



Transparency Note: Azure Face service

Updated 9/8/22

Table of Contents

What is a Transparency Note?	3
The basics of Azure Face service	3
Key terms	3
Face Functions	5
Limited Access to Azure Face service	5
Limited Access commercial use cases	5
Limited Access public sector use cases.....	6
Considerations when using Azure Face service.....	7
Characteristics, limitations, and best practices for improving accuracy	8
Defining accuracy	8
How accurate is the Face model?.....	8
Tradeoffs	9
Recognition confidence threshold tuning	9
Why choose a recognition confidence threshold less than one?.....	9
How should a recognition confidence threshold be selected?	10
Example of a scenario minimizing false positives.....	10
Example of a scenario optimizing true positives.....	10
Best practices for improving accuracy.....	11
Plan for an evaluation phase.....	11
Face size.....	11
Face orientation.....	12
Control image capture environment.....	12
Plan for variations in subject appearance and behavior	13
Design the system to support human judgment.....	14
Use multiple factors for authentication	15
Use the latest models.....	15
Use the provided quality attributes to provide user feedback and determine if the image is good enough for facial recognition	15
Learn more about responsible AI	16
Learn more about Azure Face Service	16
Contact us	16
About this document	16

What is a Transparency Note?

An AI system includes not only the technology, but also the people who will use it, the people who will be affected by it, and the environment in which it is deployed. Creating a system that is fit for its intended purpose requires an understanding of how the technology works, its capabilities and limitations, and how to achieve the best performance.

Microsoft's Transparency Notes are intended to help you understand how our AI technology works, the choices system owners can make that influence system performance and behavior, and the importance of thinking about the whole system, including the technology, the people, and the environment. You can use Transparency Notes when developing or deploying your own system, or share them with the people who will use or be affected by your system.

Microsoft's Transparency Notes are part of a broader effort at Microsoft to put our AI principles into practice. To find out more, see [Microsoft AI principles](#).

This Transparency Note is part of our effort at Microsoft to implement our [Facial Recognition Principles](#), which set out how we approach the development and deployment of facial recognition technology. We encourage you to use the principles to guide your development efforts as you use this technology.

The basics of Azure Face service

Accessible through Azure Cognitive Services, Face detects, recognizes, and analyzes human faces in images using pre-trained machine learning models that have been developed by Microsoft. Developers can integrate Face functions into their systems without creating their own models.

When used responsibly, facial recognition is an important and useful technology that can improve efficiency, security, and customer experiences. Face is a building block for creating a facial recognition system.

Key terms

Facial verification	A "one-to-one" matching of a face in an image to a single face from a secure repository or photo to verify they are the same individual, using unmanipulated images. An example is a banking app that enables users to open a credit account remotely by taking a selfie and taking a picture of a photo ID to verify their identity.
Facial identification	"One-to-many" matching of a face in an unmanipulated image to a set of faces in a secure repository. An example is a touchless access control system in a building that replaces or augments physical cards and badges in which a smart camera captures the face of one person entering a secured door and attempts to find a match from a set of images of faces of individuals who are approved to access the building.
Facial recognition	A term that captures both Face Identification and Face Verification scenarios.
Facial detection	Finds human faces in an image and returns bounding boxes indicating their locations. Face detection models alone do not find individually identifying features, only a bounding box marking the entire face.

