Knight
Georgetown
Institute

JANUARY 2026

Age Assurance Online:

A Technical Assessment of Current Systems and Their Limitations

Eric Rescorla
Zander Arnao
Alissa Cooper
Knight-Georgetown Institute

About the Knight-Georgetown Institute

The Knight-Georgetown Institute (KGI) is dedicated to connecting independent research with technology policy and design. KGI serves as a central hub for the growing network of scholarship that seeks to shape how technology is used to produce, disseminate, and access information. KGI is designed to provide practical resources that policymakers, journalists, and private and public sector leaders can use to tackle information and technology issues in real time. Georgetown University and the Knight Foundation came together to launch the institute in 2024. Learn more about KGI at https://kgi.georgetown.edu.

Acknowledgements

This report was strengthened through feedback from a range of experts. KGI would like to thank the following reviewers: Alice Marwick, Allison Ivie, Daniel Kahn Gillmor, Dennis Jackson, Elizabeth Galicia, Eric Mill, John Perrino, Joseph Lorenzo Hall, Marshall Erwin, Martin Thomson, Mike Specter, Nick Doty, Paul Ohm, Riana Pfefferkorn, Richard Barnes, and Zeve Sanderson.

Review of the report does not constitute an endorsement of the report's conclusions or recommendations. Any errors or omissions are the responsibility of KGI.

The authors would also like to recognize the contributions of the KGI team, including Julie Anne Miranda-Brobeck, and Peter Chapman, as well as Kate Hudson, who was responsible for the figures.

Executive Summary

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, this report provides a comprehensive technical assessment of the landscape of age assurance technologies. The report assesses each of the most commonly deployed age assurance architectures and mechanisms (referred to as "age signals") against a common set of criteria: baseline accuracy, circumvention resistance, availability (which refers to the ability of eligible users to access age-restricted services), and privacy. The key findings of the technical assessment 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.

Privacy protection: 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.

Table of Contents

I. Introduction

There is broad agreement that some content and experiences on the internet are not suitable for minors. What to do about this and whose responsibility it should be is the subject of a long-running debate in the technology and policy communities.

In recent years, an increasing number of jurisdictions around the world have begun evaluating and adopting age assurance requirements of different kinds. Numerous US states, the UK, EU, Australia, and other countries have passed laws requiring that some service providers offer different experiences or restrict certain classes of content to users in specific age ranges, typically to users above a certain age. Many more jurisdictions are considering similar requirements. 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, this report provides a comprehensive technical assessment of the landscape of age assurance systems.

The report assesses two age assurance architectures:

- Server-based age evaluation
- Device-based age evaluation

These architectures can be used with multiple types of age assurance mechanisms (referred to as “age signals”). The report assesses each of the most commonly deployed age signals against a common set of criteria: baseline accuracy, circumvention resistance, availability (lack of impediments for eligible users to be able to access age-restricted services), and privacy. The age signals covered are:

- Self-declaration
- Commercial and government records (banking records, mobile network operator records, credit cards, other commercial and government records retrieved by name, email, etc.)
- Government IDs (in both physical and digital form)
- Facial age estimation
- Behavioral signals

Importantly, this report does not address the broader question of the desirability of age assurance requirements or whether any particular service or experience is age-appropriate. Rather, it takes those policy questions as a given and asks whether age assurance mechanisms effectively deliver on the stated policy objectives. The focus of the technical assessment is on whether existing and proposed age assurance mechanisms are fit for their claimed purpose and what effects can be expected from their deployment. The results in this report are intended to inform discussions and decisions about the use of age assurance, including age assurance mandates, and best practices for deployment of age assurance systems.[1]

Age assurance systems are complex to implement and deploy correctly, and the stakes of imposing widescale age assurance requirements are high. The results of this technical assessment reinforce the need for policymakers, service providers, and age verification providers to exercise care in the design and deployment of age assurance systems and requirements. The key findings of the assessment are:

<u>Multiple use cases</u>: 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. The privacy impact and circumvention incentives are not the same in all cases. Tailoring approaches and guidance to each use case is critical.

<u>Multiple age signals</u>: No single age signal is sufficient on its own. Facial age estimation is highly available but inaccurate near the age threshold. Systems based on commercial and government records are often unable to verify the ages of eligible users. Government-ID-based systems are accurate but not always available. Age thresholds below 18 are harder to deploy because minors often do not possess government-issued ID. Behavioral signals are less suitable for primary age assurance because they first require users to develop a behavioral history and cannot be used to determine a user’s age on first contact. As a result, 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 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.

<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 with their data. Systems with stronger technical privacy guarantees—device-based enforcement and zero-knowledge proofs—are possible but not widely deployed.

[1] Some of the material in this report is inevitably technical, but it is intended to be broadly accessible to readers without advanced knowledge of computer science. Where some of the discussion strays into more technical topics, as with zero-knowledge proofs, the report summarizes the conclusion for the lay reader and provides more technical material in appendices.

Circumvention: 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.

Taken together, this report's 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 technical characteristics of different age assurance mechanisms vary 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.

The remainder of this report is organized as follows. Section II defines terminology. Section III explains age assurance fundamentals. Section IV discusses the policy landscape. Section V explains the assessment methodology. Sections VI and VII assess age assurance architectures and age signals respectively. Section VIII synthesizes the key findings of that assessment. Section IX addresses the broader impacts of widespread deployment of age assurance. Section X concludes.

II. Terminology

This section provides a brief overview of the relevant terminology used in this report.

Age assurance is a broad term meant to capture age estimation, age inference, and age verification systems.

Age estimation systems attempt to estimate a user's age to within some level of error, e.g., from a still picture or from a live video interaction.

Age inference[2] systems make use of records tied to a person's identity to assess whether a user is within a given age range.

Age signal refers to a piece of information which is used in the process of age assurance to help determine or estimate a person's age.

[2] Although the term "age inference" is widely used, it can create confusion because "inference" (or "AI inference") can also refer to the process of using an AI model to classify a specific input, as is done in age estimation.

Age verification systems attempt to ascertain with precision whether a user fits within a given age range. Note that these mechanisms may determine the user's exact age but then output a yes or no response for the target age range.

Age verification providers (AVPs) are entities that provide the service of age verification mechanisms or age assurance for their customers, such as service providers.

Application Programming Interface (API) is an interface used by one piece of software to talk to another, either on the same device or across a network.

App integrity refers to a set of remote attestation mechanisms that allow app vendors to verify that a given user device is running an unmodified copy of their app.

Availability is to the extent to which the eligible user population will be able to use a given age assurance system.

Baseline accuracy is the accuracy of an age assurance system under normal conditions (i.e., when it is not under attack).

Closed device is a computing device where the user has limited or no ability to control the software that runs on it.

Credential issuer is an entity which issues identity credentials, whether physical or digital.

Digital IDs are a form of identification which is stored on a computer or mobile device rather than on a physical card or booklet. Digital IDs are often the digital counterpart of a form of physical ID such as a driver's license.

Digital signature is a cryptographic technique that allows a user to "sign" a piece of data using their private key in such a way that the signature can be "verified" by anyone with the public key as reflecting the same data value and being generated by the private key.

Eligible age range is the range of ages where users are eligible to access a service or experience.

Enforcer is the role in an age assurance system that uses the results from the evaluator to determine which content and experiences a user is eligible for.

Evaluator is the role in an age assurance system that determines whether a user is within a given age range.

Hashing is a one-way transformation on data which converts an arbitrary amount of data into a characteristic shorter value.

Injection attack is an attack on a remote video or still-image authentication system where the attacker bypasses the camera and provides their own video stream or image.

Internet Protocol (IP) address is the numeric network-level identity for a device on the internet, allowing it to be reached by other devices.

Mobile Driver's License (mDL) is a digital ID based on a person's driver's license.

Open device is a computing device that allows the user to control the software that runs on it.

Presentation attack is an attack on a remote video or still-image authentication system where the attacker manipulates the input to the camera, such as by holding a photo up to the lens or wearing a mask.

Private key refers to a cryptographic key which can be used together with the corresponding public key. The private key is kept secret.

Public key refers to a cryptographic key which can be used together with the corresponding public key. The public key can be safely distributed.

Random errors are errors that do not necessarily occur consistently even under similar conditions, such as a user who is sometimes estimated as 18+ and sometimes as 17.

Remote attestation is a process whereby a computing device demonstrates to some other device that it is running a specific piece of software in a specific untampered configuration.

Secure element is a processor or part of a processor in a computing device that is designed to resist attacks, such as exfiltrating secret data, including from an attacker who otherwise controls the computing device. Secure elements are often used for cryptographic key storage.

Service providers provide websites, mobile applications, games, or other services or content.

Systematic errors are errors that repeatedly occur under similar conditions, such as a user who is 18+ being consistently estimated to be 17.

Virtual private network (VPN) is a technology for concealing the IP address of a device by forwarding it through another computer with its own different IP address.

Zero-knowledge proofs (ZKPs) are a cryptographic technology which allows one party to prove knowledge of a statement without revealing any other information. These can be used as part of a privacy-preserving authentication system.

III. Age Assurance Fundamentals

This section provides an overview of the fundamentals of age assurance.

A. Use Cases

At a high level, the relevant use cases for age assurance can be grouped into two primary categories:

- **Safer defaults.** Providing minors with a tailored set of default settings, such as restricting the use of personalized feeds or notifications during certain hours. The majority of these cases are associated with social media and AI chatbots.

- **Blocking.** Entirely restricting access to certain classes of content or experiences, such as adult content or gambling, though in some cases also including social media. Many of these cases are associated with services which are entirely restricted, but there can be cases where only part of a service is restricted, for example a website that hosts generally available content but also hosts some age-restricted content.

These categories are not always sharp but can be helpful in understanding the requirements for various scenarios, as discussed throughout this report.

B. Identity, Pseudonymity, and Anonymity

Services vary widely in the extent to which they require users to identify themselves. There are at least three common scenarios:

- **Anonymous usage.** The vast majority of websites and some apps do not require users to identify themselves at all. For example, general purpose search sites such as Google allow users to access the site with no name or account, although many of these sites allow accounts to be created. Users concerned about technical identifiers being used to trace their identities (IP address, cookies, etc.) can use widely available anonymization tools for further privacy protection if they choose.

- **Account required.** Some services require users to create accounts but do not require the user to provide a verifiable name. In some cases, these services require that the user provide a reachable contact email address or phone number, and they may even ask for a name or birthdate, but do not require these identifiers to be verifiable. Example services in this category include social media sites such as X and Facebook[3] and dating apps such as Tinder and Grindr.[4]

- **Verifiable identities.** Finally, some services require users to provide verifiable identities, for example by requesting that they upload a picture of a government ID. For example, inside the US, the DraftKings gambling site requires the user to provide a full name, address, and date of birth, and may require them to provide their social security number or upload a picture of their ID to successfully verify their identity.[5]

Note that in many cases, services allow access with a low level of self-identification but provide an enhanced experience (e.g., a badge indicating that the user is verified) if the user conducts a more thorough self-identification process.[6]

Even if users do not directly identify themselves to a site, they may already be sharing significant amounts of information. For example, while Instagram does not require users to verify their names, users frequently upload pictures of themselves. Similarly, users frequently upload pictures of themselves to dating apps and some apps such as Tinder require pictures of the user's face. Moreover, routine usage of many sites involves disclosing personal information, both with explicitly personal sites (medical, dating, etc.) and general purpose sites such as search engines.

Depending on the amount of identity data already shared with the service, users may have different privacy expectations. For example, if the user is expecting to access a site anonymously, then requiring them to identify themselves for age assurance purposes has a larger additional impact on their privacy than if they already had to provide a government ID to use the site at all.

In many cases, legal requirements for safer defaults are aimed at services where an account is already required, because the motivation is to present a different set of defaults to users once they have logged into the service. Legal requirements for blocking have been applied across different kinds of services that support the full range of identification options from anonymous to verifiable.

C. Reference Architecture

This section provides a reference architecture for a typical age assurance interaction. 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.

[3] Facebook nominally requires their users to "provide for your account the same name that you use in everyday life," but does not generally require them to verify that name. Meta, "Terms of Service."

[4] Some services will also require age assurance for users they believe to be under 18.

[5] DraftKings, "Why am I being asked to verify my identity? (US)."

[6] Tinder, "ID + Photo Verification." Note that in some regions Tinder requires a face scan as a liveness check, but does not require government ID. See Tinder, "FAQ Mandatory Liveness Check."

Age assurance systems involve four key actors:

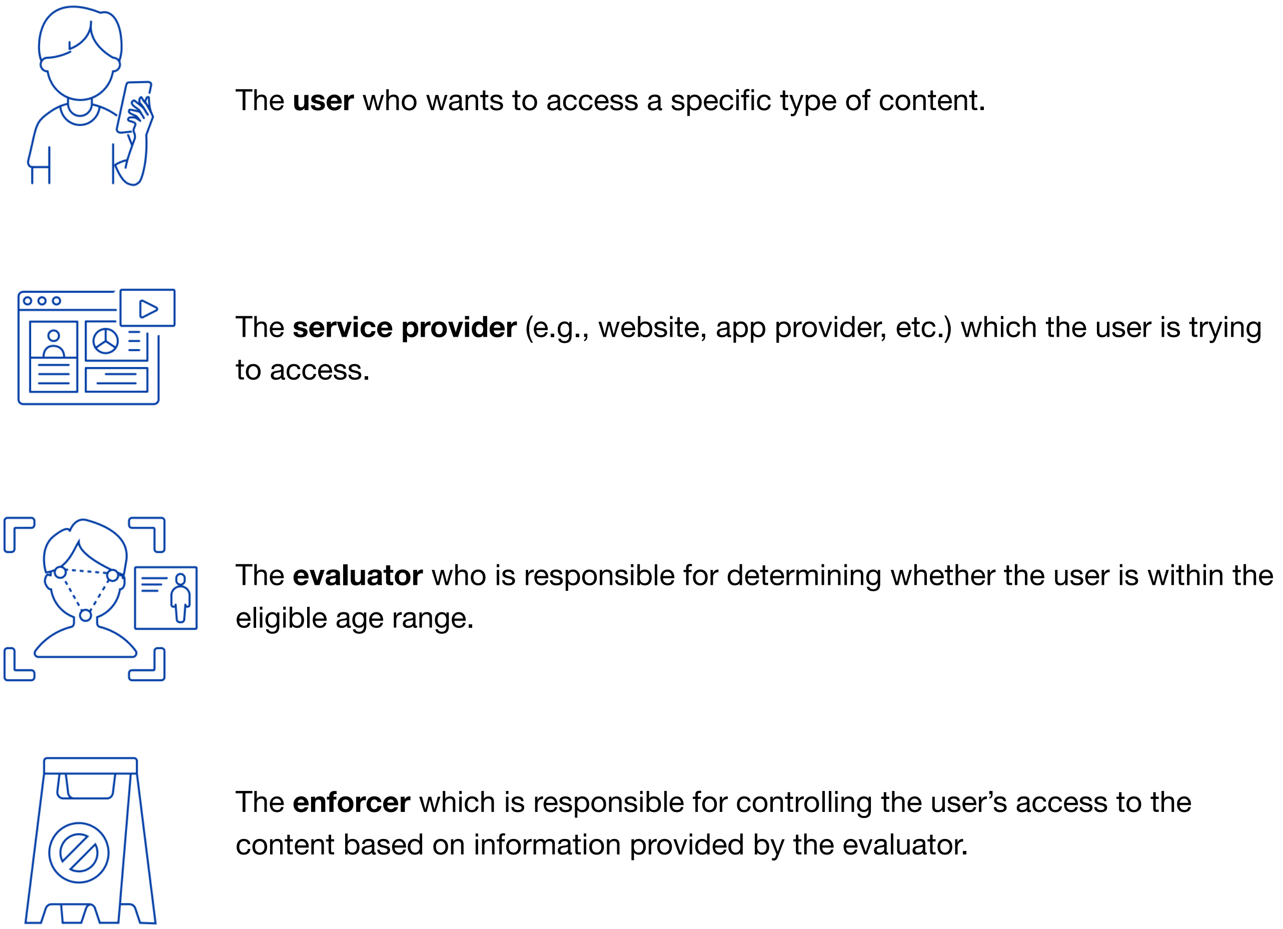

The **user** who wants to access a specific type of content.

The **service provider** (e.g., website, app provider, etc.) which the user is trying to access.

The **evaluator** who is responsible for determining whether the user is within the eligible age range.

The **enforcer** which is responsible for controlling the user's access to the content based on information provided by the evaluator.

Figure 1. Age assurance roles.

The evaluator and enforcer roles can be implemented in a number of locations, for instance on the user's device, or on separate internet services. This report uses the names of the roles rather than the concrete entities that are implementing them in order to retain clarity about where functions are being performed.

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

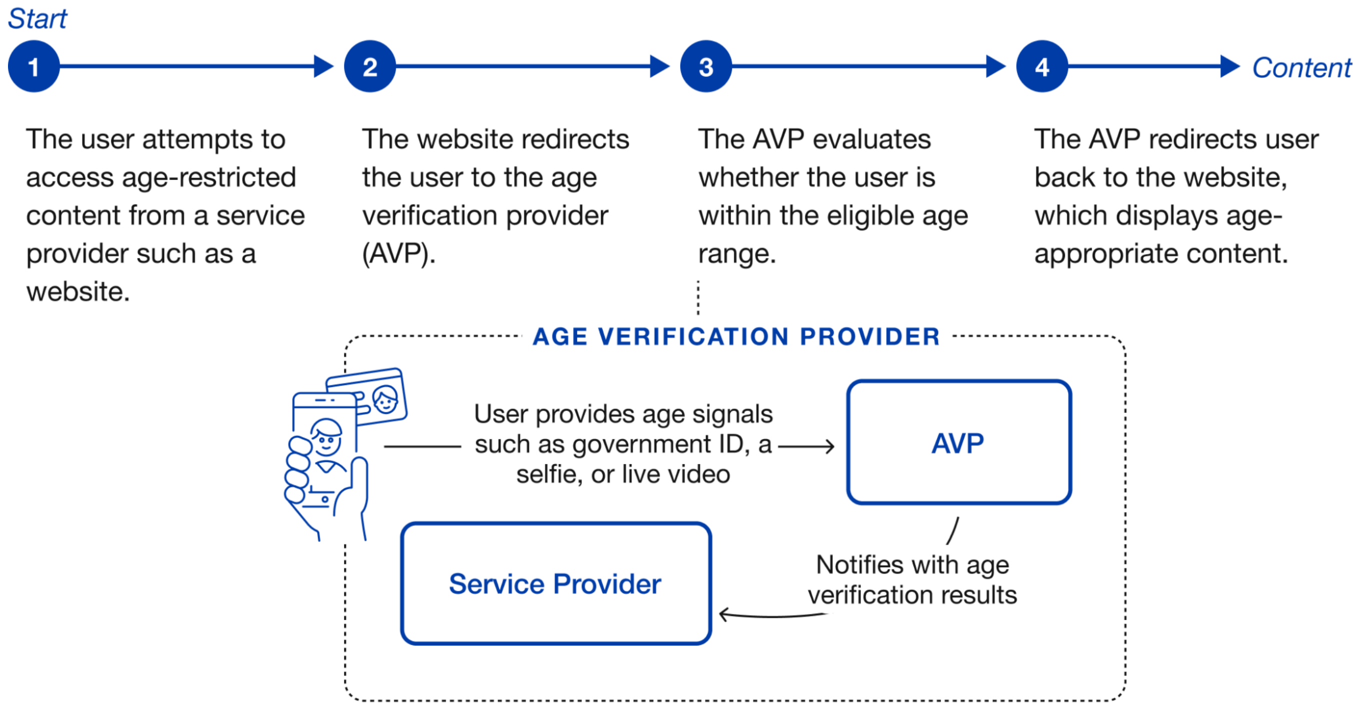

1 The user attempts to access age-restricted content from a service provider such as a website.

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 2. A typical web-based age assurance system.

The process starts with the user visiting the **service provider** which has some age-restricted content or which is designed to provide the user with an age-appropriate experience. The service provider redirects the user to their selected **age verification provider (AVP).** Depending on the user interface, this may be largely invisible to the user, appearing as just a transition to a different web page and only secondarily to another entity. The AVP requests that the user provide some *age signals*, such as government ID, a selfie, or live video (the following section covers the primary age signals in use).The AVP then evaluates those signals and based on them determines whether the user meets the required age criteria and notifies the service provider accordingly. The AVP then redirects the user back to the service provider so that the user can view the content.

The figure below shows a typical age assurance architecture for mobile apps:

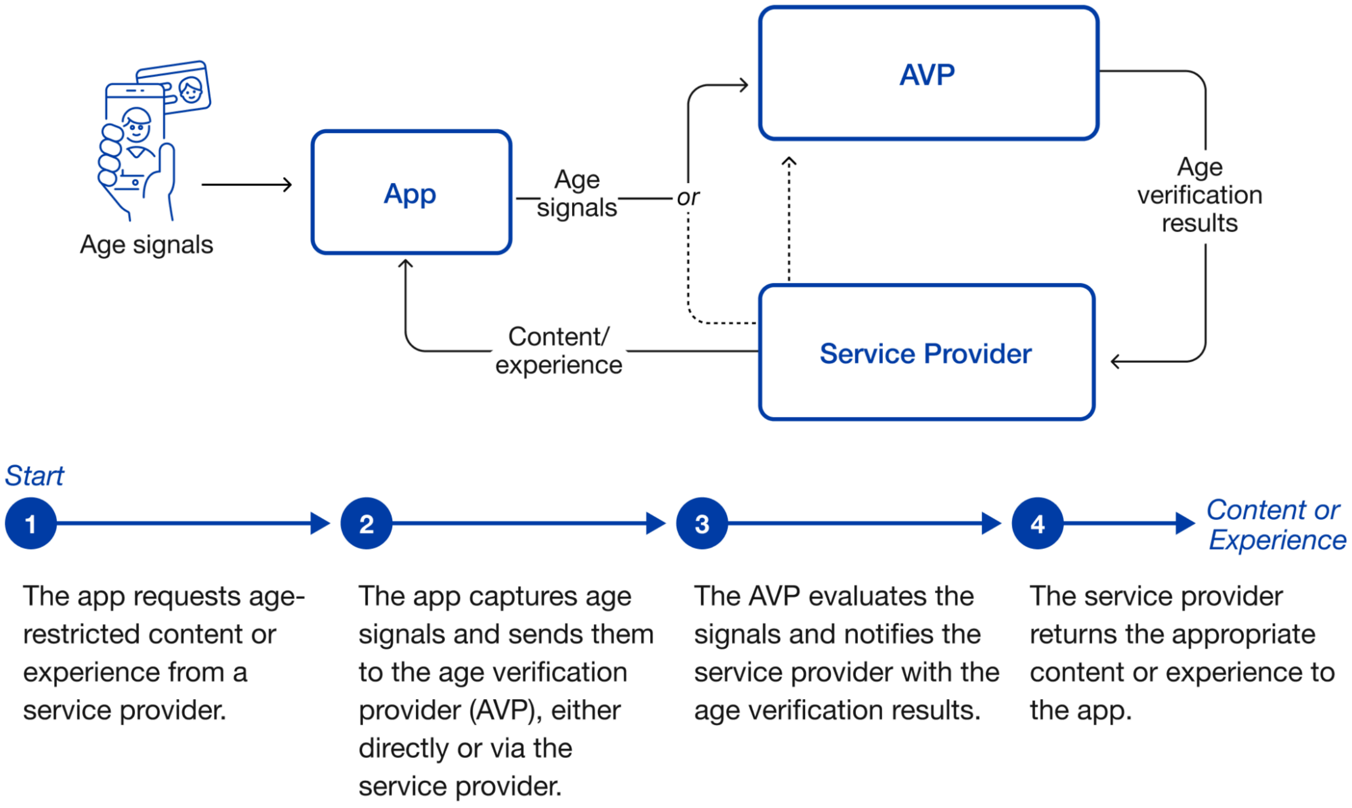

1. The app requests age-restricted content or experience from a service provider.
2. The app captures age signals and sends them to the age verification provider (AVP), either directly or via the service provider.
3. The AVP evaluates the signals and notifies the service provider with the age verification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

Figure 3. A typical app-based age assurance system.

This architecture also performs all the evaluation and enforcement on the server, but instead of having the web browser visit the AVP, the app captures the appropriate age signals and sends them to the AVP directly or to the service provider who can forward them to the AVP. The AVP notifies the service provider of the age result, and the service provider can provide the appropriate content or experience, either on the server or in the app.

Importantly, the question of what kinds of signals the user presents to demonstrate their age is largely orthogonal to the architecture of the age verification system and in fact usually does not need to be known by the enforcer or the service provider because the evaluator just returns a yes or no answer.[7]

It is also possible to have evaluation and/or enforcement happen on the device. Section VI discusses the various possible architectures and their implications.

[7] The situation is more complicated in cases where the system wants to enforce a maximum age as well, because the service provider needs to know when users age out of a given range.

D. Age Ranges

In most regimes where age enforcement is desired, the actual question is not the user's precise age but rather whether their age falls into a specific age range. One common case requires users to prove that they are at or above a certain age. The lower bound in these cases varies according to the legal regime and features of the service. Some example lower bounds include: 13 (COPPA, EU GDPR minimum), 16 (GDPR default), 17 (proposed in COPPA 2.0 and Kids Online Safety Act (KOSA)), 18 (common age of majority), and 21 (minimum age to purchase alcohol in the US).

In some cases, services are tailored to specific age ranges, implying the need to enforce both minimum and maximum ages. For example, Instagram aims to provide a suite of safer defaults and protections for users aged 13-17, with users aged 13-15 requiring parental consent to override some of these settings.[8] Roblox provides differentiated feature access for users under 13, 13-17, and 18 and over.[9] Similarly, the social media site Yubo operates two communities, one for under 18 and one for 18+.[10] Effectively enforcing these ranges requires the ability to detect when someone is over a certain age as well as when they are under a certain age. This is more challenging for a number of age assurance mechanisms, and creates the additional problem of knowing when a user who was previously eligible ages out of the permitted range, which requires either storing the user's approximate age or frequently revalidating their eligibility.

IV. Policy Landscape

The past decades have seen vigorous debate about age assurance requirements, along with the more recent broad introduction of those requirements across many jurisdictions. This section briefly surveys these existing policy debates and the landscape of relevant policy activity.

Debates on age assurance have centered on exactly when the law should require evidence of a user's age and how this process should unfold. This includes arguments that age assurance should be proportionate to risks and tailored to the specific design and affordances of individual platforms.[11] Given this desire for flexibility, many analysts call for a spectrum of acceptable methods that enable multiple options for users in different contexts,[12] while some champion particular approaches and technologies.[13]

[8] Instagram, "Instagram Teen Accounts"; Instagram, "Introducing Instagram Teen Accounts."

[9] Kaufman, "Revolutionizing Digital Connection"; Kaufman, "Roblox Announces Ambitious Plan to Expand Age Estimation to All Users."

[10] Yubo, "Staying Safe on Yubo."

[11] 5Rights Foundation, *But how do they know it is a child?*; Brennen and Perault, *Keeping Kids Safe Online*; Digital Trust & Safety Partnership, *Age Assurance*; Grosshans, "Comments."

[12] 5Rights Foundation, *But how do they know it is a child?*; Grosshans, "Comments"; Hales, "Re: Advanced Notice of Proposed Rulemaking."

[13] Hogg and Swarztrauber, *On the Internet, No One Knows You're a Dog*.

The potential for mandatory age verification to restrict lawful speech was a critical concern during the formative years of the policy framework shaping the web,[14] and those concerns remain hotly debated today. Wide scale mandates inevitably have the potential for collateral effects on speech and expressive activities of both adults and minors, and questions about the balance of interests, chilling effects, and the appropriateness of mandates versus parental or user controls continue to be contested. Some advocates have long held that parental controls[15] are a less restrictive alternative to mandatory age assurance,[16] but current policy debates reveal frustrations with limited uptake and effectiveness.[17]

Among civil society voices, data misuse and collection represent a common concern about age assurance. Because a digital service often must receive sensitive information to infer or verify a user's age, these systems might inadvertently weaken anonymity and chill expression online.[18] Significant disagreement exists as to the extent of these problems. To some, age assurance threatens to fundamentally destabilize data privacy and security on the internet,[19] while others are more confident that advances in technology can ameliorate these risks.[20] Commentators in all cases stress that any mandate must minimize the processing and retention of personal information, and some suggest relying on trusted third-parties to conduct age assurance.[21]

Another source of concern is the real-world performance and feasibility of age assurance systems. Many question whether existing technologies can deliver accurate and equitable outcomes, given the risk of misclassification and uneven effects across social groups.[22] Other concerns center on the financial burden and uncertainty surrounding expectations created by age assurance requirements, particularly for smaller services.[23] To help address these issues, observers have discussed the need for standardization and quality benchmarks to support trustworthy and interoperable approaches to age assurance.[24]

[14] Supreme Court of the United States, "Ashcroft v. ACLU"; Supreme Court of the United States, "Reno v. ACLU"; United States, "Communications Decency Act"; United States, "Child Online Protection Act."

[15] The technical controls imposed by these systems are similar to those contemplated as part of age assurance systems, but they are different from age assurance in a number of important respects. First, parental controls are opt-in, with the parent or administrator having to install them, whereas age assurance systems are either opt-out (with parental consent) or mandatory (no parental consent). In order to ensure this, age assurance systems require third party validation of age eligibility based on uniform standards, whereas parental controls systems allow the administrator of the device to determine whether controls should be in place based on their judgement of what is appropriate for each minor.

[16] Barthold et al., "Brief of Amici Curiae."

[17] Family Online Safety Institute, "Parental Controls for Online Safety are Underutilized, New Study Finds"; Tenbarge, "Fewer than 1% of parents use social media tools to monitor their children's accounts, tech companies say."

[18] Forland et al., *Age Verification*; Goldman, "The 'Segregate-and-Suppress' Approach to Regulating Child Online Safety."

[19] Ibid.

[20] Hogg and Swarztrauber, *On the Internet, No One Knows You're a Dog*; Tutor, *Age Verification*.

[21] Brennen and Perault, *Keeping Kids Safe Online*; CNIL, *Online age verification*; Grosshans, "Comments"; Hales, "Re: Advanced Notice of Proposed Rulemaking"; Forland et al., "Age Verification"; Tutor, *Age Verification*.

[22] Brennen and Perault, *Keeping Kids Safe Online*; Forland et al., "Age Verification"; Stockwell and Powell, *Age Assurance Technologies and Online Safety*.

[23] Brennen and Perault, *Keeping Kids Safe Online*.

[24] 5Rights Foundation, *But how do they know it is a child?*; Brennen and Perault, *Keeping Kids Safe Online*; Hogg and Swarztrauber, *On the Internet, No One Knows You're a Dog*.

Despite these disagreements, lawmakers in many jurisdictions have begun imposing age assurance requirements, producing an uneven and rapidly evolving policy landscape.

A. United States

In the 1990s, the United States (US) attempted to regulate the ability of minors to access certain categories of content. First, the Communications Decency Act (CDA) restricted speakers and online services from purposely sharing obscene or indecent content with minors,[25] but this legislation was quickly struck down by the Supreme Court in 1997.[26] In response, Congress passed the Child Online Protection Act in 1998, which barred commercial websites from hosting content harmful to minors unless they implemented age verification;[27] however, this law was blocked from enforcement and later invalidated in 2004.[28]

The Children's Online Privacy Protection Act (COPPA),[29] also passed in 1998, restricts the use of the personal information of children under 13, but does not require that service providers verify the ages of their users. As a consequence, while service providers may limit use by under 13 children either via their terms of service or explicit age gates, they have not historically needed to verify user ages beyond self-declaration. Although federal legislation has been introduced to require age assurance,[30] there are no federal requirements in this area at the time of this writing.

As of this writing, over 25 US states have passed some form of age assurance requirements.[31] This legislation falls into two broad categories:

- Requiring service providers to exclude minors from material that the state deems harmful to minors, such as pornography.
- Requiring social media services to provide a different default experience to minors and/or to obtain parental consent for minors to use the service at all.

In addition, Louisiana, Texas, and Utah have imposed requirements for app stores to determine the user's age and to require parental consent for new app installation.[32] California AB 1043,[33] effective January 1, 2027, requires operating system providers to request-–but not verify-–a user's birthdate and make it available to applications via an application programming interface (API). Applications are

[25] United States, "Communications Decency Act."
[26] Supreme Court of the United States, "Reno v. ACLU."
[27] United States, "Child Online Protection Act."
[28] Supreme Court of the United States, "Ashcroft v. ACLU."
[29] United States, "Children's Online Privacy Protection Act."
[30] See, e.g., United States, "S.737 - SCREEN Act"; United States, "S.2714 - CHAT Act"; United States, "S.3062 - GUARD Act."
[31] Age Verification Providers Association, "US state age assurance laws for social media"; Age Verification Providers Association, "US State age verification laws for adult content"; Free Speech Coalition Action Center, "State Age Verification Laws."
[32] Louisiana, "Act No. 481"; Texas, "App Store Accountability Act"; Utah, "App Store Accountability Act."
[33] California, "AB 1043."

required to request age information and use it for compliance with other laws, but AB 1043 does not itself require content, feature, or experience restrictions for minors.

These age assurance mandates also implicate an emerging trend of states passing laws regulating how minors interact with “companion chatbots,” i.e., generative AI products designed with human-like features such as names and the ability to make friendly conversation.[34] These laws create general safeguards for users and heightened protections for minors.[35] For instance, California SB 243, signed into law in 2025, requires operators of companion chatbots to notify users that they are interacting with an AI system, and to implement certain safety protocols to reduce risks like self-harm when users are known to be underage.[36] Age assurance requirements might affect when such provisions of “safer default” laws for companion chatbots are triggered.

The trend of states passing age assurance laws is likely to continue. Many states which have not passed such laws are actively considering legislation that would implement age assurance for both safer defaults and blocking.[37] Polling suggests that large majorities of US residents support requiring parental consent before minors can use social media as well as age assurance for all users.[38]

Importantly, this trend will continue to be shaped by ongoing litigation related to age assurance. Though the Supreme Court recently blessed a Texas state law requiring online distributors of pornography to verify the ages of users in the landmark *Free Speech Coalition v. Paxton* case,[39] the extent to which age assurance requirements are permissible under the First Amendment for other services such as social media remains an open question. As such, various state laws imposing such requirements are being challenged in court,[40] creating uncertainty in the medium term.

