



AgeVerif Analysis of the Technical Report by *AI Forensics*

Table of content

1. Independence of the Age Verification Provider from the Pornographic Platforms	1
2. Clarification on Promotional Articles and SEO Practices	2
3. Reliance on state-of-the-art solutions calibrated to avoid false positives.....	2
4. Use of Open-Source Models and Technical Rationale.....	3
5. Clarification on System Operation and Model Evaluation	4
6. Clarification on Fairness and Testing Methodology	5
7. Limiting the Risk of Circumvention	5
8. Conclusions	6

Disclaimer

We acknowledge the publication of the AI Forensics report concerning the AgeVerif system and appreciate any external analysis that aims to improve transparency and accountability in age assurance technologies.

AgeVerif remains open to constructive criticism and dialogue that can help enhance the quality, reliability, and fairness of the service we provide.

We recognize that the report raises some pertinent technical and governance questions, several of which we are already in the process of addressing or reviewing internally.

Unfortunately, we were only made aware of the report's existence through media coverage, rather than being contacted directly prior to its release, a dialogue that we would have welcomed.

Nevertheless, we view the report as a valuable risk assessment document that can contribute positively to ongoing improvements and cooperation between regulators, researchers, and industry stakeholders.

1. Independence of the Age Verification Provider from the Pornographic Platforms

AgeVerif is operated by **PlanetSeason Lda**, a company registered in Portugal that develops and maintains the technology independently.

The adult entertainment platforms mentioned in the report are operated by **EvoFill Ltd**, a separate company incorporated in Cyprus.

While the same individual serves as director and shareholder of both PlanetSeason Lda (AgeVerif) and EvoFill Ltd, all decisions regarding product development, compliance, and technical strategy are made entirely within the AgeVerif project. The two entities have separate objectives, accounting, and infrastructure. EvoFill Ltd has no influence over AgeVerif's technical or operational choices.

AgeVerif was created to address a clear need for a free or low-cost verification system, enabling smaller and free websites to comply with age protection laws without facing unsustainable costs. For many publishers, having access to a reliable and privacy-focused verification service at minimal cost is essential to remain compliant and to keep their business viable.

The fact that the director has experience in managing online content platforms is public information. This background helps ensure that AgeVerif's technology meets the real-world expectations of website operators and regulators. It allows the system to provide integration methods that are practical, efficient, and aligned with how publishers actually operate, which ultimately improves compliance and the protection of minors online.

The corporate information referenced in the report is publicly available in official company registries. This is not hidden, and its visibility demonstrates our commitment to transparency.

It would have been easy to use nominee directors or layered holding structures, as some other providers do, but we chose not to. The same individual's name appears openly in all filings precisely because AgeVerif has nothing to conceal.

This structure has also been clearly disclosed to regulators, including during exchanges with the French regulator Arcom, which raised no objection to the governance model.

What matters under regulatory frameworks such as the DSA or the SREN law is that the age verification provider operates with functional independence, data separation, and privacy integrity, all of which AgeVerif fully ensures.

AgeVerif Analysis of the Technical Report by AI Forensics

AgeVerif remains open to any independent audit that regulators may wish to conduct regarding its governance or data separation.

The AI Forensics report itself confirms that **there is no personal data shared** between AgeVerif and any adult websites. The report acknowledges that **no selfie or video data is ever transmitted**, as all processing takes place locally on the user's device.

After completing a verification, users can choose whether to save their proof of age or not. Those who wish to save it have two options: they can link it to an email address for convenience or store it anonymously through a passkey. The passkey option is by far the most popular, as it allows users to reuse their verification instantly without sharing any personal information with AgeVerif or the websites.

AgeVerif follows a privacy-first by design approach and fully complies with European data protection standards. The system has successfully passed the **IEEE 2089.1-2024 certification** through **ACCScheme**, confirming its adherence to recognized international standards for age assurance and data minimization.