Facial redaction	Redaction enables modification of images or videos in order to blur or block faces of individuals.
Bounding box	A box drawn around the location of a face in the photo in response to Face Detection calls.
Template	A face template is a unique set of numbers generated from an image that represents the distinctive features of a face. The images themselves – whether enrollment or probe images (see below) – are not stored by Face, and the original images cannot be reconstructed based on a template.
Enrollment	Enrollment is the process of enrolling photos of individuals for template creation so they can be recognized. High-quality photos yield high-quality enrollment templates.
Person ID	When a person is enrolled to a verification system used for authentication, their template is also associated with a primary, randomly-generated identifier ¹ called the Person ID that will be used to determine which template to compare with the probe image (see below).
Probe image	A probe image is an image submitted to a facial recognition system to be compared to enrolled individuals. Probe images are also converted to probe templates to compare to enrolled templates. All images are immediately deleted once they are converted to templates
Recognition confidence score	When a probe image is queried to the Face Verification or Identification API, the API will return a recognition confidence score of whether two faces match in the range of [0, 1], such as 0.6. This is not the same as the percent likelihood that two faces match (i.e., a 0.9 recognition confidence score does not mean there is a 90% chance that the two faces match).
Recognition confidence threshold	The minimum confidence score required to determine whether two faces belong to the same person based on the recognition confidence score. For example, if the confidence threshold is 0.5 and the recognition confidence score returned from a probe image query is 0.6, then the two faces are considered a match.
Candidate list	For Face Identification scenarios, a candidate list is the list of faces with scores above the recognition confidence threshold.

¹ Face does not store primary identifiers, such as customer IDs, alongside facial templates. Instead, Microsoft associates stored facial templates with random GUIDs or globally unique identifiers. System developers can associate the GUID generated by Microsoft with an individual's primary identifier to support verification of that individual.

Face Functions

Face Detection (“Detection” / “Detect”) answers the question, “Are there one or more human faces in this image?” Detection finds human faces in an image and returns bounding boxes indicating their locations. Face detection models alone do not find individually identifying features, only a bounding box. All other functions are dependent on Detection: before Face can identify or verify a person (see below), it must know the locations of the faces to be recognized.

Face Detection Attributes: The Detect API can also optionally be used to analyze attributes about each face using additional AI models, such as pose and facial landmarks like eye or nose position. The attribute functionality is completely separate from the verification and identification functionality of Face. The full list of attributes is described in the service [concepts](#). The values returned by the API for each attribute are predictions of the perceived attributes and are best used to make aggregated approximations of attribute representation rather than an individual assessment.

Face Verification (“Verification” / “Verify”) builds on Detect and addresses the question, “Are these two images of the same person?” Verification is also called “one-to-one” matching because the probe image is compared to only one enrolled template. Verification can be used in identity verification or access control scenarios to verify a picture matches a previously captured image (such as from a photo from a government issued ID card). For more information, see the [Verify API](#) reference documentation.

Face Identification (“Identification” / “Identify”) also starts with Detect and answers the question, “Can this detected face be matched to any enrolled face in a database?” For this reason, identification is also called “one-to-many” matching. Candidate matches are returned based on how closely the probe template of the detected face matches each of the enrolled templates. For more information about facial identification, see the [Identify API](#) reference documentation.

Face Find Similar (“Find Similar”) also builds on Detect and searches for similar looking faces from all enrollment templates. For more information, see the [Find Similar API](#) reference documentation.

Face Group (“Group”) also builds on Detect and creates smaller groups of faces that look similar to each other from all enrollment templates. For more information, see the [Group API](#) reference documentation.

For more information on functions of Azure Face service, see the [Face documentation](#).

Limited Access to Azure Face service

Azure Face service is a Limited Access service, and registration is required for access to some features. To learn more about Microsoft’s Limited Access policy visit aka.ms/limitedaccesscogservices. Certain features are only available to Microsoft managed customers and partners, and only for certain use cases selected at the time of registration. Note that facial detection and facial redaction use cases do not require registration:

Facial detection: Detect the locations and attributes of faces for accessibility (such as [Seeing AI](#)) or modern productivity.

Facial redaction: Redact or blur detected faces of people recorded in a video to protect their privacy.

Limited Access commercial use cases

The following use cases are approved for commercial contexts:

Facial verification for identity verification to grant access to digital or physical services or spaces. Such verification may be used for opening a new account, verifying a worker, or authenticating to participate in an online assessment. Identity verification can be done once during onboarding, and repeatedly as someone accesses a digital or physical service or space.