In addition, state Attorneys General have initiated lawsuits related to the age assurance practices of major platforms under existing consumer protection laws. For instance, in 2023, a group of plaintiffs representing state Attorneys General as well as individuals and school districts filed complaints against Meta, Snap, ByteDance, and Google alleging that their failure to verify ages (among other design choices) amounts to deceptive and unfair business practices,[41] with parallel litigation initiated individually by several Attorneys General.[42]

[34] Gluck, “Understanding the New Wave of Chabot Legislation.”
[35] Ibid.
[36] California, “SB 243.”
[37] See, e.g., Minnesota, “SF 2105”; North Carolina, “HB 301.”
[38] Anderson and Faviero, “81% of adults - versus 46% of teens - favor parental consent for minors to use social media.”
[39] Supreme Court of the United States, “Free Speech Coalition v. Paxton.”
[40] See, for instance, United States District Court for the Northern District of California, “Order Granting in Part and Denying in Part Motion for Preliminary Injunction”; United States District for the Western District of Arkansas, “Memorandum Opinion and Order.”
[41] United States District Court for the Northern District of California, “Plaintiffs’ Second Amended Master Complaint”; United States District Court for the Northern District of California, “Order Granting in Part and Denying in Part Motion for Preliminary Injunction.”
[42] Clark County District Court, “Complaint and Demand for Jury Trial”; First Judicial District Court of New Mexico, “Plaintiff’s Complaint for Abatement and Civil Penalties and Demand for Jury Trial.”

B. European Union

The EU generally employs a risk-based approach to implementing age assurance technology[43] and has passed three laws—the General Data Protection Regulation (GDPR), Audiovisual Media Services Directive (AVMSD), and Digital Services Act (DSA)—governing legal requirements to do so.

The GDPR broadly restricts the processing of personal data in the EU and, together with the DSA, implicitly requires age assurance when risks are deemed to be high.[44] Article 8 of the GDPR specifies that consent to data processing can only be given when a child is over the age of 13 to 16 (varying across Member States) or by those with parental responsibility when a child is under the age of 13.[45] Data controllers must make “reasonable efforts” to verify that the data subject giving consent is over 16 or holds parental responsibility for a child under 13.[46]

The AVSMD specifies how Member States should legislate with respect to audiovisual media, including online video-sharing platforms.[47] Together, Article 6(a) and 28(b) specify that legislation by Member States should mandate that video-sharing platforms take “appropriate measures” to protect minors, including the “strictest measures” for the content deemed most harmful, such as gratuitous violence and pornography.[48]

In practice, these measures include regulating the time in which harmful content is shown, requiring it to be labelled, and verifying the age of users before it can be accessed.[49] Though EU Member States have adopted different approaches when transposing the AVMSD into national law, some explicitly mandate age verification.[50] For instance, Austria’s Audiovisual Media Services Act obliges video-sharing platforms to employ age verification or comparable access controls when hosting categories of harmful content like pornography.[51] Because the AVSMD is set to be revised in 2026, and protections for minors have been highlighted as a specific priority,[52] it is possible that these age assurance requirements could be further strengthened.

The DSA likewise employs a risk-based approach. Article 28 of the DSA requires “providers of online platforms accessible to minors to put in place appropriate and proportionate measures to ensure a high level of privacy, safety, and security of minors, on their service.”[53] The European Commission has produced guidelines[54] on how platforms can comply with these regulations, which include preventing

[43] Livingstone et al., “Children’s Rights and Online Age Assurance Systems.”

[44] Ibid.

[45] European Union, “General Data Protection Regulation,” Article 8.

[46] Ibid.

[47] European Union, “Audiovisual Media Services Directive.”

[48] Ibid.

[49] Wukovits, *Transposition of the 2018 Audiovisual Media Services Directive*.

[50] Ibid.

[51] Better Internet for Kids, “Austria”; KommAustria, “Guidelines for the promotion of self-regulatory bodies for the protection of minors adopted.”

[52] European Union, “Revision of the Audiovisual Media Services Directive.”

[53] European Union, “Digital Services Act,” Article 28.

[54] European Commission, “Guidelines on measures to ensure a high level of privacy, safety and security for minors online.”

minors from accessing content deemed age-inappropriate (e.g., pornography or gambling content) as well as providing safer defaults for minors (e.g., preventing unsolicited interaction by other accounts, disabling features that "contribute to excessive use").

While the guidance does not require platforms to use a specific form of age assurance, it specifies a set of criteria for effectiveness[55] that age assurance mechanisms should meet. The EU is in the process of building a government ID-based age verification system based on the EU Digital Identity Wallet,[56] scheduled for deployment in 2026. In the interim, the EU has sponsored the development of an EU Age Verification Solution[57] which uses a similar set of technologies,[58] although it was not available in the iOS or Android app stores at the time of this writing.[59]

The EC has opened investigations into a number of online platforms[60] to determine whether they are complying with their obligations under DSA Article 28. The results of these investigations will shed additional light about whether age assurance will be interpreted as required and under what circumstances under European law.

These measures accompany a variety of initiatives related to age assurance and other elements of platform product design that are growing in prominence at the time of this writing. The European Parliament issued a report in late 2025 calling for greater action to protect minors online, including a minimum age for social media platforms and targeted regulation of addictive features.[61] At the national level, various EU member states (e.g., France and Germany) are imposing or enforcing laws requiring digital services that host adult and other harmful content to verify the ages of their users.[62] Several EU member states are considering social media bans for youth.

C. United Kingdom

The UK Online Safety Act (OSA) imposes a set of duties on a broad range of online services, directed at protecting minors from content which is deemed harmful to children.[63] These duties are enforced by Ofcom as the independent regulator.

Services that publish or produce pornography (part 5 services) are required to use age assurance. User-to-user services (part 3 services) must conduct one or more assessments (a "children's access assessment" and then potentially a "children's risk assessment") and, depending on the outcome, may be required to implement age gating for some or all parts of their service in order to prevent minors from encountering certain types of content deemed harmful.

[55] European Commission, "Guidelines on measures to ensure a high level of privacy, safety and security for minors online."
[56] European Commission, "A digital ID and personal digital wallet for EU citizens, residents and businesses."
[57] European Commission, "EU Age Verification Solution."
[58] European Commission, "Operational, Security, Product, and Architecture Specifications."
[59] European Commission, "Installing the App."
[60] European Commission, "Commission opens investigations to safeguard minors from pornographic content under the Digital Services Act."
[61] European Parliament, "Children should be at least 16 to access social media, says MEPs."
[62] Desmarais, "The age verification era."
[63] Ofcom, "Protecting Children from Harms Online."

Ofcom does not require a specific age assurance mechanism but instead requires the use of "highly effective age assurance measures." Ofcom has published criteria for age assurance mechanisms and provides a list of mechanisms which are capable of being highly effective if implemented correctly.[64] It is the service provider's responsibility to ensure that their mechanisms are in fact highly effective.

The initial rollout of age assurance requirements in the UK has encountered a number of challenges, including sites failing to deploy age assurance,[65] a surge in downloads of Virtual Private Networks (VPNs),[66] the publication of an easy circumvention technique for Discord,[67] and the subsequent disclosure of the personal information (names, email addresses, credit card numbers) for around 70,000 users from Discord's age assurance customer service system due to a data breach of that system.[68] Recently, there have been proposals in the UK to expand age assurance mechanisms by restricting use of VPNs to adults[69] and to require device manufacturers[70] to "make it impossible for children to take, share or view a nude image."[71] With the exception of parents providing their child's age to digital services, UK residents have expressed discomfort with most methods of age assurance.[72]

D. Australia

Australia's Online Safety Act (2021) designates the eSafety commissioner as the regulator for online safety. The commissioner has published a set of age assurance requirements, including:

- A "Social Media Minimum Age" of 16[73] for accounts on social media platforms, which took effect on December 10, 2025 (per a 2024 amendment to the Online Safety Act).
- The "Phase 2" industry codes, which require platforms to restrict access to a variety of types of material judged harmful to minors (pornography, self-harm, simulated gambling, violence

[64] Ofcom, "Quick guide to implementing highly effective age assurance."

[65] Ofcom, "Investigation into 4chan and its compliance with duties to protect its users from illegal content"; Ofcom, "Investigation into AVS Group Ltd's compliance with the duty to prevent children from encountering pornographic content through the use of age assurance"; Ofcom, "Investigation into the provider of xgroovy.com's compliance with the duty to prevent children from encountering pornographic content through the use of age assurance."

[66] McMahon, "VPNs top download charts as age verification law kicks in."

[67] Ridley, "Brits can get around Discord's age verification thanks to Death Stranding's photo mode, bypassing the measure introduced with the UK's Online Safety Act. We tried it and it works—thanks, Kojima."

[68] Peters, "Discord customer service data breach leaks user info and scanned photo IDs."

[69] Parliament of the United Kingdom, "Children's Wellbeing and Schools Bill."

[70] Starmer, "A bold new mission."

[71] This text comes from a blog post by UK Prime Minister Keir Starmer, and does not provide details about what is being contemplated technically. If read literally, this proposal would extend far beyond any existing device-based age assurance mandate, because those mandates do not (for instance) prevent messaging apps from being used to transmit nude pictures or web browsers from viewing nude pictures on sites which do not implement age assurance. Absent a complete proposal, it is difficult to evaluate the feasibility of this approach, but as stated it would require far more invasive measures than discussed in this report in order to restrict devices in this fashion.

[72] Lai, "Age assurance and online safety."

[73] eSafety Commissioner, "Social Media Minimum Age Campaign."

instruction)[74]. These requirements apply both to adult sites and to more generic services such as search engines and phase in between December 2025 and March 2026.

In order to comply with these requirements, providers are required to restrict access to the relevant content and experiences unless users have demonstrated their age.[75] Enforcing these restrictions requires adult sites[76] and social media platforms to implement age assurance.[77] Other providers may need to implement age assurance in some conditions. For example, search engines must implement age assurance for account holders.[78] In addition, Australia has conducted an Age Assurance Technology Trial[79] and concluded that age assurance is practical.

E. Relevant International Standards

Multiple standards development organizations have developed standards in the area of age assurance. These standards provide an overall framework for implementation of age assurance, including privacy, security, and benchmarks for accuracy. They can be used as the basis for certification of age assurance systems. They do not provide detailed specifications for how to conduct age assurance using any specific mechanism, however.

These key international standards are:

- Institute of Electrical and Electronics Engineers
  - IEEE 2089.1-2024: IEEE Standard for Online Age Verification[80]
- International Standards Organization/International Electrotechnical Commission
  - ISO/IEC 27566-1: Information security, cybersecurity and privacy protection — Age assurance systems — Part 1: Framework[81]
  - SO/IEC 27566-2: Information security, cybersecurity and privacy protection — Age assurance systems — Part 2: Technical Approaches and guidance for implementation (working draft)[82]
  - ISO/IEC 27566-3: Information security, cybersecurity and privacy protection — Age assurance systems — Part 3: Approaches to Analysis and Comparison (committee draft)[83]

[74] Specifically, class 1C and class 2 material. eSafety Commissioner, "Consolidated Industry Codes of Practice for the Online Industry (Class 1C and 2 Material) Head Term."

[75] In the case of search, search engines must perform age assurance.

[76] eSafety Commissioner, "Consolidated Industry Codes of Practice for the Online Industry (Class 1C and 2 Material) Head Term."

[77] eSafety Commissioner, "Social media age restrictions."

[78] eSafety Commissioner, "Schedule 3 - Internet Search Engine Services Online Safety Code."

[79] Age Assurance Technology Trial, *Part A*.

[80] IEEE, "IEEE 2089.1-2024."

[81] International Organization for Standardization, "ISO/IEC 27566-1:2025(en)."

[82] British Standards Institution, "ISO/IEC NP 27566-2."

[83] International Organization for Standardization, "ISO/IEC CD 27566-3."

The IEEE and ISO/IEC specifications are to some extent complementary and to some extent overlap. Parts 2 and 3 of the ISO/IEC specification family are under development.[84]

V. Assessment Methodology

This 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 in Section V.B and uses them to examine each age assurance signal and architecture in turn. Because age assurance is an adversarial setting, in which the service cprovider is trying to restrict the access of some set of users (often in order to comply with some law or regulation), the assessment is performed in the context of a threat model, which defines both the capabilities of the various actors and their potential objectives, such as a minor accessing services and experiences for which they are not authorized or a service provider gathering more information about a user than the user wishes to reveal. This threat model is laid out in Section V.A.

Following the threat model, Section V.B explains the general assessment criteria used to examine each age assurance signal and architecture: baseline accuracy, circumvention resistance, availability, and privacy.

While this report does not address the broader question of the desirability of age assurance requirements or whether any particular service or experience is age appropriate, other policy questions naturally form part of the technical assessment, however. In particular, age assurance systems can be misused for surveillance or to selectively disenfranchise specific individuals for reasons other than age, and these issues are considered in the assessment where relevant.

A. Threat Model

The first step in the evaluation of any security mechanism is to define what is known as a “threat model.” A threat model defines both the security properties that the system is intended to guarantee and the capabilities and objectives of the “adversary.” While in a typical security setting the adversary might be depicted as a hacker or criminal, in the age assurance setting the adversaries are the ordinary members of the ecosystem: users, service providers, age verification providers, etc., who have conflicting incentives and objectives. Age assurance systems must be evaluated from the perspective of two threat models: the service provider’s threat model, and the user’s threat model.

[84] These standards supersede British Standards Institution PAS 1296:2018: Online age checking. Provision and use of online age check services. Code of practice, which was withdrawn January 5, 2026. See Age Check Certification Scheme, “Comparison Guide”; British Standards Institution, “PAS 1296:2018.”

1. Service Provider's Threat Model

The basic security property that the service provider is attempting to ensure (or is required to attempt to ensure) is that users are provided services consistent with their actual age. In order to accomplish this, the service provider needs to ensure that:

1. Users cannot access restricted content and experiences without undergoing age assurance.
2. When users' ages are assessed they are placed within the correct ranges.
3. Users cannot substitute someone else's age assessment for their own.

From the perspective of the service provider, the adversary is any user who does not fall within the age range required to access specific content, features, or experiences and wants to have said access. Stated more directly: the minor is considered the "adversary" from this perspective*.* This perspective is why age assurance mandates generally do not consider self-declaration of age sufficient for compliance.[85] If minors were not expected to try to circumvent the age gate, then such a declaration would be sufficient without *any* age assurance mechanism.

An additional consequence of this threat model is that the system needs to deny users access to the content or experiences in question by default until their age can be sufficiently verified. Otherwise, users could trivially circumvent age assurance by refusing to cooperate with the age assurance process (e.g., by providing a corrupted and unreadable picture for facial age estimation).

Unlike conventional access control mechanisms, age assurance systems are not expected to perfectly exclude all ineligible users, but rather to be "highly effective". Expectations about the level of effectiveness vary, but the general understanding is that some minors will be able to access age-restricted content and experiences.

a. Adult Assistance

It is also possible that the minor will have the assistance of someone who is within the eligible age range, whether a parent or carer, or simply an older friend. Parents and carers[86] frequently knowingly allow or even actively assist minors in misrepresenting their age to social media services,[87] and it seems likely that in some cases they would also assist minors in bypassing age assurance, or at least not act to prevent minors from doing so.[88] In addition, many minors will have older friends who might be willing to assist them.

[85] In some cases, it is possible to declare that one is a minor without evidence, thus receiving the default experience, but to need to undergo age assurance to get the adult experience. New York, "Stop Addictive Feeds Exploitation (SAFE) for Kids Act."

[86] This report uses the term "parents" for "parents and carers" for brevity reasons.

[87] boyd et al., "Why parents help their children lie to Facebook about age"; Ofcom, "Children's Online User Ages Quantitative Research Study."

[88] A recent Australian poll found that 34% of parents would be "likely to find a way to/help their child find a way to still use social media." Wilson, "1 in 3 parents will help kids get around teen social media ban, government privately warned."

The extent to which others will be willing to assist a minor in circumvention may depend on the type of content or service in question. Some parents who are willing to assist their children in circumventing age assurance to access social media may not be willing to do so for other classes of blocked content. In contexts where the same age restrictions exist for multiple kinds of content or services, minors may be able to persuade an adult to assist them with age assurance for one service but then leverage that to access another. For example, in some proposals for device-based age assurance, the device is persistently configured in either a restricted or unrestricted mode, so a minor could ask an adult to remove restrictions to allow them to access social media but then use the device to consume adult content.

b. Technical Sophistication

Minors will span a wide range of technical sophistication. At one end, younger children may be somewhat technically unsophisticated and by and large will not be able to mount complex technical attacks. By contrast, some older teens may be very technically sophisticated and will have programming–and even hacking—skills, as well as access to tooling developed by others specifically for circumvention, such as AI image generation (“deepfakes”) tools. Moreover, the history of computer security shows that once tools are developed they rapidly become commodities that are available to less-skilled users.[89]

c. Device Type

The range of devices that users will be using to access services will vary widely. The most important factor is the degree to which the user has control over loading software onto the device. At one end of the spectrum, the user might have a closed device such as an iPhone which is configured so that the user cannot load any software on it that is not approved by Apple. At the other end of the spectrum, they might have a PC loaded with an open source operating system such as Linux and have complete control over the software. In between these two extremes, there are many cases where devices are partially closed, such as devices with parental controls enabled or that are centrally administered by schools.

The more control the user has over their device, the more options they have for circumventing age assurance. Even server-based age assurance systems are often reliant on the behavior of the user’s device, for instance to prevent “injection” attacks where a minor transmits AI-generated video that is allegedly of their face.

2. User’s Threat Model

From the user’s perspective, the primary property they wish to ensure is access to content and services for which they are eligible[90] while maintaining their privacy.

[89] In the specific case of the Australian Social Media Minimum Age, children are already sharing TikTok videos with information on circumvention techniques. Wilson, “How Australian teens are already planning to dodge the social media ban.”

[90] Users may also wish to bypass age assurance, but as discussed above, that desire is out of scope for this report.

a. Access

Age assurance creates the threat that users will be excluded from activities in which they should be included, e.g., because they are unable to establish their age. This exclusion could be either targeted or incidental. For example, the government could attempt to exclude some classes of users from the underlying systems on which age assurance relies, (e.g., governments restricting specific individuals from having bank accounts in-country, accessing financial services, or even accessing specific online services). Even if the provider has no direct incentive to exclude users, it may do so incidentally merely because it is inconvenient or expensive to make age determinations more accurately. For example, a provider could rely exclusively on mechanisms which require users to provide a facial image, thus excluding users who cover their faces for religious reasons.

b. Privacy

Age assurance systems present a direct privacy threat to the extent that the signals that are used by an age assurance system to determine a user's age are inherently identifying, whether directly so, as in the case of credit card numbers or government documents, or indirectly so, as in the case of face-based age estimation systems. This data collection presents a potential privacy issue on its own, as even systems which require account creation often do not require the user to identify themselves with their real name or image, let alone their age and address. This information can potentially be misused in many ways (identity theft, doxxing, extortion, etc.) and presents a special risk when children's data is collected, as recognized by COPPA and other legislation.

Even in contexts where the user provides no identity information, age assurance mechanisms might enable cross-site or cross-application tracking. For example, if a user needs to routinely demonstrate their age in order to browse the web, then an age assurance service may be in a position to correlate the user's activity across multiple sites and applications, by observing the user as they engage with each site in turn and are required to demonstrate their age.[91] This information could then be misused by the age assurance service or disclosed in a data breach. Moreover, in many cases it is possible to identify a specific pseudonymous user solely from their browsing history,[92] thus tying the profile of the user's behavior to the user's inferred identity.

This risk is highest in cases where the user might ordinarily be anonymous, as when browsing sites on which they are not logged in, especially in private browsing modes or when using a VPN. The risk may be lower in cases where the user has already provided some identity information, for instance if they had to provide their name and email address when creating an account on the service. Even in these cases, the privacy risk may not be zero, as users may have provided only an email address and a false name or no name at all, and potentially even used a temporary address, as provided by Firefox Relay[93] or Apple's Hide My Email services;[94] if the user has to provide additional information (e.g., their name,

[91] This is not conceptually different from the situation with federated authentication systems such as logging in with Google authentication, but widespread age assurance requirements have the potential to increase the number of web activities where such an authentication mechanism is required.
[92] Bird et al., "Replication."
[93] Firefox, "Protect your identity with secure phone and email masking."
[94] Apple, "How to Use Hide My Email with Sign in with Apple."

photo, or even a picture of their identification document), then this increases the information that the service provider has about the user and which might subsequently be abused. The risk is also higher in cases where the services and experiences being accessed are sensitive, as with pornography or LGBTQ content.[95]

This form of activity correlation is more effective if there is a small number of large age verification providers, as each provider will see a correspondingly larger fraction of the transactions. Recent research found significant concentration, with the top five AVPs collectively covering over 70% of the US age assurance market, and the largest service, Yoti, being used in over 60% of websites in Texas and Georgia that used an AVP.[96]

These privacy risks are inherent in the collection of information by the evaluator, but the level of risk depends on the evaluator's behavior. If the evaluator has strong security practices and deletes the data promptly, the risk is lessened; if it shares the information with others, retains the information for a long period of time, or has weak security practices, then the risk is increased. Even if the evaluator itself is well-intentioned, it may be subject to a data breach which reveals the user's information; this risk is exacerbated if the user's information is stored beyond the minimum time necessary to perform age assurance.

c. Adversary Capabilities

From the perspective of the user, most other entities on the internet are the adversary, because they are in a position to inappropriately exclude the user and/or breach their privacy.[97] This includes the service provider and its contractors, AVPs and their contractors, any other entities that the AVP consults in order to perform age assurance, and any third party who is able to obtain the AVP or service provider's records, including hackers, or the government.

Importantly, the user has no visibility into the functioning of the systems of any of these entities. This means that any information that the user provides to them (e.g., their name, photo, etc.) is potentially subject to misuse, which may not be visible to the user. For example, if the user provides their name to the AVP, which stores it in its database, and then suffers a breach, this is all invisible to the user. Moreover, these entities might cooperate to misuse the user's data, subject to whatever technical and policy controls are in place.

Note that generally, even when service providers or AVPs directly install software on the user's device, they will typically only have limited capabilities, and so will not be able to directly learn personal information about the user unless the user gives them permission. The security of the device itself and of generic software (e.g., web browsers) is outside of the scope of this report. In general, compromise

[95] For example, some PornHub users were recently extorted after a breach of the Mixpanel analytics provider revealed their viewing history. Abrams, "PornHub extorted after hackers steal Premium member activity data."
[96] Minocha et al., "Papers, Please."
[97] From the perspective of a minor attempting to circumvent the system, these entities are also the adversary because they are in a position to (correctly) restrict the minor's access.

of the device will lead to severe compromise of user security and privacy, as will malicious behavior by the device or operating system vendor.

B. Assessment Criteria

This section describes the criteria used to assess age assurance systems:

- **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.

These criteria are discussed in more detail below.

1. Baseline Accuracy

The primary function of an age assurance mechanism is to *accurately* distinguish between users who are within the eligible age range and those who are not. In this report, the term *baseline accuracy* refers to accuracy under conditions where the user is not trying to actively deceive the system, such as by showing a fake ID or a picture of someone older. The case where the user is trying to deceive the system is discussed in the next section.

Some age assurance systems (e.g., simple self-declaration) have negligible baseline error rates because the user can nearly always enter their age correctly if they choose to. However, other age assurance systems, especially age *estimation* systems, inherently have some level of error. There are a number of ways to characterize the error rate of this kind of system, but at a high level an age assurance system either grants or denies access and therefore it is natural to consider two values:

- The *false reject*[98] rate, representing the fraction of users within the eligible age range (e.g., over 18) who are not permitted to access the system, for instance.
- The *false accept* rate, representing the fraction of users outside the eligible age range (e.g., under 18) who are permitted to access the system.

It is common for systems to have a tradeoff between false rejects and false accepts. For example, age assurance systems based on biometric age estimation may internally produce a probability distribution representing the likelihood that the user is a certain age; this can readily be used to compute the

[98] The statistics and testing literature uses a number of terms for error rates, including sensitivity versus specificity, false positive versus false negative, and type I versus type II errors. These can often be hard to interpret because of confusion about whether a “positive” result leads to acceptance or rejection. This report uses the terms “false reject” and “false accept” for clarity.

system's estimate of the probability that the user is above (or below) a given age threshold.[99] Because the system needs to ultimately either accept or reject the user, it is necessary to translate this probability distribution into a yes or no answer. This is often done by selecting the highest probability age and then asking whether it exceeds some threshold, e.g., "is the highest probability age over 18?"

For any given system, selecting the right threshold involves determining the relative importance of false acceptance and false rejection. For instance, if a very low threshold is used, then borderline cases will be accepted, thus leading to low false reject rates but high false accept rates. Conversely, if a high threshold is used, then borderline cases will be rejected, leading to high false reject rates but low false accept rates. The exact nature of this tradeoff curve is determined by the technology in use, with better technologies allowing lower joint false accept/false reject rates. Importantly, error rates are not necessarily uniform. For example, some facial age estimation techniques are poorer at estimating the age of people of African heritage[100] or of females.[101]

a. Repeatability

Age assurance systems are susceptible to two kinds of error:

- **Systematic error:** Some age assurance systems consistently produce the wrong answer for the same person. For example, for facial age estimation systems, because people's apparent ages vary and some people look far older or far younger than typical, some users will be consistently categorized incorrectly (whether falsely accepted or falsely rejected). Similarly, systems that estimate age based on measured user activity are subject to systematic errors because some users have activity that looks more like activity typically associated with older or younger users.

- **Random error**: Age assurance systems can produce inconsistent results for the same user, for instance because of slightly different camera angles or lighting. The result is that if a user attempts age assurance and is rejected, they might be accepted if they tried again. Conversely, a user who is accepted might be rejected on some subsequent occasion.

Most systems will have both kinds of error to varying degrees. Because the system must ultimately either reject or accept a given user at a given time, these types of errors interact. For example, a user who is 15 but who on average appears to the system as 17 may sometimes be accepted as over 18 due to random errors.

The level of random error contributes to whether a system can *repeatably* produce the correct answer. The level of repeatability is especially important for age assurance systems because a user who is rejected may choose to try again. If an age estimation system is subject to random errors, a minor who

[99] Note that this value cannot be directly translated to the false accept and false reject probabilities because it depends on the underlying distribution of age ranges (the "base rate"). As an example, if the age threshold is set at 200 years, then all "accept" results will be false accepts because this exceeds the maximum age of the population.

[100] Oladipo et al., "Face Age Estimation and the Other-race Effect."

[101] Hanacek, *Face Analysis Technology Evaluation (FATE) Age Estimation & Verification*.

is rejected can simply try again—changing lighting, angles, or expressions—until they get a false accept result. A system with high repeatability (even if imperfect) is more secure because it will consistently reject the same user, preventing them from exploiting random errors.

The obvious way to manage random errors is to keep track of which users have tried and failed to pass age assurance and limit multiple attempts. However, this creates a new set of privacy issues because it requires tracking user behavior, and specifically the behavior of users who are likely to be children.

2. Circumvention Resistance

Baseline accuracy is concerned with accuracy in non-adversarial settings, but some users will try to deceive the system. If some users were not trying to circumvent the age gate, then the currently common age gates that merely ask the user to self-declare their age would be sufficient.

Different age assurance systems are subject to different forms of circumvention, so it is not practical to have a single metric for circumvention resistance. The assessment focuses on adversary capabilities (where the adversary is a user who is outside the eligible age range), the difficulty of circumvention, the likelihood of success, and the ability to scale/commoditize circumvention techniques.

3. Availability

Even a perfectly accurate age verification system may exclude some users who fall within the eligible age range because, for a variety of reasons, those users are not able to demonstrate their age to the satisfaction of the system. The term "availability" is intended to include both impediments which fall under the classic definition of "accessibility" to users with physical impairments and users who are unable or unwilling to engage with a given age assurance mechanism. For example:

- Age estimation systems based on selfies or live video of users may be inaccessible to users without a camera.
- Age verification systems based on government ID may be inaccessible to those without government ID.
- Age verification systems based on commercial transactions may be inaccessible to users with limited commercial records.

As discussed in the remainder of this report, essentially all practical age assurance systems will not be available to some class of users. Which class is affected depends on the nature of the system. One approach to addressing this challenge is to give the user a choice of multiple age assurance methods, thus increasing the possibility that at least one method will be available to them.

4. Privacy

The final criterion for assessment is the privacy impact of age assurance mechanisms, specifically, the degree to which any other actor in the ecosystem learns any information they would not learn in the absence of age assurance.

For privacy assessment purposes, all age assurance mechanisms can be compared against the scenario where the service provider learns only that the user is within the eligible age range, and no other party learns any additional information about the user. This is the minimal amount of information disclosure possible while still preserving the age assurance function.[102] However, many age assurance systems involve substantially more information disclosure, frequently including the user's identity. In order to evaluate the privacy of age assurance systems, the assessment asks what additional information is disclosed and the extent to which it enables the forms of attack described in Section V.A. This information falls into three main categories:

- Personal information about the user, such as their identity or image.
- What content and experiences the user is engaging with
- Linking user activity across services

Each of these is discussed below.

a. Personal Information

As noted above, many age assurance systems involve the user disclosing some level of personal information in order to demonstrate their age. In order to evaluate the risk, this assessment first considers which information the user discloses and which entities learn it; for example, the information might be disclosed just to the age verification provider or also to a service provider. Second, it considers what uses that information might be put to, especially when coupled with other information that these entities might already have, such as the user's email address or phone number.

b. Content and Experiences

The second question to consider is disclosure of the content and experiences the user desires to access. In many—though not all—cases, service providers will already know this information for a specific user, especially if the users have accounts, but third parties such as age verification providers or the sources they consult for age assurance will not already have this information. If the age verification provider learns both the user's activity and their identity, this has clear privacy implications, especially if the activity is sensitive in nature (e.g., consuming pornographic content), but even activity which appears innocuous to one user may not be seen that way by another user in another context.

[102] If a system effectively excludes users who are not within the target age range, then a user's use of a service allows the service provider to infer that they are within the target range.

c. *Linking User Activity*

Even if the user's identifying information is not revealed during age assurance, the process of age assurance may create records which could be used to link up different periods of user activity, whether using the same service or different services. This linkage could happen because the user provides some form of consistent identity (e.g., an email address) or because there is some mechanism for remembering the user for future visits (see Section VI.A.5). If coupled with learning the user's activity, this type of linkage may allow the opportunity to create a profile of user behavior, and, if coupled with any instance of providing the user's identity, may allow linkage of the whole profile to that identity.

d. *Ephemeral versus Long-Term Records*

As noted above, age assurance will often require collecting information about users. If the various components of the age assurance process retain this information, then this becomes a new source of privacy risk for users, in that it essentially becomes a database of which services a user has accessed. Such a database is not only an attractive target for attackers but is also a new source of data for government surveillance, whether by legal process or otherwise, even though that surveillance is not otherwise necessary for age assurance. Some recent age assurance mandates have contained requirements for data deletion[103] but not all do so.[104]

e. *Technical versus Policy Controls*

There are two approaches to addressing the risk to user privacy of age assurance systems:

- **Technical controls** which prevent some or all of the components of the age assurance system from learning information about the user.
- **Policy controls** which allow components of the age assurance system to learn about the user but restrict how the system can use that information or require the system to anonymize or delete it.

In general, technical controls are stronger than policy controls. In many cases, users can verify for themselves that technical controls are in place, whereas with policy controls the user must trust that some other parties are behaving correctly, often in cases where the user has no prior relationship with that party or reason to trust them. For example, consider the case of showing ID to purchase alcohol: if the clerk just visually inspects the ID, then they might be able to remember names occasionally, but it is not practical to record everyone who purchases alcohol. By contrast, if the ID is scanned, then the purchaser has no way of knowing what that data is used for or how long it will be retained. Many age assurance mechanisms are more like this second class of system in that personal information is provided to the AVP and then the user has to trust the AVP to handle it correctly.

A special concern for policy controls is that even if an entity *ordinarily* complies with policies, the entity might not do so under exceptional conditions such as requests by government agencies or data

[103] New York, "Stop Addictive Feeds Exploitation (SAFE) for Kids Act."
[104] Age Assurance Technology Trial, *Part A*, A 51.5.

breach. For example, if an age verification provider keeps records of each age assurance transaction and the site for which age was verified, then compromise of the AVP may compromise user privacy.[105] Note that some AVPs have policies[106] that explicitly permit them to disclose personal data to governments even when they are not obligated to do so.[107]

It is not uncommon for systems to have a combination of both types of controls: for example, many web-based age assurance systems have a separate age verification provider which learns the user's identity but has policy controls intended to protect that identity. The service provider, by contrast, does not learn the user's identity but merely that the user is within the eligible age range.

VI. Assessment of Age Assurance Architectures

Age assurance can be implemented using a variety of architectures. These 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)

This section considers these architectures in turn.

A. Server-Based Age Evaluation and Enforcement

The most widely deployed architecture is to perform all age assurance functions on the server side. Figure 4 below shows the typical architecture, where 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.

[105] Apthorpe et al., "Online Age Gating."

[106] For example, Incode's policies state, "We may disclose any personal data that we collect when required *or permitted* by law, such as to law enforcement agencies, courts, regulatory agencies and others, including to comply with valid legal process" (emphasis ours). Incode, "Privacy Policy." By contrast, Yoti and VerifyMy have more limited policies that only describe sharing with governments when required by law, although they do not explicitly state that they will require legal process. VerifyMy, "Privacy Policy"; Yoti, "Yoti Age Verification Service - Privacy Notice."

[107] For example, the Kids Web Services age assurance provider shares information with governments under a range of conditions. Kids Web Services, "Privacy Policy."

