JANUARY 2026

Age Assurance Online:

A Technical Assessment of Current Systems and their Limitations

Overview

Eric Rescorla
Zander Arnao
Alissa Cooper
Knight-Georgetown Institute

I. Introduction

In recent years, an increasing number of jurisdictions around the world have begun evaluating and adopting age assurance requirements of different kinds. Collectively, these moves represent a major change from how online services have been accessed over many decades, and they implicate a variety of important concerns and values for consumers, both adults and youth.

Age assurance technologies are complex systems that are being deployed on a wide scale on the internet for the first time. To help policymakers, service providers, independent experts, and users better understand how these systems work and their tradeoffs, *Age Assurance Online: A Technical Assessment of Current Systems and Their Limitations* provides a comprehensive technical assessment of the landscape of age assurance systems.

This overview document summarizes the key findings from the report, offering policymakers, service providers, and independent experts a high-level view of how age assurance systems function in practice and the tradeoffs they present.

It introduces the two principal age assurance architectures: server-based evaluation, in which service providers evaluate a user's age, and device-based evaluation, in which age evaluation and enforcement occur on the user's device or operating system, and explains how commonly used age assurance mechanisms (or "age signals") function. The overview also outlines the report's assessment methodology and summarizes the properties of both the architectures and the mechanisms across the key criteria of baseline accuracy, circumvention resistance, availability, and privacy.

II. Key Findings

The key findings of *Age Assurance Online* are as follows:

Multiple use cases: There are multiple use cases for age assurance, each with different requirements and challenges. These use cases largely fall into two main categories: (1) *safer defaults* for general-purpose services such as social media, AI chatbots, short-form video, gaming, and search, and (2) *blocking* access to specific content or services, especially adult-oriented services such as gambling or pornography.

- **Safer defaults are designed to provide users with an experience deemed more age-appropriate.** For instance, service providers might restrict the use of personalized feeds or of notifications during certain hours. These use cases typically involve the user having a long-term relationship with the service, allowing the service to adapt in response to user behavior. Because the user often has to identify themself to use these services, there may be a perceived decreased need for anonymity in age assurance, although safer defaults use cases exist where services allow pseudonymous or anonymous access. Minors may have less incentive to circumvent age assurance in safer defaults cases if the defaults do not adversely affect their experience of the service.

- **Some content and experiences may be blocked entirely for minors.** Some services are determined—often by law or regulation—to be adults-only. These use cases may support access by unidentified users without accounts and the expectation is that service providers block underage users with no previous history of interaction. Even in cases where accounts are required, users may wish to remain pseudonymous or anonymous, including for the purposes of age assurance. Minors may be more motivated to circumvent age assurance in these cases if it prevents them from accessing content or experiences that they want.

Multiple age signals: **No single age signal is sufficient on its own.** All existing age signals (self-declaration, commercial and government records, government IDs, age estimation) suffer from either accuracy or availability issues. In order to deploy a practical and effective age assurance system, any practical age assurance system needs to support multiple age signals so that users who are unable to successfully demonstrate their age with one signal can use another signal. Because the privacy properties of age assurance systems vary greatly and many of the most privacy-preserving designs are also not highly available, allowing the user to select a more private signal if available will protect user privacy more than requiring the user to try signals in a predetermined order.

- **Facial age estimation is highly available but inaccurate near the age threshold.** Anyone whose device has a camera can use facial age estimation, but it cannot reliably distinguish whether a user is just above or just below the age threshold and so must reject users who are not clearly older than the threshold.

- **Government-ID-based systems are accurate but not always available.** Systems based on government-issued ID provide accurate information about a legitimate user's eligibility based on their birthdate. However, many users do not have government-issued IDs; this is especially true of minors.

- **Behavioral signals are less suitable for primary age assurance.** Some service providers use user behavior to detect potential minors based on their patterns of usage. These systems may be usable as a backup mechanism but are less suitable for primary age assurance because they cannot determine a user's age on first contact.