Facial identification for touchless access control to enable an enhanced experience using facial recognition, as opposed to methods like cards and tickets. This can reduce hygiene and security risks from card/ticket sharing/handling, loss, or theft. Facial recognition can assist the check-in process for accessing sites and buildings, such as airports, stadiums, offices, and hospitals.

Facial identification for personalization to enable ambient environment personalization with consent-based facial recognition that enriches experiences on shared devices. For example, hot desk screens and kiosks in the workplace and home can recognize you as you approach to provide directions to your destination or jumpstart hands-free interaction with smart meetings devices.

Facial identification to detect duplicate or blocked users to control or prevent unauthorized access to digital or physical services or spaces. For example, such identification may be used at account creation or sign-in or at access to a work site.

Limited Access public sector use cases

The following use cases are approved for the public sector:

Facial verification for identity verification to grant access to digital or physical services or spaces. Such verification may be used for opening a new account, verifying a worker, or authenticating to participate in an online assessment. Identity verification can be done once during onboarding, and repeatedly as someone accesses a digital or physical service or space.

Facial identification for touchless access control to enable an enhanced experience using facial recognition, as opposed to methods like cards and tickets. This can help reduce hygiene and security risks from card/ticket sharing/handling, loss, or theft. Facial recognition can assist the check-in process for accessing sites and buildings, such as airports, stadiums, offices, and hospitals.

Facial identification for personalization to enable ambient environment personalization with consent-based facial recognition that enriches experiences on shared devices. For example, hot desk screens and kiosks in the workplace and home can recognize you as you approach to provide directions to your destination or jumpstart hands-free interaction with smart meetings devices.

Facial identification for preservation and enrichment of public media archives to identify individuals in public media or entertainment video archives for the purposes of preserving and enriching public media only. Examples of public media enrichment include identifying historical figures in video archives or generating descriptive metadata.

Facial identification to assist law enforcement or court officials in prosecution or defense of a criminal suspect who has already been apprehended, to the extent specifically authorized by a duly empowered government authority in a jurisdiction that maintains a fair and independent judiciary, *OR* to assist officials of duly empowered international organizations in the prosecution of abuses of international criminal law, international human rights law, or international humanitarian law.

Facial identification to respond to an emergency involving imminent danger or risk of death or serious physical injury to an individual.

Facial identification for purposes of providing humanitarian aid or identifying missing persons, deceased persons, or victims of crimes.

Considerations when using Azure Face service

The use of Azure Face by or for state or local police in the U.S. is prohibited by Microsoft policy.

The use of real-time facial recognition technology on mobile cameras used by law enforcement to attempt to identify individuals in uncontrolled, “in the wild” environments is prohibited by Microsoft policy. This includes where police officers on patrol use body-worn or dash-mounted cameras using facial recognition technology to attempt to identify individuals present in a database of suspects or prior inmates. This policy applies globally.

Avoid use of facial recognition or detection technology to attempt to infer emotional states, gender identity, or age. Microsoft has retired general-purpose facial detection capabilities that were used to classify emotion, gender, age, smile, hair, facial hair, and makeup. General-purpose use of these capabilities poses a risk of misuse that could subject people to stereotyping, discrimination, or unfair denial of services. These capabilities will be carefully restricted to select accessibility scenarios such as those provided by [Seeing AI](#).

Avoid use for ongoing surveillance of real-time or near real-time identification or persistent tracking of an individual. Ongoing surveillance is defined as the tracking of movements of an identified individual on a persistent basis. Persistent tracking is defined as the tracking of movements of an individual on a persistent basis without identification or verification of that individual. Face was not designed for ongoing surveillance or persistent tracking of an individual and does not work on large-scale real-time camera streams. In accordance with our [Six Principles for Developing and Deploying Facial Recognition Technology](#), the use of facial recognition technology for the ongoing surveillance of individuals by law enforcement should be prohibited except in narrow circumstances and only with adequate protections for individual civil liberties and human rights.

Avoid use for task-monitoring systems that can interfere with privacy. Face's probabilistic AI models were not designed to monitor individual patterns to infer intimate personal information, such as an individual's sexual or political orientation.

