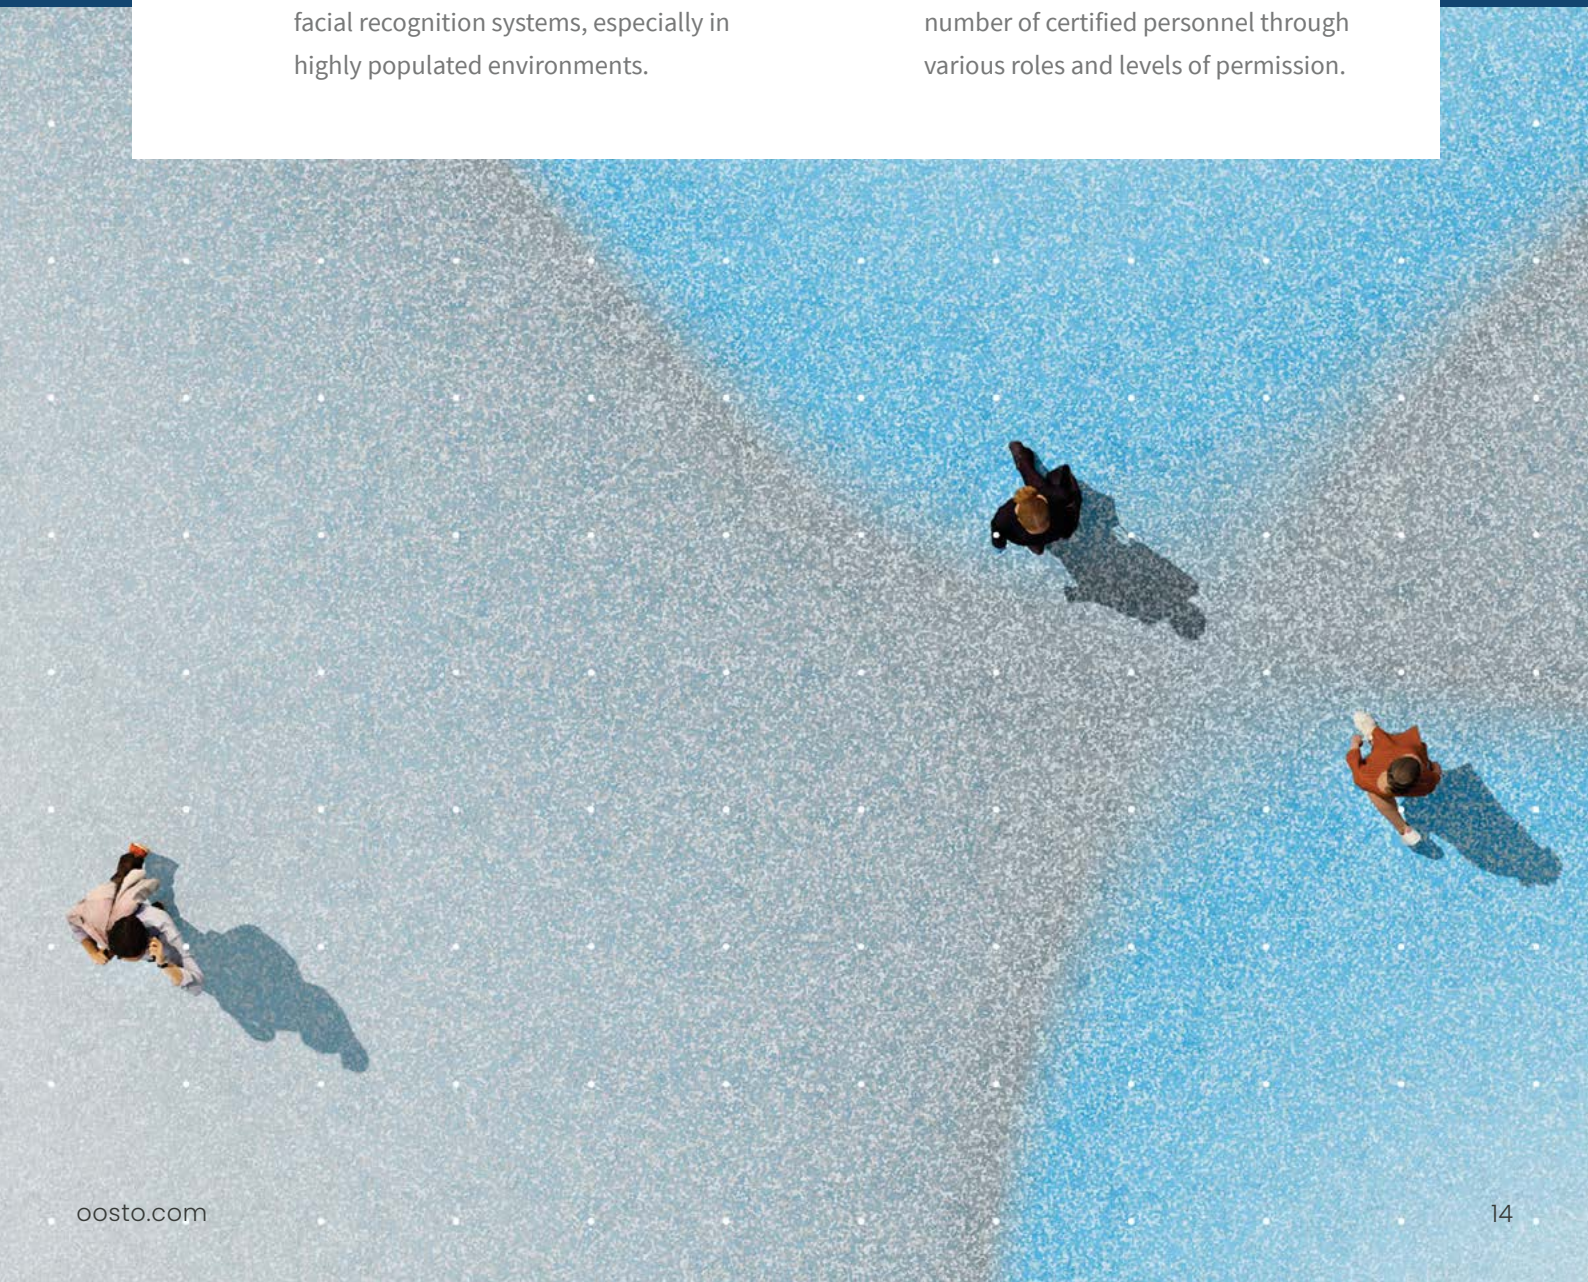

How Oosto Addresses these Challenges

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.

Data Capture & Storage

Oosto does not collect or share user data on our servers. In fact, we do not even have access to the data captured by our customers. The watchlists are created by our customers and not furnished by Oosto. We do not provide any such data, including scraped images from third-party sources (e.g., Google images, Instagram or LinkedIn). Moreover, Oosto does not use our customer's data to train our neural networks, nor do we share our data with our customers. There is a clear line in terms of who owns the data, how that data is used and where that data is stored.

Ethics Review Board

Oosto employs an internal Ethics Board which reviews every potential sales opportunity to ensure that our technology is being used for ethical purposes. The Ethics Board also considers the regional differences and compliance mandates to ensure that our technology falls within those designated guard rails. If the prospective customer is a government or law enforcement agency, we carefully review the use case and will not allow our technology to be deployed unless it meets our strict guidelines for ethical use.

Large Data Sets

Oosto 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 the vast majority of ethnicities, skin colors, races, and genders from around the world. This gives us a significant capability to address bias since our training data includes large representative populations of different demographics.

Watchlist Creation & Privacy

Unlike other popular forms of facial recognition, Oosto does not start with a pre-enrolled database of photos. Oosto 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 particular needs, which vary widely. Such lists may consist of known bad actors, authorized employees or even VIPs.

Oosto 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 which improves our ability to accurately match customers/visitors against the customer's predefined watchlist. The Oosto system actually 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."