- **Age thresholds below 18 are harder to deploy.** An age threshold below 18 (e.g., 16) requires minors to prove their age, but many minors who are close to the threshold will not have government ID. In many cases, parental consent or declaration will be the most practical option for age assurance below age 18.

- **Parental consent is difficult to establish.** In some cases, it will be possible to verify that an individual is over 18 and asserts that they are the parent of a child, but this is different from actually establishing that they are the parent. It is particularly challenging to verify parental consent while simultaneously protecting the privacy of both the adult and the minor.

<u>Privacy protection</u>: The most commonly deployed age assurance approaches present privacy risks, even though more privacy-protective approaches are possible and becoming more widely available. The most common age assurance systems require the user to either directly identify themself by name, email, or phone number, or to provide the age verification provider (AVP) with an image of their face. This forces the user to trust the AVP not to misuse their data and to protect their data from breach or disclosure even though the user may have no prior relationship with the AVP and no real alternative options if they wish to access the desired content or experiences. These risks are especially acute in cases where age thresholds below 18 are in use and minors are asked to demonstrate their age. Systems with stronger technical privacy guarantees are possible but not widely deployed.

- **Most widely deployed age assurance architectures require the user to trust the age verification provider (AVP).** When the AVP is separate from the service provider, the AVP learns the user's identity and the service provider they are trying to access, but not necessarily the specific content from the provider they are trying to access. The service provider only learns whether the user is in the eligible age range. However, there are no technical mechanisms preventing the AVP and service provider from colluding to match up the user's identity and activity. The user has no way of assuring this is not happening.

- **The most private age assurance systems are based on device-based enforcement or zero-knowledge proofs.** Both of these systems check the user's age on the device. With device-based enforcement, software on the device prevents the user from accessing restricted content or experiences. Zero-knowledge proofs use advanced cryptography to prove to sites and services that the user is in the eligible age range without revealing their identity. In both cases, neither the AVP nor the service provider learns the user's identity at all, with the result that the user need not trust either the AVP or the service provider with their data.

<u>Circumvention:</u> All age assurance systems are vulnerable to circumvention. It is not technically feasible to build an age assurance system which would prevent all minors from accessing restricted content or experiences without also blocking large numbers of adult users.

- **Server-based enforcement on the web can be circumvented by virtual private networks (VPNs).** Servers must know in which jurisdiction a user is located in order to enforce the right policy; this determination is often based on the user's IP address (especially on the web). VPNs – which are commonly used for accessing a variety of services without disclosing the user's IP address – allow users to appear to be in a jurisdiction which does not require age assurance. Some jurisdictions may attempt to restrict VPNs, which would have widespread negative security and privacy consequences for the large number of existing VPN users. VPNs are less effective with mobile apps, which can directly query the user's location, subject to user permission.

- **Device-based enforcement can be circumvented by obtaining a non-enforcing device.** Deployment of device-based age assurance on mobile devices is relatively straightforward, as most apps are installed through vendor-provided and controlled app stores which could be readily updated to restrict the use of non-compliant apps. It is less practical to require that desktop devices perform device-based age assurance, because software and operating system installation is less tightly controlled.

- **Many age assurance mechanisms allow a minor to cooperate with an adult to circumvent age assurance.** For example, an adult could buy a device for a minor and unlock it for the minor or let the minor use their credit card for credit-card-based age assurance. In some cases parents might assist minors in circumventing age assurance, but minors might also turn to older peers. Preventing this form of attack would require biometrically verifying the user at each intervention, which intensifies privacy and friction issues.

- **For many age assurance mechanisms, anti-circumvention relies on the fact that most mobile devices are closed systems.** Open systems on which the user can install software of their choice make circumvention easier, for instance by allowing users to bypass the camera and send forged "deepfake" video or by ignoring device-based enforcement. This is a larger issue for desktop devices than for mobile because most mobile devices are already largely closed ecosystems.

