

TECHNICAL REPORT: AGEVERIF

AI FORENSICS



Credits

Research & Report: Dr Paul Bouchaud.

Graphic & Brand Design: Denis Constant | Ittai Studio <https://ittai.co>

All other content (c) AI Forensics 2025 | <http://aiforensics.org/>

AI Forensics is funded by core grants from [Open Society Foundations](#), [Luminate](#), [Omidyar Network](#), and [Limelight Foundation](#).

Email : info@aiforensics.org

Social Media: [Linkedin](#) | [Bluesky](#)

Table of Content

Executive Summary:.....	3
Background on AgeVerif:.....	4
AgeVerif's system.....	7
Age Estimation Model.....	10

Disclaimer

This technical report is based on information publicly available as of September 8th, 2025. We acknowledge that manual errors may have occurred, and we do not claim to have carried out an exhaustive analysis of all existing age verification mechanisms. Finally, at no point does this report provide, attempt to provide, or purport to offer a legal assessment of non-compliance with applicable regulations.

Executive Summary:

With the stated aim of protecting minors online, the Digital Services Act (DSA), in force since August 2023, requires platforms accessible to children to implement safeguards against harmful content. In France, the SREN law of May 2024 reinforced these provisions. On August 28, 2025, six pornographic platforms deployed age verification systems following Arcom's formal notices. The French regulator Arcom stipulates that pornographic websites are "*free to choose the solutions for protecting minors that they wish, provided that they comply with the technical requirements*" set out in its guidelines¹. For instance, Tukif[.]com & reference-sexe[.]com rely on the solution developed by AgeVerif².

The report inspect AgeVerif in relation to four regulatory requirements:

I. Independence of the age verification provider from the pornographic platform

Public records seem to show that *AgeVerif* (operated by *E-Borealis Lda* until 2024 and *PlanetSeason Lda* thereafter) has one of its directors in common with the company operating *Tukif & reference-sexe*, two of the pornographic websites using *AgeVerif*'s system.

II. Reliance on state-of-the-art solutions calibrated to avoid false positives

Based on our research, *AgeVerif* relies on an open-source model released in 2019 that was not trained to distinguish minors from adults. The model was trained on a dataset containing only ~7% of minors. Our own evaluation of the underlying model suggests that around 10% of minors might be granted access. We observe that *AgeVerif* appears to apply country-specific thresholds on the estimated age: 20 years in France, 25 in Germany, 23 elsewhere.

III. Non-discrimination across population groups (Charter of Fundamental Rights)

Our evaluation of *AgeVerif*'s age estimation underlying model against fairness benchmarks reveals risks of major disparities: we estimate that minors with dark skin tones may be 3 times more likely than white minors to be misclassified as adults, while South and East Asian adults may be about twice as likely as white adults to be wrongly denied access. The datasets (*released by academic teams for non-commercial use only*) on which the age estimation model was trained contain only 6% individuals with dark skin tones.

IV. Limiting the risk of circumvention of the technical solution

It would appear that *AgeVerif* have implemented a feature that allows age verification to be bypassed simply by adding a cookie, without a specific value, named "discl" (on *Tukif*) or "noagvf" (on *reference-sexe* and by default on other websites).

¹<https://www.arcom.fr/se-documenter/espace-juridique/textes-juridiques/referentiel-technique-sur-la-verification-de-lage-pour-la-protection-des-mineurs-contre-la-pornographie-en-ligne>

² <https://www.ageverif.com>

Background on AgeVerif:

Creation: The domain *AgeVerif.com* was last registered in December 2020³ by an unidentified party. Archived versions of the website from September 2021⁴ reveal that since around March 2021, the site was owned and operated by *E-Borealis Lda*, a company registered in Portugal, which continued to run the service until at least early December 2024.

