

AN INTRODUCTION TO AGE ESTIMATION



→ paravision.ai

Trusted Vision AI

Introduction

Whether buying goods like alcohol or tobacco, accessing activities like gambling or gaming, or logging in to critical services like banking or telecommunications, age is a critical factor. This has been true in the physical world for many years, and with the rapid proliferation of fully online services, it is now highly relevant in the digital world. In fact, due to the inherently remote nature of online services, it is arguably more relevant than ever.

While understanding age may be critical, getting a clear understanding of age poses unique challenges for digital services. Traditionally, if you walk into a bar, the bartender will get a basic understanding of your age, only asking for ID (i.e., age verification) if you “look under 30.” But how does this work in a digital domain? If someone goes to order alcohol for delivery online, how does a service provider know that it isn’t an underage child using their parent’s credit card? Alternatively, how does a service provider assess age for a service that is youth-focused (such as online video gaming), but where there may be age limits for groups not covered by government ID cards? And how can access to age restricted services be safeguarded while also protecting privacy?

No doubt, the solution to these questions can involve a variety of layers, including identity document checks. Among these, automated age estimation (also commonly referred to as “age detection”) is emerging as a powerful tool to address some of these key questions about age verification in digital identity. Powered by the latest Vision AI technology, modern age estimation software can predict age with a high degree of fidelity, including across demographics such as age, gender, and race. In combination with liveness detection technology that helps to ensure a submitted photo is of a real, live person (as opposed to a printed photo or someone in a mask, for example), automated age detection can perform the same function as the bartender mentioned above has done for years.

In this white paper, we will explore Paravision Age Estimation, the company’s proprietary age estimation technology, including how it works and the rigorous, privacy-protecting ethical standards used for development and deployment. We will detail how Paravision is building age detection technology that is easily deployable, usable, and performant.

Table of Contents

Terminology	4
Use Cases	5
Ethical Considerations	7
Ethical Development: Datasets	7
Ethical Development: Demographic Performance	7
Ethical Deployment: Privacy	9
Ethical Deployment: Use Case Analysis	9
How Paravision Age Estimation Works	10
Face Image (Selfie) Capture	10
Age Estimation	10
Accuracy Analysis	11
Mean Absolute Error (MAE)	11
Demographic Analysis Overview	11
Comparison with Human Accuracy	12
Conclusion	13

Terminology

Age Estimation Technology

Software which assesses images and determines the estimated age of faces in the field of view. Age estimates are characterized in terms of accuracy by MAE and as described below, and are typically reported as a certain number of years old. Also commonly referred to as “Age detection” or “Age verification” technology.

Image Quality Metrics

Software which assesses images and determines if they are of sufficiently high quality to perform a given analysis (which in this case is age estimation). If an image is of low quality (for instance, low resolution, poorly lit, or blurry) it may not be possible to accurately estimate the age of faces in the field of view.

Mean Absolute Error (MAE)

This is the primary metric for reporting the accuracy of age estimation software. It is calculated by determining the absolute error of a given sample—how many years older or younger the system is reporting a person to be as compared to their actual age—and calculating the mean (average) of a statistically significant number of samples. MAE is reported in years, i.e., the average number of years of error that is expected for a given piece of software. The use of absolute error is critical here as otherwise, if a system had positive and negative errors, they could cancel each other out. It should also be noted that mean is just that, and not, for instance, maximum. While a system may have a mean absolute error rate of for instance 2 years, some samples will show error rates of less than 2 years and some samples will show error rates of more than 2 years.

False Positive Rate (FPR)

This a measure of the frequency with which an age estimation system reports that a person who is actually of age X or lower is of age T or higher. For instance, FPR could detail the likelihood that a person of age 15 or lower is reported as being age 18 or higher. In this case, “False Positive” refers to the system falsely delivering an age estimate that is higher than (i.e., positive relative to) the required threshold. In the “We ID under 30” example, a False Positive would be falsely judging that a person who is under 21 appears to be greater than 30.

Use Cases for Age Estimation

Age estimation is highly relevant across various industries and applications, catering to diverse use cases:




01
Online Content Restrictions

Age estimation technologies play a crucial role in ensuring age-appropriate access to online content. Platforms hosting sensitive or explicit material such as gaming or adult entertainment websites can utilize age estimation to verify users' ages before granting access.