Taken together, these findings illustrate the inherent tradeoffs that characterize all currently available age assurance approaches. Different use cases place different demands on accuracy, availability, privacy, and resistance to circumvention, and no single mechanism excels across all of these dimensions on mobile and desktop. The suitability of different age assurance mechanisms varies significantly depending on whether the goal is to provide safer defaults or to block access entirely, and implementation choices—including whether evaluation and enforcement occur on servers or devices—have substantial implications for user privacy, system security, and dependency on closed device ecosystems. The technical assessment in this report illuminates how age assurance systems function in practice and the consequences that can be expected from their deployment.

III. Reference Architecture

This section provides a reference architecture for a typical age assurance interaction across both web-based and mobile-app based age assurance systems. Not all uses of age assurance will follow precisely this pattern but this architecture lays out the basic functions that need to be performed and provides context for the rest of the report.

The figure below shows a typical interaction with a web-based age assurance system.

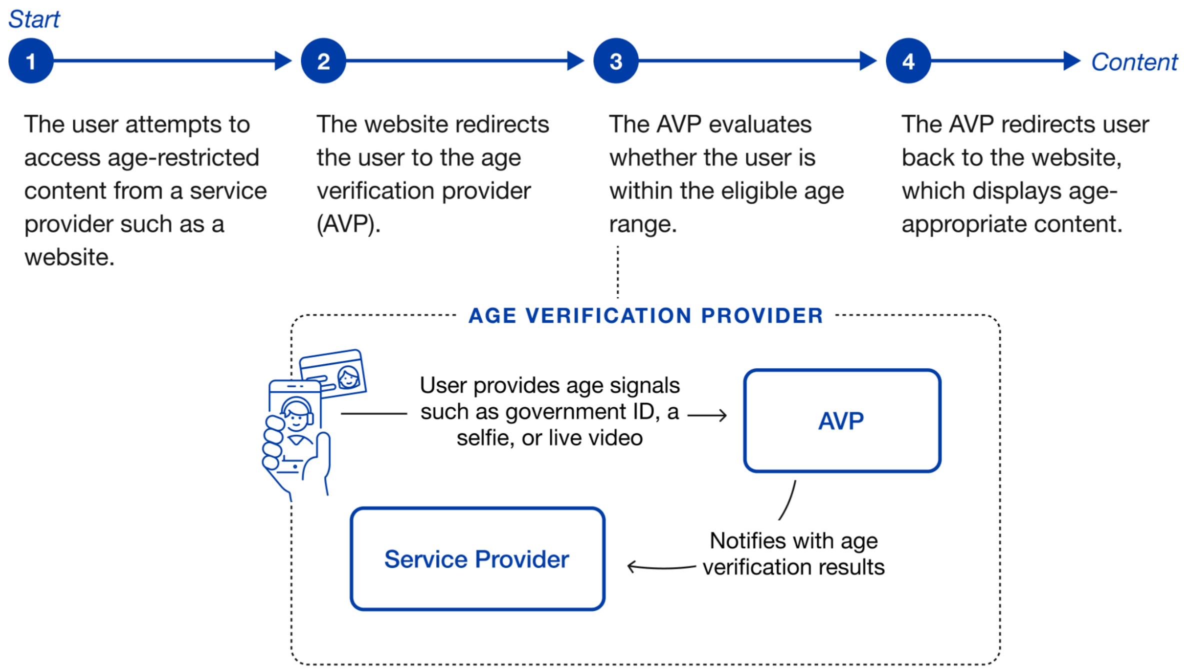

1. The user attempts to access age-restricted content from a service provider such as a website.
2. The website redirects the user to the age verification provider (AVP).
3. The AVP evaluates whether the user is within the eligible age range.
4. The AVP redirects user back to the website, which displays age-appropriate content.

User provides age signals such as government ID, a selfie, or live video

Notifies with age verification results

Figure 1. A typical web-based age assurance system.

The figure on the following page shows a typical age assurance architecture for mobile apps.