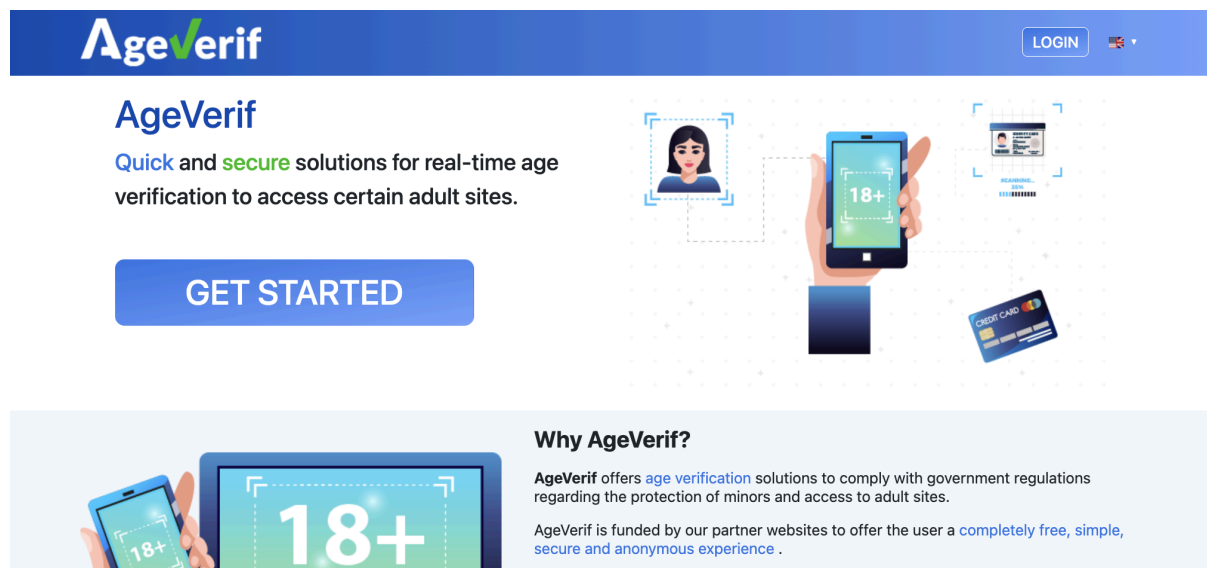


Figure 1: [AgeVerif.com](https://www.ageverif.com) website in September 2021

Promotion: On September 21, 2021, a low-key website *kalistasolutions.fr* published a promotional article praising AgeVerif's technology, titled "Welcome to the cutting edge of age verification technology"⁵. The piece sought to reassure readers about privacy concerns, claiming:

"With AgeVerif, you don't have to worry about your privacy being compromised. All verification methods take place on your own phone or computer. None of the data is collected or stored on their servers."

Just a week later, on September 29, 2021, another website, *hachette-litteratures.com* (bearing similarity with the French publishing group *Hachette Livre* whose official website is *hachette-collections.com*), published a similar piece, presenting AgeVerif as a "crucial" system, especially "as more and more countries adopt stricter laws requiring your website to ensure that no underage users can access it"⁶.

³ <https://who.is/whois/ageverif.com>

⁴ <https://web.archive.org/web/20211019043503/https://www.ageverif.com/terms-and-conditions>

⁵ <https://kalistasolutions.fr/bienvenue-au-sommet-de-la-technologie-des-systemes-de-verification-de-lage>

⁶ <https://hachette-litteratures.com/comment-utiliser-un-systeme-de-verification-de-lage-comme-un-seo-pro>

These websites appear to have been part of a search engine optimization (SEO) strategy. Their role was likely to boost the visibility of AgeVerif content in search results by embedding promotional articles among otherwise generic content filled with trending keywords. For example, before publishing the AgeVerif article, *kalistasolutions.fr* had only published two technology-related articles:

- “AMD, Intel, and Qualcomm call on the US to reduce its dependence on foreign chips” (February 7, 2021)
- “The majority considers Qualcomm to be the best processor brand for smartphones” (February 25, 2021)

Upon inspection of the other articles published by these two sites, we found references to the following domains:

pornograttuit.stream, mvideoporno.xxx, reference-sexe.com, pornofrancais.xxx, lebon.porn, hammerporno.xxx, & sexy-parade.com

According to their terms of services, these different pornographic websites are owned and operated, as of September 2025, by the same company *EvolFill LTD*, based in Cyprus.