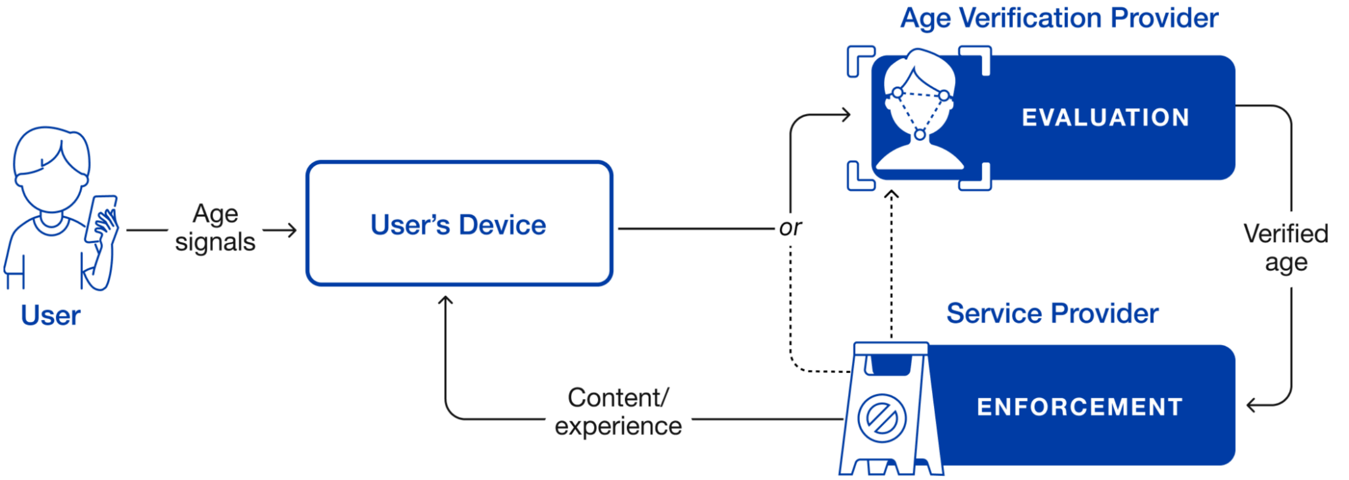

Figure 4. A typical server-based age assurance architecture.

The user experience is as follows:

1. The user attempts to access the service provider, either in their web browser or in an app.
2. The service provider then prompts the user to demonstrate their age.
3. The user provides their age signals to their device (e.g., shows their ID or turns on their camera).
4. The device provides the age signals to the AVP for evaluation, either directly or via the service provider.
5. The AVP provides the results of that evaluation to the service provider.
6. The AVP redirects the user back to the service provider for enforcement.
7. The service provider provides whichever service experience is age-applicable based on the results of the evaluation.

This architecture is deployable without requiring any changes to web browsers or mobile devices, as it makes use of only existing deployed technologies. See Appendix A for a description of deployment on the web.

The AVP is free to distribute evaluation and enforcement functions between the user's device and the server. For instance, instead of uploading the user's face image to their servers for processing, an AVP could perform facial age estimation (see Section VII.D) on the device. As a practical matter, most of the age signals discussed in this report require some server-side processing, but it is still possible for the AVP to do some processing on the device. For example, Yoti's web-based age assurance system uses code running in the user's browser to capture the user's face and then uploads the image to the Yoti server for age estimation.

Mobile apps can store the user's age eligibility status in the app. In the web case, age eligibility status could be stored using web storage technologies (e.g., cookies), but some users will use private browsing modes or delete cookies, especially for adult sites, thus erasing any age eligibility information. In these cases, the service provider may need to re-verify the user's age at each new interaction, which can be a source of increased friction for users. This friction degrades user experience and may also decrease the chance that the user will choose not to use the service, which is not desirable from the service provider's perspective. In addition, if the service provider pays a vendor to provide age assurance on its behalf,[108] then there may be additional costs associated if users need to re-verify. This gives the service provider an incentive to avoid repeated transactions by remembering the user, for instance by requiring the user to make an account.

Moreover, because each service provider is independently responsible for enforcement, they must also independently arrange for evaluation, which may require users to repeatedly undergo age assurance with each new service provider. If two service providers share the same AVP, the AVP can remember the user—potentially by asking them to make an account—and allow them to bypass repeated age assurance. This reduces user friction to some extent, though may still require some user interaction, e.g., to authorize reuse of their age eligibility information. Section VI.A.5 discusses various technical mechanisms for managing repeated interactions and their privacy properties.

1. Baseline Accuracy

In addition to the specific accuracy properties associated with whichever age signals are in use (discussed throughout Section VII), in some cases server-side enforcement can have an accuracy challenge associated with geolocating the user. Services designed to comply with different age restrictions in different jurisdictions must first determine the user's location in order to know what age range to enforce.

For web-based service providers, one common way of identifying the user's location is to collect the user's Internet Protocol (IP) address and use an IP-based geolocation service to attempt to infer the user's geographic location. This is inherently a somewhat inaccurate process because IP address allocations do not line up neatly with jurisdictional boundaries.In general, IP-based geolocation has high accuracy at the country level, with providers claiming accuracy above 99%.[109] Estimates of accuracy at the state level inside the United States vary more widely. For example, Digital Element estimates 98% accuracy, while Maxmind estimates 80% accuracy. There is limited independent research on this topic, but a 2021 study[110] found generally high accuracy for fixed broadband addresses and lower accuracy for mobile broadband, although in some cases this is due to services clustering the addresses for a given region (e.g., all of Newark). Regardless, any server-side enforcement using IP-based geolocation will have some level of inherent error.

Mobile apps have more options in terms of determining the relevant jurisdiction. An important first question is when the relevant jurisdiction is determined. The two main alternatives here are to adopt

[108] The New York State Office of the Attorney General estimates a cost of around $.05/assurance method at scale. See Office of the New York State Office of the Attorney General, "Notice of Proposed Rulemaking."
[109] Digital Element, "An Executive's Guide to IP Geolocation"; MaxMind, "Geolocation Accuracy."
[110] Saxon and Feamster, "GPS-Based Geolocation of Consumer IP Addresses."

the policy of the jurisdiction where the user originally activated the account or device or to adopt the policy of the current location. In the former case, users may be incorrectly classified if they move between jurisdictions. In the latter, users might find themselves suddenly subject to age assurance even for apps or sites they were previously using, which is more challenging for operating systems and service providers to handle correctly.

Mobile apps can use IP-based geolocation but this is not the only option. In general mobile devices are able to determine the user's physical location with high reliability via a combination of GPS, visible Wi-Fi access points, and information about the mobile network. For privacy reasons, both Android and iOS require the user's permission before allowing apps to get the user's location. In principle the user can decline to share their location with mobile apps, but evaluators might require it as part of the age assurance process.

In addition, the mobile device has access to the user's language, app store account location, and mobile configuration (including mobile number) and current mobile network. Depending on which policy the app is using for determining jurisdiction, some or all of these indicators may be useful. However, for privacy reasons, not all of this information is available to mobile apps, especially on iOS, or may require the user permission for apps to access. Moreover, this information may not be sufficient to provide location information at finer granularity than the country level, which presents an issue in the United States and elsewhere where policies vary by state or province.

The process of determining jurisdiction precisely would be made easier if mobile devices were to offer new APIs which only disclosed the user's general location down to the relevant level of jurisdiction (e.g., country or state), as the evaluator does not need the user's precise location in order to determine the appropriate age range. These APIs could be made available to apps which did not have permission to access the user's precise location—though might still require user permission—and thus reduce the privacy risk to users.

2. Circumvention

As with baseline accuracy, the main circumvention risk beyond circumvention associated with specific age signals relates to the user spoofing their location to appear to be in a jurisdiction that does not require age assurance. IP-based geolocation can be circumvented by the use of virtual private networks (VPNs), which make the user appear to have a different IP address than their true IP address. A number of VPN services explicitly offer the ability to have an IP address in specific jurisdictions,[111] can be used to evade geographic restrictions on content (e.g., to watch streaming sports content which is only available in certain regions).[112] In the case of age assurance, a VPN allows the user to evade age restrictions in their own region by appearing to be located in a different region. This is especially relevant for pornography sites, which are largely accessed through the web due to app store restrictions.

[111] ExpressVPN, "VPN Servers"; NordVPN, "Thousands of ultra-fast VPN servers across 178 locations."

[112] Dutkowska-Zuk et al., "How and Why People Use Virtual Private Networks."

There is extensive evidence that users respond to server-side enforcement of age restrictions by using VPNs.[113] Much of this usage is likely to be adults who would be able to pass an age assurance screen but do not wish to do so (e.g., for privacy or convenience reasons).[114] However, it is also a route for ineligible users to bypass age assurance mechanisms: a 2021 survey found that 46% of UK 16- and 17-year-olds had used a VPN or Tor browser.[115]

Experience with VPNs to circumvent copyright and content licensing restrictions shows that it is possible for services to identify and block users connected via VPNs to some extent, for instance by blocking traffic which comes from the IP addresses known to be associated with VPN providers.[116] However, this blocking is necessarily imperfect because the VPN providers are incentivized to evade blocking.[117] In the age assurance setting, services may be incentivized to allow VPN-connected users on their platforms because more usage means more revenue.

Jurisdictions might choose to restrict the use of VPNs or to require that VPN operators perform age assurance for all of their customers. If these restrictions are regulatory, the jurisdiction may need to enforce them against extraterritorial VPN providers (this is also an issue for extraterritorial service providers). This can be particularly challenging with anonymity networks such as Tor,[118] which are designed to resist censorship. It is also possible to technically block VPNs and other anti-censorship technologies, as China, Russia, and other countries have opted to do. In these cases, there is an arms race between blocking and evasion, with the result that some technically sophisticated users are able to evade the blocking.[119] In addition, some technical blocking techniques can create collateral damage in the form of blocking of non-VPN usage.[120]

VPNs are a less effective mechanism for spoofing location for mobile apps because mobile apps can, in some cases, query the user's location directly rather than relying on the IP address, as discussed in Section VI.A.1. A motivated user could still attempt to falsify their location, for instance by modifying their app, using GPS spoofing software,[121] or buying a hardware GPS spoofer.[122] Because mobile devices use multiple location signals besides GPS (e.g., distance from cellular towers and nearby Wi-Fi access points), vendors should be able to detect hardware-based GPS spoofing in many cases.

[113] Bradshaw, "VPN use surges in UK as new online safety rules kick in"; Castro, "'VPNs are not kryptonite of age assurance'"; Cyber Security Intelligence, "VPN Demand Surges As British Online Safety Law Takes Effect"; Datta, "VPN surge won't stop France's fight against porn, vows its digital minister"; Lang et al., "Do Age-Verification Bills Change Search Behavior?"

[114] See Baroness Kidron in Parliament citing Ofcom. Parliament of the United Kingdom, "Online Safety Act 2023: Virtual Private Networks"; "Age Verification Providers Association, "No, UK porn use was not halved by age verification."

[115] Thurman and Obster, "The regulation of internet pornography."

[116] See, e.g., GeoComply, "Helping stop geo-piracy and location fraud with award-winning VPN and proxy detection"; Spur, "Advanced detection of anonymization and threats."

[117] Khan et al., "Stranger VPNs."

[118] Tor, "Browse Privately. Explore freely."

[119] Wu et al., "How the Great Firewall of China Detects and Blocks Fully Encrypted Traffic."

[120] Sommese et al., "Disrupting the Internet in the name of copyright."

[121] Singh, "Top iOS Location Changer Apps in 2025."

[122] GPSPATRON, "GNSS Spoofing Scenarios with SDRs."

To defend against the location signals being altered on the device, app servers can use “app integrity” mechanisms provided by mobile operating systems, such as Apple’s App Attest[123] and Android’s Play Integrity.[124] These mechanisms allow the device to generate a cryptographically protected data value (an “assertion”) about the app which can then be verified at the server.[125] This allows the server to ensure that the correct app is running on an unmodified device. It is not clear whether current app integrity mechanisms can detect all forms of software-based location spoofing, but mobile OS vendors could readily extend them to do so. If vendors were to offer new dedicated APIs for determining jurisdiction based on location, as described in Section VI.A.1, these APIs would be similarly resistant to circumvention. These defenses are not available to websites relying on device-based geolocation APIs because there is no mechanism for verifying the application integrity of web browsers or that the user has not configured the browser to return a false result.[126]

A number of other location indicators are susceptible to user manipulation. As described above, app integrity can be used to determine that the app is reading the correct information from the device, but the user might be able to change the device configuration. For example, the user can change their locale and app store location and could potentially get a SIM from another location, thus making their carrier appear to be elsewhere. These tactics may affect the user experience of their device in undesirable ways or require extra effort to revert back to their home locale. It is not clear how difficult it is for apps to defend against these forms of circumvention, for instance by querying for these indicators repeatedly to make it more difficult for a user to temporarily change their location to evade age assurance.

An additional challenge for server-side enforcement is that in many cases the target services will be located outside of the jurisdiction doing the regulation. In some cases, the target service will still do substantial business in the relevant jurisdiction and so it will be possible to compel compliance. In other cases, this may be more difficult.[127] By contrast, when enforcement takes place on devices, it is more practical to require restrictions because the devices are physically present in the jurisdiction, and usually sold there.

3. Availability

Because server-based age assurance uses existing, standardized web technologies and APIs, it is in general highly available. Users of commercially available web browsers and mobile apps who can access the relevant services and are able to send the underlying age assurance signals will in general be able to use this age assurance architecture.

[123] Apple, “Establishing your app’s integrity.”

[124] Google, “Play integrity and signing services.”

[125] Apple, “Validating apps that connect to your server.”

[126] See, e.g., ilGur, “Change Geolocation”; Kumar, “How to Change or Fake Location in Chrome, Edge, or Firefox.”

[127] For example, the US-based site 4Chan has been under investigation by Ofcom but has refused to cooperate and sued Ofcom in US court in an attempt to prevent enforcement. Vallance, “4chan launches legal action against Ofcom in US.”

4. Privacy

Because the server sees the user's age signals and in most cases those signals either provide the user's identity directly (as with showing a government-issued ID) or indirectly (as with facial age estimation), this creates a privacy risk for the user. When the user's identity is associated with the age-restricted experience, the service provider can create a record of that user's access. This information may then be sold, leaked, or accessed via legal process.

When the service provider relies on an AVP for the evaluation function, the two entities learn different information about the user. The AVP learns the information about the user's identity that is provided by the age signals as well as the identity of the service that the user is trying to visit. The AVP does not necessarily learn about specific content, features, or accounts that the user is trying to access (although if only some features are restricted, then the AVP can make some inferences).[128]

In contrast, if the user's device sends age signals to the AVP without allowing the service provider to see them,[129] the service provider does not learn information about the user's identity other than age-related information. The service provider learns the results of the AVP's evaluation. Depending on the design, this might mean the user's age, estimated age, or whether their age is within the eligible age range. If the service provider runs their own age assurance or captures the age signals and sends them to the AVP, it learns both sets of information, resulting in greater privacy exposure. Users may not know whether the service provider is receiving their age signals, especially in the app context. The app captures the age signals and invisibly transmits them to back-end infrastructure, which could be operated either by the service provider or the AVP. In the web case, the user may be on the AVP's website when they send their age signals, so in principle it is easier to determine that the age signals are being sent directly to the AVP.

The user is reliant on the policy controls the AVP and service provider have in place to keep user identity information and service access records separate, but users are unable to verify for themselves whether those policies are being enforced nor if/when their data has been disclosed contrary to policy (as in a breach). In some jurisdictions, these policies may be subject to government audit, for instance to ensure that they are deleted in a timely fashion, but this is also not directly verifiable by the user.

If a mobile app uses the device's precise location to determine the appropriate jurisdiction, then this potentially reflects a risk to user privacy as the app will learn the user's precise location. This risk is especially serious if the user is performing age assurance from their home, as the evaluator would then learn their home address. In principle the user can decline to share their location with mobile apps, but evaluators might require it as part of the age assurance process. This privacy risk is reduced in cases where the app uses other indicators of jurisdiction that provide coarser grained location information.

[128] The HTTP "Referer" header indicates the site that the user came from, but the default policy of "strict-origin-when-cross-origin" does not reveal the specific page. Moreover, many age gates are on the site's front page, and therefore do not disclose which parts of a site a user is interested in. This does not, however, prevent the service provider from intentionally revealing the user's behavior to the AVP.

[129] In the web context, this would typically involve the user being redirected to the AVP's site. In the app context, the app could send signals directly to the AVP.

5. Repeated Interactions

If the user makes an account with the service provider, either for the purpose of not having to prove their age each time they access the service or for some other reason, this permits the service provider to track user behavior and may also require the user to reveal their contact information (e.g., email address) even if the age assurance mechanism itself does not reveal their identity. The result is the creation of a corpus of data about the behavior of individual users that could be misused by the service provider, sold, or subject to breach or government legal process. Moreover, the need to create an account is a source of friction for users, both because it is extra effort and because they may not wish to reveal additional information. If the user makes an account with the AVP, this friction presents similar challenges with respect to the AVP. In general, service providers have an incentive to avoid user friction as it may cause users to choose not to visit their sites.

In the normal web context, cookies can be used to persist user state between browsing sessions, even if the user does not make an account. Because cookies also allow the website to track the user between visits, privacy-preserving browsing modes which delete cookies after the browsing session are commonly used to visit adult sites, limiting the utility of storing age status in a cookie for friction reduction for those sites. Several providers are now using a secure login technology called passkeys[130] to retain state without using cookies. The OpenAge[131] initiative has developed a passkey-based solution called AgeKey, which Meta and Snap have announced they will be using.[132] Yoti[133] has developed a similar mechanism. Passkeys do not reduce friction entirely, however, because unlike cookies they require effort to set up and user confirmation for each interaction.

If passkeys or cookies are used to directly persist user state to the service provider or AVP, then this permits those entities to build an activity profile for the user: each time a service provider needs to verify the user's eligibility, it contacts the AVP, which uses the passkey/cookie to determine if the user has already performed age assurance, with the result that the AVP can link up the user's activity. In the AgeKey model, the AVP does not directly store the user's state but instead works with a server operated by OpenAge which retains the user's eligibility state tied to the user's passkey. On repeat interactions, the service provider redirects the user to the AVP, which in turn redirects the user to the OpenAge server, where they use their passkey. The OpenAge server then notifies the AVP of the user's eligibility status. As a result, the AVP is not able to use the passkey to link up multiple visits by the same user.

While this approach provides improved privacy there are still multiple approaches for linkage. First, if the user is not using a VPN or other IP concealment technology, the AVP may be able to use the user's IP address to link up multiple interactions. Second, if OpenAge and the AVP collude, together they can

[130] FIDO Alliance, "Passkeys."
[131] OpenAge Initiative, "OpenAge."
[132] Bradshaw, "Meta adopts new age-check system to meet global child safety laws"; Snap, "Implementing Australia's Social Media Minimum Age Law."
[133] Trotman, "Introducing Yoti Keys."

build a profile of the user's interactions. As with other remote interactions, the user is forced to trust the AVP and OpenAge, even though they have no real relationship and were not chosen by the user.

On mobile devices, it is also possible to use a mobile app to store age verification results for future use. The properties of this kind of app would vary depending on how those results are stored and delivered to the service provider. If the app is used to store the age verification result and thus bypass a repeated age assurance process on the AVP, then it potentially allows the AVP to link up repeated interactions, as with the options discussed above. More privacy-preserving approaches such as zero-knowledge proofs are also possible.

6. Case Study of Server-Based Architecture: Yoti

Yoti is a UK-based age verification provider.[134] As of this writing, Yoti supports 11 separate age assurance mechanisms,[135] including facial age estimation, credit card verification, and email-based age estimation.

a. Basic Age Assurance

In order to use Yoti, the service provider first registers with Yoti and creates a service provider account. The service provider has two options for how to configure the age assurance user flow:

- The service provider can collect the user's information itself and query Yoti's API for an age evaluation (accept/reject) response.
- The service provider can redirect the user to Yoti's site for age assurance (see Appendix A.B for technical details).

In the latter case, the site will also use Yoti's API to indicate to Yoti which age assurance mechanisms to use for a given interaction. When the user arrives at Yoti's site, they will see a screen similar to the one pictured in Figure 5 below.

[134] Yoti, "Adult Content Age Verification."

[135] These include Facial age estimation, Digital ID wallet, Document, Age token, Credit card, Yoti Keys, Mobile, LA wallet, Electronic ID, Social security number, Database, and Email age estimation. Yoti, "Age Verification."

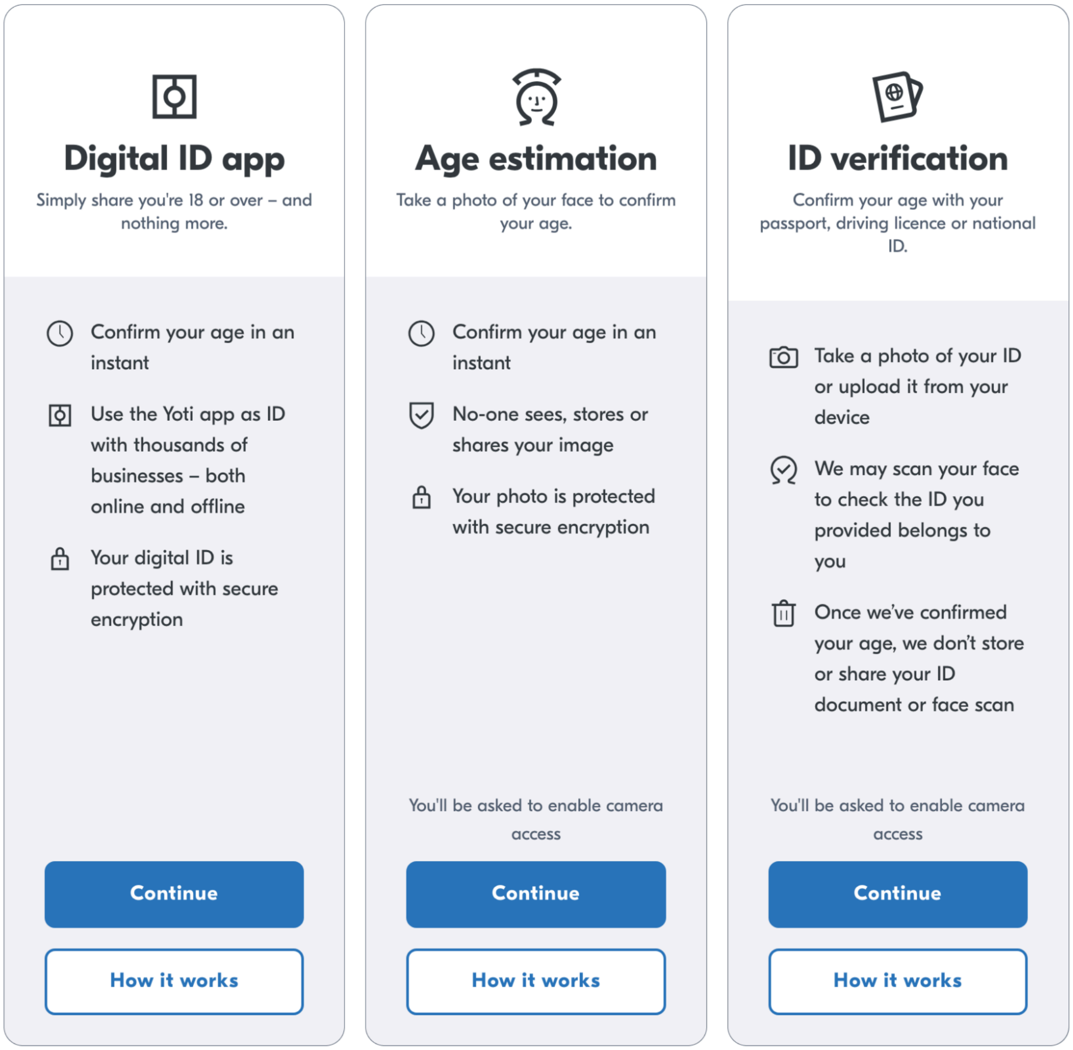

Digital ID app

Simply share you're 18 or over – and nothing more.

- Confirm your age in an instant
- Use the Yoti app as ID with thousands of businesses – both online and offline
- Your digital ID is protected with secure encryption

Age estimation

Take a photo of your face to confirm your age.

- Confirm your age in an instant
- No-one sees, stores or shares your image
- Your photo is protected with secure encryption

You'll be asked to enable camera access

ID verification

Confirm your age with your passport, driving licence or national ID.

- Take a photo of your ID or upload it from your device
- We may scan your face to check the ID you provided belongs to you
- Once we've confirmed your age, we don't store or share your ID document or face scan

You'll be asked to enable camera access

Figure 5. Example Yoti age assurance signal selection screen.[136]

b. Repeat Interactions

Yoti also supports a mechanism to allow users who have already established their age to skip repeated interactions.[137] This mechanism, known as "reusable tokens," works by having Yoti create a digitally signed object storing the user's age information in the user's browser. When creating the age assurance session for a given user, service providers can indicate whether they accept reusable tokens and what types of tokens they accept (based on which age assurance mechanisms were used). If the user has a matching token, then the user is admitted without having to perform another age assurance transaction.

Users can also create a "Yoti account" which allows them to share their age assurance results across multiple devices by logging into the account on a new device. In addition, Yoti supports the use of

[136] Yoti, "Yoti Developer Documentation."
[137] Trotman, "Introducing Yoti Keys"; Yoti, "Tokens."

passkey-based authentication which allows for tokens to persist even if the user is using a private browsing mode which does not persist cookies beyond the lifetime of the browsing session. Otherwise, the user will have to re-establish their age even if they have previously done so.

Yoti also offers a Yoti ID app[138] which allows users to register their information with Yoti to create a "reusable digital ID". They can subsequently use the app to demonstrate their age to service providers, for instance by scanning a QR code or clicking on a link.

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). The device only needs to perform age assurance once, no matter how many services the user engages with, which reduces a source of friction for users.

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).[139]
- 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).

The first of these options is only viable for settings in which apps are "all-or-nothing," such as with innocuous apps (e.g., a calculator app) or apps where all the content is age-restricted (e.g., hookup apps). However, many apps need more fine-grained enforcement because they can be used to access both age-restricted and non-age-restricted content or experiences. For example, in jurisdictions where minors can use social media apps, but only with safer defaults, the service provider must either offer multiple apps (potentially one for each minimum age) or condition the app's behavior on age eligibility information it receives from the device.

Device-based age assurance architectures have been proposed by Meta[140] and porn site operator Aylo,[141] and form the basis of legal requirements for app stores under laws passed in Texas,[142] Utah,[143] and California.[144] These requirements vary substantially between states, and there does not yet appear

[138] Yoti, "Yoti ID is your secure Digital ID."
[139] This is already the case with some parental controls systems, such as iOS and Android parental controls. See Apple, "Use parental controls to manage your child's iPhone or iPad"; Google, "Manage your child's Google Play apps."
[140] Hutchinson, "Meta Calls for New Legislation That Would Force App Stores to Implement Age Restrictions."
[141] Aylo, "Aylo response to Ofcom consultation on Guidance for service providers publishing pornographic content."
[142] Texas, "App Store Accountability Act."
[143] Utah, "App Store Accountability Act."
[144] California, "AB 1043."

to be broad consensus between jurisdictions on the form of requirements for device-based age assurance.

As shown in Figure 6 below, when device-based evaluation is used, enforcement can happen either on the device or on the server, or on a combination of the two.

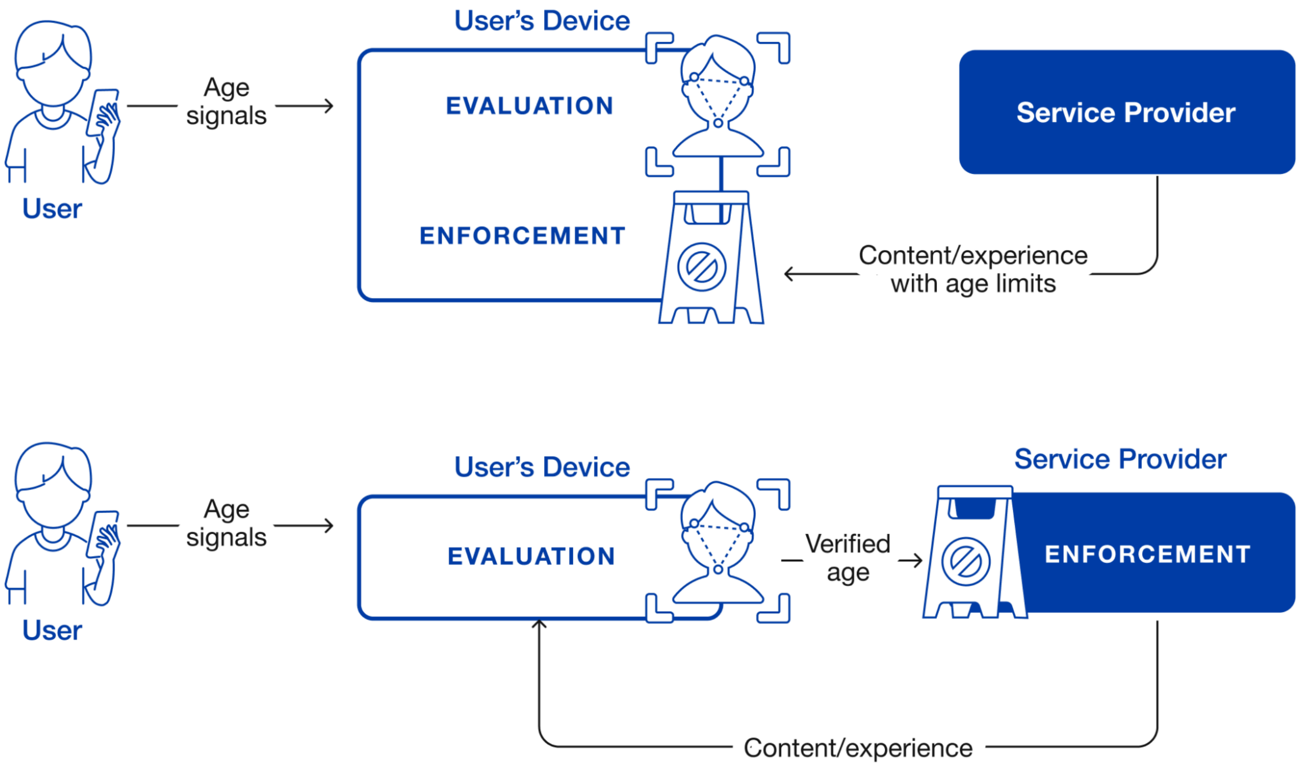

Figure 6. Two models for device-based evaluation. In device-based enforcement, the service provider offers content or experiences labeled with age limits, and the device determines whether to allow the user access to the content or experience based on the user's verified age. In server-based enforcement, the device sends the user's verified age to the service provider, which provides the appropriate content or experience.

1. Enforcement for apps

Once the app knows the user's age eligibility it can then provide age-appropriate content or experiences to the user. The precise implementation details of content delivery may vary from app to app. In particular, the app can locally select which content is appropriate based on the user's age (device-based enforcement), or tell the service provider about the user's age eligibility so that the service provider can provide appropriate content from the server side (server-based enforcement), or some combination of the two. Because the service provider also operates the app, these internal details are up to the service provider.

2. Enforcement for web browsers

Web browsing—including any app with an in-app browser[145]—is a special case because the browser can be used to access content which is not affiliated with the provider of the browser.
In the web browsing case the device and server need to cooperate to provide enforcement, as follows:[146]

- **Device-based enforcement:** Sites can self-identify as providing age-restricted content by sending an indicator such as the "Restricted to Adults" (RTA) label.[147] The browser would then be responsible for blocking the relevant content.
- **Server-based enforcement:** The browser sends an indicator of the user's age eligibility to the server. The website is then responsible for providing the appropriate content or experience from the server side.

Device-based enforcement on the web is less effort for the server, because the server only needs to signal the age-restricted status of individual pages. However, it is also less flexible than server-based enforcement because on its own the browser can only block content, not offer content conditional on the user's age.[148] If more flexible behavior is desired, it is most likely easier to have the browser send an age indicator to the site and allow the site to provide the correct experience.

3. Responsibility for Enforcement

Because there are a number of places where age enforcement can occur, this raises the question of which entities are made responsible (if any) for ensuring that enforcement happens.

For "all-or-nothing" mobile apps, where the intent is to prevent apps from being installed or used unless users are age-eligible, app stores (whether associated with the operating system or operated by a third party) may be a convenient place to locate enforcement. Once the device has established the user's age range, the app store can then block installation or execution of any apps which are age-restricted.

App stores are less suited for enforcing age assurance for apps which require finer-grained enforcement. App stores could require service providers to perform their own age enforcement as a condition of appearing in the app store, and app stores could attempt to verify that service providers

[145] Many apps that are not themselves web browsers contain in-app browsers that allow the user to view web content. See, e.g., Instagram, "Edit Instagram's in-app browser settings"; Meta, "About the in-app browser for Facebook and Instagram."
[146] It is also technically possible for the browser to analyze site content or rely on an external content filtering list. This approach is much more challenging, as it requires the browser vendor or operating system vendor or some third party to assess whether each site on the internet is suitable for minors in each jurisdiction. This is inevitably expensive and error-prone given the large number of websites available. Research on existing list-based, user-side filtering systems has found problems with both "overblocking," where content that should not be blocked is, and "underblocking," where content is inappropriately blocked. See Mathewson, "Schools Were Just Supposed to Block Porn"; Tutor, *Parents' Survey*.
[147] Association of Sites Advocating Child Protection, "What is the RTA Label?"
[148] It is technically possible for the site to send two different versions of its site to the web browser along with JavaScript code which reads the age assurance status of the user and displays the correct version. This is conceptually the same as server-based enforcement even though it is technically happening in the browser. It is also possible for the site to probe the browser to determine whether age restrictions are in place and then condition its subsequent behavior on the result.

do so as part of the app store review process. In practice, accurately and comprehensively testing how all apps (millions, in some cases) enforce age assurance would be challenging for app store providers, and requires the app store to verify the correct implementation of software which it did not write and of which it may have only a limited understanding. In addition, app vendors could attempt to evade review by providing an app which performs age assurance correctly while under test but then not when fielded.

In the case of browsers or apps with in-app browsers, locating enforcement in the app store would mean that the app store would need to check that these apps participate in age enforcement, but it would also be necessary for any age assurance mandate to require that sites participate in enforcement, either by providing the appropriate label or by adjusting their behavior based on a browser-provided indicator, as discussed above. Because many desktop users install software outside of the operating system app store and current platforms only provide minimal oversight of desktop applications, the situation on desktop is more complicated, as described in Section VI.B.5.b.

Levying direct age enforcement requirements solely on app vendors—and especially browser vendors— is likely to be unworkable in practice because it is straightforward to build a noncompliant browser in another jurisdiction and make it globally available for download. For instance, it is trivial to build a new browser based on open source web browsers such as Chrome or Firefox; a vendor could readily create such a browser which did not do age enforcement, leaving authorities with the problem of identifying the vendor and compelling them to comply. By contrast, it is possible to block such a noncompliant product in the app store, whether reactively or proactively.

4. Baseline Accuracy

The major consideration for assessing the accuracy of device-based age assurance is the accuracy of the device's age evaluation. This accuracy is minimally determined by the underlying age assurance mechanism. In general, because age assurance need happen only once, device-based evaluation may make it more feasible to use higher-friction age assurance mechanisms.

As with apps that use server-based enforcement, there is an open question about whether the user's present location or location at initial activation should be used for the purposes of determining age assurance; either choice is compatible with device-based enforcement.