Avoid use in protected spaces. Protect individuals' privacy by evaluating camera locations and positions, adjusting angles and regions of interest so they do not overlook protected areas such as restrooms.

Avoid use in environments where enrollment in identification or verification is not optional. Protect individuals' autonomy by not planning enrollment in situations where there's pressure to consent.

Avoid use where a human in the loop or secondary verification method is not available. Failsafe mechanisms, e.g., a secondary method being available to the end user if the technology fails, helps to prevent denial of essential services or other harms due to false negatives.

Carefully consider use in schools or facilities for older adults. Face has not been heavily tested with data containing minors under the age of 18 or adults over age 65. We recommend that customers thoroughly evaluate error rates for any scenario in environments where there is a predominance of these age groups.

Carefully consider use for healthcare-related decisions. Face provides probabilistic results like face detections, attributes, and recognitions. The data may not be suitable for making healthcare-related decisions.

Carefully consider use in public spaces. Evaluate camera locations and positions, adjusting angles and regions of interest to minimize collection from public spaces. Lighting and weather in public spaces such as streets and parks will significantly impact the performance of the spatial analysis system, and it is extremely difficult to provide effective disclosure in public spaces.





Characteristics, limitations, and best practices for improving accuracy

Because Face is a building block for creating a facial recognition system to which other building blocks must be added, it is not possible to provide a universally applicable estimate of accuracy for the actual system you are planning to deploy. Companies may share accuracy as measured by public benchmark competitions, but these accuracies depend on details of each benchmark methodology and therefore won't be the same as the accuracy of a deployed system.

Ultimately, system accuracy depends on a number of factors, including the technology and how it is configured, environmental conditions, the use case for the system, how people to be recognized interact with the camera, and how people interpret the system's output. The following section is intended to help you understand key concepts that describe accuracy in the context of a facial recognition system. With that understanding, we then describe system design choices, how they influence accuracy, and reference metrics.

Defining accuracy

The accuracy of a facial recognition system is based on a combination of two things: how often the system correctly matches a person who is enrolled in the system and how often the system correctly finds no match for a person who is not enrolled. These two conditions, which are referred to as the "true" conditions, combine with two "false" conditions to describe all possible outcomes of a facial recognition system:

True positive or correct match 	The person in the probe image is enrolled and they are correctly matched.
True negative or correct reject 	The person in the probe image is not enrolled and the system finds no match
False positive or incorrect match 	Either the person in the probe image is not <i>enrolled</i> but is matched to an <i>enrolled</i> person OR the person in the probe image is enrolled but is matched with the wrong person.
False negative or incorrect reject 	The person in the probe image is enrolled, but the system finds no match.

The consequences of a false positive or a false negative vary depending on the purpose of the facial recognition system. The examples below illustrate this variation and how choices you make in designing the system affect the experience of those people who are subject to it.

How accurate is the Face model?

Measuring the accuracy of facial recognition technology is a very difficult problem and methodologies vary across the industry. To learn about our commitment to Fairness and improving the accuracy of our AI systems, review the [Responsible AI Resources](#).

You can use the [Fairness Assessment Sample Notebook](#) to assess the fairness of face verification on your own data. It is a Jupyter notebook using the Fairlearn python package.

Tradeoffs

Recognition confidence threshold tuning

The purpose of this section is to help you understand how system configuration influences system accuracy and the trade-off between false positives and false negatives.

The recognition confidence threshold influences system accuracy and the trade-off between false positives and false negatives. It is not related to confidence intervals.

Recognition confidence score	A recognition confidence score describes the similarity between a probe template and an enrolled template. Recognition confidence scores range from 0 to 1. High recognition confidence scores indicate that it is more likely that the two images are of the same person.
Recognition confidence threshold	A recognition confidence threshold is a configurable value between 0 and 1 that determines the recognition confidence score required to be considered a positive match.