02
E-Commerce and Delivery Services

Age verification is essential for purchasing age-restricted products such as alcohol, tobacco, or certain medications in e-commerce settings. By integrating age estimation solutions, online retailers can verify the age of customers before completing transactions, ensuring legal compliance and mitigating the risk of selling restricted items to minors.



03
Social Media and Dating Platforms

Social networking and dating platforms employ age estimation to enforce age restrictions and enhance user safety. By verifying users' ages during account registration or content access, these platforms can prevent underage individuals from accessing adult-oriented content or interacting with older users, thus fostering a safer online environment.



04

Gaming and Entertainment

Age estimation technology can be used in the gaming industry to enforce age ratings and parental controls. Gaming platforms can employ age verification to restrict access to age-inappropriate games or features, ensuring that minors are not exposed to content unsuitable for their age group.



05

Education and Online Learning

Age estimation technology can be utilized in educational platforms to ensure age-appropriate access to learning materials and resources. By verifying students' ages, educational institutions can tailor educational content and experiences to specific age groups, enhancing learning outcomes and engagement.



06

Fraud Detection and Prevention

Age estimation technology can also contribute to fraud detection and prevention efforts within financial services. By analyzing users' facial characteristics and comparing them with other data sources, age estimation can detect inconsistencies that may indicate fraudulent behavior, such as using stolen identities or forged documents to access financial accounts.

Overall, age estimation technologies can power a wide range of solutions across industries, enabling organizations to enforce age restrictions, comply with regulations, and enhance user safety and privacy in various online and offline settings.

Ethical Considerations

The development of any AI-based technology or any technology that is involved in identity-related application demands a rigorously ethical approach. Paravision has publicly documented and rigorously adheres to a set of AI Principles¹ which states that it will ethically develop and conscientiously sell its technology. This ethical approach applies to age estimation technology in three main ways:



Ethical Development: Datasets

Paravision Age Estimation has been developed with the use of over 800,000 images. These are divided internally into separate datasets for training (over 600,000 images) and benchmarking (over 200,000 images). Paravision has obtained appropriate consents for all of the images used for training and benchmarking. Neither the training nor benchmark data use any operational data taken from any service provider. Paravision itself is a technology provider, not a service provider, and does not have access to partner or user operational data.



Ethical Development: Demographic Performance

Paravision has benchmarked its Age Estimation software across multiple demographics, including self-reported age, gender, and race, and this report includes an overview of these different factors, demonstrating strong consistency across the spectrum of demographics. While our approach is based on extensive analysis, we will continue to assess leading thinking and techniques for understanding and mitigating bias, and will update this approach over time whenever useful as part of an ongoing commitment to understand and mitigate differential performance (i.e., algorithmic bias).

Age

Age itself poses challenges due to the varied facial features and aging patterns across different age groups. Put differently, as people age, the differences in appearance across a fixed number of years grows smaller (e.g., there is more visual difference between 12 and 18 year olds than 32 and 38 year olds). So it is important to look at how age estimation performs across age groups. In addition, certain age ranges are more critical than others due to policy and societal norms (e.g., the difference between 12 and 18 year olds is more important than 32 and 38 year olds). For that reason, in Paravision's Demographic Analysis, we've highlighted analysis of target age ranges relevant to specific use cases, such as gaining access to social media or purchasing alcohol.

Gender

Following best practices, Paravision relies on self-reported gender rather than using external human or AI characterization of gender appearance. We are constantly looking for opportunities to improve gender inclusion and representation in a balanced way across training and benchmark datasets.

¹ <https://www.paravision.ai/ai-principles/>

Race, Ethnicity, and Skin Tone

Following extensive and repeated analyses by NIST FRVT / FRTE², as well as the U.S. Department of Homeland Security Science & Technology Directorate’s Biometric Technology Rally³, Paravision is well-attuned to both the importance of demographic analysis as well as the different approaches to addressing the critical consideration of race, ethnicity, and skin tone. There is no single approach to addressing this consideration: NIST often relies on nationality as it has broad access to such information. DHS S&T uses a combination of self-reported race and color-calibrated photometers, which is appropriate given the in-person, high-touch nature of their testing. Others have used the Monk Skin Tone Scale⁴, which is viewed as the most modern and sophisticated method for determining skin tone.