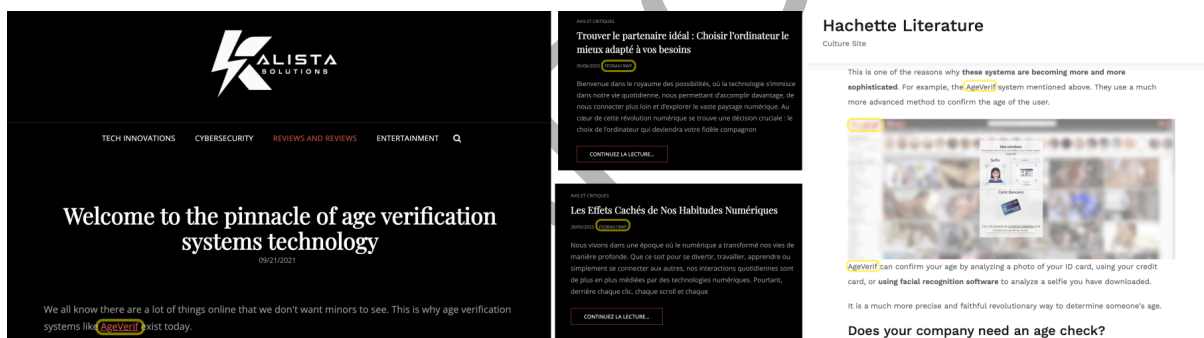


Figure 2: Articles promoting AgeVerif (Left & Right) & published by “Fedraxwp19” (center). Key Elements are highlighted in yellow.

In parallel, we note that *Tukif.com* was operated, prior to 2024, according to its terms of services, by *Fedrax Lda*, a company registered in Portugal. Interestingly, some of the articles on *kalistasolutions* were authored⁷ under the name “*fedrax19wp*” (with “wp” likely referring to WordPress, the content management system used by the site), see Figure 2.

⁷ See archived articles: <https://ghostarchive.org/archive/bsiLl>

Corporate Registry: According to public records, *E-Borealis Lda*, the company operating *AgeVerif.com* until 2024, was registered in Portugal (fiscal code: 510001408) in 2011. Its director, Marco, happens to also be the director of *Fedrax Lda* (Fiscal code: 511285469), which operated the pornographic site *Tukif.com* until 2024.

As of September 2025, *AgeVerif.com* is owned and operated by *PlanetSeason Lda* (fiscal code: 518490831). Public records list its director, Jérôme, as being also the managing director of *NWS*, a Luxembourg-based holding company that owned 50% of *Fedrax Lda*.

In April 2024, *EvolFill LTD* was incorporated in Cyprus (Registration number: HE 458856), with Jérôme as its director.

Figure 3 summarizes the network of companies, holdings, directors, and managing directors connected to AgeVerif and Tukif, as inferred from public records.