5. Circumvention

The security of device-based evaluation is rooted in the security of the device itself, which is responsible for securely evaluating the user's age and taking appropriate action, whether that is communicating it to apps or restricting some apps from executing. Circumventing device-based age assurance may be possible if a user can modify the device's configuration or software. These methods of circumventing device-based evaluation are not available on closed platforms where the device operating system's behavior can be assured.

All device-based evaluation systems are potentially circumventable if an adult assists the minor by unlocking the device for them: the adult performs the age assurance process–either on a device they purchase for the minor or on a minor-provided device–and then provides the device to the minor, who uses it as usual. If age assurance only needs to be performed once—one of the advertised benefits[149] of device-based age assurance-–then circumvention can be achieved through this single age check. Devices could be configured to require the user to demonstrate their age regularly or to demonstrate that they are the same user who demonstrated their age via biometric comparison. Both of these approaches involve additional friction for adult users.

With device-based age assurance, the device can use the user's location directly to make jurisdictional decisions. As with server-based age assurance for mobile apps, a minor might attempt to circumvent age assurance by changing the apparent location of the device by modifying apps, using location spoofing software, or GPS-spoofing hardware. The operating system can implement similar defenses in this case.

The other circumvention properties of device-based evaluation depend on the platform type and enforcement mechanism.

a. *Mobile Apps*

In situations where app stores are responsible for ensuring that apps conduct age enforcement, the main circumvention path for mobile devices would involve users obtaining apps from outside the app store. In cases where an app is age-restricted, loading apps from outside the app store would permit bypassing any app store-based enforcement.

iOS devices are designed so that all users, regardless of age, are only permitted to install apps with Apple's permission. Inside the US, users are restricted to Apple's app store. The EU permits alternative app stores but at the time of this writing apps must still be "notarized" by Apple, which provides a point of control for which apps can be installed.[150] Both ordinary app store listing and notarization require some review by Apple.[151] In principle, Apple could use that review to enforce compliance. In either case, Apple could forbid installation of noncompliant apps.

Android devices also have an app store (in the US and EU most frequently the Google Play Store). As with the iOS app store, Google requires content review,[152] but it is possible to "sideload" apps outside of the confines of the Google Play Store, and to install third-party app stores. Google has announced plans[153] to require developers to register in order to make apps for certified Android Devices.[154] Similar technical mechanisms could be used to require that apps comply with age restrictions, and third-party app stores could be required to enforce similar restrictions.

[149] Jackson, "Who Bears the Burden?"
[150] Apple, "Update on apps distributed in the European Union."
[151] Apple, "App Review Guidelines."
[152] Google, "Providing a safe and trusted experience for everyone."
[153] Frey, "A new layer of security for certified Android devices."
[154] Android, "Hundreds of partners ship Play Protect certified phones and tablets."

Unlike iOS devices, it is possible to “root” Android devices to install third-party operating systems (frequently forks of Android), which might not enforce age assurance restrictions, thus allowing circumvention. It is possible to manufacture new Android devices to resist rooting, much as Apple devices do today, but that would not prevent rooting of current devices. Rooting requires a modest degree of technical sophistication and so may not result in widespread circumvention even in the absence of such restrictions.

Even if alternative app stores and rooting are technically possible, In the case of non-browser apps, service providers can address circumvention because the service provider can perform app integrity checks to ensure that the user is running a valid version of the app which correctly performs age assurance. However, this countermeasure does not work for websites, because browser integrity cannot be verified remotely by the site provider.[155] Preventing circumvention for browsers requires preventing users from installing or running browsers that do not enforce age restrictions.

b. Desktop Devices

The primary circumvention challenge for device-based evaluation on desktop is the user’s ability to install browsers which do not enforce age restrictions, and thus access age-restricted websites (web browsers rather than apps are the dominant access method for internet services on desktop). All popular desktop operating systems allow users to install new software of their choice and users routinely do so. For example, approximately 70% of desktop users[156] worldwide use the Chrome browser and a majority of those users downloaded and installed it themselves. Both Windows and macOS implement “code signing” systems which are intended to authenticate software before users install it.

Developers for Windows can sign their code either using the Windows Trusted Signing Service[157] or with their own certificates.[158] Users can still download and run unsigned executables, but by default the operating system will warn them and require them to explicitly override the warning. Unlike on mobile platforms, Microsoft’s code signing system does not involve content review. Apple nominally requires that code both be signed by an authorized developer and scanned[159] for malicious code. Software which passes these checks can either be notarized (signed) or distributed in the app store. These checks are automated and do not require manual review of policy compliance. As with Windows, macOS users can still install unauthorized code but must reconfigure the operating system to allow it.[160] In both cases it is possible for a developer to download source code for an application and build and install it without having it signed.

[155] See Web Environment Integrity. Wiser et al., “Web Environment Integrity Explainer.”
[156] Statcounter, “Desktop Browser Market Share Worldwide.”
[157] Microsoft, “What is Artifact Signing?”
[158] Microsoft, “List of Participants.”
[159] Apple, “Notarizing MacOS software before distribution.”
[160] Apple, “Updates to runtime protection in MacOS Sequoia.”

Preventing loading non-compliant browsers on desktop devices would require a number of changes to current practice.The operating system would need to be modified to prevent overriding signing/authorization checks unless the user had undergone age assurance. Operating system vendors would need a reliable way to identify and act against signed but noncompliant software, either by reviewing software upfront or revoking software later found to be noncompliant. Both approaches present challenges: upfront review would be an expensive new task for desktop operating system vendors and revocation would be easy to evade by developers who can re-register under a different name if their original software is revoked. These restrictions would also be incompatible with the use of a desktop device as a software development platform, as developing software inherently requires the ability to run it without external review.

An additional challenge is that desktop users can readily install one of the many versions of the popular open source Linux operating system, which does not enforce any restrictions on what users can install and does not have any central vendor who could enforce such restrictions. While it is technically possible[161] to prevent users from loading Linux on their computers without parental consent, standard desktop devices are not configured to prevent it. For these reasons, even if future desktop devices were configured to prevent users from loading software of their choice, for the foreseeable future there would be a large population of desktop devices which do not enforce these restrictions.

In some cases it may be sufficient to require age assurance only for mobile devices, in which case none of these circumvention paths on desktop would be material. In cases where age assurance is expected to be enforced for desktop uses, desktop devices represent a straightforward path for circumvention. Taking the strictest anti-circumvention posture implies that users would have to be forbidden from developing and installing their own software because they might use that freedom to develop and install browsers that disable the age restriction features. Allowing adults to permanently unlock devices for minors would mitigate this concern, but also provides a roadmap for circumvention because it provides minors with legitimate—although potentially pretextual—reasons to request unlocking which would then allow access to other content and experiences which the parent did not anticipate. Section VIII.E.6 discusses the broader implications of this type of restriction.

Finally, desktop machines are frequently shared in the home setting.[162] Depending on the precise configuration of the machine—specifically whether different users have different accounts—a minor may have the same settings as an adult, which would enable circumvention. Circumvention of this kind may become more widespread in cases where age assurance becomes common; for example if age assurance is required for social media, then a child may be able to leverage a parent's age assurance both for social media and for restricted content such as pornography.

[161] Microsoft, "Secure boot."
[162] Cisco, "Actions Speak Louder Than Words"; Lucchesi, "A Family Affair."

6. Availability

Because device-based age assurance is not widely deployed, for some time there will be a significant number of non-updated devices which do not perform age assurance and therefore are able to access age-restricted content. The duration of the transition period depends on a number of factors, including how aggressively older devices are updated and the extent to which it is considered acceptable to deny services entirely to non-updated devices. For example, operating system vendors might only provide age assurance on new operating system releases, even though many users are still on old operating systems, especially on Android.[163] In the case of mobile apps, service providers can either require age assurance (thus "stranding" any users who are unable or unwilling to update their devices)[164] or permit access on older operating systems (thus enabling users of those devices to bypass age assurance). In the case of browsers, service providers will generally not be able to determine whether devices are upgraded, thus lessening the effectiveness of age assurance during this transition period.

7. Privacy

The privacy properties of device-based evaluation depend on two factors:

1. Who is responsible for evaluating the user's age in order to configure the device.
2. Which signals are used to establish the user's age.

If the device is doing evaluation directly (e.g., via a digital credential or perhaps on-device facial age estimation), then the privacy risk is relatively low, assuming that the vendor is trusted. In the more likely case where the user's age is evaluated by the device or operating system vendor (using the types of remote age assurance techniques described in Section VII), or at the point of sale, then the privacy properties depend on the information revealed to the vendor and depend on the specific age signal. Unlike server-based architectures, the device vendor need not learn which services and experiences the user is accessing.[165] In addition, the evaluator learns that this specific user is interested in having age restrictions removed from their device.

In general, the more common this request is, the less information it leaks. For example, if the clerk asks the user for their ID at the point of purchase and automatically configures the device appropriately, or the device or operating system vendor routinely asks for this information as part of system registration, then the decision to establish one's age has only modest privacy implications in terms of identifying those who want to access age-restricted content. In other words, in a world where every adult purchaser of a device is subject to an age check, the fact that any given user completed the age check does not provide much indication about that user's engagement with age-restricted

[163] Statcounter, "Android Version Market Share Worldwide"; Statcounter, "Mobile & Tablet iOS Version Market Share Worldwide."

[164] In some cases, updating may be impossible if the vendor has not provided updates for older hardware. In other cases, it may be merely disruptive.

[165] Devices may independently report this information to the device vendor, but it is not necessary for age assurance.

content. However, such a system incurs a privacy cost on everyone who wants to purchase a device, as well as users of existing devices, if age assurance is retroactively imposed on them.

In addition, if the device allows websites to learn the user's age status—for instance by sending that information unsolicited or adding a web API that allows the site to query it—device-based evaluation can leak the user's age status to websites, even those which are not age-restricted. Operating systems and browsers can prevent sites which do not need the user's age status from learning it, for instance by prompting the user before revealing their age status, at some additional cost to user friction. If enforcement happens locally in the browser, then preventing this type of leakage means it would be necessary to ask the user before loading each age-restricted site.

C. 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.

VII. Assessment of Age Signals

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

- Self-declaration
- Commercial and government records (banking records, mobile network operator records, credit cards, other commercial and government records retrieved by name, email, etc.)
- Government IDs (in both physical and digital form)
- Facial age estimation
- Behavioral signals

Each of these signals is covered in detail below.

A. Self-Declaration

The most basic age signal is simple self-declaration, 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. For example, Figure 7 shows the age gate for the Jack Daniel's whiskey site:[166]

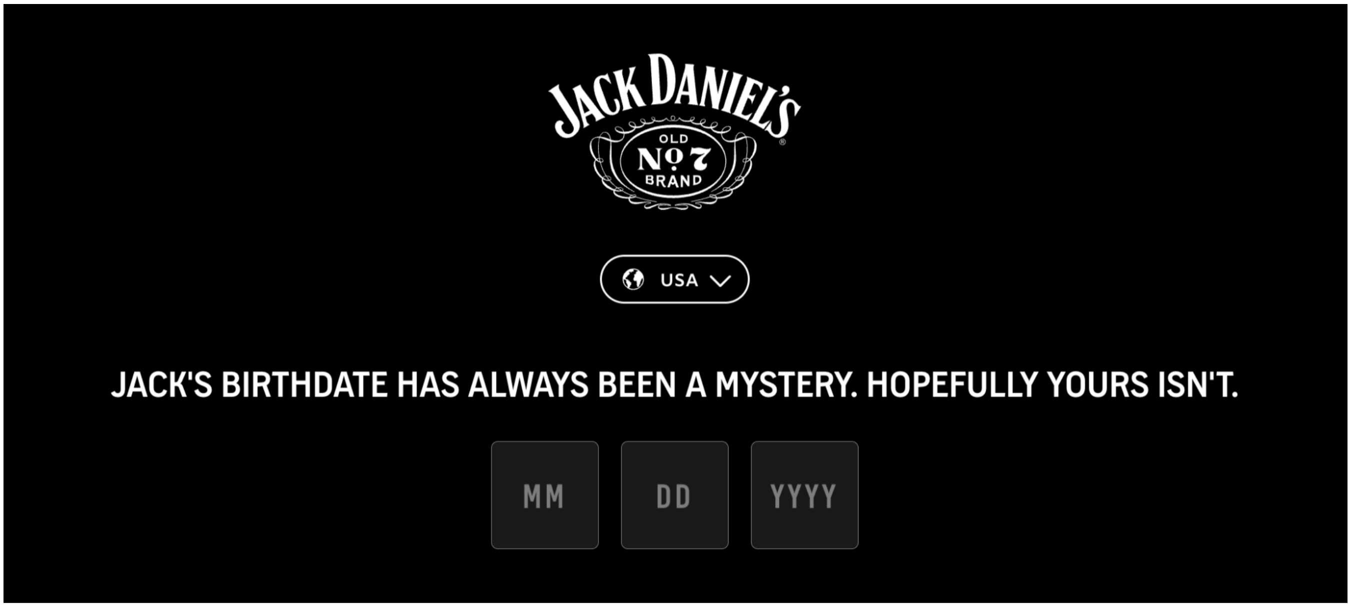

JACK'S BIRTHDATE HAS ALWAYS BEEN A MYSTERY. HOPEFULLY YOURS ISN'T.

Figure 7. An example self-declaration age-gate, from Jack Daniel's site.

[166] Jack Daniel's, "Jack's Birthdate Has Always Been a Mystery. Hopefully Yours Isn't."

Self-declaration is a very common form of age gating. For example, a recent OECD report[167] found that three of the top pornography sites (PornHub, xHamster, and xVideos) only required users to assert that they are over the relevant age and that OnlyFans only required users to enter a birthdate, although in some cases sites are changing their requirements in response to new regulations such as the UK Online Safety Act. An independent analysis of e-cigarette websites and major social media platforms found that they relied almost exclusively on different forms of self-declaration to age-gate their services.[168]

Although self-declaration is in common use, it relies entirely on user honesty; for this reason, most age assurance mandates require a stronger signal such as the ones discussed below.[169] However, self-declaration is common and can serve as a comparison point for more effective technologies.

1. Baseline Accuracy and Circumvention Resistance

If users do not misrepresent themselves, self-declaration provides an accurate signal of the user's eligibility, whether it's precise age (if birthdate is provided) or minimum age (if a "tick-the-box" mechanism is used). However, it is trivial for the user to simply lie about their age by clicking the right button or providing a false age or birthdate, as required. While some sites have mechanisms for continuously monitoring users for age-appropriate behavior,[170] there is no specific mechanism for preventing deception with simple self-declaration, often by large amounts. Multiple studies have found that minors frequently misrepresent their age in self-declarations. For example, in research commissioned by Ofcom, two-fifths of 8-12 year-olds with social media accounts had user ages of at least 16 and a third of 8-17s had a user age of at least 18.[171] Research by Pew[172] and Australia's eSafety commission[173] shows a similar pattern. Notably, in many cases (77% in the Australian report) parents or carers assisted children in setting up those accounts.

2. Availability

Self-declaration is universally available: the interface for providing age information is just a simple input field, so it is readily accessible both in apps and in web forms. There is no real chance that any user who wants to access content will be excluded because of the inability to pass the age gate, although some users might choose not to provide their birthdate and hence be excluded.

3. Privacy

The privacy properties of self-declaration depend on what is being requested from the user. If the user simply has to declare that they are above the relevant age, this is the minimum amount of information

[167] OECD, *Age Assurance Practices of 50 Online Services Used by Children*; Ofcom, "Quick guide to implementing highly effective age assurance."
[168] Dhesi and Apthorpe, "Measuring the Prevalence and Variety of Online Age Gates"; Eltaher et al., "The Digital Loophole."
[169] European Commission, "Guidelines on measures to ensure a high level of privacy, safety and security for minors online," 22.
[170] OECD, *Age Assurance Practices of 50 Online Services Used by Children*.
[171] Ofcom, "Children's Online User Ages Quantitative Research Study."
[172] Lenhart et al., "Part 3."
[173] eSafety Commissioner, *Behind the screen*.

that can be disclosed in order to restrict access to those in the eligible age range. If the user is required to provide their birthdate, this may be sufficient along with other otherwise non-identifying to identify the user.[174]

As noted above, there are a number of contexts in which a person's name combined with their date of birth are used as authenticators, such as for medical services in the US. If the user makes an account, they will generally be required to provide their name and/or email address from which their name can be derived. The combination of name and birthdate may allow the evaluator to learn enough to impersonate the user to a third party; in particular this is the type of information that would be needed for identity theft.

4. Assessment Summary for Self-Declaration

	Baseline Accuracy	Circumvention	Availability	Privacy
Self-Declaration	High if the user is honest.	Easy.	Ubiquitous.	High. Moderate if birthday is requested.

B. Commercial and Government Records

A broad class of age assurance mechanisms—often referred to as "age inference"—uses commercial and government records tied to a user's identity.[175] 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 commercial and government records.

1. Banking Records[176]

In many if not most jurisdictions, banking customers are required to prove their identity in order to open an account, as part of Know Your Customer programs. As part of this process, the customer will provide personal identification that includes their date of birth, allowing the bank to determine the customer's age. An age evaluator can use these bank records to determine whether a user is above the required age.[177] The basic process here is similar to services that allow users to "Sign in with Google" (or Facebook or Apple) on third-party websites:

1. At the website the user is trying to access (the evaluator), the user chooses the option to verify via bank records and selects their bank.
2. The user is redirected to the bank website.
3. The user logs into the bank using their ordinary credentials.

[174] For example, there are cases where the ZIP code, birth date, and sex are sufficient to identify an individual. See Sweeney, "k-anonymity."
[175] These mechanisms are sometimes referred to as "age inference" methods.
[176] In Ofcom's taxonomy, this mechanism is sometimes referred to as the Open Banking API.
[177] OneID, "How OneID's open banking-powered identity verification services help boost productivity for small businesses."

4. The user gives the bank consent to share information about their age (e.g., that they are over 18) with the evaluator.
5. The bank informs the evaluator that the user is within the eligible age range.

a. Baseline Accuracy

Because the bank knows the user's precise age, the baseline accuracy of banking-based mechanisms is very high. The main source of potential inaccuracy is if the bank has made errors in authenticating the customer when the account was opened. These errors can of course occur, but banks have a strong incentive to authenticate customers correctly for regulatory reasons.

b. Circumvention

Because age verification is tied to the ability to authenticate to the bank account, the main form of circumvention is for a user outside the eligible age range (e.g., a minor) to obtain access to the banking account of a customer who is within the eligible age range (e.g., an adult). In general, it seems unlikely that adults outside of a user's family will want to share their banking credentials with a minor. However, there are a number of scenarios where a minor might nevertheless be able to circumvent this form of age assurance:

- A parent or other adult shares their banking credentials with a minor for the purpose of passing the age gate.
- A parent shares their banking credentials with a minor for some other purpose and the minor uses those credentials to pass the age gate.
- A parent and minor share a computer where the banking credentials are stored and the minor uses the credentials to pass the age gate.
- A minor accesses a parent or other adult's computer where the banking credentials are stored and uses the credentials to pass the age gate.
- A minor obtains a third party's banking credentials and uses them to pass the age gate.

The first three scenarios (and to some extent the fourth) rely on a parent or other adult trusting that the minor will not abuse banking credentials for financial gain. It is not uncommon for parents to leave their computers in an insecure state and trust their children. By contrast, the final scenario represents a breach in banking system security generally.

c. Availability

This mechanism requires users to have online access to a bank account. While most adults have bank accounts, a significant fraction do not. In 2023, the Federal Deposit Insurance Corporation (FDIC) found that 4.2% of US households did not have a bank or credit union account.[178] Minority households are unbanked at a higher rate, with Black, Hispanic, and American Indian/Alaska Native having rates of between 9.5 and 12.2%. This probably underestimates the fraction of adults who do not have access to an online bank account, as in some households only one member will have a bank account and not

[178] FDIC, "FDIC Survey Finds 96 Percent of U.S. Households Were Banked in 2023."

everyone who has a bank account has online access. The United States Census Bureau's Survey of Income and Program Participation recently found that, among married couples sharing a household, just over 5% feature a single individual with a bank account in 2023.[179] Therefore, it should be expected that a significant percentage of users will be unable to use this mechanism.

Minors are even less likely to have bank accounts, which represents an obstacle to using this mechanism for younger age ranges. In the US, there is no minimum age on bank accounts and approximately 60% of children have accounts (frequently opened for them by parents)[180] although not all of those children will actually have access to the account credentials, as many of those accounts were opened for children under 6. Recent research on 89 OECD countries found 62% of 15 year-olds having some kind of account at a "bank, building society, post office, or credit union,"[181] although rates vary dramatically between countries, from 90+% in Denmark to 13% in Peru.

An additional barrier to availability is that banks must have support for bank-based age assurance; if they do not already have this kind of mechanism they will need to pay to develop it, and may have only minimal financial incentive to do so.

d. *Privacy*

This mechanism is intended to conceal the user's identity from the evaluator while only disclosing whether they are within a given age range. However, as part of the process, the bank will learn that the user is attempting to establish their age and the identity of the evaluator (e.g., the AVP). The evaluator will likely learn the user's banking institution.

An additional concern with bank account-based age verification is that it creates a potential vector for phishing scams where users are deceived into entering their banking credentials on malicious websites. In this scenario, the user goes to a site requiring age verification and is then prompted to log into what appears to be their bank account but is actually a malicious site which collects their password. While the long-term solution for phishing is wide deployment of phishing-resistant authentication mechanisms,[182] those mechanisms are nowhere near universal deployment, and encouraging users to routinely log into their bank as part of another website's authentication process is likely to have a detrimental impact on security.

2. Mobile Network Operator Verification

Under the UK Code of practice for the self-regulation of content on mobiles,[183] mobile network operators (MNOs) filter internet access[184] ("content restriction filters") by default for users who have

[179] Opanasets, "Almost a Quarter of Married Couples Didn't Have Joint Accounts in 2023, Up From 15% in 1996."
[180] Wrinn and Savvy, "Youth Accounts Map a Promising Path Forward for Banking Providers."
[181] OECD, *PISA 2022 Results*.
[182] For instance, technologies like WebAuthn or passkeys. See FIDO Alliance, "Passkeys"; MDN, "Web Authentication API."
[183] British Board of Film Classification, "Mobile Content"; EE et al., "UK Code of practice for the self-regulation of content on mobiles."
[184] Note that this does not impact subscriber's ability to access age-restricted sites via non-mobile connectivity such as Wi-Fi.

not demonstrated that they are 18 or over.[185] That demonstration can take a number of forms, as described in the code of practice:

> a) at the point of mobile device sale for new customers: inspection of document containing customer's date of birth (e.g. Driving licence, Citizen Card etc.); visual check (is the customer clearly over 18?); b) "customer not present": a valid credit card transaction for the customer; age confirmation using 3rd party agencies (e.g. Experian or Dun & Bradstreet etc.); c) documents and/or process used for contract mobile phone customers, combined with a process by which customers can manage access controls.

As a result of this procedure, the MNO already knows if the user has demonstrated that they are over 18 and this can be used as a form of age verification. The process for age verification via MNO works as follows:

1. The user provides their mobile number to the evaluator.
2. The evaluator confirms that the user is reachable at the mobile number provided. The evaluator uses a reachability mechanism such as sending a link or a code to the user via SMS and the user clicks on the link or enters the code on the evaluator's site.
3. The evaluator then sends the mobile number to the MNO.
4. The MNO consults its records and returns a response indicating whether the user is 18 or over.

Note that there is no requirement for the device which is being used to authenticate the challenge to be the one making the internet connections. For instance, the user could access a site with their web browser and use that to request the code, then receive the code on their phone, and enter it into the evaluator's site; this mechanism is only intended to demonstrate control of a device which is associated with an 18+ user.

a. Baseline Accuracy

The accuracy of this form of age verification depends on the accuracy of the MNO's mechanism for verifying or estimating age (assuming they do so at all). As noted above, MNOs use a variety of mechanisms for assurance, such as driver's licenses, credit cards, etc. Discussion of the accuracy of those mechanisms is in the remainder of this section. If the MNO has accurately determined the user's age range, MNO-based verification provides a highly accurate result.

b. Circumvention

If the child has the assistance of their parent, then the parent can easily allow them to circumvent MNO-based verification. First, the adult can simply represent to the MNO that the mobile number is associated with themself rather than with the child, providing the appropriate supporting information. Second, the adult can allow the child to use their mobile number to demonstrate their age and answer

[185] Open Rights Group, "Content filtering by UK ISPs."

the challenge for the child.[186] The child may also be able to enlist an older non-parent adult such as a sibling or a friend.

If the child is unable to obtain the assistance of an adult, they may still be able to circumvent a phone-based age verification challenge if they have access to a parent's phone, either by using the phone directly or by borrowing the SIM card and inserting it into their own device.[187]

c. Availability

Mobile phone ownership is high but not ubiquitous, with 84% of people worldwide having mobile phones.[188] In 2025, 98% of 15-year-olds in OECD countries had an internet-connected phone.[189]

However, in order to make use of this mechanism, users must first establish their age with their carrier. In the UK, because phones default to a restricted filtered mode and users can demonstrate their age to have the content restriction filters for their device disabled,[190] the MNO's records can be used for online age assurance. Some users may not wish to prove their age to their carrier for privacy reasons (see below), in which case they will not be able to use this mechanism. Data about what fraction of UK users have chosen to remove content filters is not publicly available.

Outside the UK, carriers may not maintain records of user ages, in which case this form of age assurance is not applicable. Even inside the UK, this mechanism is not applicable to age ranges other than over 18, because it leverages the existing checks for users over 18. In order to use it for other age ranges, carriers would need to start collecting ages for all users, not just those who request unblocking.

d. Privacy

MNO-based age verification results in information leakage both to the evaluator and the MNO. The evaluator learns the user's phone number, personal information which can be used to retrieve the user's identity and can be used to link up the user's behavior across providers. Even if that behavior is not itself sensitive, the creation of a behavioral profile for users can present a privacy concern.

As part of the age assurance process, the MNO will learn that the user is attempting to establish their age and the identity of the evaluator. If the evaluator is an AVP, then the MNO may not learn the service provider, but if the service provider performs their own evaluation, then the MNO will learn the service provider as well.

[186] Ibid.

[187] Note that "SIM swapping" attacks in which the attacker takes over the victim's mobile number are a common form of fraud, and would enable MNO-based circumvention. It is unclear how many minors will be willing to go to this extent to access age-restricted content or experiences.

[188] Klapper et al., *The Global Findex Database 2025*.

[189] OECD, *How's Life for Children in the Digital Age?*

[190] Note that in the UK this can be done using a credit card, as discussed below.

In addition, the mobile number can be used—either directly by the evaluator or any entity with whom the evaluator shares the number—to contact the user via voice or text in the future, thus enabling marketing, spam, and/or fraud.

3. Credit Cards

In some jurisdictions (notably, the UK), credit card issuance is restricted to users over 18.[191] In these jurisdictions, the credit card can be used as a demonstration of age. The typical process is to request that the user enter their credit card information and then make a small charge or authorization (e.g., $1.00) in order to verify the validity of the credit card information. After validity is verified the transaction is abandoned or refunded, so there is no cost to the user.

a. Baseline Accuracy

As with MNO-based verification, credit card-based verification's accuracy depends on the accuracy of the authentication used by the credit card issuer to ensure that it only issues cards to people of the appropriate age. If that mechanism is accurate, then credit-card-based verification will also be accurate. In many cases, credit cards are issued remotely with no positive identity check, relying solely on government records.

b. Circumvention

Credit card-based verification has a number of trivial circumvention measures because all that is required to pass the age gate is to have a valid credit card. This implies that a child can circumvent the age gate with the cooperation of a parent or another adult. Because the credit card is not charged for any significant amount and is later refunded, the use of the credit card to authenticate one user does not necessarily preclude it being used to authenticate another user. Anti-fraud mechanisms may prevent some forms of multiple use, but it is normal for a user to demonstrate their age multiple times, for instance for two different providers. Where evaluators choose to keep records of age assurance attempts to prevent this kind of circumvention this increases privacy risk.

Even without the explicit cooperation of an adult for age assurance purposes, a child could still potentially circumvent credit card-based verification by use of a credit card borrowed for another purchase, especially of a parent.[192] It is already common for children to make use of borrowed cards for other purchases[193] and a parent might not notice the trivial charge.

c. Availability

Credit card-based verification depends on the user having a valid credit card. In the US, as of 2023, around 82% of US adults had credit cards;[194] as of 2025 about 68% of UK adults have a credit card;[195]

[191] Ofcom, "Age checks for online safety."

[192] Even identity professionals loan their credit cards to their kids. See, e.g., Andrew Chevis, Chief Executive of CitizenCard Ltd, discussing loaning his credit card to his daughter to pay for her 16th birthday party. Chevis, "UK: CitizenCard."

[193] Papandrea and Sherrier, "46% of Parents Say Their Child Used Their Credit or Debit Card Without Permission, Racking Up $500+."

[194] Government Accountability Office, *Credit Cards*.

[195] Valev et al., "Percent people with credit cards."

and as of 2021, large shares of European residents also possessed a credit card, though with substantial variance by country (e.g., approximately 57% in Germany vs. 23% in Greece).[196] The remaining adults would not be able to use this form of age assurance.

Credit card-based verification can only be used where there is a minimum age of issuance and where that age is greater than or equal to the minimum age where access is permitted. The minimum age of credit card issuance varies widely across jurisdictions. As noted above, the UK imposes an 18-year-old minimum. In the United States, the Truth In Lending Act generally requires that cardholders be 21 or older; however, if an adult is the primary cardholder, they can allow a minor to be an additional cardholder who is issued their own card. The minimum age for authorized users varies but some banks have no minimum at all.[197] Regulations vary across the EU with over half of EU states having a minimum age under 18 or no minimum with parental consent, as of the most recently available official data (2017).[198] In Canada, the minimum age is 18 or 19 depending on the province,[199] but it is possible for minors to become additional cardholders.

d. Privacy

As described above, credit card-based verification requires the user to disclose their credit card number and potentially other information such as their address. Even if the user's name is not required, this information is often sufficient for the evaluator to determine the user's identity. Moreover, repeated use of the same credit card allows for behavioral profiling of the user, even without knowing their identity.

In addition, the various components of the credit card ecosystem involved in processing the user's credit card (payment processor, issuer, etc.) learn that the user has requested age assurance. When the service provider uses a separate AVP, then the AVP need not reveal the service provider to the credit card processor, however, recent research found that Yoti's credit card verification system leaks the identity of the service provider to their payment processor (Stripe).[200]

Finally, the card transaction can leak to any other individuals who have access to the user's credit card statement or who get notified of transactions. If a user has a joint credit card—for instance with their spouse—then other cardholders may learn that the user is visiting an age-restricted site or using an age-restricted app.[201]

As with banking-based age assurance, the use of credit cards presents a security risk to the user if the service provider or age assurance provider uses the user's credit card to commit fraud. Although many

[196] Ibid.

[197] Norman, "What Age Can You Get a Credit Card."

[198] European Union Agency for Fundamental Rights, "Minimum age requirements related to rights of the child in the EU."

[199] Canada, "Choosing a Credit Card."

[200] Minocha et al., "Papers, Please."

[201] Minocha et al.'s investigation of Yoti found that they were notified by their issuing bank of a transaction from "YOTI LTD" but that the transaction did not appear in their transaction list because Yoti had configured their payment provider (Stripe) to authorize the transaction but not capture the funds. See Minocha et al., "Papers, Please."

jurisdictions have consumer protections for credit card fraud, users might not notice the charges or might be unwilling to report the fraud if the content or experiences being accessed were sensitive or embarrassing.

4. Commercial and Government Records Searches

There is a broad class of age assurance techniques which depend on searches of commercial and/or government records. At a high level, these techniques work by asking the user to provide some form of personal information, such as an email address, name and address, or (in the US) a social security number, which is then used to search a variety of commercial and government sources for evidence that the user is within the eligible age range. The personal information can either be collected by the service provider and sent to the AVP or be collected by the AVP directly.

In some cases, this search will identify the user's specific date of birth and produce a verified age. In other cases, it simply provides an estimate, for instance if an email address has been in use for a sufficiently long time or is associated with commercial transactions (e.g., utility bills) that provide evidence that the user is over 18. The specific details of this form of age assurance mechanism are typically proprietary to the age verification provider.[202]

Some AVPs which use this type of mechanism will store the results of the verification associated with the user's personal information (e.g., the email address). This allows the AVP to skip future checks for previous users by looking up their results from their personal information.[203]

a. Baseline Accuracy

There is little independent evidence of the baseline accuracy of these mechanisms. Their accuracy is dependent on the personal information provided, the records that the AVP is able to search, and the quality of records in the jurisdiction(s) where the user lives/lived. For example, VerifyMy is an AVP which uses "proprietary algorithms and external data sources" and claims a false accept rate below 1% for an age threshold of 18.[204] However, it also reports a false reject rate in excess of 10%; this is likely to be a systematic error related to the quality of available records, rather than a random one.

b. Circumvention

The primary form of circumvention for these mechanisms is for the user to provide personal information that is associated with someone else to the verifier. Without the other person's cooperation, how easy this is depends on the personal information:

- Ownership of an email address can be readily verified by sending a one-time code to the provided address.[205] A number of AVPs require verification of ownership of email addresses.[206]

[202] McConvey, "Email address age assurance is private, complaint, and simple."
[203] This does not reduce user friction, just cost and effort by the AVP.
[204] VerifyMy, *Innovative age assurance*.
[205] Or, in the case of some email providers, such as Gmail, via social login.
[206] Ibid.; Yoti, "Age Verification."

- Social security numbers are typically considered secret and so may be difficult for children to access, although they are often not considered secret *within* families and most families will receive paper correspondence (e.g., tax records) with the social security number on them.
- There is no quick and effective way to verify that a user is truthfully providing their own name, address, and birthdate over the internet. It is unclear whether these identifiers are used in practice for records checks for remote age assurance.

With the cooperation of an eligible user, these mechanisms may be circumventable. For example, an eligible user can provide their email address and answer the one-time code email on behalf of the ineligible user. This form of circumvention may be less effective for account based services if the same identifier is used for account creation and age assurance. For example, if Alice and Bob both have accounts with TikTok (authenticated via phone number), then Alice might not be able to allow Bob to use her phone number to demonstrate her age, as it is already bound to her account. However, if there are multiple identifiers that can be used for both age assurance and account creation, then Alice may be able to use one to make an account and another on behalf of a minor who wants to circumvent age assurance.

c. *Availability*

Essentially all users will be able to provide some form of personal information as input to this process, but in many cases, there will not be sufficient records to assess the user as eligible. In VerifyMy's 2024 study[207] of 102,460 email addresses, when assessed against an 18+ threshold, no age/gender cohort (over 18) had a false rejection rate under 10%, as shown in Figure 8. As a result, this mechanism used alone will exclude a large number of users.