User Consent

According to SIA's Principles for the Responsible and Effective Use of Facial Recognition Technology: "Organizations should provide reasonable notice to individuals who, by continuing a course of action, will make their image subject to facial recognition analysis by the organization, unless public safety considerations make this infeasible. Enrollment of an image in a facial recognition system for physical security, safety, fraud prevention or asset protection purposes should be guided by easy-to-understand written policies governing the criteria and human review process by which the enrollment is approved. Such implementations must also respect the reasonable expectations of privacy held by customers and individuals whose images or information are captured by security devices." You can receive consent by executing the solution on an opt-in basis and educating visitors on how their image is captured, used and permanently deleted. This will help you effectively promote privacy and customer reassurance every step of the way. In fact, this type of consent is starting to be mandated by local and regional regulations.

A new law in New York requires commercial establishments (including retail stores, places of entertainment, restaurants, food trucks, and other food and drink establishments) that use biometrics in order to identify their customers to post a clear and conspicuous sign notifying customers of the biometric collection activity.

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 Oosto, 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.

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. We also can limit the amount of time an individual remains on an organization's watchlist. This means the suspect will automatically be removed from the watchlist after the specified retention period.

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 in order to match them with images of individuals in a pre-existing database. Leading facial recognition providers deliver continuous improvements, even after it's deployed in production environments, and these new versions of the models are regularly deployed which reflect the latest ML models.

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 can expect better quality and accuracy from our neural networks with each subsequent release.

Advanced Privacy Options

Oosto offers advanced privacy settings including Face Blur and Discard Detections designed to protect the identities of innocent individuals not on the watchlist. The Face Blur option effectively blurs all faces of people (on video playback) not explicitly listed on an organization's watchlist. When this feature is activated, only individuals identified (i.e., individuals subject to the selected detection) on the watchlist are visible — all other people in the field of view of the camera are blurred. This functionality can also be applied to exported videos.

The Discard Detections mode goes a step further as it discards all face detections of non-enrolled individuals. When the Discard Detections option is activated, Oosto does 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. Importantly, operators cannot even search for these non-watchlist individuals in the system when the Discard Detections mode is activated.

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).

Real-World Conditions

Oosto's technology is able to 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).

User Roles and Permissions

Oosto 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).

Object Recognition

Oosto'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, existence of a 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.

Terms of Service

Oosto believes all technology products, including facial recognition, must only be used for purposes that are lawful, ethical and nondiscriminatory. We recognize that facial recognition has the potential to be misused if placed in the wrong hands, and that we have an inherent responsibility to ensure our technology and products are used properly. As stated in our end-user license agreement, all customers are prohibited from using the technology for inappropriate, improper or unlawful purposes.

Setting the Right Thresholds Based on the Use Case

Low Threshold
Less accurate recognitions

Medium Threshold
More accurate recognitions

High Threshold
Risk of missing-out recognitions

It may be that more false positives or false negatives occur due to the setup configurations of the system, especially with regard to the 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.

The figure above shows the impact of low, medium, and 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.

Conclusion

Public and private-sector uses of facial recognition technology are extremely varied, including identity verification, loss prevention, physical security and investigative applications. Law enforcement agencies use facial recognition to identify suspects by comparing images with criminal records and other databases. They've also used the technology to find missing children, combining facial recognition with ageing software to predict how children would look several years on and find them even when they've been missing for years.

Facial recognition software can be an effective preemptive measure against organized retail crime. Business owners use the software and security cameras to identify known or suspected thieves, and notification the technology is in use, as well as the presence of the cameras themselves work to deter theft in the first place. If a business does end up getting stolen from, the software can also help identify and track the thieves throughout a property. Facial recognition also helps improve safety and security in non-retail spaces, like airports and banks. It's been a regular part of airport security screening for years. Similar to identifying criminals that come into shops, the software has helped identify criminals and potential threats to airlines and passengers.