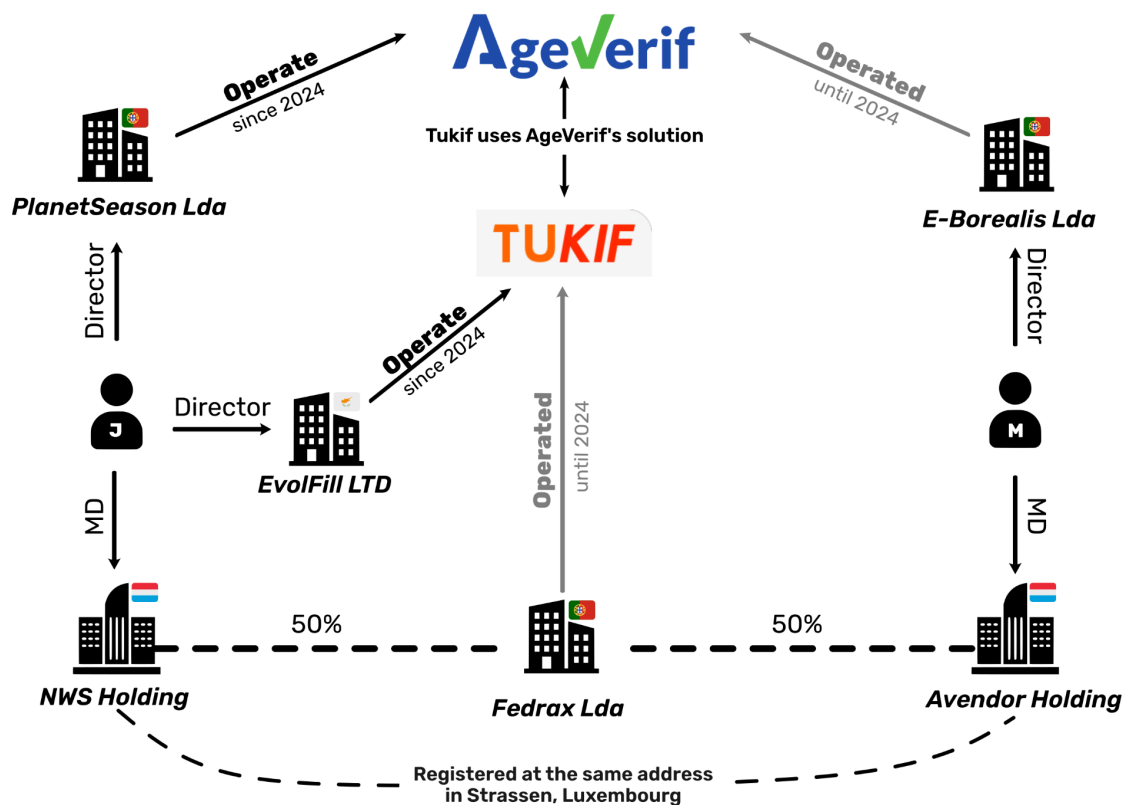


Figure 3: Network of companies, holdings, directors, and managing directors (MD) connected to AgeVerif and Tukif

AgeVerif's system

On websites that have deployed the AgeVerif solution, pornographic content is inaccessible until users verify their age. To this end, users can choose one of multiple verification methods⁸ offered by the system. As shown in Figure 4, the list of available options is determined by the website implementing the solution.

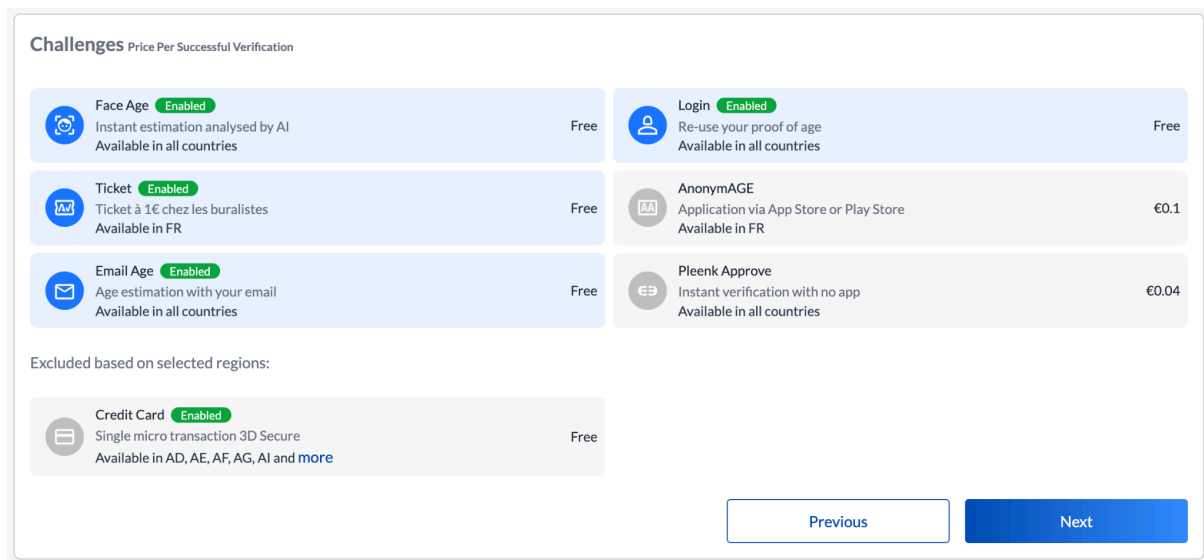


Figure 4: Difference Age Verification methods offered by AgeVerif to webmasters.

One of the preset options accessible to users without an account is the “Selfie: instant age estimation analyzed by AI”.