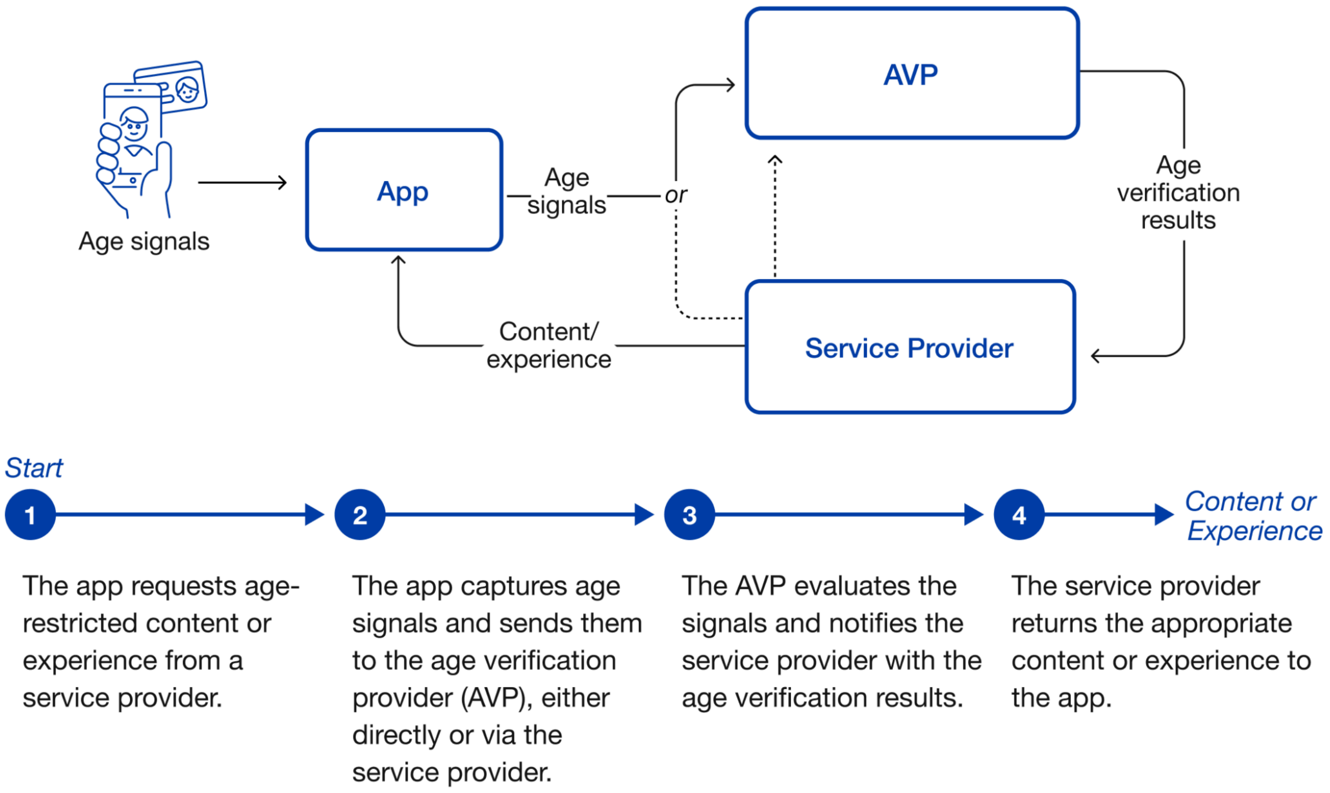

1. The app requests age-restricted content or experience from a service provider.
2. The app captures age signals and sends them to the age verification provider (AVP), either directly or via the service provider.
3. The AVP evaluates the signals and notifies the service provider with the age verification results.
4. The service provider returns the appropriate content or experience to the app.

Figure 2. A typical app-based age assurance system.

IV. Assessment Methodology

The report treats age assurance systems as security mechanisms, which are intended to grant or deny access based on properties associated with the user seeking access. It is intended to address the question of how well each system fulfills that function, i.e., does it effectively restrict access to services and experiences to users who are eligible while minimizing other negative consequences?