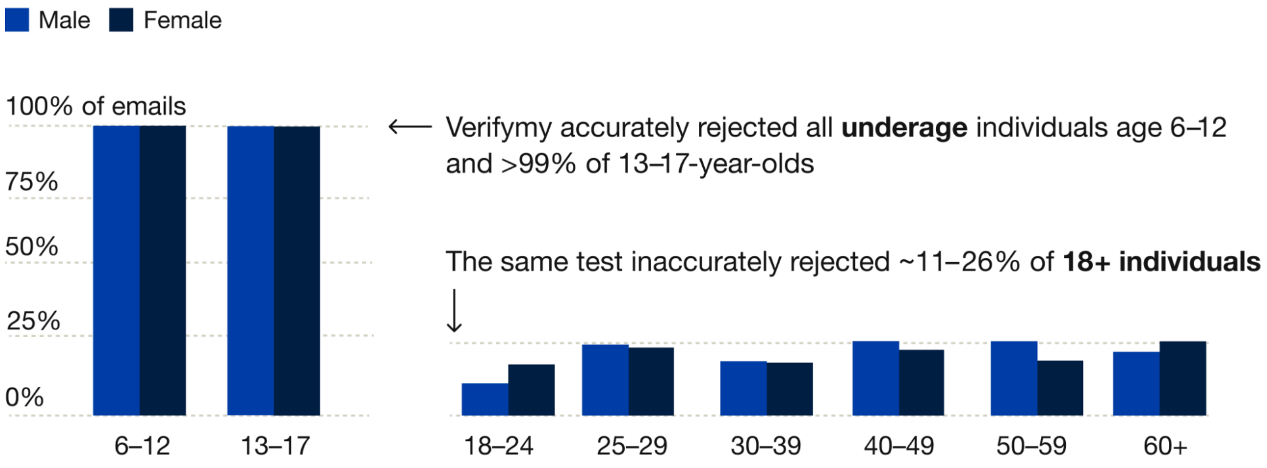

⟵ Verifymy accurately rejected all **underage** individuals age 6–12 and >99% of 13–17-year-olds

The same test inaccurately rejected ~11–26% of **18+ individuals** ↓

Figure 8. VerifyMy's reported results for email age estimation: Percentage of email addresses rejected as "unable to estimate as 18+."[208]

[207] VerifyMy, *Innovative age assurance*.
[208] Ibid.

d. *Privacy*

By definition, these mechanisms require revealing user identifiers to the evaluator. If the evaluator collects the personal information, then this allows them to learn the user's identifiers if those identifiers would not otherwise be collected by the evaluator (for account setup). If the AVP collects the personal information, then they will learn the user's identifiers and the evaluator will learn the user's age eligibility.

If the AVP stores the results of the evaluation, then this creates a semi-permanent record of the user's activity, which is especially relevant if the content or experience the user is accessing is itself sensitive. This is a general risk of AVPs, but is a particular risk with this class of age signal because the process inherently involves a persistent user identifier which can be used to store the results. In some cases, AVPs claim to protect this data with encryption[209] and hashing[210] identifiers (e.g., hashing the email address),[211] but this does not in fact protect user privacy against a malicious or compromised AVP.[212] Hashed email addresses can effectively be reversed using a brute-force attack to determine the email addresses in the AVP's database.[213] And anyone with access to the AVP's database can hash an email address and look it up at any time to determine if a specific user is in the database.

[209] Yoti, "You're in safe hands."

[210] Hashing refers to applying a one-way transformation to an input value to produce a fixed-length output which is characteristic of the input value.

[211] OECD, *Age Assurance Practices of 50 Online Services Used by Children*.

[212] Even if data is encrypted at rest before being stored by the AVP, the AVP needs to have the encryption keys in order to use it, and so a malicious or compromised AVP can use those keys to decrypt the data.

[213] Englehardt et al., "I never signed up for this!"

5. Assessment Summary for Commercial and Government Records

	Baseline Accuracy	Circumvention	Availability	Privacy
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.

C. 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).[214] Both versions are discussed below.

1. Physical IDs

Identification via physical government-issued identification documents is familiar to most people from contexts such as purchasing alcohol or cigarettes. The way these documents are used is that the the evaluator compares the photo on the document to the person in front of them and if they match, the evaluator trusts the rest of the information on the document, which may either be printed on the card or in a machine readable format such as a bar code or NFC-readable chip.[215] To state the security logic clearly:

1. The card is tamper-resistant and so the evaluator can trust that the picture and information on the card are what was intended by the issuer.
2. The picture matches the person in front of them and therefore the card is theirs.
3. Because the card binds the picture and the information on the card together, the information on the card applies to the person in front of them.

When used for remote age assurance, the user provides a digital copy of the identity card, via a photo, scan, or video. In many cases users are also required to demonstrate that they are the owner of the identity card using either a still selfie or a live video of their face. The evaluator then validates (1) the legitimacy of the identity card; (2) that the user matches the photo on the card; and optionally (3) that the user is live rather than a still image. Once all of these checks are performed, the evaluator can then use the birthdate on the card to determine whether the user's age is within the required range.

a. Baseline Accuracy

Government ID-based age assurance has essentially the same accuracy as the underlying issuance mechanism, which is in general quite high. For example, US REAL ID cards require the user to provide documentation showing "1) Full Legal Name; 2) Date of Birth; 3) Social Security Number; 4) Two Proofs of Address of Principal Residence; and 5) Lawful Status."[216] In principle, the age assurance system could incorrectly scan any printed information on the card, but existing computer vision techniques have high accuracy as long as the image is of reasonably high quality. If the card has machine-readable features such as a bar code or NFC, then these can usually be read without error.[217]

[214] International Organization for Standardization, "ISO/IEC 18013-5:2021."

[215] For example, REAL ID-compliant US driver's licenses have PDF 417 bar codes on them, containing much of the relevant information in the document, and ICAO 9303 passports contain NFC-readable information. See International Civil Aviation Organization, "Doc 9303"; International Organization for Standardization, "ISO/IEC 15438:2015."

[216] United States Department of Homeland Security, "REAL ID Frequently Asked Questions."

[217] In part, because they are designed with features to enhance readability under poor conditions.

b. *Circumvention*

Like all age assurance mechanisms that involve capturing information via the user's camera, government-ID mechanisms are subject to two main avenues of technical attack:

- **Presentation attacks** in which the attacker manipulates the camera input, for instance by taking a video of an image of someone else.
- **Injection attacks** in which the attacker hijacks the camera feed and provides any content of their choosing.

Injection attacks are a more powerful and general technique because they provide the attacker with complete control over the input, but in some cases they can be prevented by the use of trusted devices such as mobile phones. Appendix B provides more information on both forms of attack.

Whichever attack modality is used, there are two main approaches to circumventing a physical ID-based age assurance system:

- Providing a fake identification card.
- Providing a valid identification card that does not belong to the user.

These are discussed in more detail below.

Fake Identification Cards

The obvious attack on any ID-based system is for the user to use a fake identification card. The appearance of valid ID cards is well-known, and there are existing services which will manufacture cards of the user's choice.[218] These cards will have the user's image and a valid age, so detecting circumvention relies on detecting that the card is fake, typically based on using the physical security features on the card (e.g., holograms, microprinting, etc.), which were designed for in-person use and may not be effective remotely.[219] Some AVPs claim to detect these features remotely, but little independent data on effectiveness is available at the time of this writing.[220] It may also be possible to mount an injection attack where a wholly virtual card is transmitted to the verifying system.

If the evaluator is able to query the credential issuer's database and obtain the contents of the ID, including a photo, then these physical security features are no longer necessary, as the credential just becomes a lookup key. However, some jurisdictions do not allow private parties to query their databases. For example, the US Driver's Privacy Protection Act restricts the disclosure of "highly restricted personal information" on a driver's license, such as the subject's photograph.[221] In jurisdictions where private parties can query the database for the user's details but not their photo,

[218] Cox, "Inside the Underground Site Where 'Neural Networks' Churn Out Fake IDs."

[219] ANSSI and BSI, "Remote Identity Proofing ANSSI-BSI Joint Release."

[220] Notably, the US Department of Homeland Security (DHS) has declined to publish measured rates, but instead recommends a 10% false accept rate, suggesting that substantially better performance is difficult to achieve. Howard et al., *A Quantitative Framework for Evaluating Remote Identity Validation Systems*.

[221] United States, "18 U.S. Code § 2721."

then the security features are still relevant because the user might substitute their own photo in a card for a valid user. Evaluators may also query non-governmental databases for the validity of ID cards.

Non-Matching Cards

In many scenarios, there is no need for the provider to know the user's name, and so if the user is not asked to provide a selfie or self-video, then there is no real barrier to using another person's card. In this case, it most likely suffices to have an appropriate card scan in order to mount an attack. In other scenarios, such as account creation, the user may be requested to provide a name and the provider may refuse to create an account with a non-matching name.

If the user provides a selfie or a self-video, then the evaluator attempts to match the photo on the card against the provided image of the user. In practice, this matching will almost certainly involve some kind of algorithmic verification rather than human verification in order to reduce cost. Attacks on matching are similar to those on facial age estimation, including injection and presentation attacks (see Appendix B).

c. Availability

The availability of this mechanism for older teens and adults is tied directly to the availability of government-issued IDs, which varies substantially across jurisdictions. On the low side of availability, a 2023 survey estimated that 9% of American adults do not have a valid driver's license.[222] Non-white adults are overrepresented in the population without licenses, as are those with lower socioeconomic status and younger adults (18-29). This last group presents a special problem in settings where age estimation is implemented with a policy where the user is asked to prove their age with an identification card if they appear close to the age limit, as this is the hardest group to accurately estimate their eligibility, and they are the most likely to lack ID. On the high side of availability, many countries have a mandatory national identity card. In these countries, it is reasonable to expect near-universal availability for older teens and adults.

In many jurisdictions, this mechanism cannot be used to verify the ages of children or young teens, because they are not required or able to obtain government-issued identification. Even in countries which have mandatory ID, that ID may not be mandatory for those under 18. For these reasons, government ID cards are often not suitable for demonstrating that a user is a minor (as distinct from requiring ID to demonstrate that a user is 18+), especially for eligible age ranges below 18, as a large fraction of minors will have no suitable identification.

An additional availability challenge with government IDs is that the evaluator may need to be able to process many different types of ID, both because there are multiple types of ID within a jurisdiction (passport, driver's license, social insurance card), and because there may be users with IDs from multiple jurisdictions (e.g., other states in the US or other countries in the EU). If an evaluator chooses to limit the types of ID they accept, this may limit availability.

[222] Rothschild et al., *Who Lacks ID in America Today?*

Some users who are technically able to participate in age assurance may be unwilling or unable to provide the AVP with a facial image, for instance for privacy or religious reasons. In addition, users with visible facial differences may not be able to authenticate via government ID-based systems if the system cannot confirm that their face matches their credentials.[223]

d. Privacy

Using a photographic ID to demonstrate age inherently reveals the user's identity to the entity performing the evaluation, as well as other information that is on the card, such as a precise birthdate, national identity number, etc. In the US, the REAL ID Act of 2005 requires that compliant driver's licenses contain the user's sex and home address, and some US states contain multiple other kinds of data. For example, a California driver's license also includes the height, weight, eye color, and whether the driver is an organ donor. In addition to the information on the ID itself, if the user is required to provide a selfie or self-video, this provides a current image of the user, thus making it harder for them to deny their activity.

If the evaluator retains the ID information used by the user to demonstrate their age, the resulting store of data becomes a privacy risk, as it could be compromised by an attacker, subject to legal process, or otherwise leak. This risk is not theoretical: in one 2025 incident, information provided to Discord for age assurance was compromised.[224] Driver's licenses, selfies, and other identity information collected by the Tea app was hacked and leaked in 2025,[225] as were driver's licenses collected by the TeaOnHer app.[226]

In addition, if the user provides their photo, it can later be used for other purposes, such as AI-based deepfake or "nudification" systems. Note that this is a potential concern even in cases where the user might otherwise be providing their photo to the service provider, as with social media services; if the service provider uses a third-party age verification provider, that AVP gains access to the user's photo that it would not have had in the absence of age assurance.

2. Digital IDs

Digital IDs are meant to be the digital equivalent of physical ID cards. Instead of physical tamper-resistance features, they rely on cryptography to ensure that the information they are intending to convey is authentic and unaltered when it is sent over the internet.

There has been increasing deployment of fully digital identification mechanisms such as Mobile Driver's Licenses (mDLs) in many US states, the EU Digital Identity Wallet,[227] and EU Age Verification App.[228] These digital ID systems are generally based on a standardized technical specification,

[223] Facial Equality International, "Facial Recognition and the Facial Difference Community."
[224] Peters, "Discord customer service data breach leaks user info and scanned photo IDs."
[225] Wise, "Tea encouraged its users to spill. Then the app's data got leaked."
[226] Silberling and Whittaker, "TeaOnHer, a rival Tea app for men, is leaking users' personal data and driver's licenses."
[227] European Commission, "A digital ID and personal digital wallet for EU citizens, residents, and businesses."
[228] European Commission, "EU Age Verification Solution."

ISO/IEC 18013-5.[229] These mechanisms can be used for remote age assurance and have some superior—though still imperfect—technical properties when compared to physical credentials.

In order to translate a physical credential into the digital domain, the physical security features are replaced with digital security features. Specifically, the credential is protected using a cryptographic technique called a "digital signature." A digital signature is used with a pair of cryptographic keys, a "private key" and a "public key." The private key is known to the signer and is used as an input to the digital signature function along with the document; the public key can be known by anyone and is used to verify the signature. As long as the digital signature associated with a document (in this case a digital identity document), can be verified correctly the verifier can be confident that the holder of the private key signed it and that the document has not been tampered with. Thus, a digital credential can be created by taking an existing physical credential and having the credential issuer (e.g., the Department of Motor Vehicles) sign the data on it.

Digital credentials can be used like physical credentials by including a biometric such as a picture. The verifier can then ask the user to provide a selfie and compare it to the biometric. However, they can also be used in a more privacy-preserving fashion using a cryptographic technique called "selective disclosure," in which each individual attribute (e.g., name, birthdate, address, etc.) is separately cryptographically protected and can be individually revealed,[230] thus allowing the user to prove their age without revealing their identity or their face.

In order to prevent reuse of digital credentials by others, many digital credential systems "bind" the credential to a given device. Device binding works by the device having its own public-private key pair, with the private key stored in a "secure element" on the device, so that the user cannot extract it and share it with others ("cloning"). In order to use the credential, the user provides the verifier with (1) the signed credential, (2) a proof that it knows the device private key, and (3) the revealed values of the relevant attributes (e.g., date of birth).

In many cases (US mDLs, EU Digital Identity Wallet), the digital identity is a digital version of an existing physical credential. In the typical case, the user has a physical credential and can load it onto their mobile phone by taking a photo of the credential and then using the phone camera to provide a selfie or self-video that demonstrates that they are the subject of the credential. The phone then communicates with the original credential issuer (e.g., the relevant Department of Motor Vehicles) to get a digital version of the credential which is automatically loaded on the device.

A digital credential can also be a separate credential as in the case of the EU Age Verification App, in which the user proves their age with another mechanism (e.g., by connecting to the credential issuer and showing their ID or performing facial age estimation) and then is provided with a digital credential

[229] International Organization for Standardization, "ISO/IEC 18013-5:2021."

[230] In an ISO 18013-5 credential, selective disclosure is provided by signing over a list of hashes where each hash is computed over a single attribute and a random value. In order to disclose a value, the user provides the attribute and its associated random value, and the verifier can compute the hash itself and verify that it matches the value in the credential.

that carries only age information. The digital credential would be loaded onto the device and used similarly to a digital credential based directly on a physical credential.

When a user has a digital ID stored on their device, apps on the device can directly query the digital ID for age verification purposes, as shown in Figure 9 below:

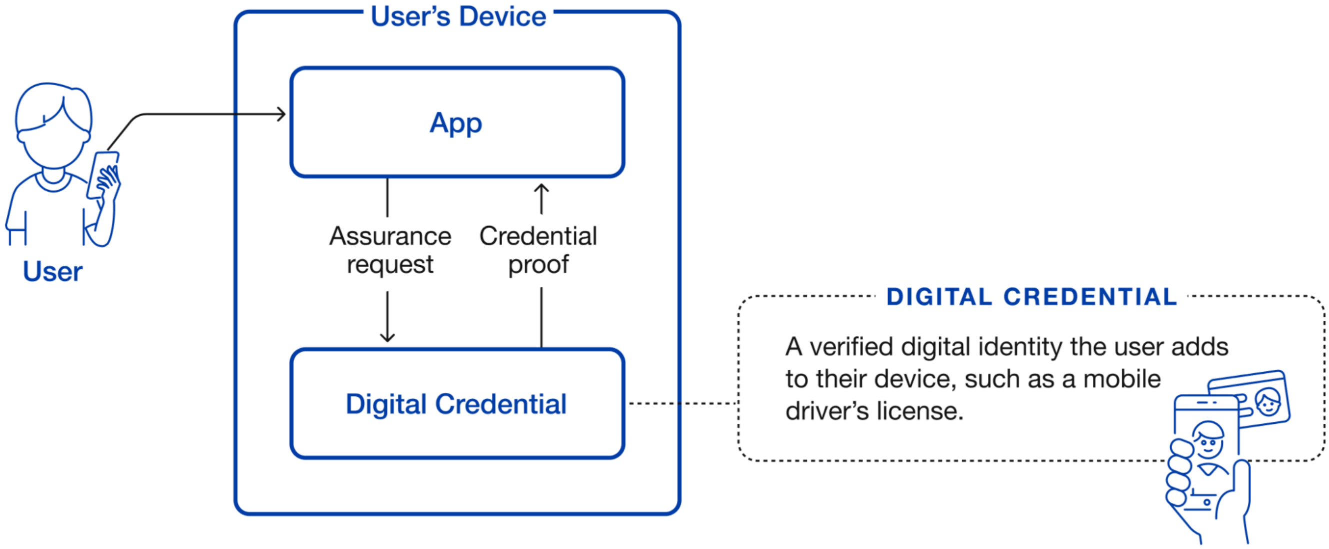
User's Device
App
User
Assurance request
Credential proof
Digital Credential
DIGITAL CREDENTIAL
A verified digital identity the user adds to their device, such as a mobile driver's license.

Figure 9. App-based age assurance using a digital credential.

Prior to allowing the user to access age-restricted content or experiences, the app would query the digital credential via operating system application programming interfaces (APIs). The APIs could either provide cryptographically verified age information or just attest that the OS has verified the age information on the credential.[231] Either way, the app then knows that the user is in the eligible age range and can safely provide the requested content or experience.

In the case of a user accessing a website rather than an app, the user's age information stored in the digital ID can be conveyed to the website via the W3C Digital Credentials API, which serves an analogous purpose to the operating system APIs in the web context.[232] A web-based age assurance flow using a digital credential is shown in Figure 10 below.

[231] The app has to trust the operating system in any case, so it need not necessarily verify the credential itself.
[232] Caceres et al., "Digital Credentials."

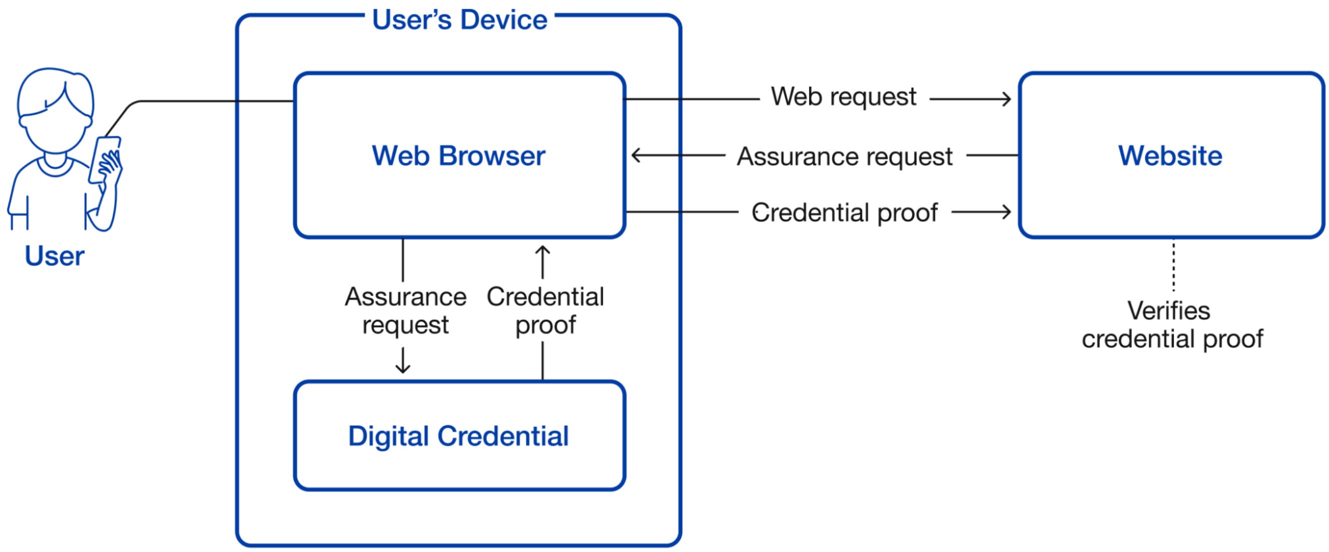

Figure 10. Web-based digital credential verification.

If the user is browsing the web on their mobile device which has the credential enrolled, the browser can talk directly to the on-device wallet and use it to produce a proof of the validity of the credential and the user being in the eligible age range. That proof is then sent to the service provider, which verifies the proof and provides the appropriate content or experience. It is also possible to use the digital ID for age verification when the user is on a desktop browser, either via desktop-to-mobile or by having the browser display a QR code that the user can capture on their mobile device, allowing it to send the proof of validity to the evaluator directly, as shown in Figure 11 below.

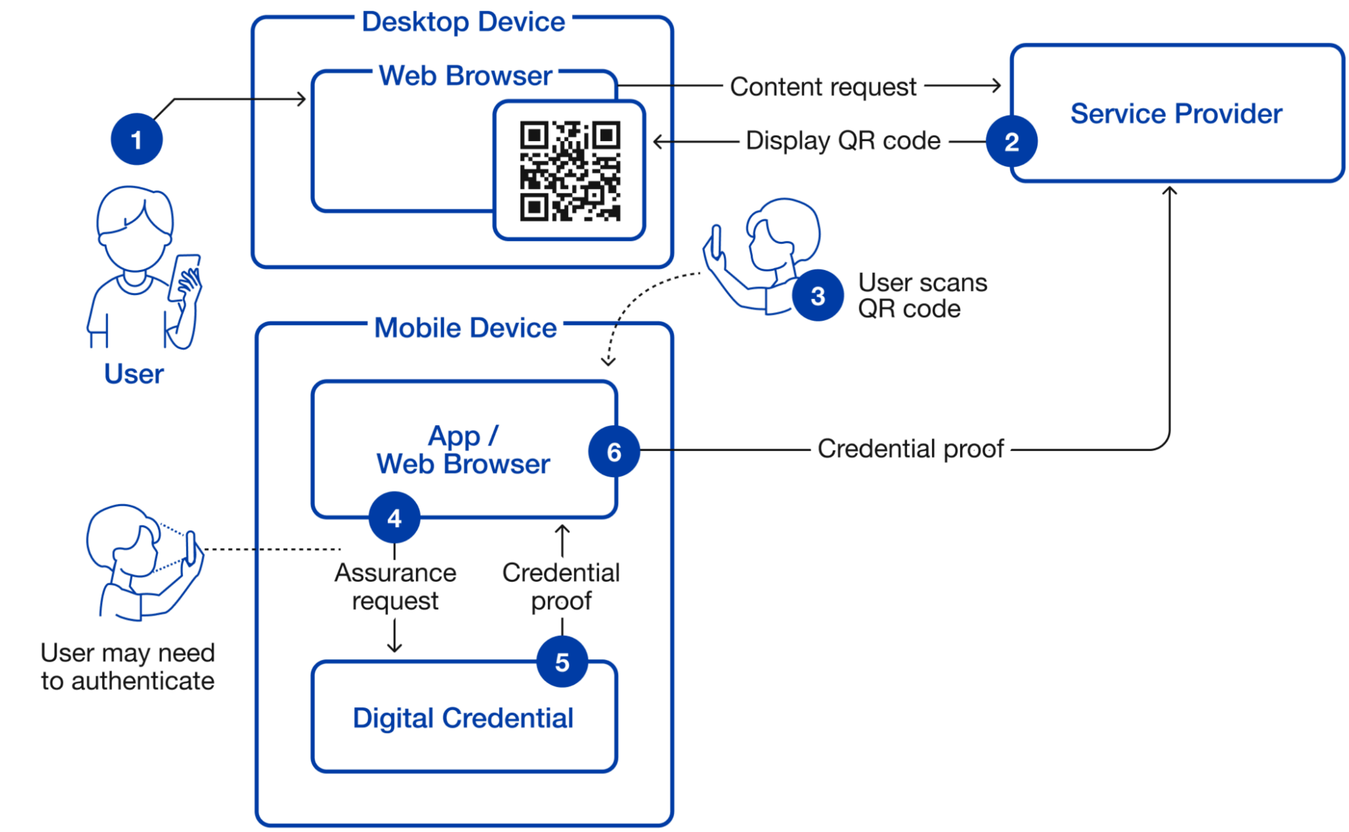

Figure 11. Digital-credential-based age assurance with a desktop browser and mobile credential.

a. *Baseline Accuracy*

The baseline accuracy of digital IDs should be the same or better than that of physical IDs in the age assurance context. Digital IDs have the potential to be significantly better, as they remove the need to process potentially flawed images with visual artifacts, optical distortion, etc. The digital credential itself can be transmitted without error, and when the user initially enrolls their physical credential, the enrollment process can ensure they provide a high quality image and verify the details in the credential against the original database. Moreover, any given credential issuer only needs to process its own physical credentials and does not need to be prepared to handle credentials from a variety of jurisdictions.

b. *Circumvention*

Because the digital ID is digitally signed, it is not subject to attacks which attempt to produce a fake digital ID. The evaluator can directly verify the signature and hence the validity of the credential.[233] If

[233] Note that it is technically possible for users to attempt to enroll in the digital ID system with a fake physical ID. However, enrollment only needs to happen once, and digital ID providers are typically government entities that can check government records at enrollment time to determine whether the physical ID that the user presents was legitimately issued by the relevant governmental authority. Therefore, successfully using a fake ID to obtain a fake digital ID will be challenging for any user as long as the digital ID provider cross-checks government records at enrollment time.

the evaluator requires face matching via the presentation of a selfie or self-video, then the circumvention situation is similar to physical IDs, though without the need to account for errors introduced in the image capture process for the physical credential.

If the evaluator does not require face matching, then circumvention resistance depends on the security of the device. Because the private key is stored in the secure element, users should not be able to copy (clone) the credential onto another device without first breaking the secure element. However, the user might give or loan their device to someone outside the eligible age range. Defense against this attack depends on how the device/app enforces that only authorized users can use the device to authenticate. The general pattern for these systems is that the user authenticates to the device at enrollment time and then must re-authenticate when presenting their credential for age assurance. The EU Age Verification App uses a passcode for this purpose whereas Apple uses biometric authentication (TouchID or FaceID) with a fallback to a passcode in accessibility modes.[234]

If the device does not do biometric authentication, then it is straightforward for an eligible user (e.g., a parent or older friend) to enroll the device and then provide it to another ineligible user (the minor) on a temporary or permanent basis, just by providing the PIN. If the system does biometric authentication, then the eligible user must be present whenever age assurance is requested in order to provide their biometrics. In principle, if face-based biometrics (e.g., FaceID) are used, the device/app could attempt to match the user's biometrics to those found on the identity card, but this introduces a new source of user friction and potential false rejections if the user's appearance has changed. It is not clear if the existing systems do this. Moreover, it would be a significant expansion of the core purpose of these systems on mobile devices, which is to verify that the current user is the authorized user of the device, not to match to any real-world identity.

In general, this type of digital credential is unlikely to provide a strong defense against cooperation with an adult who is willing to share their credential. However, as sharing their credential might allow the minor to impersonate the adult in other settings as well, that may constitute some deterrent, as the adult might be willing to let a minor leverage their ID to demonstrate age but not to impersonate them in other cases where ID is required.

c. *Availability*

While mDLs are increasingly common, they are far from ubiquitous. For example, as of this writing, 14 US states support mDLs in Apple Wallet.[235] While some EU member states issue mDLs, the EU-wide mDL is planned to be introduced in 2030.[236] The EU Age Verification App can ingest an existing identity card or passport and exchange it for a digital credential, which may allow for wider availability. In general, any of the age assurance mechanisms described above can be translated into a digital

[234] Apple, "Add your driver's license to Apple Wallet."
[235] Apple, "ID in Wallet."
[236] European Parliament, "Modernising EU driving rules to increase road safety."

identity and used as described in this section, but this requires setting up a new issuer infrastructure.[237]

The anti-cloning mechanisms used by ISO/IEC 18013-5-style credentials depend on the device key being stored in a secure element. This means that users will need to have a device manufactured by one of a small number of vetted vendors such as Apple or vendors selling Google-certified Android devices. Some users may not have access to these devices or may prefer to use a device that does not have these security features;[238] these users will not be able to use digital IDs. Note that in the web case the user may not need to actually be using the secure device for browsing, as it is possible for the secure device to cooperate with an insecure device, as discussed above.

Digital credentials present a special availability concern in the form of “de-credentialing:” because digital credentials can be centrally managed (refreshed, revoked, etc.) it is possible for the credential issuer to disable a given user’s credential, thus potentially preventing them from accessing specific forms of content and experiences. The effectiveness of this form of attack depends on which experiences are subject to age assurance and the availability of other forms of age assurance; if age assurance requirements are widespread, then users might experience barriers to access, or at least be forced back into forms of age assurance with inferior privacy properties.

d. Privacy

If the digital credentials system does not include a selective disclosure feature, then the privacy properties of digital credentials systems are similar to those with physical IDs. If the digital credentials system includes a selective disclosure feature, as with ISO 18013-5 credentials, the user need not disclose all the details of their identity to the evaluator but can demonstrate only that they are in the eligible age range. ISO 18013-5 also supports “minimum age” attributes (e.g., “user is at least 18”) so that the user need not even disclose their birthdate to the evaluator, but only that they are within the eligible age range. However, there are still a number of privacy challenges, as discussed below.

Excessive Requests

Although selective disclosure systems allow users to disclose only the minimum necessary information, it is possible for evaluators to ask for additional information that they do not need,[239] especially if the same credentials and APIs are used not just for age assurance but also other remote identification functions, which might require information such as the user’s name and address. This risk can be mitigated by having the user approve which information is disclosed, or, in some cases, by having the user’s device or software restrict what can be requested, as Apple’s system does (see Section VII.C.2.e).

[237] See Bellovin, “Privacy-Preserving Age Verification—And Its Limitations.”

[238] For instance, because they want their device to be completely under their control rather than subject to the control of the manufacturer.

[239] Hancock and Collings, “Zero Knowledge Proofs Alone Are Not a Digital ID Solution to Protecting User Privacy.”

Linkage

A selective disclosure system improves privacy compared to a physical identity card-based system by preventing the evaluator from learning any information about the user other than the specific attributes that are disclosed. However, the credential is a unique user identifier that can be used to link together different user transactions. Consider the case where the user uses their credentials twice, once at an adult site to prove they are over 18, and once at the airport to prove that their name matches their boarding pass.

The table below shows the information disclosed in each scenario:

Scenario	Disclosed
Adult site	signature, age >= 18
Airport	signature, name

The credential is the same in both cases, even if different attributes are revealed, making it possible to link the two transactions.[240] Specifically, the airport and the adult site can collude to allow the adult site to learn the user's name even though the user did not disclose their name to the adult site. More generally, evaluators can collude to determine the union of all of a user's disclosed attributes for a single credential.

The privacy risk can be mitigated by having the issuer give the user multiple credentials with the same information so that the user can use a separate one for each transaction. This prevents evaluators from linking up individual transactions because the signed blocks will be different.

Even if the user is provided with multiple credentials and uses a different one per transaction, this does not prevent the evaluator from colluding with the digital ID issuer to track the user: because the issuer knows which credentials were issued to which users, the relying parties and the issuer can share information to determine which users engaged with which evaluators.

There are a number of plausible scenarios in which this could happen. In the context of age verification, the issuer is typically a government entity that issues ID cards. The government could obtain records from the evaluator (e.g., the AVP), thus learning which digital IDs were used to access age-restricted services. The government could then link those records to its own records that map the digital IDs to a specific person, thereby identifying which specific users accessed which age-restricted services.

[240] Specifically, the "mobile security object" (MSO) and signature are the same for each transaction with the same credential.

Enrollment

Enrollment of a physical credential onto a device requires disclosing the physical credential to the device and potentially to the device manufacturer.[241] The credential issuer only learns that you enrolled a physical credential, which will be a common behavior as many people will choose to have digital credentials. If age verification is required for each visit to age-restricted sites (as advised under the EU DSA Article 28 guidelines on the protection of minors, for example),[242] then the user will need to retrieve batches of credentials frequently, which indicates to the credential issuer that the user may be accessing age-restricted content or experiences. The privacy-sensitivity of this indication depends on which content or services are age-restricted.

If the user enrolls an age verification-specific credential, such as the one being built in the EU, then the situation is similar to that with a third-party AVP: the credential issuer learns whatever information the user supplies to demonstrate their age.

e. *Case Study: Apple Digital Certifications API*

Apple has proposed a generic system for remote verification using digital credentials that also permits age assurance[243] based on mDL-style digital credentials. This system is a fairly straightforward implementation of the selective disclosure mechanism described in Section VII.C.2, with the web interface provided by the W3C Digital Credentials API,[244] thus allowing the user to remotely authenticate to a website.