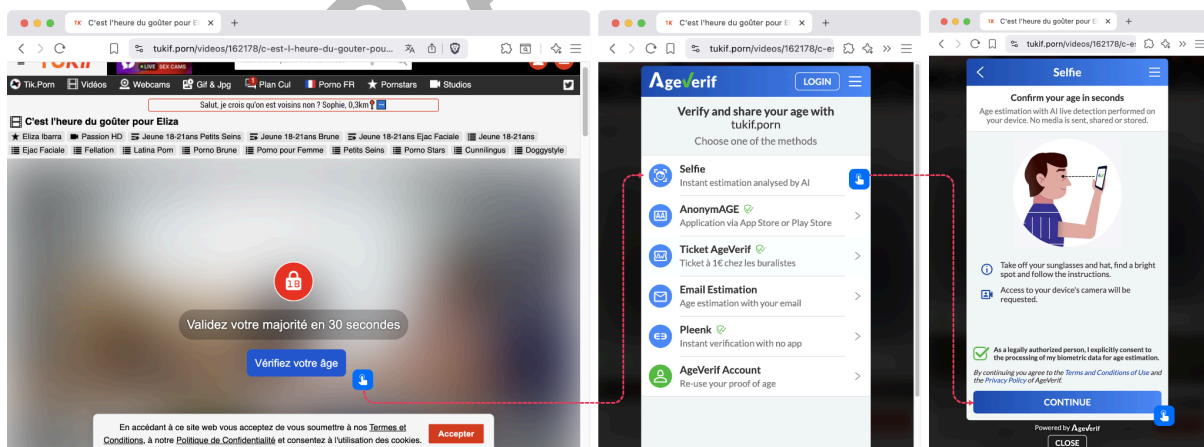


Figure 5: AgeVerif 'age verification' deployed on tukif.com

⁸ The verification methods offered by AgeVerif appear to vary across websites. For example, the 'AnonymAGE' option is available on *tukif.com* but not on *ixxx.com*, even though both rely on the AgeVerif system.

AgeVerif states that this estimation is performed locally on the user's device: “No media is sent, shared or stored”. Inspecting the network traffic, we can confirm this claim: as of September 2025, the webcam stream is **not** transmitted to AgeVerif or any third party.

Instead, AgeVerif sends three different computer vision models to the user's device to enable the age estimation. These models are deployed using TensorFlow.js, a JavaScript library that allows web pages to load and run machine learning models directly in the browser.

We observed that AgeVerif's “Face Age” estimation unfolds as such: a first model (tiny_face_detector_model) detects the user's face in the webcam stream, while another model (face_landmark_68_tiny_model) identifies facial landmarks such as the eyes and nose, see Figure 6. AgeVerif then prompts the user to move their head left, right, and so on, as a *liveness* check. Finally, the system estimates the person's age by calculating the median of the predicted ages across different frames (through the age_gender_model).

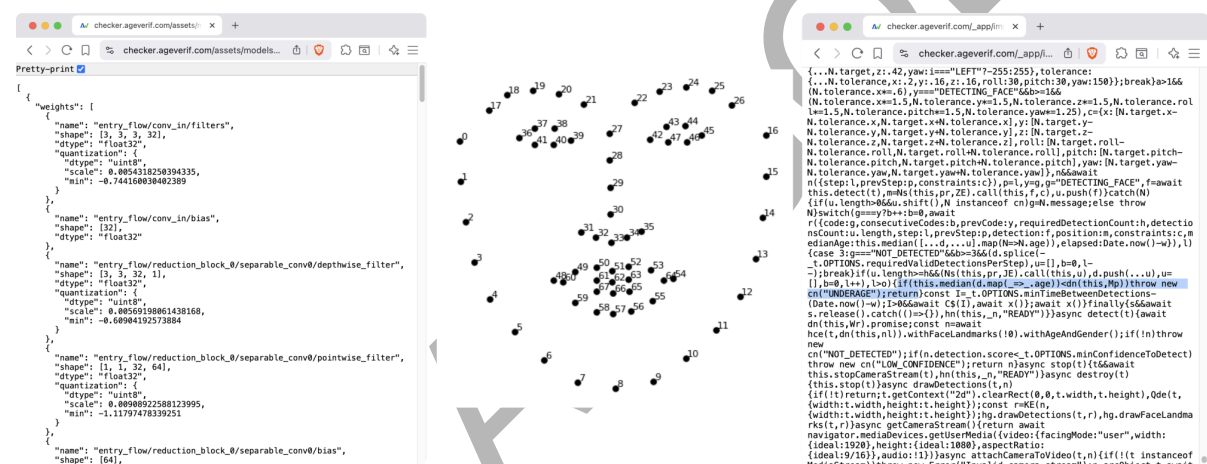


Figure 6: (Left) Configuration of the model used to estimate user age, (Center) facial landmarks, (Right) AgeVerif logic that compares the estimated age against a threshold.

User estimated age is then compared with a hard-coded minimum age threshold; if it falls below this threshold, the face is deemed *underage* and access to the pornographic platform is denied.

Interestingly, rather than applying a universal threshold, AgeVerif appears to have manually set country-specific thresholds. In particular, we observed the following thresholds: {"UK": 25, "DE": 25, "FR": 20}, with a default value of 23 years old for other countries⁹.

⁹ Amusingly, during a manual test using a VPN to log in from the United Kingdom, AgeVerif identified our country as “GB” rather than “UK”. As a result, the 25-year-old threshold was not applied, and the fallback value of 23 years was used instead.