In our analysis, Paravision has elected to use self-reported race as opposed to skin tone or nationality. We believe that relying on self-reporting for race rather than any other type of characterization follows best practices while minimizing the introduction of unintended biases. Why did we use race? Nationality is not an option, as Paravision does not have access to this type of data (which may be more easily obtained by government agencies). Paravision has also performed extensive testing with the Monk Skin Tone Scale, but has found inconsistencies both with self-reported values (when subjects are given a printed reference card) as well as extraction of data (i.e. “apparent” skin tone) from images. The latter is problematic in that lighting and camera quality can both dramatically affect perceived skin tone. Fundamentally, the goal is to deliver consistently-performing technology across wide demographic swaths, and we believe that approach used is successful in that regard. Race is broken down in this report into Black, South Asian, East Asian, and white. Here, the definition of Black and white follows conventions set forth by the U.S. Census Bureau, but Asian has been broken down into South Asian (i.e., representing race from the Indian subcontinent) and East Asian (all other Asian populations).

Environment, Usage, and Camera Technology


While demographic performance is vital for assessing age estimation software, it’s also important to note that environmental factors like lighting, camera quality, and facial positions affect performance. For this reason, it is critical to consider these factors in both the development of age estimation technology as well as in the image quality metrics used during a capture process. Paravision’s approach to ensuring sufficient capture quality is discussed in the section of this document titled “Face image (selfie) capture.”



2. <https://www.paravision.ai/news/paravision-shines-in-latest-nist-frvt-11-rankings/>

3. <https://www.paravision.ai/news/paravision-face-recognition-delivers-top-performance-in-latest-dhs-biometric-test/>

4. <https://blog.google/products/search/monk-skin-tone-scale/>




Ethical Deployment: Privacy

In the realm of age estimation technologies, one pivotal consideration is the preservation of user privacy. It's paramount to offer solutions that provide accurate age verification without compromising individuals' personal data, whether for age assurance in applications used by children or privacy-sensitive applications used by adults. AI-powered facial age estimation prioritizes privacy by enabling age checks without the need to process customer data extensively, by delivering age estimation based on a facial image only, without other personal information. That being said, even though Age Estimation can help with user privacy by only using facial images for age assurance, it's important to note that face images used for age verification purposes are considered Personally Identifiable Information (PII) and are therefore should always be treated with extreme care.

Paravision Age Estimation does not “learn” from any operational data from Paravision Partners; there is no training from the images or results in partners' systems. Rather, Paravision collaborates with partner applications in an offline manner to identify opportunities for improved performance and either update technology or datasets to close performance gaps. The system is then re-trained offline based on properly-consented, carefully curated datasets which are fully independent of partner services. Any improvements are deployed just as a standard software update. In addition, it is worth reiterating that Paravision is not a SaaS provider: Paravision software is delivered to partners and installed, deployed, operated from within the partners' environments, to which Paravision has no access.

Finally, while Paravision delivers world-class face matching technology in support of positive identity verification, this technology is in no way bound to Paravision Age Estimation. In other words, while partner solutions can use a combination of face matching (i.e. selfie-to-driver's licensing identity checks) and age estimation, there is no need to do so. For age-sensitive applications where privacy is valued and identity verification is inappropriate or unnecessary, Paravision Age Estimation can be used in a fully stand-alone mode (or in conjunction with Paravision Liveness and Deepfake Detection, which also do not check biometric identity). Age Estimation is ephemeral (face images or other identity do not need to be stored) and depending on service provider policy, does not need to be linked to any sort of identity or transaction repository.



Ethical Deployment: Use Case Analysis

As a part of its AI Principles, Paravision conducts a detailed use case review prior to allowing its technology to be deployed with a new partner and / or new use case. Paravision has strict limits of the use of its technology by region and application, maintaining a goal of supporting safety, security, strong identity, and positive customer experiences while protecting against inappropriate or unethical use. With years of experience handling the deployment of face recognition technology, which is by its nature highly sensitive, Paravision has a strong framework to assess ethical use of Age Estimation and protect against potential abuses.