In order to examine this question, the report lays out a set of general assessment criteria and uses them to examine each age assurance architecture and signal in turn. The criteria used to assess age assurance systems in the report are:

- **Baseline accuracy**: the accuracy of the system in the absence of any attempts by the user to circumvent it.
- **Circumvention resistance:** the degree to which the system resists attempts by users to establish an age different from their true age.
- **Availability:** the degree to which the system will be usable by the eligible population.
- **Privacy:** the degree to which use of age assurance by a user reveals information that would not be accessible without the use of age assurance.

V. Age Assurance Architectures

Age assurance architectures can be broadly divided into two principal categories based on who is responsible for evaluating the user's age:

- The service provider (conventionally referred to as "server-based" architectures)
- The device or operating system vendor (conventionally referred to as "device-based" architectures)

A. Server-Based Age Evaluation and Enforcement

The most widely deployed architecture is to perform all age assurance functions on the server side. The diagram below shows the typical architecture, in which the service provider contracts with a third-party AVP, which performs evaluation, with the service provider performing enforcement on the basis of the evaluation results.

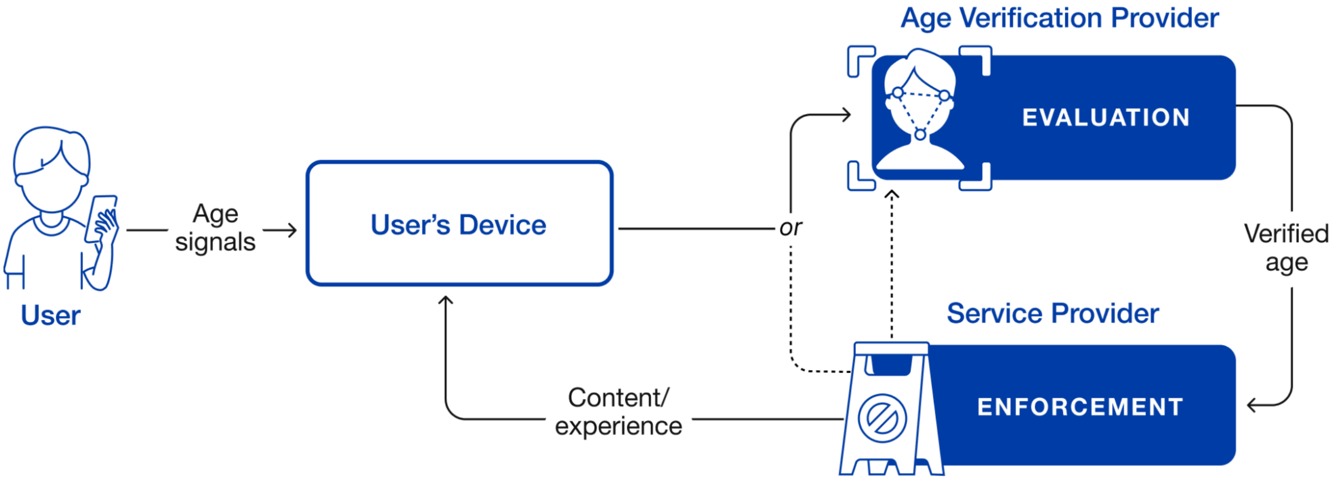

Figure 3. A typical server-based age assurance architecture.

B. Device-Based Age Evaluation

It is also possible to perform age assurance on the device. In this scenario, the device operating system would be responsible for acquiring the appropriate age signals and performing age assurance. At the end of this process, the device would then know whether the user was within the eligible age range (and potentially the user's exact age).

Once the device knows the user's age eligibility, there are two main options available for restricting access to age-restricted content and experiences:

- The device can prevent users from installing or running apps which access restricted services or experiences (for blocking use cases).

- The device can make the user's age status available to apps via an operating system API, and the apps then perform age enforcement (for blocking or safer defaults use cases).