While setting a higher threshold than the legal age—i.e., requiring users to be classified by the system as at least 20/23/25 years old to access content rated for 18+ years—is motivated and actually recommended, for instance, by the German Commission for Youth Media Protection¹⁰, we cannot explain why these specific values were chosen nor why they vary as a function of the country of visit. However, such differential thresholds may result in discriminatory outcomes based on the user's country of connection. For instance, a user who appears to be 24 years old would be granted access in France but denied in Germany.

Bypass:

We observed a mechanism in AgeVerif's solution that allows it to bypass age verification.

Specifically, a function named `verified()` decides whether to initiate age verification based on:

- whether age verification is required for the page,
- whether the user already has a valid “verification cookie” stored in the browser, or
- whether a special “bypass cookie” is present.

In the configuration used on the site *tukif*, if a user's browser contains a cookie called “discl” (see Figure 7), the age verification step is skipped and access is granted. On other sites, the default name for this bypass cookie is “noagvf”.

By manually adding such a cookie we successfully manage to bypass the age verification process:

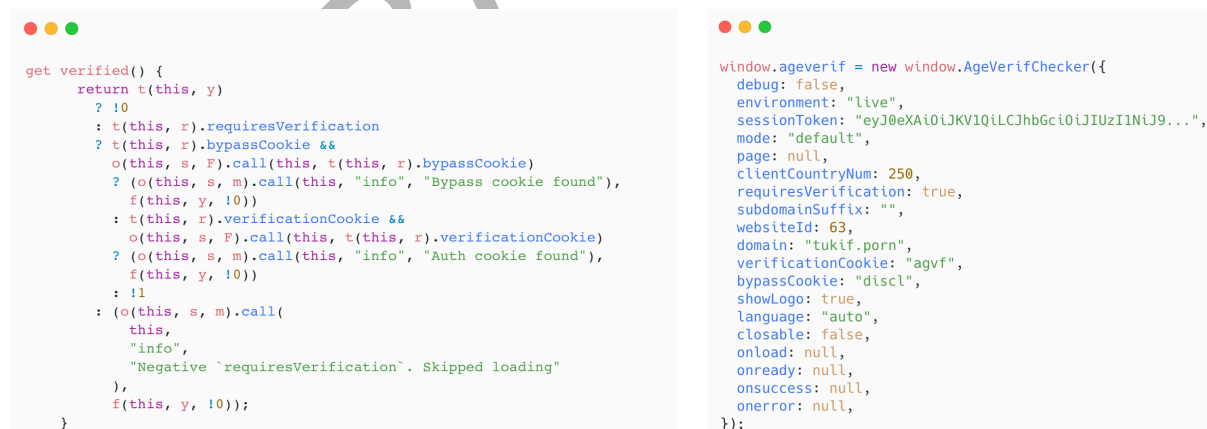


Figure 7: (Left) AgeVerif `verified` function (Right) Configuration used on Tukif

¹⁰ <https://www.kim-online.de/pressemitteilungen/altersverifikation-persona-gocam/>

Age Estimation Model

AgeVerif does not disclose any details on its website about how age estimation is performed, merely claiming that it relies on “*AI-powered biometric analysis*”¹¹.

By searching on GitHub for the combination of the three models (tiny_face_detector_model, face_landmark_68_tiny_model, and age_gender_model), we identified a popular open-source JavaScript library. Comparing the cryptographic hash of the model weights confirmed that **AgeVerif uses the version of the model released in 2019 by this open-source library.**

Because our analysis focuses solely on the implementation of this model by AgeVerif, we have chosen not to name the library. The library and its models were developed and released in good faith by an open-source developer, without any suggestion that they should be used in sensitive or commercial contexts. Our concern lies with the companies that choose to deploy such tools without due consideration of their limitations and implications.

Model Training: The open-source developer provides further information about how the model was trained. In particular, we learn that the model was trained on a combination of datasets, the largest (public) ones, in number of faces, being IMDB, Wiki, MegaAge, UTK.

Age and Gender Recognition Model

The age and gender recognition model is a multitask network, which employs a feature extraction layer, an age regression layer and a gender classifier. The model has a size of roughly 420kb and the feature extractor employs a tinier but very similar architecture to Xception.

This model has been trained and tested on the following databases with an 80/20 train/test split each: UTK, FGNET, Chalearn, Wiki, IMDB*, CACD*, MegaAge, MegaAge-Asian. The * indicates, that these databases have been algorithmically cleaned up, since the initial databases are very noisy.