The overall workflow is shown in Figure 12 below:

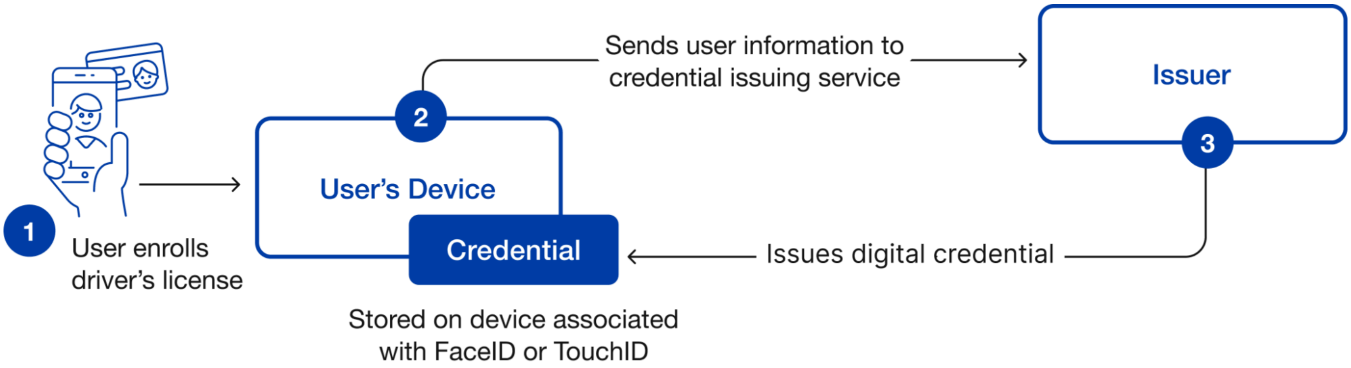
Sends user information to credential issuing service
Issuer
2
3
User's Device
1
User enrolls driver's license
Credential
Issues digital credential
Stored on device associated with FaceID or TouchID

Figure 12. Enrollment of a digital credential on an Apple device.

[241] Apple, "IDs in Apple Wallet."
[242] European Commission, "Guidelines on measures to ensure a high level of privacy, safety and security for minors online," 22.
[243] Apple, "Verify identity documents on the web."
[244] Caceres et al., "Digital Credentials."

The process starts with the user loading their mDL into the device.[245] They can load the mDL into either Apple Wallet or a third-party wallet, but this discussion focuses on the Apple Wallet case.[246] As part of this process, the user is asked to take views of their face from multiple angles in order to ensure that they are the person associated with the ID. This process only has to be done once.[247] The user also has to authenticate via FaceID or TouchID.

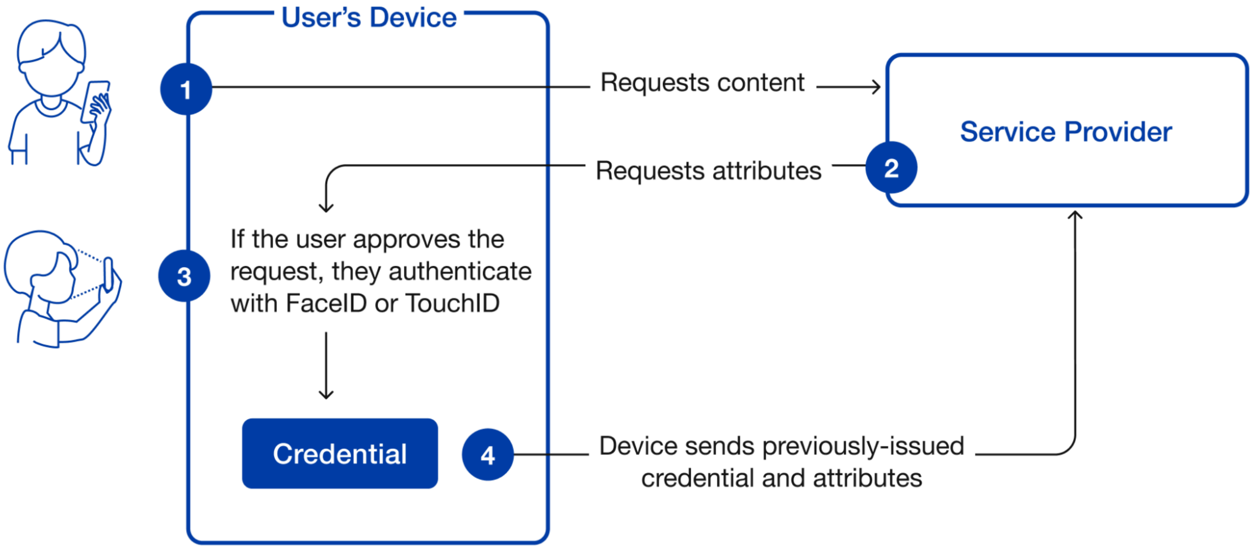
User's Device
1
Requests content
Service Provider
Requests attributes
2
3
If the user approves the request, they authenticate with FaceID or TouchID
Credential
4
Device sends previously-issued credential and attributes

Figure 13. Use of a digital credential for age assurance on an Apple device.

As shown in Figure 13, when the user later visits a website, that site can use the Digital Credentials API to request the desired attributes. The browser then queries the device for authentication. The device prompts the user about whether they want to reveal the requested attributes. When the user approves, they have to authenticate again in order to demonstrate that it's the same person who enrolled the device. Assuming the user consents, the device provides a verifiable response back to the browser. The browser provides the response back to the site, which then can verify the response and check the relevant attributes.

It is also possible to authenticate via this API on a desktop browser as long as the user's iPhone has a digital credential. This works by default in Safari but macOS provides APIs[248] that allow other browsers to do the same thing. Outside of macOS, Apple provides a QR-code-based mechanism for initiating the authentication on the phone, but does not seem to have published protocol specifications that would allow for the same seamless experience as on macOS.

[245] Apple, "Add your driver's license to Apple Wallet."
[246] See, e.g., Maryland Department of Transportation, "How to add your Maryland Mobile ID to Apple Wallet" for a video demonstration of the process.
[247] See Apple, "IDs in Apple Wallet" for some discussion of the privacy properties.
[248] Apple, "IdentityDocumentWebPresentmentController."

User Binding

Because the device requires the user to authenticate, this system provides a measure of binding to the user even if the site does not request the user's photo. Only the user who enrolled the mDL is able to use it to authenticate. Note that when using TouchID, this system does not require that the user who provided the fingerprint for TouchID be the same user who provided the mDL. In other words, person A could enroll their mDL on person B's mobile phone. It should be technically possible for the device to match FaceID against the mDL, but it is unclear whether Apple actually does this. Apple also appears to allow the user to bypass the biometric checks entirely if the user has accessibility features enabled.[249]

Privacy

Apple's system attempts to preserve privacy by retrieving batches of credentials, each with its own device key. The idea is that each credential in the batch is used only once, so that evaluators cannot link up multiple interactions when they see the same credential used multiple times[250] (see Section VII.C.2.d for a more comprehensive discussion). This does not prevent the credential issuer from linking up transactions, but such linkage would require the evaluator's cooperation (either voluntarily or legally compelled).

In addition, Apple requires evaluators to register with Apple Business Connect[251] and to obtain a signing certificate that is used to authenticate the evaluator's requests for remote user authentication. As part of this registration, the evaluator needs to document what attributes it will be requesting and why it needs them. This list is enforced at user authentication time, so that the evaluator cannot ask the user for extra attributes. This partly addresses concerns about sites asking for unnecessary attributes,[252] but at the cost of restricting the set of evaluators to those approved by Apple.

At present, it is unclear whether Apple applies restrictions to the types of evaluators it will accept as it has done with which apps can appear in the app store. For example, the Apple App Store guidelines explicitly forbid pornography and "apps that encourage consumption of tobacco and vape products, illegal drugs, or excessive amounts of alcohol."[253] It is possible Apple could opt to forbid evaluators from using evaluation to enable these types of activities, thus restricting the usability of digital credentials for these service providers.

3. Digital IDs with Zero-Knowledge Proofs

The privacy risks discussed above can be addressed by replacing the direct presentation of a digital ID with a zero-knowledge proof (ZKP). A ZKP is a cryptographic technique that allows one party (the prover) to convince another party (the verifier) that a statement is true, without revealing any other information beyond the fact that the statement is true. In this case, the user has a valid digital ID, but

[249] Apple, "Add your driver's license to Apple Wallet."

[250] The device may still reuse credentials if it runs out and cannot contact the credential issuer in time.

[251] Apple, "Your business. Open on Apple apps."

[252] Hancock and Collings, "Zero Knowledge Proofs Alone Are Not a Digital ID Solution to Protecting User Privacy."

[253] Apple, "App Review Guidelines."

instead of presenting it to the evaluator, they instead use a ZKP to prove that they have a digital ID with the eligible age range, with the user acting as the prover and the evaluator acting as the verifier. The advantage of this technique is that it removes the ability of evaluators and credential issuers to link multiple presentations of the same user's ID.

Google has deployed a version of a ZKP-based system on Android and it is in use on a limited scale.[254] Complete specifications are not available at the time of this writing. ZKP support has also been proposed for use with the EU Digital Identity Wallet and the EU Age Verification Solution.[255] See Appendix C for a more detailed discussion of the technical details of ZKPs.

The overall properties of ZKP-based systems are largely similar to those of ordinary digital ID systems, with a few exceptions as discussed below.

a. *Circumvention*

As with ordinary digital ID-based systems, the security of ZKP-based digital ID systems is tied to the security of the private key associated with the credential. In most cases, this key will be stored in a secure element on the user's device to prevent extraction and "cloning" of the credential. However, the consequences of extracting the private key are more severe with ZKP-based systems because the improved privacy properties make it harder to detect excessive use (e.g., many users performing age assurance with the same credential). "ZK-rate limiting" technologies are available to limit credential reuse, but are comparatively new and untested.

b. Privacy

As noted above, the privacy properties of ZKP-based systems are superior to normal digital ID systems because they do not permit linkage of multiple ID presentations. This prevents service providers from colluding to track users, as well as the credential issuer. Otherwise, the privacy properties are similar to those systems.

[254] Stapelberg, "It's now easier to prove age and identity with Google Wallet."
[255] European Commission, "EU Age Verification Solution."

4. Assessment Summary for Government IDs

	Baseline Accuracy	Circumvention	Availability	Privacy
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.

D. Facial Age Estimation

A number of age assurance systems are based on facial age estimation, in which 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.[256] 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. In practice, the results can be provided in one of three ways:

- As an estimate of the user's most likely age.
- As an estimate of whether the user is over or under a threshold age (potentially with a specified probability).
- As a distribution of estimates of the probability of the user being each age within a range.

1. Baseline Accuracy

Facial age estimation is inherently inexact, especially for ages close to the threshold. Even ignoring individual variation in appearance, people's appearances do not change significantly on their birthdays from the way they were the day before. This presents a problem because whether to allow a user to access an age-gated service or experience is a binary decision. Any service using facial age estimation for age assurance must therefore accept a certain minimum and nontrivial amount of error, either in the form of false rejects, false accepts, or both.[257]

The minimum error rate depends on the accuracy of estimation. There has been extensive research on the accuracy of facial age estimation technologies. The most comprehensive study, NIST's Face Analysis Technology Evaluation (FATE) Age Estimation & Verification[258] project (2025) measured the performance of 33 facial age estimation systems for still images. The results show high error rates, with the best-performing algorithms having a mean absolute error (MAE) value of 2.7 years on a set of

[256] Other physiological estimation techniques are also possible. Needemand's Borderage is a gesture-based technology which estimates age from hand movements. See Borderage, "AI technology based on medical research." This technique does not appear to be in wide use and there is little independent research available. The Australian Age Assurance Technology Trial found a false acceptance rate of 14.29% with seven 17-year-old subjects and an 18-year threshold, suggesting that a significant buffer age would be needed. The false rejection rate for users around 18 years old was not published, but as with facial age assurance, the need to use a buffer age suggests that this technique is not sufficient on its own. Age Assurance Technology Trial, "Needemand." Because gesture-based age assurance does not require an image of the user's face, it is likely to have superior privacy properties to facial age estimation. The circumvention and availability properties are likely to be roughly similar to facial age estimation. The company Privately provides voice-based age estimation technology, based on having the user read a sentence. As with gesture-based estimation, there is little independent research on voice-based age estimation, and Privately's reported results focus on settings with very large buffers (e.g., accurately detecting 13-14-year-olds as minors). The privacy properties of voice-based systems are likely to be better than facial age estimation because voice samples labeled with identities are less widely available. The circumvention and availability properties are likely to be roughly similar to facial age estimation. See Age Verification Providers Association, "Privately"; Yürüten, "VoiceAssure."

[257] Literature suggests that facial age estimation can be more error-prone for certain demographic groups, including women and the elderly, especially when these biases are not accounted for in the design and deployment of age assurance systems. For evidence from the empirical literature, see, e.g., Ganel et al., "Biases in human perception of facial age are present and more exaggerated in current AI technology"; Panić et al., "Addressing Demographic Bias in Age Estimation Models through Optimized Dataset Composition"; Puc et al., "Analysis of Race and Gender Bias in Deep Age Estimation Models."

[258] Hanacek, *Face Analysis Technology Evaluation (FATE) Age Estimation & Verification*, Figure 7.

images taken of visa applicants.[259] Using a higher quality image set captured on mobile phones, Yoti cites better accuracy with an overall MAE of 1.1[260] for 18-year-olds versus NIST's 2.63 result using lower quality images.[261] Australia's independent report on Yoti's system found an MAE of 1.0 for the same age cohort.[262]

Figure 14 below may be helpful in understanding the context. It shows the fraction of users of various ages who are estimated by Yoti as being over 18 for Yoti, as measured by the Australian Age Assurance Technology Trial:

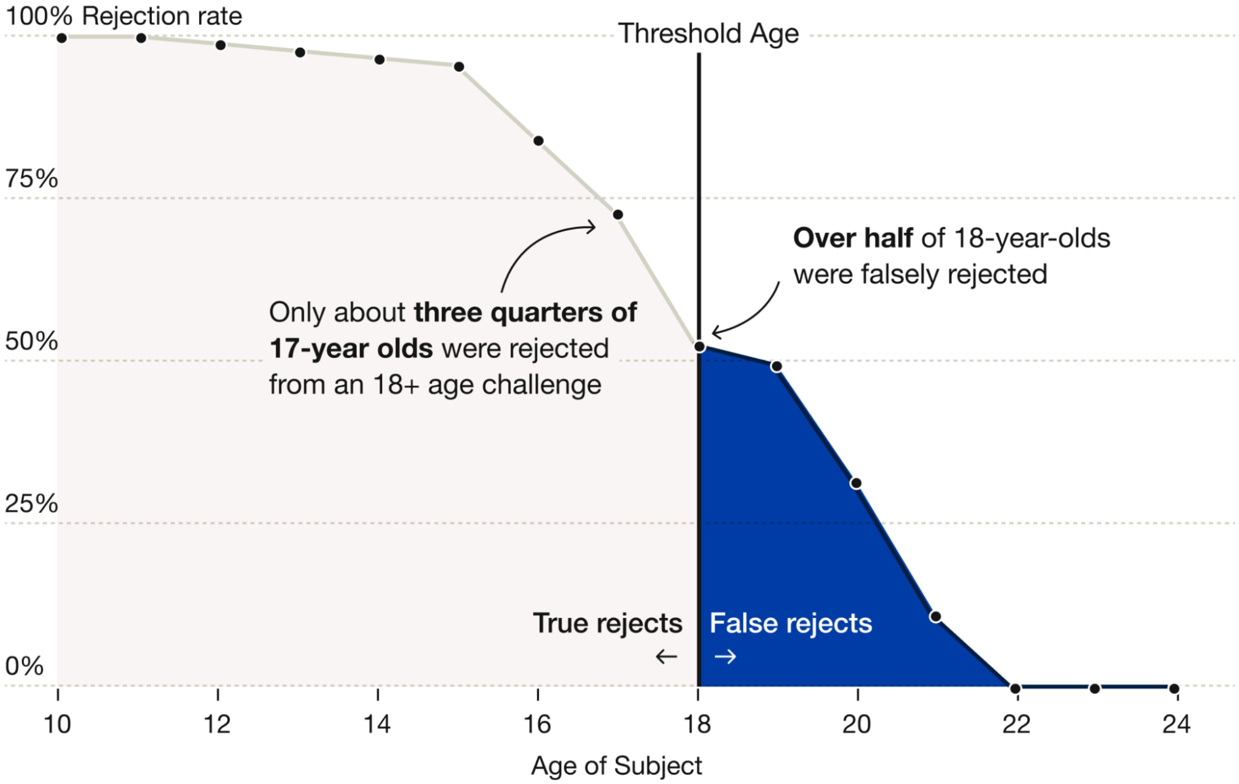

Figure 14. With an 18-year-old threshold, 25% of users aged 17 would be falsely accepted, and over 50% of users aged 18 would be falsely rejected. Using a higher threshold would decrease the false acceptance rate but increase the false rejection rate.

There are a number of options to deal with this inherent inaccuracy:

- Requiring that users' estimated ages be significantly above the threshold (e.g., 21 for an 18-year-old threshold), thus falsely rejecting a large number of people whose true ages are above the threshold for the sake of minimizing false accepts.

[259] This is an improvement from 2014, where the best algorithm had a MAE value of 4.27 years. Ibid.
[260] Age Check Certification Scheme, *Age Estimation Test Report*.
[261] Result of evaluation by the Age Check Certification Scheme. See Yoti, "Yoti Facial Age Estimation."
[262] Age Check Certification Scheme, *Age Estimation Test Report*.

- Allowing users whose ages are estimated to be at or above the specific age threshold, thus falsely rejecting many people who appear younger than their true ages and falsely accepting many people who appear older than their true ages.
- Allowing users whose ages are estimated to be significantly below the threshold, thus falsely accepting many people whose true ages are below the threshold for the sake of minimizing false rejects.

In cases where age assurance is intended to ensure that users are above a certain age, the first approach is the most common, requiring that users appear to have an age that is equal to the target age plus some buffer, such as 2 years (described by the Australian Age Assurance Technology Trial as "typical")[263] or 3-5 years (recommended by Yoti for "highly regulated sectors" such as adult content, gambling, alcohol, and tobacco).[264] This means that *by design* any user in the buffer range will be rejected, even if they appear older than the target age.

As a result of these issues, where facial age estimation is in use, it needs to be used with a "buffer" where the system attempts to determine whether the user is well outside the eligible age range (e.g., 25 for a target of 18) and if not, prompts them to demonstrate their age using a second, more accurate mechanism.

Facial age estimation systems also show performance variation in error rate for different demographics, with accuracy being higher for lighter skinned users. In NIST's results, error rates were generally-–but not always—higher for women than men.[265] Users with visible facial differences may also not be able to achieve accurate age estimation. While little data is available on facial differences and age estimation, this is a known issue with face recognition systems.[266] If a user is from a demographic group where there is a higher than normal error, then the system will be less available, and will have to fall back to some other method of assurance.

It is potentially possible that video-based systems would perform better. No comprehensive studies of the accuracy of video-based facial age estimation are available.

2. Circumvention

The most obvious form of circumvention for age estimation is for the user to just try again, potentially with a different hairstyle, glasses, etc. Because the repeatability of facial age estimation is fairly low, this is likely to be successful with subjects whose apparent age—at least according to the age estimation system—is close to the threshold.

[263] Ibid., 11.
[264] Yoti, "Yoti Facial Age Estimation," 8.
[265] Hanacek, *Face Analysis Technology Evaluation (FATE) Age Estimation & Verification*, Table 6.
[266] Facial Equality International, "Facial Recognition and the Facial Difference Community."

As an example, Figure 15 below shows the variation in estimated age for each frame of a 60-second video clip of a single subject, age 58:

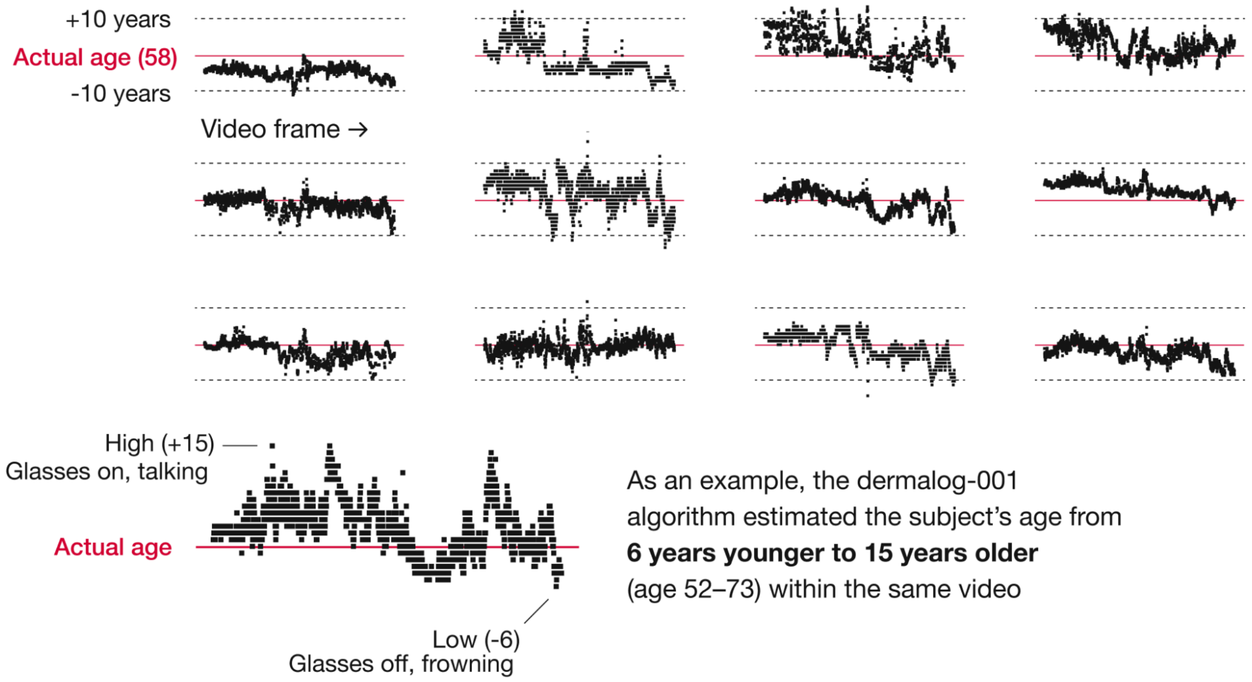

Figure 15. Each graph shows age estimates for one of 13 different algorithms applied to 1688 frames of Patrick Grother, a 58-year-old man, extracted from a 60-second video captured using an Android phone.[267]

It seems likely that many users on the margin of an eligible age range would be able to be successful with multiple attempts.

The obvious way to defend against this form of circumvention is for the evaluator to keep track of which users have attempted to demonstrate their age. However, this has privacy implications because it requires the evaluator to retain records of users who have *failed* age assurance. This is particularly problematic because users who have failed age assurance are more likely to be children.

As with other video and selfie-based systems, facial age estimation systems are vulnerable to both presentation and injection attacks, as described in Section VII.C.1.b While there is less specific research on the vulnerability of age estimation systems to these kinds of attacks than on remote identity verification systems, there has been at least one high-profile example of such an attack on the age verification system used by Discord, which is an interactive system in which the user has to provide stills in various facial expressions in order to demonstrate "liveness."[268] The attacker was able to fool it with a presentation attack using a video game character displayed on a device and held in front of the camera. The character could be told to adopt specific facial expressions to simulate

[267] Hanacek, *Face Analysis Technology Evaluation (FATE) Age Estimation & Verification*.
[268] Ridley, "Brits can get around Discord's age verification thanks to Death Stranding's photo mode, bypassing the measure introduced with the UK's Online Safety Act. We tried it and it works—thanks, Kojima."

liveness. This is an unsophisticated attack that anyone could launch using readily available tools on the internet. This should raise concerns about how well these systems have been designed to resist circumvention in practice.

3. Availability

The minimum technical requirement for a facial age estimation system is that the user's device has a camera. However, without a technical mechanism that prevents the user injecting their own video, age estimation is very susceptible to injection attacks (see Appendix B.2. If input source integrity is required, then use of facial age estimation will be restricted to users with verifiable hardware, which practically speaking means iOS devices and Android devices which are compatible with Google Play Integrity.[269]

Some users who are technically able to participate in age assurance may be unwilling or unable to provide the AVP with a facial image, for instance for privacy or religious reasons.

4. Privacy

Although the evaluator does not learn the user's name by virtue of using facial age estimation, they capture an image of the user, which presents a number of privacy threats. First, the evaluator may be able to use commercially available facial recognition systems to learn the user's identity from their photo.[270] As discussed above, this is a privacy threat if the user is attempting to remain anonymous but less so if they are accessing a service where they have a lower privacy expectation, such as a social media service. In the former case, the picture might also be used as evidence that the user accessed sensitive content; this becomes a more severe risk if the user has to provide their name for some other reason, for instance, if they fail age assurance and the user has to fall back to banking records.

As with other photo-based systems, if the user provides their photo, it can later be used for other purposes, for instance AI-based deepfake or "nudification" systems (see Section VII.C.1.d).

5. Assessment Summary for Facial Age Estimation

	Baseline Accuracy	Circumvention	Availability	Privacy
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.

[269] Google, "Play integrity and signing services."
[270] There are commercially available systems such as Vigilant Solutions, PimEyes, FaceCheck.id, and ProFaceFinder that offer this service.

E. 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. Many companies have publicly disclosed information about their use of behavioral signals for the purpose of age estimation:

- **Character.ai**'s age assurance technology examines a number of signals, including "login info, activity on the platform, and some signals from third parties," to estimate whether a user is a minor.[271] When this is the case, users must then verify their ages to access mature content and features.[272]
- **Google** assesses the age of users across its services through a "variety of signals already associated with a user's account" like search history on Google and YouTube watch history.[273] For accounts identified as likely belonging to minors, Google disables personalized advertising,[274] activates watch time reminders and restricts some types of content recommendation on YouTube,[275] and imposes content filters on Gemini, among other measures.[276]
- **Meta** estimates the age of Facebook and Instagram users through an "adult classifier" that is trained on signals such as "profile information" (e.g., when the account was created) and "interactions with content."[277] Meta uses this age estimation to proactively classify accounts as "Teen Accounts"[278] on Facebook and Instagram, which provide enhanced privacy settings by default, limit users' ability to see mature content, and cannot be changed without parental consent.[279]
- **OpenAI** predicts the age of ChatGPT users through a "combination of behavioral and account-level signals," such as time since signup, long-term usage patterns, and stated age.[280] When the age prediction model infers that a user is a minor, their account is defaulted to an experience designed to limit access to sensitive content, such as graphic violence and sexual roleplay.[281]

[271] Character.ai, "Age Assurance."
[272] Ibid.
[273] Beser, "Extending Our Built-In Protections to More Teens on YouTube"; Brooks, "Ensuring a safer online experience for U.S. kids and teens."
[274] Ibid.
[275] YouTube, "Building content recommendations to meet the unique needs of teens and pre-teens."
[276] Google, "Guide your child's Gemini Apps experience."
[277] Finkle et al., "How Meta uses AI to better understand people's ages on our platforms"; Meta, "Introducing New Ways to Verify Age on Instagram."
[278] Meta, "Working With Parents and New Technology to Enroll More Teens Into Teen Accounts."
[279] Facebook, "How Teen Accounts work on Facebook"; Instagram, "Instagram Teen Accounts Will Be Inspired by Movie Ratings for Ages 13+"; Instagram, "Introducing Instagram Teen Accounts"; Meta, "We're Introducing New Built-In Restrictions for Instagram Teen Accounts, and Expanding to Facebook and Messenger."
[280] OpenAI, "Our approach to age predictions."
[281] Ibid.; OpenAI, "Updating our Model Spec with teen protections."

- **Reddit** uses various signals, including users' posts, comments, and subscriptions, to estimate the age of users.[282] When a user is estimated to be under the minimum age as required by local law or Reddit's terms of service, they are required to verify their age through a third-party provider.[283]
- **Roblox** says that it is "constantly evaluating user behavior to determine if someone is significantly older or younger than expected."[284] For Roblox users whose behavioral signals deviate from their stated age, Roblox plans to require them to re-verify their ages. Age verification on Roblox is required to access the chat feature, and is used to limit the age range of users allowed to chat with one another.[285]
- **TikTok** employs age estimation using a "range of signals" (e.g., profile images and references to upcoming birthdays) to ensure that a user's stated age is correct.[286] TikTok reportedly plans to implement this process in more regions globally.[287]

Thus far, age estimation based on behavioral signals has primarily been deployed by platforms with large existing user bases where the dominant mode of engaging with the service is for a user to obtain an account or profile and retain that account or profile for consistent use over time.[288] This mode is what makes behavioral age estimation possible, because the platforms' features and value proposition incentivize or encourage users to maintain a persistent identity on the platform, in order to obtain personalized recommendations through sustained engagement, build a following, share their real identity, or save information to their profile or account.

1. Baseline Accuracy

There is little independent research on the accuracy of this type of behavioral age estimation. To obtain reliable accuracy information, platforms would need to share with evaluators the data and algorithms used to estimate ages, or conduct their own accuracy testing and subject it to independent auditing.

However, there is an inherent source of inaccuracy in that the user must establish a baseline amount of activity on the site in order for the system to form an estimate of their age. As a consequence, the user will have the default level of access for some period of time, which means minors would default into experiences designed for adults or the general population, or adults would default into experiences designed for minors. YouTube defaults all users who self-declare as 18 or over into its unrestricted

[282] Reddit, "Why is Reddit asking for my age?"
[283] Ibid.
[284] Roblox, "Roblox Requires Users Worldwide to Age-Check to Access Chat."
[285] Ibid.
[286] TikTok, "An update on our work to provide teens with age appropriate experiences."
[287] Ibid.
[288] See, e.g., Beser, "Extending Our Built-In Protections to More Teens on YouTube"; Brooks, "Ensuring a safer online experience for U.S. kids and teens"; Character.ai, "Age Assurance"; Finkle et al., "How Meta uses AI to better understand people's ages on our platforms"; OpenAI, "Our approach to age predictions"; Reddit, "Why is Reddit asking for my age?"; Roblox, "Roblox Requires Users Worldwide to Age-Check to Access Chat"; TikTok, "An update on our work to provide teens with age appropriate experiences."

experience, and later moves users into teen experiences if it gathers behavioral signals that cause it to infer that the user is a teen.

2. Circumvention

No information is publicly available about the circumvention of age estimation based on behavioral signals.

For existing users on a platform with extensive histories of engagement and a desire to consume content or otherwise engage in ways that may reveal their age cohort, these systems may be difficult to evade. However, for users willing to open new accounts that have no baseline behavioral signals associated with them, or for services where account creation is not required, circumvention is more straightforward. Opening a new account, using private browsing, or using a VPN can prevent the service from building a baseline of the user's behavior.

3. Availability

This age assurance mechanism is only available to services that collect sufficient behavioral data to provide indications of age. On many services, user behavior will not necessarily correlate with or provide any indication of age. Where this form of age assurance is available to services, the availability to users is essentially the same as the availability of the service itself. Any user who can access the service may be subject to behavior-based age assurance.

4. Privacy

The large-scale services that have announced the use of behavioral signals for age estimation are relying on behavioral data for age assurance that they already collect for other purposes. However, users may not expect that data which was gathered for one reason (to provide the service) will be used for another (age estimation).

If services begin to require additional data collection specifically for the purposes of age assurance, this could introduce additional privacy risks. For example, if services begin requiring users to create accounts in order to access restricted experiences, this can impact privacy because the account information may identify the user directly and the user will be unable to access age-restricted experiences in a way that prevents the site from building a profile of their activity.

Because this age assurance mechanism relies on observations of user behavior over time, reliance on this mechanism could serve as a deterrent to deleting behavioral data. At present there is insufficient public information about which behavioral data is being used for age assurance and how long it is stored to understand whether this mechanism is leading to richer profiles of users being stored for longer. Behavioral signals may also implicate data protection laws like the EU General Data Protection Regulation (GDPR).[289] In September 2025, the European Data Protection Board adopted draft guidelines on the interplay between the DSA and GDPR, which reference the ways in which age

[289] European Union, "General Data Protection Regulation."

assurance goals of the DSA may qualify as a legal basis for processing of personal data under the GDPR.[290]

5. Assessment Summary for Behavioral Signals

	Baseline Accuracy	Circumvention	Availability	Privacy
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. Key Findings

While the landscape of age assurance mechanisms is very complicated, it is possible to draw some general conclusions about age assurance as a whole.

A. 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.

User expectations around privacy vary based on the use case, depending on both the existing privacy properties of the user's relationship with the service provider and the level of sensitivity of the content or experience. If the user currently accesses a site anonymously or pseudonymously, especially with privacy-enhancing technologies such as VPNs, then their concerns over the privacy implications of age assurance are likely to be higher than if they have created an account and/or provided potentially sensitive information such as their picture or messages. If the user wishes to conceal that they are seeking a given type of content or experience—for instance, pornography—then privacy concerns are likely to be heightened compared with content and experiences which are less sensitive, such as nighttime notifications or algorithmic feeds.

The user's desire to circumvent age assurance may also vary depending on the use case, and the extent to which lack of access to age-restricted content and experiences adversely affects the user's experience. Individual users will make different judgements about what constitutes an adverse impact, but in general the more desirable an age-restricted experience is, the more motivation users will have

[290] European Data Protection Board, "Interplay between the DSA and the GDPR."

to circumvent age assurance. Additionally, a minor may find it more or less practical to persuade an adult to assist them with circumvention based on the type of content or experience being sought.

B. 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. None of the existing signals are both sufficiently accurate and available to avoid excluding large numbers of eligible users when used alone. For example:

- While facial age estimation is technically available to nearly all users, because it is not able to accurately distinguish users who are near (but below) the age threshold and users who are near (but above) the age threshold, it will necessarily either include a large number of ineligible users or exclude a large number of eligible users.

- Age verification based on government IDs (whether physical or digital) is highly accurate but in many jurisdictions, some users will not have access to any form of ID or be unwilling to disclose it to access a given service or experience. Thus, a system which requires a government ID will exclude many eligible users.

Any system which does not accept multiple signals will reject a large number of users. As a consequence, common practice is to offer users multiple signals. This can take the form of offering the user a choice of which signal to use or of starting with one signal but falling back to others, as with systems which try to estimate the user's age but will require the user to provide proof of their age if the system is not able to make a determination with high confidence.

In most cases, individual age signals will not provide a conclusive determination that the user is outside of the eligible range: rather they will be unable to confirm that the user is within the range. For example, an age estimation system might estimate that a user is 17 but this estimate has low certainty for the same reason it is unsafe to assume that a user who appears 18 actually is not an old-looking 17-year-old. As a result, if a user fails to demonstrate their age with one signal, they usually need to be invited to try another, rather than just being excluded. Because the privacy properties of these 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.