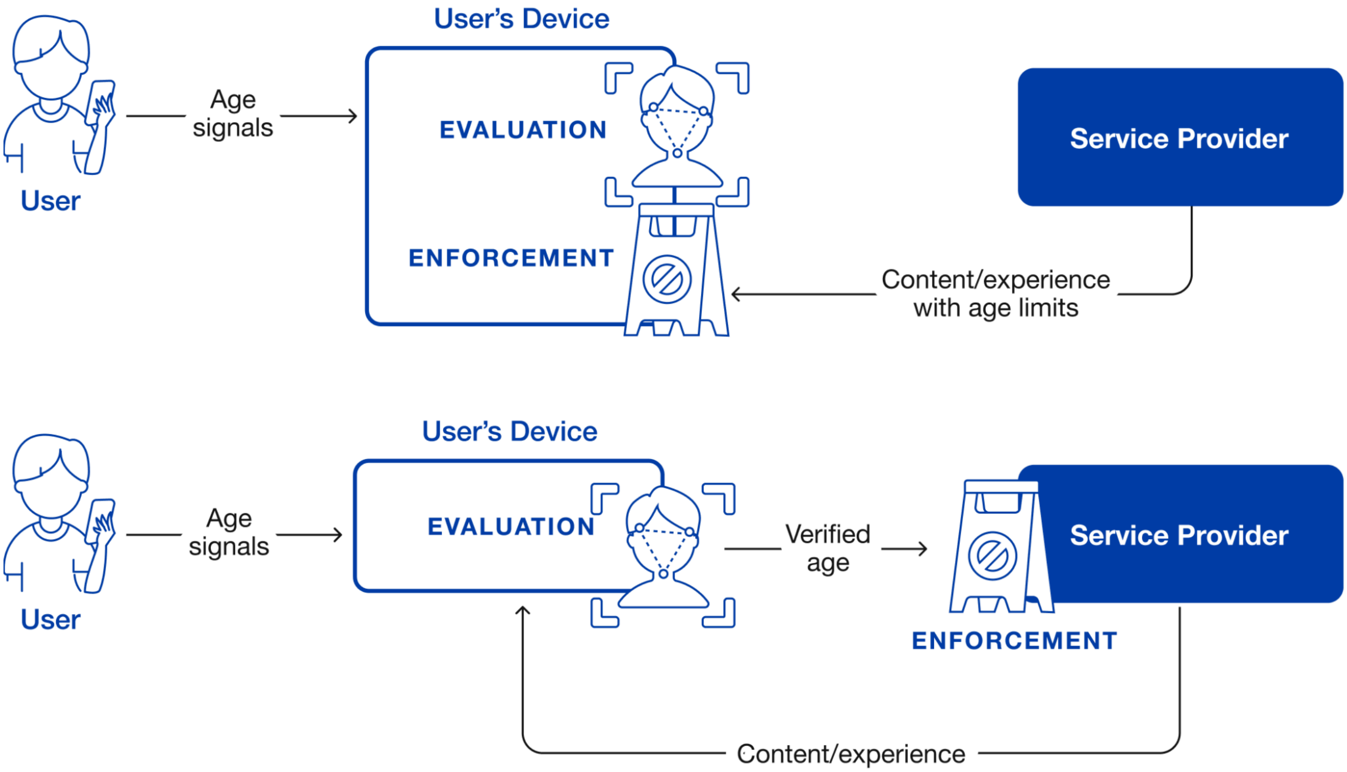

Figure 4. Two models for device-based evaluation. In device-based enforcement, the service provider offers content or experiences labeled with age limits, and the device determines whether to allow the user access to the content or experience based on the user's verified age. In server-based enforcement, the device sends the user's verified age to the service provider, which provides the appropriate content or experience.

VI. Age Signals

Current age assurance systems rely on a variety of different signals to evaluate the age of the user:

- **Self-declaration.** This is the most basic age signal, in which the user is asked to represent that they are over a given age ("Yes, I am over 18"), or, sometimes, to provide their birthdate.

- **Commercial and government records** (banking records, mobile network operator records, credit cards, other government and commercial records retrieved by name, email, etc.). A broad class of age assurance mechanisms—often referred to as "age inference"—uses government and commercial records tied to a user's identity. These mechanisms attempt to leverage pre-existing commercial relationships in which the user had to prove their identity and age. The relevant records in this category can be sorted into four groups: banking records, mobile network operator status, credit cards, and other government and commercial records.