Total Test Results

Total MAE (Mean Age Error): **4.54**

Total Gender Accuracy: **95%**

Figure 8: Description of the age estimation open-source model used by AgeVerif.

Most of them being openly available online, we report in Table 1 its curator, license, number of faces, fraction of minors depicted¹², fraction of non-white and fraction of Black individuals¹³.

¹¹ <https://www.ageverif.com>

¹² Relying on the dataset labels.

¹³ When skin tone was not reported in the dataset labels we inferred using Deepface, as such those results are estimates.

Dataset Name	Curator	Licence	#Faces	Fraction of Minors	Fraction of Non-White	Fraction of Black
IMDB	ETH Zurich	Academic research purpose only	460k	5%	~23%	~6%
WIKI	ETH Zurich	Academic research purpose only	62k	3%	~39%	~8%
MegaAge	CUHK	Non-commercial research purposes only	41k	35%	~54%	~3%
UTKFace	U. of Tennessee	Non-commercial research purposes only	20k	21.4%	~55%	~21%

Table 1: Datasets over which was trained the model used by AgeVerif

First of all, we note that the datasets on which the model (commercially used by AgeVerif) was trained were released by academic research teams for **non-commercial research purposes only**.

Additionally, we observe a significant imbalance in terms of skin tone: only 27% of the training samples are of non-white individuals, and **only 6% of Black individuals**. In addition, despite AgeVerif using the model to distinguish minors from adults, **only 7% of training samples are faces of minors**.

Despite the significant imbalance in terms of skin tone and the minimal representation of minors in the training dataset, we have **no indication** that any particular care was taken during training to **mitigate racial or age biases**.

The model appears to be trained to minimize the Mean Age Error (MAE) —the average difference between a person’s real age and the predicted age— **without any specific emphasis on distinguishing between minors and adults**.

Importantly, we emphasize that the open-source developer did not train the model for the specific purpose of distinguishing between minors and adults. The responsibility therefore should lie with *AgeVerif*, which chose to deploy a model that can be constructed as ill-suited to this application.

Model Evaluation: Considering these elements, we deemed it appropriate to evaluate the performance of the model deployed by AgeVerif.

To this end, we considered two ethnically diverse datasets of faces labeled with age: UTK¹⁴ (on which the model was partially trained) and FairFace¹⁵ (released after the model's training).

For each face we estimate their age using the model used by *AgeVerif* and compare it with the face declared real age.

1) UTK: Estimating the ages of faces in the UTK dataset seems to result in an average error of 5.3 years, which aligns with the performance advertised on the project page (5.25 over UTK).

Additionally, we inspected the model's performance across different age ranges, comparing the Mean Age Error reported by the original developer (over combination of test datasets) with the values we computed over the UTK dataset.

These results may appear commendable for age estimation and may even compete with some commercial models. Yet AgeVerif does not use this system merely to estimate a person's age, but rather to determine whether an individual is a minor or an adult.

Age Range	0 - 3	4 - 8	9 - 18	19 - 28	29 - 40	41 - 60	60 - 80	80+
MAE (Reported)	1.52	3.06	4.82	4.99	5.43	4.94	6.17	9.91
MAE (UTK)	1.61	4.12	4.18	4.61	6.30	6.88	7.23	10.10

Table 2: Mean Age Error per age ranges.

To better visualize the results, we display in Figure 4 the fraction of faces predicted as being over 18 as a function of the person's real age; segmenting the results by skin tone.

We observe that as the real age increases, the fraction of faces predicted as over 18 also increases. However, not all individuals above 18 are classified correctly, particularly those in the 18–25 age range. Similarly, regardless of skin tone, a non-negligible fraction of minors are misclassified as over 18.