2. Clarification on Promotional Articles and SEO Practices

The report mentions early promotional articles about AgeVerif published in 2021. We confirm that some of these early communications were coordinated by the AgeVerif team as part of a legitimate visibility and information strategy, which is common in the technology industry.

These materials were factual and focused on explaining the privacy-preserving design and compliance goals of AgeVerif. There was nothing improper in coordinating or authoring explanatory articles about the system, provided the information was accurate and transparent.

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

Clarification on Age Thresholds (20 / 23 / 25 years)

The age thresholds applied by AgeVerif are not arbitrary values, but **regulatory safety margins** designed to comply with national frameworks and reduce the risk of minors being incorrectly classified as adults during biometric age estimation.

In the **United Kingdom**, the Ofcom framework under the Digital Economy Act recommends that age assurance systems apply a safety buffer above 18 years to account for model uncertainty and ensure minors are not mistakenly approved. Accordingly, AgeVerif applies a **25-year threshold** for its selfie-based verification method.

In **Germany**, following the guidance of the **Kommission für Jugendmedienschutz (KJM)**, AgeVerif applies the “**25 principle**”, requiring users to appear at least 25 years old before access is granted. This follows the KJM’s long-standing position that verification systems must include a clear safety margin (“*ersichtlich erwachsen*”) to compensate for biometric and behavioural variability.

A **20-year threshold**, previously used in France during the early testing phase, was originally selected due to the **absence of detailed technical guidelines from the local regulator, Arcom**, at that time.

Following regulatory convergence and our ACCS certification, this threshold has since been **revised to 23 years**, aligning France with the standard used in the United Kingdom and several other EU jurisdictions.

These thresholds are **regulatory compliance parameters**, not technical limitations. Their purpose is to minimize false positives (minors being incorrectly approved). False negatives, where an adult may be mistakenly blocked, present no compliance risk since any adult can always complete an alternative verification method to confirm their age.

4. Use of Open-Source Models and Technical Rationale

AgeVerif does not currently train its own age estimation models from scratch. Instead, we have thoroughly tested and benchmarked all publicly available open-source models to identify the one providing the lowest error margins for faces under 30 years old while remaining lightweight enough to run efficiently on most user devices.

This approach was selected because one of the key requirements for this verification method is that it must be free, on-device, and accessible to all webmasters who wish to integrate an age verification solution without infrastructure or licensing costs.

The open-source datasets mentioned in the AI Forensics report were used only by the original academic teams to train the model. AgeVerif never accesses or distributes these datasets. Our system simply loads the already trained model weights locally on the user’s device to perform age estimation. We do not modify, resell, or redistribute the model or its components, which keeps us fully compliant with the original research license.

This model was chosen because it meets all technical and privacy criteria and was extensively tested to ensure accuracy consistent with regulatory safety margins. AgeVerif's system uses it strictly within the user's browser environment, with no data transmission or storage on any server.

It is also important to clarify that this selfie-based verification method is offered entirely free of charge and is not exploited commercially. It exists to provide a privacy-preserving, device-based option for users.

5. Clarification on System Operation and Model Evaluation

The AI Forensics report explicitly states that their analysis was limited to the standalone age estimation model and **did not audit the overall AgeVerif system**. In practice, the AgeVerif implementation involves multiple models and logic layers that together determine the final verification outcome.

Our facial age estimation process requires the user to **position their face properly**—centered, well-lit, and facing forward - before performing a short sequence of head movements (left, right).

During this process, AgeVerif performs continuous local estimations using the open-source model referenced in the report and then calculates the **median value** of all detections.

This median is compared against the **configured age threshold** for the user's region (for example, 25 years for Germany or the United Kingdom, and 23 years for France and other countries).

This multi-frame and liveness-based approach provides a **far more accurate and reliable estimation** than evaluating a single still image, as was done by the researchers in their tests.

It ensures that transient errors caused by lighting, expression, or pose are minimized, resulting in a more stable and realistic age assessment.