Because different users will end up using different age signals, evaluating the availability of an age assurance system requires asking what fraction of users will be willing and able to use at least one of the provided age signals. Note that this is not a straightforward matter of multiplying the false reject rates of each signal, as inability to use signal A is not necessarily independent of inability to use signal B. For example, a user who conceals their face for religious reasons may not be able to use either facial age estimation or a government ID-based system. Similarly, because users who wish to circumvent age assurance can keep trying different signals until they succeed, in order to determine the overall false acceptance rate of an age assurance system, it is necessary to determine the fraction of ineligible users who will be correctly rejected by all available signals collectively.

C. Age Ranges

The most common age range in use for age assurance is 18 or over, corresponding to the common age of majority. However, there is also interest in other ranges, which present additional challenges.

1. Other than 18

A number of contexts use age ranges other than 18+. In general, an age threshold *over* 18 (e.g., the 21-year-old threshold to buy alcohol in the US) is comparatively easy to manage using the same techniques as an 18-year-old threshold. However, thresholds below 18 are more challenging. Estimation-based techniques are not able to accurately classify all users with a sufficiently low error margin to be used exclusively, thus necessitating a fallback to identity-based age verification techniques. However, many jurisdictions do not routinely issue age documents to people under 18, in which case it is not practical to accurately determine the age for these users. This also makes it difficult to make fine-grained distinctions about which content is appropriate at each age, as the available technologies do not support age assurance at this level of granularity.

In many cases, the best technique available for age thresholds under 18 will be parental attestation.This is susceptible to circumvention assisted by parents, or, depending on how well parental relationships are verified, assisted by other adults. Research in Australia found that 34% of parents/carers indicated willingness to help children evade Australia's social media minimum age of 16.[291]

2. Maximum Age

In some contexts, there is a desire to assure a *maximum* age rather than a minimum age. For example, one might want to have a space dedicated only to children, as is done by the social media services Yubo and Promly.[292] For the same reasons that it is difficult to enforce over 18 thresholds, it is difficult to precisely enforce maximum ages under 18. Age estimation techniques can be used to exclude older adults, who will not be able to pass a check designed for under 18 even if it is set with a conservative threshold designed to avoid excluding actual children (favoring false accepts over false rejects). For

[291] Australian Government Department of Infrastructure, Transport, Regional Development, Communications, Sport and the Arts, *A summary report on developmental research to inform a Social Media Minimum Age campaign*.

[292] See Promly, "Promly"; Yubo Team, "How Yubo Pioneered 100% Age Verification to Set a New Standard for Trust & Safety on Social Media."

example, if the age threshold is set to 21 (to try to allow anyone under 18 in), then most 30-year-olds will still not be able to pass. However, this would still have a high false accept rate.

An additional challenge with maximum ages is that unlike minimum ages, users can be eligible at one time and then ineligible at some later time. If users are required to frequently re-perform age assurance, this is not an issue, but if the result of age assurance is stored, for example to reduce friction, then it is necessary to store not just a binary yes/no but also the user's age or estimated age, in order to know when a previously eligible user becomes ineligible.

D. Privacy Protection

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 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.

The privacy properties of age assurance systems vary widely, both in terms of the amount of information about the user that needs to be shared in order to evaluate a user's eligibility and the parties with whom that information is shared. Most age assurance systems share enough information with the evaluator to allow the evaluator to determine the user's actual identity. This includes:

- The user's actual name (as on a government ID);
- Another identifier such as an email address, phone number, or social security number; or
- The user's face.

The first two of these can be mapped directly to the user's identity. The user's face can be used with facial recognition systems to identify the user.

The vast majority of age assurance systems currently in use have the server perform the evaluation of user eligibility, whether by the service provider or by a separate AVP. In this case, the server usually learns (1) the user's identity and (2) the service provider that they are trying to access. The user's privacy exposure is entirely dependent on administrative and policy controls; in principle the server can retain the user's identity and/or distribute or sell it. Depending on the jurisdiction, there may be laws or regulations restricting the use of this information, but the user is unable to determine that the server is complying with these or with its own privacy policy. It is possible to use this type of system as a surveillance mechanism by recording the personal information of users who request age assurance. This can happen prospectively by requiring evaluators to disclose that information to authorities or retrospectively if the evaluator already keeps such logs.

In systems where eligibility is evaluated on the device—whether directly via some form of digital ID or via the manufacturer/vendor making the determination-–then it is possible to technically restrict the amount of information that is revealed about the user. When enforcement also happens on the device, privacy can be provided by limiting the information apps can request, e.g., by only permitting them to ask whether a user is in the appropriate range and limiting which ranges they can request.

When evaluation happens on a server, the most private option is zero-knowledge proof systems which only reveal that the user meets the specified age criteria without revealing anything else about the user. When properly designed, these systems do not permit linkage between multiple demonstrations of eligibility, even with the collusion of the authority that issued the credential.

When a new credential is used for each age assurance transaction, selective disclosure systems prevent evaluators from linking a user's activities together on their own, although linkage is still possible with the cooperation of the credential issuer.

Even in cases where the evaluator does not learn the user's identity, the initial act of obtaining the credential can be privacy risk: if the primary reason for digital ids to access sensitive content such as pornography, as with age assurance-specific credentials, then when the user obtains a digital ID, the credential issuer can infer that they are interested in sensitive content. If there are other innocuous reasons to establish one's age or identity-–e.g., to use a mobile driver's license at the airport—then the privacy issue is lessened.

E. Circumvention

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.

The importance and prevalence of circumvention may differ depending on the specific use case. In blocking use cases where the intent is to prevent minors from accessing certain content or experiences entirely, minors may be more motivated to circumvent age assurance if it prevents them from accessing content or experiences that they want. In safer defaults cases, minors may have less incentive to circumvent age assurance if the defaults do not adversely affect their experience of the service. In some safer defaults proposals, minors are allowed to access age-restricted experiences with parental consent,[293] in which case parental assistance in circumvention would not be considered problematic.

This section examines the major methods of circumvention.

[293] New York, "Stop Addictive Feeds Exploitation (SAFE) for Kids Act."

1. VPNs

For web users, if age restrictions are evaluated on the server and there is significant variation in the requirements for age assurance, then it is straightforward for users to use a VPN to appear to be in a jurisdiction which does not require age assurance. There is already ample evidence that users respond to age assurance requirements by using VPNs and VPN providers are not themselves required to perform age assurance. Some jurisdictions[294] have considered restricting use of VPNs or requiring VPN providers to employ age assurance for their users.[295] In jurisdictions which use technical restrictions to block VPNs, this has led to an arms race[296] between VPN users and network-based restrictions, with the result that some technically sophisticated users can evade blocking.[297]

Mobile apps may be able use system APIs to determine the correct jurisdiction without being fooled by VPNs. These APIs may not be available in all cases or may require user permission. If apps use precise geolocation APIs to determine jurisdiction, this increases privacy risk to users, as their location may be sensitive, for instance if they are at home.

2. Credential Reuse

Any age assurance mechanism which does not involve directly authenticating the user at the time of visit is subject to attacks in which an ineligible user makes use of an eligible user's identity to establish their age. This can be done either with or without the eligible user's intentional assistance. Examples of such assistance include:

- Establishing an account and sharing the password.
- Enrolling a device as their own and then sharing it with the ineligible user.
- Sharing their external credentials (credit card number, SSN, etc.).
- Purchasing an unlocked device if enforcement is on devices.

Many parents will be skeptical of or hostile towards age requirements. For example, it is quite common for parents to allow their children to have social media accounts even if the child does not meet the service provider's minimum age to have an account: around 1/3 of Australian parents expect to assist their children in bypassing the Social Media Minimum Age. Similarly, just as it is not uncommon for over-21 students to purchase alcohol for under-21 students in the US, it is likely that many under-18s will know 18-and-overs who can assist them in circumventing age restrictions. The broader age assurance mandates are, the easier it is for children to find legitimate-appearing reasons for adults to help them circumvent those restrictions. For example, if enforcement happens at the device level, children may be able to persuade their parents to unlock the device in order to allow them to access social media, with the side effect of allowing them to access adult content.

[294] See, e.g., Parliament of the United Kingdom, "Children's Wellbeing and Schools Bill."

[295] Note that VPNs are widely used to connect remote workers to enterprise networks. This type of VPN is distinguishable from privacy-oriented consumer VPNs at a regulatory level but not at a technical level in the network.

[296] Wu et al., "How the Great Firewall of China Detects and Blocks Fully Encrypted Traffic."

[297] Another approach is to require service providers to perform age assurance for any user who appears to be behind a VPN. See Tutor, *Age Verification*. However, this has the disadvantage that the service provider does not know which age threshold to apply.

Even without intentional assistance, it is often possible to take advantage of another person's eligibility, such as if a parent leaves their device unlocked, their email open, or a child is able to obtain their parent's credit card number or social security number. The types of credentials used for age assurance are often not considered sensitive within families, even when compared to login credentials.

Requiring age assurance at every interaction makes this kind of credential reuse more difficult, because the ineligible user will need long-term access to the relevant credential. However, such a requirement also increases the impact on eligible users.

3. Presentation and Injection Attacks

All of the mechanisms based on remote user biometrics, whether facial age estimation or selfie/ID matching are vulnerable to presentation attacks and injection attacks. It is likely that injection attacks will continue to improve, with the primary defense being some form of remote mechanism for determining that input is not coming from an untrusted source. App integrity mechanisms are already available with mobile apps but not at all on the web and are unlikely to be available in the foreseeable future.

4. Device-Based Enforcement

Device-based enforcement mechanisms are harder to circumvent when the device is trusted but depend on the evaluator trusting the device. For example, VPNs can be used to circumvent server-based enforcement on the web, as described above, but do not impact device-based enforcement. However, if the user is able to obtain an untrusted device, they may be able to modify it to circumvent device-based enforcement mechanisms. In order to provide effective circumvention resistance against moderately sophisticated attackers, it would be necessary not only to restrict the behavior of existing applications but also to restrict the use of these devices for software development,[298] which would be highly disruptive.

Some proposals for device-based enforcement restrict only mobile devices, whereas others cover both mobile and desktop/laptop devices. Requiring device-based enforcement for desktop devices would be much more disruptive than for mobile only, both because desktop software is currently subject to minimal policy enforcement on both Windows and macOS and because the use of alternative operating systems such as Linux is very common in both software engineering and services environments. Requiring age enforcement before purchasing Linux-compatible hardware or being able to install Linux would be a major change for many enterprises, as well as for ordinary Linux users.

If only mobile devices are restricted, then device-based enforcement will be more suitable in settings where untrusted devices are harder for users to obtain (e.g., for younger children) or where they are not close substitutes for trusted devices (e.g., settings where desktop/laptop devices are less useful,

[298] If software development is not restricted then users can simply download and build any open source web browser, thus bypassing device-based restrictions.

such as social media), and less suitable for cases where the user is capable, motivated, and where desktop devices are practical (e.g., adult content), though they still present some barrier in those cases.

5. Open vs. Closed Systems

All of the age assurance mechanisms described in this report are easier to circumvent on an open device in which the user can install software of their choice. Some mechanisms, such as facial age estimation, are more vulnerable to attacks–in this case, injection attacks–when the user has an open device.

It is important to distinguish between systems which rely on adults being able to obtain closed devices in order to demonstrate their age and systems which rely on minors not being able to get open devices. For example, deployed digital ID-based systems require that the user have a trusted device in order to demonstrate their age,[299] but do not preclude them otherwise using an open device. Some of these systems allow the user to use an open device in connection with a closed device. Adults who do not have access to a closed device may not be able to establish their age using this method, but are otherwise able to do as they please. By contrast, with device-based age assurance, a minor who can obtain an unrestricted device can bypass age assurance entirely. As a result, effective device-based enforcement requires making access to such devices prohibitively difficult for most minors, with collateral effects on adults who want open devices.

6. Limitations of Circumvention Resistance

An important policy challenge is balancing the costs of anti-circumvention measures against the incremental reduction in the amount of circumvention. In particular, many technologies which enable circumvention have high levels of legitimate use. For example:

- VPNs are widely used as a privacy measure to prevent both the local network and remote systems from monitoring user behavior.
- Open operating systems such as Linux are widely used for software development and dominate the server market.
- Many users and families share devices and accounts for both convenience and economic reasons.

While it may be possible to somewhat reduce circumvention by restricting some of these legitimate uses, such restrictions will not entirely eliminate circumvention, nor is it possible to do so without severely restricting both the internet and consumer computing devices. Even in China, which tightly restricts internet usage, research has found that attempts to reduce youth internet gaming usage have not successfully done so,[300] and children frequently borrow their parents' IDs to evade age

[299] It is possible to have a system which does not require a trusted device, with somewhat weaker anti-circumvention properties, as described in Appendix C.

[300] Zendle et al., "No evidence that Chinese playtime mandates reduced heavy gaming in one segment of the video games industry."

assurance.[301] This suggests that it is necessary to be realistic about the level of achievable circumvention resistance rather than striving for perfection.

IX. Broader Impacts

Widespread age assurance will have effects extending beyond those on minors who wish to access age-restricted content. This section examines those impacts.

A. Deterrent Effects on Adults

In addition to making it difficult for minors to access age-restricted material, one of the main consequences of enforcing age restrictions is to make it more difficult for adults to access age-restricted content and services.[302]

Increasing the difficulty of accessing age-restricted content and experiences is likely to have a deterrent effect, even for eligible users. For example, polling in the UK shows low acceptance rates for age assurance, with less than 15% of adults saying they would be likely to be willing to show “any proof of age (e.g. a photo/ video, photographic ID, using banking information, digital ID wallets etc)” to access pornography sites.[303] Willingness to participate in age assurance varies by age signal, with email being the most acceptable (56%) and banking information and credit cards being the least acceptable (18% and 17% respectively).

For users who are willing and able to establish their age, the age assurance process itself is an additional burden on the user, who must perform some additional action in order to complete the age assurance process. In the best-case scenario, this is just a matter of clicking through some set of consent screens (as with the Apple Digital Credentials flow). But in many cases, it is significantly more cumbersome, requiring the user to turn on their camera, show ID, or go through an additional authentication process.

When device-side enforcement is used, the user likely only needs to complete the age assurance process once because the device can remember the user’s status. With server-side enforcement, the user will most likely need to complete it repeatedly, potentially every time they want to access age-restricted content. Digital credential type mechanisms—whether of the selective-disclosure or zero-knowledge-proof variety—in principle can be set up to automatically establish the user’s age on each access. However, none of the current implementations do so, and they instead require the user to re-authenticate with either a biometric or a PIN, adding friction.

For most account-based services, such as social media services, the user can be expected to establish their age when they create their account and rarely thereafter. However, for

[301] Yimeng, “Parents help children dodge time limits on online games.”
[302] For some advocates of age assurance this is an intended rather than unintended consequence.
[303] Pedley, “Britons back Online Safety Act’s age checks, but are sceptical of effectiveness and unwilling to share ID.”

non-account-based services or services where the user chooses not to make an account (e.g., for privacy reasons) the situation is more complicated.

To reduce friction in web browsing use cases, either the service provider or the age verification provider (if separate) can remember the user's age status by storing a cookie on the browser. This has privacy implications for the user, as cookies can be used to allow the site to track the user across visits and to link it to the user's identity. Many users already browse adult sites in private browsing modes which do not store state, so it seems likely they will be unwilling to create an account or even allow persistent tracking.[304] Passkey-based mechanisms like AgeKey attempt to bridge this gap and reduce friction somewhat, but it is too soon to tell whether they will be acceptable to users.

It is difficult to estimate precisely the size of the impact of age assurance on legitimate adult access to age-restricted content or experiences. Adult sites that have introduced age restrictions report low rates of completion of the age assurance process and significantly reduced traffic volumes, with less than 10% of users completing age assurance.[305] Aylo reported an 80% reduction in traffic from Louisiana after the introduction of required age assurance.[306] It is highly unlikely that this reduction was solely attributable to excluding minors. Moreover, jurisdictions which have introduced age restrictions have seen large increases in interest in and use of VPNs. However, it is possible that if age assurance requirements become more commonplace, more users will be willing to go through the age assurance process.

Finally, once requirements for age assurance have been set for one category of content and services, it is easier to expand requirements to other categories of content, because it does not require adding new technical mechanisms, but only adding new categories of content which need to require age assurance.

B. Parental Consent

In some cases, age assurance requirements are paired with parental consent provisions. For example the New York SAFE for Kids Act requires social media services to restrict certain features for users under 18 unless there is explicit parental consent for the user of these features.[307]

Unfortunately, it is challenging to demonstrate that a specific individual is the parent or carer for a specific minor. In some cases, it will be possible to establish 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. Creating an accurate parental verification system for use by services generally available to the public over the internet implies the need to collect identification

[304] The EU Guidelines recommend that "adult-restricted online platforms should not allow sharing of user account credentials and thus conduct age assurance at each instance when their service is accessed." European Commission, "Guidelines on measures to ensure a high level of privacy, safety and security for minors online," 26.
[305] Free Speech Coalition, "Appendix."
[306] Iovine, "Do age-verification laws work?"
[307] New York, "Stop Addictive Feeds Exploitation (SAFE) for Kids Act."

information from both the adult and the minor, plus possible additional information to prove the parental relationship. All of the privacy considerations discussed throughout this report related to identity proofing and government records would apply to both individuals. In the absence of a system to verify parental relationships, the pathways for circumvention discussed in this report that rely on a minor collaborating with any adult would be available to circumvent parental consent.

C. Market Concentration

Because the process of age assurance often involves friction, service providers and AVPs have an incentive to reduce the need for repeated age assurance (see Section VI.A.5). Many of the mechanisms for this involve having the AVP remember that the user has demonstrated their age. This gives service providers an incentive to select a widely used AVP because this increases the chance that the user will have already demonstrated their age to that AVP.

As a result, there is pressure towards market concentration on a small number of AVPs. There is already some evidence of concentration in the AVP market, with recent research finding that the top five providers covered almost 75% of the market in Texas and Georgia, with Yoti alone representing more than half the sites.[308] In addition to typical economic concerns about market concentration, as discussed in Section V.A.2.b, concentration in the AVP market increases the risk to user privacy in case of disclosure or breach of an AVP's records.

D. Implementation Quality

This report has largely focused on the "best-case" properties of the systems under examination, which is to say the properties if they are implemented correctly. However, it is well known in the software engineering community that much software is of low quality with large numbers of defects and vulnerabilities. In these cases, the properties of the resulting systems may be significantly worse than those described above. For example:

- Implementations of facial age estimation systems may not incorporate strong liveness testing, allowing them to be easily fooled by synthetic faces, as with the Discord facial estimation system described in Section VII.D.2.
- Service providers and AVPs may not properly implement privacy mechanisms designed to prevent the AVP from learning the user's behavior and the service provider from learning the user's identity.[309]
- AVPs may retain uploaded documents and fail to secure them correctly, leading to disclosure of sensitive user data.[310]

[308] Minocha et al., "Papers, Please."
[309] Bouchaud, *Technical Report*.
[310] TrustCloud, "The AU10TIX case."

It is likely that at least some implementations of age assurance will have quality issues that will result in easy circumvention, compromise of user privacy, or both.[311] However, some designs are more susceptible to these errors than others. For example, in a ZKP-based system, the server never learns the user's identity and therefore even full compromise of the server will not result in a user privacy breach. When selecting a design, it is important to consider the impact of potential implementation errors and favor designs which are less susceptible to errors and where errors are less serious.

In many cases, service providers do not have a strong incentive to select AVPs based on implementation quality, because age assurance is an issue of regulatory compliance rather than of user satisfaction. This may lead them to select for cost and performance but not for privacy and user protection. This incentive is exacerbated if age assurance mandates specify effectiveness criteria but not user protection criteria.

X. 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.

[311] Livingstone et al., "Children's Rights and Online Age Assurance Systems."

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.

Appendix A. Web Technology Background

A. Cookies and Web Tracking

Unlike the telephone network, web browsers do not have fixed identities, which means that sites need a way to identify returning users. The basic technology for this is the “cookie”, which is simply an opaque string stored by the site which the browser provides in future requests. Cookies can be used in either “first-party” contexts where the user is visiting the site or “third-party” contexts for assets that are embedded on the site but served from a different server (e.g., ads).

When cookies were originally designed, third party cookies were tied only to the server they came from, with the result that if a user visited site A which embedded an asset from site T and then visited site B which also embedded an asset from T, T would receive the same cookie, allowing the user to be tracked between A and B. The result is that the tracker is able to build a “profile” of the user by recording the set of sites that a user visits. In many cases this profile alone is sufficient to identify the user,[312] but if the user ever visits a site where they identify themselves that has embedded the tracker, the tracker can then correlate that identity across the entire user profile.

Many modern browsers[313] provide anti-tracking features which prevent third party cookies from being used in this way. However, there are other mechanisms which can be used to correlate user behavior, including the user’s IP address or “fingerprinting” mechanisms which take advantage of specific features of the user’s browser or device.[314] VPNs or proxy technologies such as iCloud Private Relay[315] can be used to conceal the IP address and some browsers have mechanisms designed to resist fingerprinting, though their effectiveness is not entirely clear. Note, however, that traditional VPNs allow the VPN provider to build a per-customer profile of which services are in use, which can then be leaked or sold.[316] More advanced technologies such as iCloud private relay are designed to avoid creating a single point of visibility for the user’s identity and their behavior.

B. Third-Party Web Age Assurance: Web APIs and Web Hooks

This section provides some detail on how age verification can be deployed in the web context. This is modeled on how the Yoti system works but is broadly applicable. Figure 16 below provides an overview of the entire system in action. All of this process just uses standard web APIs and so will work with essentially any modern browser.

[312] Bird et al., “Replication.”
[313] Chrome is an exception in that it does not provide protection from third party cookies by default.
[314] See Vekaria et al., “SoK” for an overview of tracking technologies.
[315] Apple, “About iCloud Private Relay.”
[316] Center for Democracy & Technology, “VPNs.”

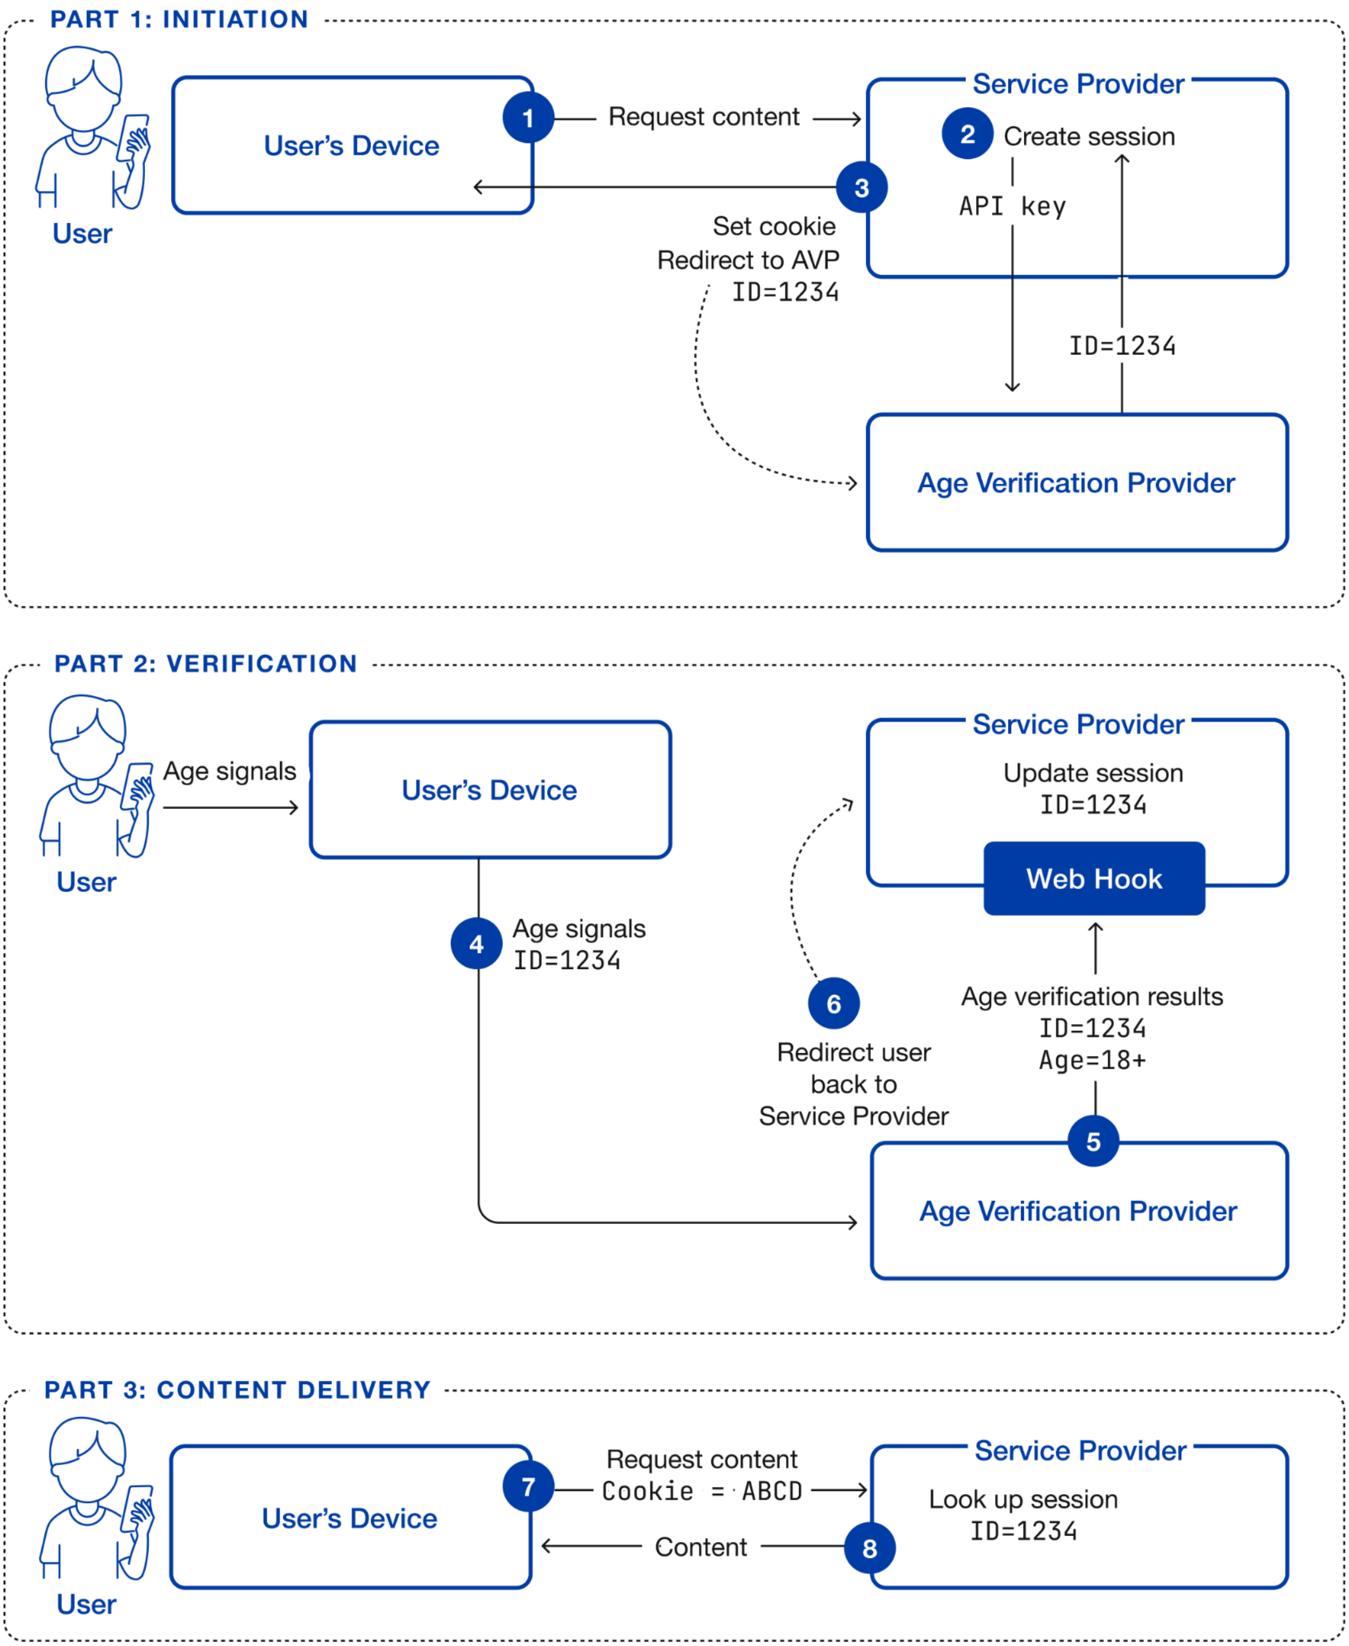

Figure 16. Age assurance on the web with a separate AVP.

The process starts with the web browser navigating to the service provider and requesting age-restricted content (step 1). The service provider recognizes that this is a new user and therefore needs an age check. The service provider contacts the age verification provider using a web API on the AVP to start an age assurance transaction. At some previous time, the service provider must have set up an account with the AVP and as part of that process, it will get an API key (a secret that it can use to authorize transactions). It will provide this to the AVP at the time the session is set up and receive a session ID for this user's transaction, in this case 1234 (step 2).

Once the age assurance transaction has been set up, the service provider sends the browser to the age verification provider, potentially with an HTTP Redirect[317] or (more likely) using JavaScript APIs to navigate the web page (step 3). The service provider provides the transaction ID to the browser so it can provide it to the AVP and at the same time stores a cookie (ABCD) on the browser so that it can recognize the browser when it comes back. The service provider stores the cookie and the transaction ID in its database for use after the browser has completed age assurance.

After step (3) is complete, the user is visiting the AVP's site. The AVP then takes the user through the process of age assurance, for instance asking them to turn on their camera and take a selfie, etc (step 4). Once the AVP is satisfied with the user's age it notifies the service provider using an API provided by the service provider (often called a "web hook") (step 5). The AVP provides the transaction ID so that the service provider knows which transaction is being approved, as there may be many concurrent transactions.

Once the AVP has notified the service provider, it then redirects the browser back to the service provider (step 6). When the browser returns to the service provider it provides the cookie that the service provider had provided in step 3 (step 7). The service provider then uses the cookie to look up the transaction in its database, sees that the age assurance completed, and so knows that it can give the browser access, and returns the requested content (step 8).

This description helps understand what each side learns during the transaction:

1. The AVP learns (1) the service provider's identity (because it's provided in the API key) and in the redirect URL and (2) whatever age signals the user sends, potentially including their identity.
2. The service provider learns the result of age assurance and what content the user engages with.

However, the AVP and the service provider can collude to correlate these two sets of knowledge based on the transaction ID. This collusion is invisible to the browser and is only enforced—if at all—based on policies.

[317] Fielding et al., "RFC 9110."

This example has shown the user as doing a new age assurance transaction, but it is also possible to skip this stage if the AVP is able to automatically verify the user's age. For example, once the user finishes the age assurance process, the AVP can set a cookie and then use that to know the user is already eligible in a future transaction. In that case, it can just redirect the user back to the service provider immediately without prompting for any age signals. Note that a cookie of this type allows the AVP to build a *profile* of the user's history on age-restricted sites. The use of such a cookie is visible to the user's browser (though not usually shown to the user).[318]

Appendix B. Attacks on Remote Face Analysis

A number of the age assurance mechanisms described in this report depend on the user providing either a selfie or a self-video,[319] which is then either matched against a biometric (picture) provided in an existing credential or analyzed directly for age estimation. The precise input provided spans a wide range of levels of interactivity, from:

- Provide a single static image (a selfie)
- Provide a single non-interactive video
- Interactively perform a set of motions (e.g., blink, turn your head, etc.) in front of a live camera.

While the analysis performed on the input varies, the general problem of assuring authentic inputs is common to both identity verification (for ID matching) and age estimation. In general, the more interactive the input, the more difficult it is to mount attacks. Typically, attacks on this kind of system are grouped into two main categories:

Presentation Attacks in which the user provides a fake input to their camera, such as by holding up a static picture to the lens or wearing a face mask.

Injection Attacks in which the user provides fake input to the device directly, bypassing the camera.

The state of attack and defense for each of these mechanisms is discussed below.

A. Presentation Attacks

Historically, most of the interest in attack detection was focused on presentation attacks. These range from unsophisticated attacks where the attacker holds up a photo to the camera to sophisticated

[318] Note that third party cookie blocking does not prevent this form of tracking, because the AVP site is "first party"; this is sometimes called "bounce tracking".
[319] For example, see Facebook, "How video selfie age verification works on Facebook."

attacks where they wear a silicone mask. These have been extensively studied in the context of remote identification, which is the same problem as ID-based age verification.

NISTIR 8491[320] reports on tests with commercial face recognition (a related problem to age assurance) against a variety of presentation attacks, some of which are not disclosed. Because there is a tradeoff between false accept and false reject rates, the standard way to report effectiveness is to set either the false accept rate or the false reject rate and then measure the false reject rate. NIST reports the false accept rate when the false reject rate is .01 and vice versa. Error rates vary widely depending on the system and the attack scenario, but the best systems exhibit both false accept and false reject rates under 10% in this setting.

NIST reports results for 6 different types of presentation attacks, with performance varying considerably between the attacks. For example, the best results for attack type 4 are a false accept rate of 13% and a false reject rate of 20%.[321] Moreover, which systems perform best varies between attacks. For example, the best system for detecting impersonation for attack type 1 does not appear in the top 10 for attack type 3. In some cases, performance on video input is superior to performance on still images, but in other cases it is worse. Nevertheless, error rates vary considerably, with the best performing systems having near zero error rates on the easiest attacks and error rates in excess of 10% on the strongest attacks. Somewhat better results can be obtained by fusing multiple algorithms.

NIST does not report on “active” systems with liveness detection where the user is asked to perform certain actions for the camera (close eyes, open mouth, etc.). These are a common feature of remote face-based age assurance and identity verification systems.

B. Injection Attacks

In an injection attack, the attacker has direct control of the image/video inputs to the system rather than having to use the system inputs. A 2024 ENISA report[322] found a large recent increase in injection attacks: “Digital injection attack incidents surged during 2022, with approximately five times more frequent and sophisticated incidents than current presentation attacks”.