How Paravision Age Estimation Works


Paravision Age Estimation includes the following main steps:



01

Face Image (Selfie) Capture

Age estimation is performed on a selfie image, most frequently taken from a smartphone. The capture itself is automated, and assisted by a collection of image quality metrics that help to ensure that the captured image will be of sufficiently high quality to make an accurate age estimate. If all conditions are optimal, the software will automatically capture the face image. If not, the software will provide intuitive guidance, such as “move camera forward” (in the case that the face is too small), “hold camera steady” (in the case that the face is too blurry), or “increase lighting” (in the case that the face is under-exposed). The image quality metrics are typically deployed within the smartphone app itself to improve response time and usability. Once conditions are met, the software will capture the face and send the face image to our partner’s cloud service or other backend infrastructure for age estimation.



02

Age Estimation

Once a selfie is taken and sent to our partner’s backend systems,, Paravision’s Age Estimation software processes the image and returns an estimated age. The Paravision software is AI-based, and trained offline with a large, representative dataset. It takes an image as an input, and using a specifically designed Convolutional Neural Network, outputs an age estimate. Partners can then use this estimate within their applications and bind the result to any other relevant data to make an overall age estimation decision.

As noted previously, Paravision Age Estimation is not a SaaS offering, and face images and other data are not sent to Paravision as a part of the age estimation process. Paravision’s operational software resides on partner systems, helping to support trust and security by minimizing 3rd party data interfaces.

Accuracy Analysis

As noted above, Paravision has curated a benchmark dataset of over 200,000 images that is used for accuracy analysis. This benchmark is independent of the training dataset to ensure meaningful results. The results of this benchmark are as follows:

Mean Absolute Error (MAE), Overall

Age Range	Mean Absolute Error (MAE) Rate
6-12	1.29
13-17	1.77
18-24	2.30
25-70	3.63
6-70	2.25

Demographic Analysis Overview

In Paravision’s extensive internal benchmarks, Paravision Age Estimation software has demonstrated strong consistency across the spectrum of demographics, including self-reported age, gender, and race.

Key takeaways from this demographic analysis show:

- The distribution across race and gender within any given age bracket is quite small.
- For the age ranges of 6-12 and 13-17 years, there is no greater than a 0.67 year variation from the highest MAE to the lowest MAE across gender and race.
- When viewed across the complete 6-70 year range, the difference between the maximum MAE (2.77 years) and minimum MAE (2.38) between gender and race demographics is 0.39 years.



For more information about Paravision Age Estimation, including detailed data on accuracy and demographic performance, please reach out to info@paravision.ai or book a meeting at paravision.ai/ contact.

Comparison with Human Accuracy

The reported accuracy and target use cases discussed in this paper naturally raise the question “How does this compare to the human ability to estimate age?”

In the June 2015 edition of IEEE Transactions on Pattern Analysis and Machine Intelligence, Han et al. from the Michigan State University Department of Computer Science and Engineering published “Demographic Estimation from Face Images: Human vs. Machine Performance,”⁵ in which they determined human accuracy measured across multiple reviewers on over 5,000 images. The team determined that with basic image quality metrics applied to screen out poor quality images, **human review delivered a MAE of 4.5 - 6.6 years**, depending on the dataset (MAE was 4.7 - 7.2 years without the image quality metrics).

Alternatively stated, **human age estimation delivers error rates that are 2.0 - 2.9X higher than the Paravision age estimation technology.**



5. <https://ieeexplore.ieee.org/document/6920084>

In Conclusion

Paravision Age Estimation represents a thoroughly modern solution, where ethics, accuracy of core technology, and usability are all combined to deliver a solution that can be deployed at scale with confidence:

- The technology has been developed with properly consented datasets, without the use of operational data, has been benchmarked by extensive datasets fully independent from training data, and delivers consistent performance across gender, race, and age.
- While it can be used as a part of a larger identity verification solution, Paravision Age Estimation works in a fully stand-alone manner, allowing for high accuracy age verification without requiring a linkage to identity information.
- Meanwhile, the resulting technology is deployed on partner systems and not as a Paravision SaaS offering, minimizing risk and maximizing protections for users.

Paravision Age Estimation supports our partners and respects the rights of end users while setting a new standard for technology in support of safe, frictionless customer experiences for people of any age.



Trusted Vision AI

For more information or to schedule a demo, please contact us at:

info@paravision.ai