¹⁴ <https://doi.org/10.1109/CVPR.2017.463>

¹⁵ <https://doi.org/10.1109/WACV48630.2021.00159>

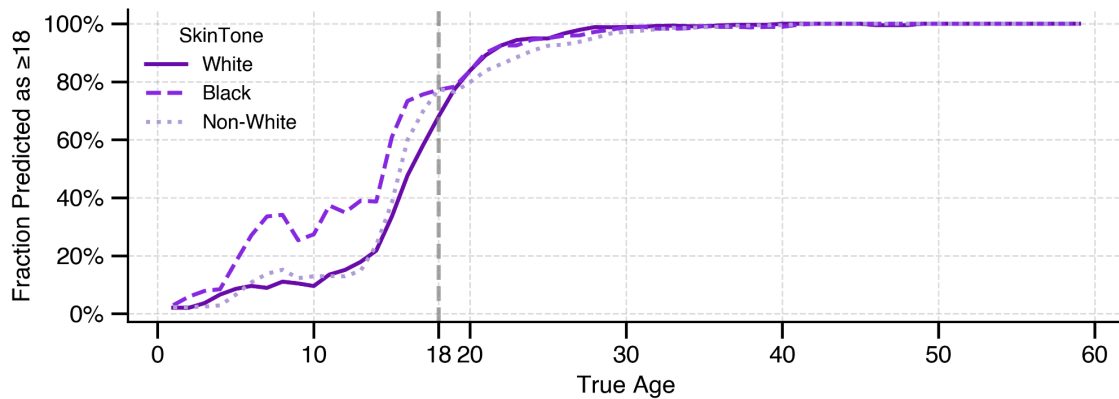


Figure 9: Proportion of faces from the UTK dataset predicted as over 18 years old, as a function of the person's true age, segmented by skin tone.

More importantly, we observe disparities based on skin tone: Black minors younger than 12 years old are more frequently predicted as over 18 compared to white minors.

For clarity, we report in Figure 4-6 the fractions of minors and adults allowed/blocked by the system, segmented by skin tone and considering the different country-dependent thresholds used by AgeVerif.

As expected, as the threshold on the estimated age increases, the fraction of minors being blocked also increases, from 91% when the threshold is set at 20, to 96% when it is set at 25. Complementarily, the fraction of adults being blocked rises as well, from 9% at a threshold of 20 to 27% at a threshold of 25.




Figure 10: Confusion matrix over the UTK dataset, threshold on estimated age set at 20 years

Importantly, we observe **significant differences based on skin tone**. Considering the case of France, where the threshold on estimated age is set at 20, we estimate that the AgeVerif model **allows access to 11% of white minors but 27% of Black minors**. This threefold bias persists across other thresholds: for example, in Germany (threshold set at 25), 4% of white minors are allowed access compared to 12% of Black minors.




Figure 11: Confusion matrix over the UTK dataset, threshold on estimated age set at 23 years




Figure 12: Confusion matrix over the UTK dataset, threshold on estimated age set at 25 years

2) FairFace: FairFace is a dataset of over 100k images, curated from the *YFCC-100M Flickr* dataset to provide a balanced representation of age, gender, and seven ethnic groups: *White, Black, Indian, East Asian, Southeast Asian, Middle Eastern, and Latino*.

Faces are labeled in age ranges. Unfortunately, one range spans 10-19 years, which partially overlaps with the legal threshold of 18 years in most European countries. To unambiguously assess performance on underage individuals, we therefore focus on the 3-9 and 20-29 year age ranges.

Our results, displayed below in Figure 13, show that, with the threshold set at 20 years, around 14% of faces of individuals aged 3–9 with Black skin tone are granted access to pornographic content by AgeVerif, compared to only 6% of white individuals in the same age range.




Figure 13: Fractions of faces from the FairFace dataset in the 3-9 age range that are granted access, shown by ethnic origin and threshold on estimated age used for allowing/blocking.

Conversely, we observe that, based on the FairFace dataset, 34% of East Asian and 30% of Southeast Asian adults aged 20-29 are classified as being under 20 years old, compared to 18% of white adults. As a result, the model seems to discriminatorily restrict access for legitimate adults seeking to access pornographic content.





Figure 14: Fractions of faces from the FairFace dataset in the 20-29 age range that are denied access, shown by ethnic origin

Limitations: Several limitations of the performed analysis should be noted. First, our previous results concern only the machine learning model used by AgeVerif to estimate a person's age from their face. We did not audit the overall system. In particular, the AgeVerif solution involves two additional models: one to detect the presence of a face and another to identify facial landmarks. Since these models operate on a live webcam stream, evaluating whether a face is detected in a single frame provides limited insight into the overall capability of the system, which is why we did not explore them further.

Additionally, we relied on a public dataset of facial images annotated with skin tone and age. Such annotations may contain errors, and the dataset itself is unlikely to be representative of the conditions under which the model is actually used. Indeed, it is reasonable to assume that the AgeVerif solution, deployed on pornographic websites, will more often be used at night, under poor lighting conditions, factors that could further degrade performance.

We generated AI-generated faces of minors with dark and light skin tones, captured at night under pool lighting conditions, and observed qualitatively similar biases.



Hence, despite these limitations, our analysis raises legitimate concerns about the potential for discriminatory behavior in the deployed solution.