When using the [Verification](#) function for authentication, if the recognition confidence score between the probe template and the enrollment template associated with the primary identifier is at least as high as the recognition confidence threshold, Face will indicate that the probe image represents the person presenting identification.

When using the [Identification](#) function, it can be useful for a person to review a list of candidates ranked by recognition confidence scores to make a final determination. Face customers can choose how many candidate templates that reach the recognition confidence threshold will be returned in ranked order of similarity to the probe template. These matches are referred to as a "*candidate list*". Face will only return candidates with recognition confidence scores at least as high as the recognition confidence threshold. When no templates have recognition confidence scores that reach the recognition confidence threshold, no matches are returned.

Why choose a recognition confidence threshold less than one?

Setting a recognition confidence threshold lets you balance the errors between false positives and false negatives to best address your specific scenario. The overall accuracy of the system is unlikely to be 100%, and when the recognition confidence threshold is set to 1, the strictest value, all errors that occur will be false negatives: the system will return "no match" because the submitted probe template will not perfectly match any enrolled templates.

If the recognition confidence threshold is set to 0, then any probe template will match any enrollment template. Because recognition confidence scores are affected by the quality of the probe and the enrollment images, a lower recognition confidence score can indicate poor quality images, rather than less similarity between people in the images. When doing Identification, if the recognition confidence threshold is set too high, the system may not return enough candidates to find the true match.

On the other hand, a low recognition confidence threshold may return low quality matches and can reduce the efficiency and accuracy of the humans reviewing the matches.

Face has a default recognition confidence threshold of 0.5, which is a balance applicable to many identity verification applications, but you can change the recognition confidence threshold to suit each application.

How should a recognition confidence threshold be selected?

The best recognition confidence threshold for your system is based on:

- The system's purpose,
- The impact of false positives and false negatives on the people who will be subject to facial recognition,
- Whether the final judgments are made by a human, and
- How the whole system, including the experience design, supports resolution of errors.

Before selecting a recognition confidence threshold, Microsoft recommends that you, as a facial recognition system owner, collect accurately labeled evaluation data on site to determine how the recognition confidence threshold affects the achievement of your goals and affects people subject to and interpreting the output of the system.

Accurately labeled data can be compared to the output of the system to establish the overall accuracy and error rates, and the distribution of errors between false positives and false negatives. This accurately labeled evaluation data should include adequate sampling of people with diverse characteristics who will be subject to recognition so that performance differences can be understood, and corrective action taken. Based on the results of this evaluation, you can iteratively adjust the recognition confidence threshold until the trade-off between false positives and false negatives meets your objectives.

Example of a scenario minimizing false positives

Facial recognition can help users log on to access controlled applications such as a banking app. A false positive in this scenario reduces customer security because it results in an incorrect match, while a false negative could prevent the customer from accessing their account. Because the purpose of the system is security, false positives must be minimized and as a result, most errors will be false negatives (account access fails). The application developer should set a high recognition confidence threshold to minimize false positives. Because a higher confidence threshold will create more false negatives, system owners can provide a fallback mechanism, like pushing a notification to the customer's phone with an access code. The customer's experience may be less efficient in this case, but account access is not blocked, and security is prioritized.

Example of a scenario optimizing true positives

Applications with deeply involved human review may want to provide more matches because humans can manually eliminate false positives. For example, consider a photo gallery application that surfaces possible photos of the user. In this case, the application builder would choose a lower recognition confidence threshold number, leaving the user (who is also the human reviewer in this case) with room to eliminate false positives (surfaced photos that are not of the user).

Here is a list of expected theoretical false positive rates for a given confidence score based on a dataset for the recognition_03 recognition model; there may be variations in real life:

Recognition confidence threshold	False Positive Rate
0.1	1 in 10
0.2	1 in 100
0.3	1 in 1,000
0.4	1 in 10,000
0.5	1 in 100,000
0.6	1 in 1,000,000
0.7	1 in 10,000,000
0.8	1 in 100,000,000
0.9	1 in 1,000,000,000

Best practices for improving accuracy

Facial recognition technology is improving and many systems, including Microsoft's Face, can perform well even when conditions are not ideal. However, these are specific actions that you can take to ensure the best-quality results from your facial recognition system.