6. Clarification on Fairness and Testing Methodology

The fairness conclusions presented in the AI Forensics report are based on tests performed using **AI-generated faces**, rather than real individuals. This approach introduces major methodological flaws, as such generated images are not authentic human faces, but synthetic representations created by another AI system. These faces reflect the biases and inaccuracies of the generation model itself and cannot be considered reliable data for evaluating real-world age estimation performance.

In practice, the AgeVerif system operates under live conditions where users must meet strict criteria for **lighting, positioning, and movement** before any result is accepted. Several additional checks are performed to ensure that only genuine, well-lit, and properly aligned facial data is analyzed. These safeguards are precisely designed to reduce uncertainty and prevent circumvention.

Had the researchers tested the complete AgeVerif system rather than isolating a single open-source component, they would have observed these multiple validation layers and the significant difference in accuracy between a static image test and the live verification process actually deployed on partner websites.

7. Limiting the Risk of Circumvention

The AI Forensics report notes a bypass mechanism based on a cookie named "discl" or "noagvf". This was an early legacy feature requested by some partner sites to prevent AgeVerif from loading by setting a cookie. It was never publicly disclosed or documented and only a technically experienced person would be able to find and exploit it.

To remove any doubt we have **completely disabled and removed this bypass mechanism** from production.

Separately, AgeVerif has several safeguards to limit circumvention, even for experienced users with access to developer tools:

- A successful local verification is valid for a maximum of one hour and is not shared between websites.
- Copying a verification cookie from one site to another does not grant access on the receiving site.
- Tampering with the AgeVerif interface using browser developer tools will trigger a full screen error state that prevents the underlying page from becoming visible.

These measures are intentionally conservative. Any technically skilled actor could attempt to defeat client side protections, but the obstacles above make simple cookie tricks ineffective and raise the bar for circumvention. We continually monitor attempted bypasses, apply timely fixes, and harden both client and server logic.

8. Conclusions

Clarification on Regional Availability of Verification Methods

The AI Forensics report focuses exclusively on one optional component of the AgeVerif platform: the selfie-based AI age estimation method. It does not evaluate the overall system architecture or the range of verification methods available through AgeVerif.

AgeVerif offers several verification methods, with availability depending on each country's regulatory and contractual framework.

In France, for example, only the selfie-based AI estimation method is operated directly by AgeVerif, in accordance with Arcom's current operational and licensing framework. Other verification routes, such as third-party integrations (AgeGO, AnonymAGE, Pleenk or the Ticket AgeVerif system available in licensed kiosks), are fully integrated within the AgeVerif interface but operated independently.

In other jurisdictions, website owners and users may select from a broader range of AgeVerif-integrated or external verification options.

This flexibility reflects AgeVerif's role as a compliance integrator, designed to ensure lawful access control while preserving user choice and market competition. The goal is to provide interoperable, privacy-preserving age assurance pathways that respect local regulations and maintain double anonymity, rather than enforcing any single proprietary solution.

By analysing only the open-source model used in the optional selfie estimation method and ignoring the rest of the platform, the AI Forensics report presents an incomplete and partial view of how AgeVerif operates in practice.

The omitted components, including certified or third-party verification methods, are central to ensuring compliance, redundancy, and user choice across jurisdictions.

Closing Remarks

We acknowledge the strong interest that AI Forensics has shown in the technical aspects of selfie-based age estimation. While their report is critical, we believe that constructive engagement between researchers and solution providers can lead to genuine progress in the field of privacy-first age assurance.

The model currently used by AgeVerif represents the state of the art in open-source age estimation for on-device applications. It allows local, private processing without transmitting personal data, offering a clear privacy advantage compared with most commercial systems.

AgeVerif remains open to collaboration with academic and independent research groups, including AI Forensics, to further benchmark and enhance open-source age estimation models.

As technology evolves, we plan to contribute to the development of new, purpose-built neural architectures that continue to advance privacy-preserving and inclusive age verification worldwide.