- **Government IDs.** Government-issued identity documents such as driver's licenses and passports can be used for age verification. Presently, this mostly involves remote presentation of the physical card but it is increasingly possible to use digital forms of identification such as mobile driver's licenses (mDLs).

- **Facial age estimation.** In deployments of this mechanism, the user supplies a selfie or a self-video (potentially interactively) and the evaluator uses artificial intelligence or machine learning algorithms to estimate the user's age. As suggested by the name, these are *estimation* systems which do not provide an exact age but rather a probability distribution about the user's age.

- **Behavioral signals.** Some services have deployed age estimation technologies that infer the user's age or age range based on the behavior they observe of the user on the service. These systems use a wealth of data about how users interact, the content and accounts they engage with, the demographic information they provide, and other factors to infer users' ages or age cohorts.

The tables on the following pages provide summaries of how the above architectures and age signals perform in practice, providing a comparative view across baseline accuracy, resistance to circumvention, availability, and privacy.

VII. Assessment Summary

A. Assessment Summary for Age Assurance Architectures

	Baseline Accuracy	Circumvention	Availability	Privacy
Server-Based Evaluation and Enforcement	Depends on underlying age signals. Applying the correct jurisdictional policy depends on the server being able to determine the user's location.	Vulnerable to location spoofing via VPNs and to injection attacks on untrusted devices (for apps) and on the web generally.	High if untrusted devices are acceptable. Much lower if trusted devices are required to prevent injection attack.	Evaluators frequently learn information about the user, which can be abused.
Device-Based Evaluation	Depends on how the device determines the user's age.	Depends on whether the user can obtain an unlocked device or get an adult to obtain one for them. Circumvention is easier on desktop.	Device-based enforcement only restricts behavior on devices which are configured to enforce restrictions. Mobile app users on non-upgraded devices may be excluded.	Service providers do not learn anything other than that the user is in the eligible age range. Any user who wants an unrestricted experience must undergo age assurance.

B. Assessment Summary for Age Signals

	Baseline Accuracy	Circumvention	Availability	Privacy
Self-Declaration	High if the user is honest.	Easy.	Ubiquitous.	High. Moderate if birthday is requested.
Commercial and Government Records				
Banking Records	High.	Easy with access to an adult's account. Difficult otherwise.	Depends on having a bank account. A significant fraction of adults do not. Low availability for below 18s.	Evaluator does not learn the user's identity, but bank learns about age assurance. Evaluator learns the user's banking institution.
Mobile Network Operator Verification	Depends on the MNO's procedures for verifying age.	Easy with cooperation of an adult or temporary access to an adult's phone.	Only available in jurisdictions that impose default restrictions on mobile phones. Not practical for under 18s.	Evaluator learns the user's mobile number.
Credit Cards	Depends on issuer's procedures for verifying age.	Easy with cooperation of an adult or temporary access to an adult's credit card.	Only available in jurisdictions where credit cards are age-restricted. Depends on having a credit card, which a significant number of adults do not. Low availability for under 18s.	Evaluator learns the user's credit card number and usually postal code, and may learn the user's name and address if payment processor requires it.
Other Commercial and Government Records	Unknown. Reported false reject rates in excess of 10%.	Depends on the identifying information used. For birthdate, address, and SSN, fairly easy. Email address or mobile number verification is easy to circumvent with assistance of an adult, difficult otherwise.	Depends on quality of records. Reported false reject rates in excess of 10% suggests that this may be low.	Evaluator learns the user's identity or a proxy for their identity such as email address. Stored records are difficult to anonymize.