Plan for an evaluation phase

As discussed in selecting a recognition confidence threshold, before a large-scale deployment or rollout of any facial recognition system, Microsoft strongly recommends that system owners conduct an evaluation phase in the context where the system will be used and with the people who will interact with the system.

You should work with your analytics and research teams to collect ground truth evaluation data to:

- Establish baseline accuracy, false positive and false negative rates.
- Choose an appropriate recognition confidence threshold to meet your objectives.
- Determine whether the error distribution is skewed towards specific groups of people.

This is likely to be an iterative process with adjustments to sensor position, lighting, and other factors that influence accuracy, as discussed in this section. This evaluation should reflect your deployment environment and any variations in that environment, such as lighting or sensor placement, as well as ground truth evaluation data that represents the diversity of people who will interact with your system.

In addition to telemetry data, you may also want to analyze feedback from the people making judgments based on the system output, satisfaction data from the people who are subject to recognition, and feedback from existing customer voice channels to help tune the system and ensure successful engagement.

Face size


Making sure the face in an enrollment or probe image is sufficiently large is critical to high quality matches.

Faces are detectable when their size is as small as 36x36 pixels, but for best performance Microsoft recommends a minimum size of 200x200 pixels with at least 100 pixels between the eyes. Note that higher image resolution increases latency, but [there are ways to optimize latency](#). The maximum size allowed is 4096x4096.



Face orientation

Faces looking away from the camera may not be detected. Faces should be oriented looking towards the camera within 35 degrees for both the pitch (head tilt towards the front or back) and the yaw (head rotation to the left or to the right); the roll (head tilt to the left or to the right) doesn't matter.



Control image capture environment

Lighting and camera calibration

Pay attention to how well the details of people's faces can be seen in images.

- Capture images in appropriate lighting conditions. Is the lighting too bright, too dark? Are faces backlit? Is there too much light from one side and not enough from the other? When possible, place sensors away from areas with extreme lighting.
- Is the lighting adequate to accurately capture the details of people's faces with different skin tones?



Backgrounds

- Strive for neutral, non-reflective backgrounds. Avoid backgrounds containing faces, for instance where there are pictures of people displayed, or where people other than the person to be recognized are prominent in the photo.

Sensor placement and maintenance

- Position sensors at face level to best capture images that meet the quality specifications.
- Ensure sensors are regularly checked for dust, smudges, and other obstructions.

Plan for variations in subject appearance and behavior

Facial occlusions

Facial recognition works best when the person's entire face is visible. Faces may be partially or entirely occluded for a variety of reasons, including:

- Religion: Head wear that covers or partially obscures faces.
- Personal Protective Equipment (PPE): PPE such as face protective masks that cover or partially obscure faces
- Weather: Garments like scarves wrapped across the face.
- Injury: Eye patches or large bandages.
- Glasses: Very opaque glasses and pinhole glasses (other glasses and lenses should be fine).
- Personal style: Bangs over eyebrows, baseball caps, large facial tattoos, etc.

Enrolling occluded faces can result in errors. While it's possible for facial recognition to verify a face that has occlusions, we recommend the following mitigations for addressing occlusion challenges, in addition to sensor placement:


- A fallback method, such as a non-biometric alternative, is critical. For some people, the fallback may be the option they consistently use.
- Pay attention to challenges that people face during evaluation and deployment to identify remediations that work best for your environment.
- Use the [Recogniton 04 model](#) in the Verify, Identify, Group, and FindSimilar APIs which can recognize people wearing face coverings (surgical masks, N95 masks, cloth masks).



Significant changes in appearance

Dramatic changes in appearance, like the removal of a full beard or many years passing between enrollment and probe images (for adults), or even short periods of time between photos of children, can result in errors.