But, facial recognition has also been the subject of plenty of controversy, especially in the aftermath of the shooting of George Floyd, which drew attention to racially motivated police brutality in the U.S. and issues of trust between law enforcement and communities. That's why it's so crucial to understand the underlying technology and how their facial recognition process actually works. It's important to understand how AI driven technology and deep learning algorithms were developed and the efforts your solution providers are taking to minimize demographic bias. It's equally important to know how the solution can improve over time and helps make accurate identifications, which ultimately minimizes unnecessary encounters with police.

At Oosto, we are committed to integrating the highest level of transparency and ethics into the development and performance, our facial recognition products and related processes, as we work with our customers to ensure they are used effectively and accurately in ways that benefit society.

Hopefully, this eBook helped to define essential elements of ethical facial recognition and the key considerations for deploying a fair, accurate and unbiased solution which adheres to those principles.

About Oosto

Oosto is the world's leading developer of facial, body, and object recognition platforms, powered by cutting edge artificial intelligence, machine learning, and deep neural networks. Oosto believes all technology products, including facial recognition, must only be used for purposes that are lawful, ethical and nondiscriminatory.

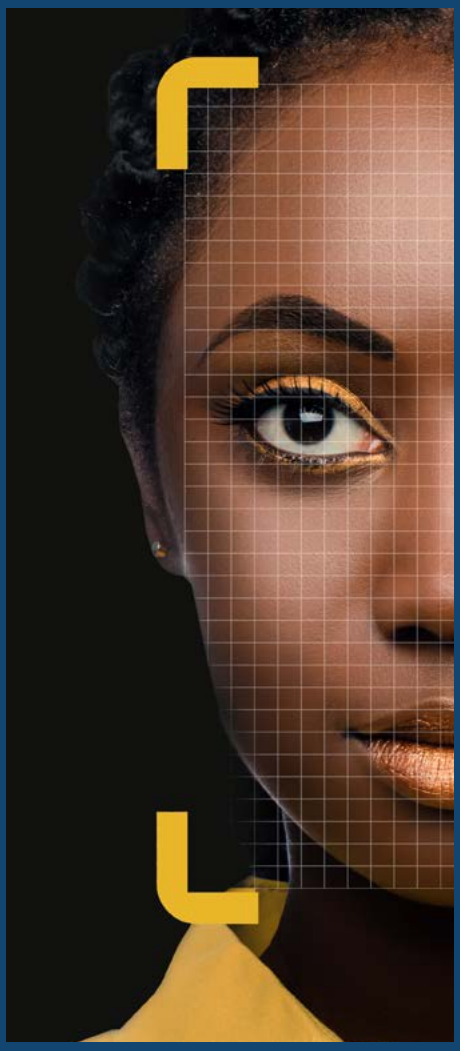

Our Commitment to Ethical Facial Recognition

Commitment to the ethical use of artificial intelligence sits at the heart of everything we do at Oosto. From designing balanced data sets and creating game-changing privacy features to supporting commonsense policies surrounding use of AI-driven technologies, we understand that as pioneers of responsible facial recognition, we are accountable for laying the foundations for a safe and ethical future.

Our Six Principles of Artificial Intelligence Ethics

Oosto understands the great value that its technology and systems can provide to society. At the same time, we recognize that powerful technology has the potential to be misused if placed in the wrong hands. We have an inherent responsibility to ensure that our technology and products are used properly. Accordingly, Oosto has adopted the following six principles for ethical facial recognition:

Fairness:

Our software must be deployed in a manner that reflects a commitment to treat all people fairly.

Transparency:

We shall communicate the capabilities and limitations of our software to our respective partners and customers.

Accountability:

We shall ensure the operation of our software is subject to human control, specifically for uses that may affect people in consequential ways.

Non-Discrimination:

Our software must not be used for unlawful discrimination.

Notice and Consent:

We instruct our partners and customers to provide adequate notice and secure consent in the deployment of our software.

Lawful Surveillance:

We advocate for lawful surveillance and will not allow the deployment of our software in scenarios that we believe will undermine this risk.