Government IDs				
Physical IDs	High.	Users may acquire a fake ID or attempt to use a borrowed ID. Remote attack detection is difficult.	Depends on prevalence of the underlying credential. In jurisdictions where IDs are not mandatory, significant fractions of adults do not have them.	Evaluator learns the user's identity as well as other personal information such as address. Evaluators may be able to misuse face image if provided.
Digital IDs	High.	Depends on the security of the device. May be possible for an adult to enroll their ID in a minor's device or allow their device to be used for a one-time age assurance.	Depends on prevalence of the underlying credential. Also requires a device which can enroll that credential for age assurance, which is not currently available in most jurisdictions.	Only reveals the user's age eligibility and not identity. Allows for linkage with the assistance of the credential issuer. May allow for linkage between evaluators if credentials are reused.
Digital IDs with zero-knowledge proofs	High.	Same as for Digital IDs.	Same as for Digital IDs.	Only reveals the user's age eligibility and not identity.
Facial Age Estimation	Many users in the eligible age range are rejected.	Depends on the implementation. Vulnerable to presentation attacks and very vulnerable to injection attacks.	Requires a device with a camera. If trusted devices are required to prevent injection attacks, then cannot be used on the web.	Evaluator learns the user's face. May be able to use this to identify the user or misuse it in other ways.
Behavioral Signals	Unknown.	Unknown. Opening a new account or using privacy tools can prevent creation of a behavioral profile.	High. Challenging to use for primary age assurance because it cannot provide results for new users.	Requires storing and retaining a profile of user behavior, even if the provider does not already do so.

VIII. Conclusion

In recent years, an increasing number of jurisdictions around the world have begun evaluating and adopting age assurance requirements of different kinds. Collectively, these moves represent a major change from how online services have been accessed over many decades, and they implicate a variety of important concerns and values for consumers, both adults and youth.

Age assurance is not a single technology but a suite of technologies which must be used together in combination in order to build an age assurance system. Understanding the properties of these technologies is essential both to deploying effective systems and crafting effective age assurance requirements.

The first and most important consideration is what use cases age assurance is intended to serve. The requirements for an age assurance system which is intended to prevent minors from accessing content are different from the requirements for a system which is intended to ensure that minors have safer defaults. Distinguishing age ranges below 18 is also more difficult because estimation methods are imprecise and those under 18 often do not have ID which establishes their precise age.

Because all existing age signals either have high error rates or exclude significant fractions of the population, any practical age assurance system needs to support multiple age signals. This allows users who are unable to establish their age via one signal to “fall back” to another signal. A common design is a “waterfall” in which users are presented with a low-friction signal such as facial age estimation and then ask users who are unable to establish their age with that signal (e.g., because they are close to the age threshold) to use a more precise but higher friction signal such as showing ID.

Just as age signals have different error rates, they also have different privacy properties. The most commonly deployed signals effectively disclose the user’s identity to the age verification provider or service provider. This concern is of lesser importance in cases where the user discloses their identity anyway (e.g., to make an account on a social media service), and greater importance in other cases when users have a prior expectation of anonymity. There are two emerging approaches which have superior privacy properties: zero-knowledge proofs based on government IDs and device-based age assurance. Zero-knowledge proofs can be deployed in parallel with existing age signals, allowing users with compatible devices and software to enjoy superior privacy properties. Alternately, device-based age assurance allows users to establish their age to the device manufacturer without having to reveal personal information to services with whom they have no existing relationship.

Minors may be motivated to circumvent age assurance if it prevents them from accessing content or experiences that they want. All age assurance systems are vulnerable to circumvention in one form or another. It is not practical to prevent all circumvention without also restricting devices and networks in ways that would have severe detrimental impacts on many legitimate uses of the internet. Many systems allow a minor to cooperate with an adult to evade age assurance. In cases where adults view age restrictions as illegitimate, they may be more likely to assist minors in circumventing them.

Finally, there is an important tradeoff between openness and security. Because open systems are more vulnerable to circumvention than closed systems, there is an inherent tension between policies that are designed to give users more control of their own devices and those which are designed to prevent minors from accessing certain content and experiences. There are inherent tradeoffs between the level of circumvention resistance and the degree to which adult users' ability to control their own devices and experiences is restricted.

Age assurance technologies are complex systems that are being deployed on a wide scale on the internet for the first time. Understanding how these systems work, along with their capabilities and limitations, is essential to making good decisions about the use of these emerging technologies.