- In addition to supporting a fallback method, designing the user experience to support immediate reenrollment following a recognition failure can improve user satisfaction.
- Facial recognition systems are generally less accurate for children. Microsoft recommends using Face for recognition of people over 18. Facial recognition can be especially challenging with people 13 and younger.



Motion, blur, and extreme expressions

Blurry faces caused by fast movement, or even extreme expression (like yawning widely with eyes closed), makes recognition challenging. To address these challenges:

- Design the user experience so people understand how to provide high-quality images.
- Create an environment where people naturally face the camera and slow down. Otherwise movement could result in blurry images that would be difficult to recognize.
- Provide clear instructions for how people should behave during recognition (eyes open, mouth closed, stand still, etc.).



Biometric twins

Twins, family members, and other people who look very similar to each other will be difficult for facial recognition systems to distinguish from one another. This is another reason to support a fallback method.



Design the system to support human judgment

In most cases, Microsoft recommends using Face's facial recognition capabilities to support people making more accurate and efficient judgments rather than fully automating a process. Meaningful human review is important to:

- Detect and resolve cases of misidentification or other failures.
- Provide support to people who believe their results were incorrect.
- Identify and resolve changes in accuracy due to changing conditions (like lighting or sensor cleanliness).

For example, when using Face for admittance into a building, a trained security officer can help when the facial recognition system fails to match someone who believes they are enrolled by deciding whether the person should be admitted to the building. In this case, Face helps the security officer work more efficiently, requiring a judgment to admit someone only when the person is not recognized.

The user experience that you create to support the people who will use the system output should be designed and evaluated with those people to understand how well they can interpret the output, what additional information they might need, how they can get answers to their questions, and ultimately, how well the system supports their ability to make more accurate judgments.

Face only supports facial recognition with still images: there are no anti-spoofing countermeasures built into Face, such as depth detection or motion detection. When facial recognition is supporting human judgment and improving efficiency, this is generally not a key limitation: humans can easily detect when a person is holding up a picture to a camera.

Use multiple factors for authentication

Use Face along with one or more other identification factors when creating authentication systems, such as confirming passengers who are about to board a plane or confirming a banking transaction. As discussed above, Verification makes use of facial recognition as a second factor for identifying someone rather than a single or primary factor. Identification does not require another factor; however, Identification is a more technically difficult problem because the probe template is compared to ALL enrolled templates instead of just the template for the primary identifier associated with the probe template. It is often still possible to use other signals to support authentication when using Identification, such as narrowing the set of enrolled templates to compare by limiting the search to people who have a ticket for a specific flight. While it is not possible to choose Verification for all scenarios, Microsoft recommends Verification for uses including secure access to buildings and for other key business and security functions.

Use the latest models

Use [the latest detection and recognition models](#) in your applications. By default, the oldest detection and recognition models are used for backwards compatibility, so you need to specify the latest models in your API requests.

Use the provided quality attributes to provide user feedback and determine if the image is good enough for facial recognition

From our investigation, many issues are caused by low quality images affected by the limitations and tradeoffs described above used for facial recognition purposes. In such cases, even a human can struggle to make the correct decision. To support the capture of high-quality images, the Detect API offers recognition quality attribute which flags image quality issues that pertain to lighting, blur, occlusions or head angle in images submitted. See the face [QuickStart](#) for how to add users into a face using the quality filter and how to call Face Detect using the face client SDKs, and visit the [API Console](#) to test out the endpoint.

Learn more about responsible AI

[Microsoft AI principles](#)

[Microsoft responsible AI resources](#)

[Microsoft Azure Learning courses on responsible AI](#)

Learn more about Azure Face Service

[What is Azure Face service?](#)

Contact us

[Give us feedback on this document](#)

About this document

© 2022 Microsoft Corporation. All rights reserved. This document is provided "as-is" and for informational purposes only. Information and views expressed in this document, including URL and other Internet Web site references, may change without notice. You bear the risk of using it. Some examples are for illustration only and are fictitious. No real association is intended or inferred.

Published: 3/29/19

Last updated: 6/21/22