Injection attacks preclude certain classes of detection that are possible with presentation attacks, such as finding the borders of a photograph held up to the camera. In the most sophisticated attacks, the attacker can use AI-based “deepfake” technology to provide a completely synthetic image. In the case of age assurance, there are two relevant attacks:

- Impersonating a specific person on an ID in order to attack a remote identity verification system which is then used for age estimation.
- Making the subject appear older in order to attack an age estimation system.

[320] Ngan et al., *Face Analysis Technology Evaluation (FATE) Part 10*.
[321] Note that these are different systems.
[322] Vrachnos et al., *Remote ID Proofing Good Practices*.

In general, the first attack is likely to be more challenging for the user to carry out than the second because the user needs to match their altered image to an image in an existing legitimate ID rather than just generically appear older. A user who can make themself appear to match an existing older person will also successfully appear older.

Tools such as Deep-Live-Cam[323] or Swapface[324] for modifying one's appearance in both still and video formats are widely available. In settings where the age assurance system asks the user to upload a static image or a video, the user can use these tools to produce the desired artifact. In cases where the age assurance system asks the user to provide live video, it is possible to use virtual camera tools such as Open Broadcaster Software (OBS) Studio[325] to modify the video on the fly and feed it to the age assurance system. It is also possible to inject video via other mechanisms such as external hardware that simulates a camera.

Broadly speaking, there are two classes of defenses against injection attacks: identifying fake content and input source integrity.

C. Identifying Fake Content

The first major defense is to analyze the input (still or video) in order to detect signals of attack. These signals can vary widely, ranging from analysis of sensor noise (which may be different in a virtual camera) to analysis of the content itself for signs of being AI-generated.

As AI-based deepfake generation is comparatively new, so is detection of AI-based deepfakes. There is a significant gap in the literature in terms of assessing effectiveness: many deepfake detection products report high accuracy rates, but independent reviews of effectiveness in general report quite low accuracy, including cases where a given system performs well on one data set and badly on another.[326] In some cases, deepfake detectors will report very high accuracy on older training sets and collapse on a newer data set.[327] This is often a sign that they have been trained on data which is not representative of the broader environment. These error rates imply that there will be a large number of missed detections of fakes in deployed systems; for example the best commercial model for still images will miss about 30% of deepfakes.

Given the immaturity of this field, and rapid advancements in both deepfake generation and detection, it is difficult to make firm predictions about whether in the future it will be easier or harder to detect deepfakes. It is likely that for some time there will be an arms race in which generation techniques

[323] Deep-Live-Cam, "Deep-Live-Cam 2.0.1c."
[324] Swapface, "Easy-to-use Faceswap AI tool in the world!"
[325] OBS Project, "OBS Studio."
[326] For example, Pirogov and Artemev tested six top open detection tools with a variety of deepfake data sets for still images: only one tool, SBI, was able to consistently deliver Receive Operating Characteristic-Area Under the Curve (ROC-AUC) values in excess of .6, and its performance degraded under downscaling. See Pirogov and Artemev, "Evaluating Deepfake Detectors in the Wild."
[327] Chandra et al., "Deepfake-Eval-2024."

improve and then detection techniques improve to compensate. Given the current state of the art, this suggests that there will be a significant period of time when generation outpaces detection, even if detection eventually prevails. Moreover, there are theoretical reasons to believe that it will eventually be possible to generate undetectable fakes, both due to the limit in quality of the true input sources and the difficulty existing classifiers have with adversarial examples generated via more sophisticated techniques.[328]

D. Input Source Integrity

The other main approach to defend against injection attacks is to prevent direct injection of media by ensuring the integrity of the equipment used to capture the media. For example, if the user is providing their video from a smartphone, the evaluator might want to ensure that the input came from the built-in camera and was not modified before being transmitted to the evaluator. This assurance is usually provided via an attestation mechanism, in which the device in question has a hardware root of trust which is able to cryptographically attest to the hardware and software running on the device.

As an example, consider an age assurance app which wants to use the camera to acquire video of the user and send it back to the app vendor for processing. It is possible for some third party to build a clone of the app that replaces the true camera input with a deepfake that makes the user appear older. The service provider offering the app can use an app integrity mechanism (see Section VI.A.2) to verify that any video came from the correct camera. Mechanisms of this type provide two challenges to availability. First, they do not work over the web, which has no mechanism for establishing that the remote endpoint is unmodified.[329] As a result, it is not possible to determine whether any image or video uploaded from a web browser is coming directly from the hardware camera on the device rather than via an injection attack.[330]

Second, remote integrity attestations inherently depend on closed devices, because it is the hardware and operating system's responsibility to attest to the content of the application. For example, it is possible for users to install their own Android Open Source Project (AOSP)-based[331] operating system on their Android-based mobile devices, but doing so precludes them from being able to issue an attestation about the software they use to supply images or video.[332] Similarly, devices based on customized Android or on fully open hardware may not be able to attest. iPhone/iOS is an inherently closed platform, so this issue is less relevant there, as it is not generally practical to get an iOS device that does not run MacOS.

[328] Goodfellow et al., "Explaining and Harnessing Adversarial Examples."

[329] In 2023, Google published such a proposal, called Web Environment Integrity but later abandoned it. See Wiser et al., "Web Environment Integrity Explainer.

[330] The Coalition for Content Provenance and Integrity (C2PA) has published a set of protocols that allow cameras and camera-containing devices to cryptographically sign their output (still photos or video) as coming from a given device. In the future this may help prevent injection attacks; however, C2PA is not yet sufficiently widely deployed to cover most users. In addition, current Web browsers modify (compress) camera output for transmission, which effectively removes the C2PA annotations. See C2PA, "Content Credentials"; Mozilla, "Codecs used by WebRTC."

[331] Android, "Android Open Source Project."

[332] It is not sufficient to ensure that the user has an official device because a modified operating system can affect the behavior of the application.

Anti-injection mechanisms based on obfuscated code running on an open device are generally not effective. For example, Yoti's web-based facial age estimation system use a "Secure Image Capture" feature[333] in JavaScript which performs on-device face detection and captures on-device metadata to send to Yoti for analysis along with the user's image. This data is protected with a cryptographic key which is embedded in the device-side JavaScript. This design suffers from a number of straightforward security vulnerabilities, as documented by recent research.[334]

Appendix C. Zero-Knowledge Proofs

One of the most widely discussed mechanisms for privacy-preserving age assurance is to use zero-knowledge proofs to demonstrate the user's age. This section provides a very brief overview of zero-knowledge proof technology, sufficient to understand its application in this context.[335]

The basic idea behind a zero-knowledge proof (ZKP) is to allow one party, the prover (conventionally: *P*), to demonstrate to the verifier (*V*) the truth of some statement (*S*) without revealing any other information to the verifier. Of particular interest is what's called a zero-knowledge proof of knowledge (ZKPK), in which *P* demonstrates they know some secret value *w* (called a "witness"). Anyone who knows *w* would be able to prove the truth of *S.*

In principle, ZKPs are a generic tool, but historically many ZKPs only allowed them to prove specific narrow statements, e.g., "I know the discrete log of this value." However, in recent years, it has become practical to use ZKPs to prove *arbitrary* statements, e.g., that a given function when run on some input will have a specific output. Operationally, one can think of the statement to be proved (*S)* as a computer program that takes the witness as input, i.e., *S*(w). The prover then sends the verifier a proof *p* which shows that the prover knows a witness *w* such that if you ran *P* on *w* the output would be ***true***, but without allowing the verifier to learn *w*.

In an age assurance context, a ZKP-based system can be deployed on top of a standard digital credentials system like mobile driver's license. In typical deployments of those systems (see Section VII.C.2.e for a description of Apple's system), the credential C contains a public key K_{pub} which corresponds to a private key K_{priv} stored in the device. In order to authenticate, the device provides both the credential C and a digital signature S using K_{priv}, thus demonstrating that the user has access to the device to which the credential was issued. Because authentication requires providing the credential C, that credential can be used to track the user. This is addressed in a ZKP system by having the device produce a ZKP that it was able to produce the witness (C, S) and that C has the right attributes (e.g., that the user is over 18) but without actually showing C or S to the evaluator.

[333] Yoti, "Secure Image Capture."

[334] Minocha et al., "Papers, Please."

[335] For more information on these systems, see Thaler, "Proofs, Arguments, and Zero-Knowledge."

While ZPKs are a very powerful tool, it can be difficult to make them efficient. For example, a ZKP of the validity of a signature using the standard Elliptic Curve Digital Signature Algorithm (ECDSA)[336] is expensive with a number of ZKP systems due to the specific math of ECDSA. Unfortunately, ECDSA is commonly used to sign existing credentials such as mobile driver's licenses, which makes it challenging to reuse those credentials. A 2025 paper[337] shows how to efficiently compute ZKPs for these credentials and is the basis for Google's new ZKP-based age assurance system.[338] ZKPs need to be specifically tailored for each statement which needs to be proved, and doing so correctly is challenging. For example, if a credential changes its date format, the ZKP program will need to be modified.

As with non-ZKP digital credential systems, binding the credential to the device is intended to resist cloning attacks where an adult obtains a valid credential and then shares C, K_{priv} with a minor. These attacks are more serious with ZKPs because the evaluator never learns C and therefore is not able to detect cases where many users are authenticating with the same credential. If K_{priv} is stored in a secure element, then the attacker must break the secure element in order to perform a cloning attack. However, requiring device binding restricts the availability of ZKPs to users who have devices with secure elements precluding the use of fully open devices.

It is also possible to deter cloning attacks by use of zero-knowledge rate limiting techniques such as rate-limiting nullifiers.[339] Effectively, these techniques restrict the number of times a given credential can be used to authenticate, thus reducing the impact of leakage of K_{priv}. These techniques can be used as a backup to device binding or as a standalone security measure without device binding, although prominent deployments of ZKPs such as Google's seem likely to require device binding.[340]

[336] National Institute of Standards and Technology, "FIPS 186-5 Digital Signature Standard."
[337] Frigo and shelat, "Anonymous credentials from ECDSA."
[338] Stapelberg, "It's now easier to prove age and identity with Google Wallet."
[339] Privacy & Scaling Exploration, "Rate-Limiting Nullifier."
[340] Frigo and shelat, "Anonymous credentials from ECDSA."

Bibliography

5Rights Foundation. *But how do they know it is a child? Age Assurance in the Digital World*. 5Rights Foundation, October 2021. https://5rightsfoundation.com/wp-content/uploads/2024/09/But_How_Do_They_Know_It_is_a_Child-1.pdf.

Abrams, Lawrence. “PornHub extorted after hackers steal Premium member activity data.” *BleepingComputer*, December 15, 2025. https://www.bleepingcomputer.com/news/security/pornhub-extorted-after-hackers-steal-premium-member-activity-data/.

Age Assurance Technology Trial. “Needemand: Vendor Case Study.” Accessed January 23, 2026. https://ageassurance.com.au/v/nee/.

———. *Part A: Main Report*. Age Assurance Technology Trial, August 2025. https://ageassurance.com.au/.

Age Check Certification Scheme. *Age Estimation Test Report: Yoti*. Age Assurance Technology Trial, October 6, 2025. https://ageassurance.com.au/wp-content/uploads/2025/08/IndividualTestReport-YOTI_AE.pdf.

———. “Comparison Guide.” Accessed January 22, 2026. https://accscheme.com/wp-content/uploads/Comparison-Guide.pdf?srsltid=AfmBOoprLGouZDOWFc-y1Hyy58bkiYyrLa2OIocBYg44uDfn2Gx_z4f8.

Age Verification Providers Association. “No, UK porn use was not halved by age verification.” September 8, 2025. https://avpassociation.com/thought-leadership/no-uk-porn-use-was-not-halved-by-age-verification/.

———. “Privately.” Accessed January 28, 2026. https://avpassociation.com/member/privately/.

———. “US state age assurance laws for social media.” Last modified August 2025. https://avpassociation.com/us-state-age-assurance-laws-for-social-media/.

———. “US State age verification laws for adult content.” Last modified June 30, 2025. https://avpassociation.com/4271-2/.

Anderson, Monica, and Michelle Faviero. “81% of adults - versus 46% of teens - favor parental consent for minors to use social media.” Pew Research Center, October 31, 2023. https://www.pewresearch.org/short-reads/2023/10/31/81-of-us-adults-versus-46-of-teens-favor-parental-consent-for-minors-to-use-social-media/.

Android. “Android Open Source Project.” Accessed January 23, 2026. https://source.android.com/.

———. “Hundreds of partners ship Play Protect certified phones and tablets.” Accessed January 23, 2026. https://www.android.com/certified/partners/.

ANSSI and BSI. “Remote Identity Proofing ANSSI-BSI Joint Release.” December 20, 2023. https://cyber.gouv.fr/sites/default/files/document/ANSSI-BSI-Joint-Release_20231220.pdf.

Apple. “About iCloud Private Relay.” Apple Support. August 31, 2023. https://support.apple.com/en-us/102602.

———. “Add your driver’s license to Apple Wallet.” Apple Support. September 15, 2025. https://support.apple.com/en-us/111803.

———. “App Review Guidelines.” Apple Developer. Last modified November 13, 2025. https://developer.apple.com/app-store/review/guidelines/.

———. “Establishing your app’s integrity.” Apple Developer. Accessed January 26, 2026. https://developer.apple.com/documentation/devicecheck/establishing-your-app-s-integrity.

———. “How to Use Hide My Email with Sign in with Apple.” Apple Support. Accessed January 22, 2026. https://support.apple.com/en-us/105078.

———. “ID in Wallet.” Accessed January 23, 2026. https://learn.wallet.apple/id#states-list.

———. “IDs in Apple Wallet: privacy and security overview.” Apple Support. November 12, 2025. https://support.apple.com/en-us/118260.

———. “IdentityDocumentWebPresentmentController.” Apple Developer. Accessed January 23, 2026. https://developer.apple.com/documentation/identitydocumentservicesui/identitydocumentwebpresentmentcontroller.

———. “Notarizing MacOS software before distribution.” Apple Developer. Accessed January 23, 2026. https://developer.apple.com/documentation/security/notarizing-macos-software-before-distribution.

———. “Update on apps distributed in the European Union.” Apple Developer. Accessed January 23, 2026. https://developer.apple.com/support/dma-and-apps-in-the-eu/#notarization-for-ios-apps.

———. “Updates to runtime protection in MacOS Sequoia.” Apple Developer. August 6, 2024. https://developer.apple.com/news/?id=saqachfa#:~:text=In%20macOS%20Sequoia%2C%20users%20will%20no%20longer,software%20to%20run%20after%20reviewing%20security%20information.

———. “Use parental controls to manage your child’s iPhone or iPad.” Apple Support. December 12, 2025. https://support.apple.com/en-us/105121.

———. “Validating apps that connect to your server.” Apple Developer. Accessed January 23, 2026. https://developer.apple.com/documentation/devicecheck/validating-apps-that-connect-to-your-server.

———. “Verify identity documents on the web.” Apple Developer. Accessed January 23, 2026. https://developer.apple.com/videos/play/wwdc2025/232/.

———. “Your business. Open on Apple apps.” Accessed January 23, 2026. https://businessconnect.apple.com/.

Apthorpe, Noah, Brett Frischmann, and Yan Shvartzshnaider. “Online Age Gating: An Interdisciplinary Evaluation.” Preprint, submitted September 24, 2024. https://papers.ssrn.com/sol3/papers.cfm?abstract_id=4937328.

Association of Sites Advocating Child Protection. “What is the RTA Label?” Accessed January 22, 2026. https://www.rtalabel.org/.

Australian Government Department of Infrastructure, Transport, Regional Development, Communications, Sport and the Arts. *A summary report on developmental research to inform a Social Media Minimum Age campaign*. Australian Government Department of Infrastructure, Transport, Regional Development, Communications, Sport and the Arts, September 15, 2025.

https://www.esafety.gov.au/sites/default/files/2025-10/Social-Media-Minimum-Age-Campaign-research-report-summary-Oct2025.pdf?v=1765497600038.

Aylo. "Aylo response to Ofcom consultation on Guidance for service providers publishing pornographic content." Accessed January 22, 2026. https://www.ofcom.org.uk/siteassets/resources/documents/consultations/category-1-10-weeks/272586-consultation-guidance-for-service-providers-publishing-pornographic-content/responses/aylo.pdf?v=370047.

Barthold, Corbin K., Aaron Mackey, Ari Cohn, Ben Szoka, Elizabeth Femia, and David Greene. "Brief of Amici Curiae Electronic Frontier Foundation, Woodhull Freedom Foundation, and Techfreedom in Support of Petitioners." Supreme Court of the United States. Accessed January 21, 2026. https://www.supremecourt.gov/DocketPDF/23/23-1122/326628/20240923142310077_23-1122_Amici%20Brief.pdf.

Bellovin, Steven M. "Privacy-Preserving Age Verification—And Its Limitations." Preprint, submitted October 2025. https://www.cs.columbia.edu/~smb/papers/age-verify.pdf.

Beser, James. "Extending Our Built-In Protections to More Teens on YouTube." YouTube. *YouTube Official Blog*, July 29, 2025. https://blog.youtube/news-and-events/extending-our-built-in-protections-to-more-teens-on-youtube/.

Better Internet for Kids. "Austria - Policy Monitor Country Profile." 2025. https://better-internet-for-kids.europa.eu/en/knowledge-hub/austria-policy-monitor-country-profile.

Bird, Sarah, Ilana Segall, and Martin Lopatka. "Replication: Why We Still Can't Browse in Peace: On the Uniqueness and Reidentifiability of Web Browsing Histories." *Proceedings of the Sixteenth Symposium on Usable Privacy and Security* (2020): 488-503, https://www.usenix.org/system/files/soups2020-bird.pdf.

Borderage. "AI technology based on medical research." Accessed January 23, 2026. https://borderage.com/technology/#research.

Bouchaud, Paul. *Technical Report: AgeGO Age Verification on Pornographic Platforms*. AI Forensics, September 2025. https://aiforensics.org/uploads/AIF_report_AgeGO_porn_platforms.pdf.

boyd, danah, Eszter Hargittai, Jason Schultz, and John Palfrey. "Why parents help their children lie to Facebook about age: Unintended consequences of the 'Children's Online Privacy Protection Act.'" *First Monday 16, no. 11* (2011), https://firstmonday.org/ojs/index.php/fm/article/view/3850/3075.

Bradshaw, Tim. "Meta adopts new age-check system to meet global child safety laws." *Financial Times*, December 17, 2025. https://www.ft.com/content/8164c36c-f224-4006-bc1f-561952c119b6.

———. "VPN use surges in UK as new online safety rules kick in." *Financial Times*, July 22, 2025. https://www.ft.com/content/356674b0-9f1d-4f95-b1d5-f27570379a9b.

Brennen, Scott Babwah, and Matt Perault. *Keeping Kids Safe Online: How Should Policymakers Approach Age Verification?* The Center on Growth and Opportunity at Utah State University and the Center on Technology Policy at the University of North Carolina at Chapel Hill, June 2023. https://www.thecgo.org/wp-content/uploads/2023/06/Age-Assurance_03.pdf.

British Board of Film Classification. “Mobile Content.” Accessed January 23, 2026. https://www.bbfc.co.uk/mobile-content.

British Standards Institution. “PAS 1296:2018. Online age checking. Provision and use of online age check services. Code of Practice.” March 31, 2018. .https://knowledge.bsigroup.com/products/online-age-checking-provision-and-use-of-online-age-check-services-code-of-practice.

———. “ISO/IEC NP 27566-2: Age assurance systems — Part 2: Technical approaches and guidance for implementation.” December 7, 2024. https://standardsdevelopment.bsigroup.com/projects/9024-10663.

Brooks, Mindy. “Ensuring a safer online experience for U.S. kids and teens.” Google. *The Keyword*, June 20, 2025. https://blog.google/innovation-and-ai/technology/safety-security/age-assurance-measures-safer-online-kids-teens-us/.

Caceres, Marcos, Tim Cappalli, Mohamed Amir Yosef, and Sam Goto. “Digital Credentials.” WC3. January 13, 2026. https://www.w3.org/TR/digital-credentials/.

California. “AB 1043.” California Legislative Information. October 13, 2025. https://leginfo.legislature.ca.gov/faces/billNavClient.xhtml?bill_id=202520260AB1043.

———. “SB 243.” California Legislative Information. October 13, 2025. https://leginfo.legislature.ca.gov/faces/billNavClient.xhtml?bill_id=202520260SB243&utm_campaign=wp_the_technology_202&utm_medium=email&utm_source=newsletter.

Canada. “Choosing a Credit Card.” Last modified October 15, 2025. https://www.canada.ca/en/financial-consumer-agency/services/credit-cards/choose-credit-card.html.

Castro, Chiara. “‘VPNs are not kryptonite of age assurance’ - Age verification experts explain why governments don’t need to ban VPNs.” *TechRadar*, August 14, 2025. https://www.techradar.com/vpn/vpn-privacy-security/vpns-are-not-kryptonite-of-age-assurance-age-verification-experts-explain-why-governments-dont-need-to-ban-vpns.

Center for Democracy & Technology. “VPNs.” Accessed January 23, 2026. https://cdt.org/vpns/.

Chandra, Nuria Alina, Ryan Murtfeldt, Lin Qiu, et al. “Deepfake-Eval-2024: A Multi-Modal In-the-Wild Benchmark of Deepfakes Circulated in 2024.” Preprint, submitted March 4, 2025. https://arxiv.org/pdf/2503.02857v4.

Character.ai. “Age Assurance: What you need to know.” Last modified November 2025. https://support.character.ai/hc/en-us/articles/42828297541787-Age-Assurance-What-you-need-to-know.

Chevis, Andrew. “UK: CitizenCard - How to ID a person with no ID (Vouching).” Video, January 2025. Posted by Age Check Certification Scheme. YouTube. https://youtu.be/v2zgkRLoPI4.

Cisco. “Actions Speak Louder Than Words: Despite Claiming Security Awareness, Many Remote Workers Engage in Risky Online Behavior.” October 9, 2006. https://newsroom.cisco.com/c/r/newsroom/en/us/a/y2006/m10/actions-speak-louder-than-words-despite-claiming-security-awareness-many-remote-workers-engage-in-risky-online-behavior.html.

Clark County District Court. “Complaint and Demand for Jury Trial.” January 30, 2024. https://ag.nv.gov/uploadedFiles/agnvgov/Content/Issues/2024.01.30%20Snap%20Complaint.pdf.

CNIL. *Online age verification: a complex issue with significant privacy risks*. CNIL, September 22, 2022. https://www.cnil.fr/en/online-age-verification-balancing-privacy-and-protection-minors.

Coalition for Content Provenance and Integrity. “Content Credentials : C2PA Technical Specification.” Accessed January 28, 2026. https://spec.c2pa.org/specifications/specifications/2.3/specs/C2PA_Specification.html.

Cox, Joseph. “Inside the Underground Site Where ‘Neural Networks’ Churn Out Fake IDs.” *404 Media*, February 5, 2024. https://www.404media.co/inside-the-underground-site-where-ai-neural-networks-churns-out-fake-ids-onlyfake/.

Cyber Security Intelligence. “VPN Demand Surges As British Online Safety Law Takes Effect.” July 8, 2025. https://www.cybersecurityintelligence.com/blog/vpn-demand-surges-as-british-online-safety-takes-effect-8580.html.

Datta, Anupriya. “VPN surge won't stop France's fight against porn, vows its digital minister.” *Euractiv*, June 6, 2025. https://www.euractiv.com/news/vpn-surge-wont-stop-frances-fight-against-porn-vows-its-digital-minister/.

Deep-Live-Cam. “Deep-Live-Cam 2.0.1c.” GitHub. Last modified December 2025. https://github.com/hacksider/Deep-Live-Cam.

Desmarais, Anna. “The age verification era: Which EU countries are restricting access to adult sites?” *Euronews*, November 17, 2025. https://www.euronews.com/next/2025/11/17/the-age-verification-era-which-eu-countries-are-restricting-access-to-adult-sites.

Dhesi, Tajveer Singh, and Noah Apthorpe. “Measuring the Prevalence and Variety of Online Age Gates.” Accessed January 23, 2026. https://www.ieee-security.org/TC/SPW2025/ConPro/papers/dhesi-conpro25.pdf.

Digital Element. “An Executive’s Guide to IP Geolocation.” Accessed January 22, 2026. https://www.digitalelement.com/resources/guides/guide-to-ip-geolocation/.

Digital Trust & Safety Partnership. *Age Assurance: Guiding Principles and Best Practices*. Digital Trust & Safety Partnership, September 2023. https://dtspartnership.org/wp-content/uploads/2023/09/DTSP_Age-Assurance-Best-Practices.pdf.

DraftKings. “Why am I being asked to verify my identity? (US).” Accessed January 21, 2026. https://support.draftkings.com/dk/en-us/why-am-i-being-asked-to-verify-my-identity-us?id=kb_article_view&sysparm_article=KB0010467.

Dutkowska-Zuk, Agnieszka, Austin Hounsel, Amy Morrill, Andre Xiong, Marshini Chetty, and Nick Feamster. “How and Why People Use Virtual Private Networks.” *Proceedings of the 31st USENIX Security Symposium* (2022), https://www.usenix.org/conference/usenixsecurity22/presentation/dutkowska-zuk.

EE, Telefonica UK, Vodafone, and Three. "UK Code of practice for the self-regulation of content on mobiles." July 1, 2013. https://cdn.prod.website-files.com/5b7ab54b285dec5c113ee24d/5d5d4f65228753840d521965_uk-code-of-practice-for-the-self-regulation-of-content-on-mobiles.pdf.

Eltaher, Fatmaelzahraa, Rahul Krishna Gajula, Luis Miralles-Pechuán, Christina Thorpe, and Susan Mckeever. "The Digital Loophole: Evaluating the Effectiveness of Child Age Verification Methods on Social Media." *Proceedings of the 11th International Conference on Information Systems Security and Privacy* 2 (2025): 213-222, https://www.scitepress.org/Papers/2025/132483/132483.pdf.

Englehardt, Steven, Jeffrey Han, and Arvind Narayanan. "I never signed up for this! Privacy implications of email tracking." *Proceedings on Privacy Enhancing Technologies* 1 (2018): 109-126, https://petsymposium.org/popets/2018/popets-2018-0006.pdf.

eSafety Commissioner. *Behind the screen: The reality of age assurance and social media access for young Australians*. eSafety Commissioner, February 25, 2025. https://www.esafety.gov.au/sites/default/files/2025-02/Behind-the-screen-transparency-report-Feb2025.pdf?v=1739990594940.

———. "Consolidated Industry Codes of Practice for the Online Industry (Class 1C and 2 Material) Head Term." 2025. https://www.esafety.gov.au/sites/default/files/2025-09/Consolidated-Industry-Codes-of-Practice-for-the-Online-Industry-%28Class-1C-and-Class-2-Material%29-Head-Terms-9-September-2025.pdf?v=1763069128966.

———. "Schedule 3 - Internet Search Engine Services Online Safety Code (Class 1C and 2 Material)." 2025. https://www.esafety.gov.au/sites/default/files/2025-06/Schedule-3-Internet-Search-Engine-Services-Online-Safety-Code-%28Class-1C-and-Class-2-Material%29.pdf?v=1763069128966.

———. "Social Media Minimum Age Campaign." Accessed January 21, 2026. https://www.esafety.gov.au/about-us/industry-regulation/social-media-age-restrictions/campaign.

———. "Social media age restrictions." Last modified January 13, 2026. https://www.esafety.gov.au/about-us/industry-regulation/social-media-age-restrictions.

European Commission. "A digital ID and personal digital wallet for EU citizens, residents and businesses." Accessed January 21, 2026. https://ec.europa.eu/digital-building-blocks/sites/spaces/EUDIGITALIDENTITYWALLET/pages/694487738/EU+Digital+Identity+Wallet+Home.

———. "Commission opens investigations to safeguard minors from pornographic content under the Digital Services Act." May 26, 2025. https://ec.europa.eu/commission/presscorner/detail/en/ip_25_1339.

———. "EU Age Verification Solution." Accessed January 21, 2026. https://ageverification.dev/.

———. "Guidelines on measures to ensure a high level of privacy, safety and security for minors online, pursuant to Article 28(4) of Regulation (EU) 2022/2065." July 10, 2025. https://ec.europa.eu/newsroom/dae/redirection/document/118226.

———. "Installing the App." Accessed January 21, 2026. https://ageverification.dev/Getting%20started/app_installation/.

———. "Operational, Security, Product, and Architecture Specifications." Accessed January 21, 2026. https://ageverification.dev/av-doc-technical-specification/docs/architecture-and-technical-specifications/.

European Data Protection Board. "Interplay between the DSA and the GDPR: EDPR adopts guidelines." September 12, 2025. https://www.edpb.europa.eu/news/news/2025/interplay-between-dsa-and-gdpr-edpb-adopts-guidelines_en.

European Parliament. "Children should be at least 16 to access social media, says MEPs." November 26, 2025. https://www.europarl.europa.eu/news/en/press-room/20251120IPR31496/children-should-be-at-least-16-to-access-social-media-say-meps.

———. "Modernising EU driving rules to increase road safety." October 21, 2025. https://www.europarl.europa.eu/news/en/press-room/20251016IPR30947/modernising-eu-driving-rules-to-increase-road-safety.

European Union. "Directive (EU) 2018/1808 of the European Parliament and of the Council of 14 November 2018 amending Directive 2010/13/EU on the coordination of certain provisions laid down by law, regulation or administrative action in Member States concerning the provision of audiovisual media services (Audiovisual Media Services Directive) in view of changing market realities." November 14, 2018. https://eur-lex.europa.eu/legal-content/EN/TXT/PDF/?uri=CELEX:32018L1808.

———. "Regulation (EU) 2016/679 of the European Parliament and of the Council (General Data Protection Regulation)." April 27, 2016. https://eur-lex.europa.eu/eli/reg/2016/679/oj/eng.

———. "Regulation (EU) 2022/2065 of the European Parliament and of the Council of 19 October 2022 on a Single Market For Digital Services and Amending Directive 2000/31/EC (Digital Services Act)." 2022. https://eur-lex.europa.eu/eli/reg/2022/2065/oj/eng.

———. "Revision of the Audiovisual Media Services Directive." Accessed January 21, 2026. https://digital-strategy.ec.europa.eu/en/policies/revision-avmsd.

European Union Agency for Fundamental Rights. "Minimum age requirements related to rights of the child in the EU." Last modified October 19, 2018. https://fra.europa.eu/en/publications-and-resources/data-and-maps/minag?dataSource=MINAG_en_74851&media=png&width=740&topic=group13&question=MINAG_EMP07&plot=MAP&subset=NONE&subsetValue=NONE&answer=MINAG_EMP07&year=2017.

ExpressVPN. "VPN Servers: Choose the best VPN location for your needs." Accessed January 22, 2026. https://www.expressvpn.com/vpn-server.

Facebook. "How Teen Accounts work on Facebook." Help Center. Accessed January 23, 2026. https://www.facebook.com/help/1123135202753352.

———. "How video selfie age verification works on Facebook." Help Center. Accessed January 23, 2026. https://www.facebook.com/help/1386337538619854/.

Facial Equality International. "Facial Recognition and the Facial Difference Community: 2024 Survey Results." 2024. https://faceequalityinternational.org/FEI_2024_survey_results.pdf.

Family Online Safety Institute. “Parental Controls for Online Safety are Underutilized, New Study Finds.” May 28, 2025. https://fosi.org/parental-controls-for-online-safety-are-underutilized-new-study-finds/.

FDIC. “FDIC Survey Finds 96 Percent of U.S. Households Were Banked in 2023.” November 12, 2024. https://fdic.gov/news/press-releases/2024/fdic-survey-finds-96-percent-us-households-were-banked-2023.

FIDO Alliance. “Passkeys.” Accessed January 22, 2026. https://fidoalliance.org/passkeys/.

Fielding, R., M. Nottingham, and J. Reschke. “RFC 9110: HTTP Semantics.” Internet Engineering Task Force. June 2022. https://www.rfc-editor.org/rfc/rfc9110.html.

Finkle, Erica, Sheng Luo, Christine Agarwal, and Dave Fryer. “How Meta uses AI to better understand people’s ages on our platforms.” Meta. *Tech at Meta*, June 22, 2022. https://tech.facebook.com/artificial-intelligence/2022/6/adult-classifier.

Firefox. “Protect your identity with secure phone and email masking.” Firefox Relay. Accessed January 22, 2026. https://relay.firefox.com/.

First Judicial District Court of New Mexico. “Plaintiff’s Complaint for Abatement and Civil Penalties and Demand for Jury Trial.”f December 5, 2023. https://www.nmag.gov/wp-content/uploads/2024/01/2023-12-05-NM-v.-Meta-et-al.-COMPLAINT-REDACTED.pdf.

Forland, Sarah, Nat Meysenburg, and Erika Solis. *Age Verification: The Complicated Effort to Protect Youth Online*. New America, April 23, 2024. https://www.newamerica.org/oti/reports/age-verification-the-complicated-effort-to-protect-youth-online/.

Free Speech Coalition. “Appendix: About Age Verification Data.” Accessed January 23, 2026. https://www.freespeechcoalition.com/appendix-about-age-verification-data.

Free Speech Coalition Action Center. “State Age Verification Laws: Laws in Effect.” Accessed January 21, 2026. https://action.freespeechcoalition.com/age-verification-resources/state-avs-laws/.

Frey, Suzanne. “A new layer of security for certified Android devices.” Android. *Android Developers Blog*, August 25, 2025. https://android-developers.googleblog.com/2025/08/elevating-android-security.html.

Frigo, Matteo, and abhi shelat. “Anonymous credentials from ECDSA.” Cryptology ePrint Archive. Last modified December 20, 2024. https://eprint.iacr.org/2024/2010.

Ganel, Tzvi, Carmel Sofer, and Melvyn A. Goodale. “Biases in human perception of facial age are present and more exaggerated in current AI technology.” *Scientific Reports* 12 (2022), https://www.nature.com/articles/s41598-022-27009-w.

GeoComply. “Helping stop geo-piracy and location fraud with award-winning VPN and proxy detection.” Accessed January 22, 2026. https://www.geocomply.com/industries/media-entertainment/.

Gluck, Justine. “Understanding the New Wave of Chabot Legislation: California SB 243 and Beyond.” Future of Privacy Forum. November 4, 2025. https://fpf.org/blog/understanding-the-new-wave-of-chatbot-legislation-california-sb-243-and-beyond/#:~:text=As%20more%20states%20consider%20how,include%20protections%20tailored%20to%20minors.

Goldman, Eric. “The ‘Segregate-and-Suppress’ Approach to Regulating Child Online Safety.” *Stanford Technology Law Review* 28 (2025): 173-232, https://law.stanford.edu/publications/the-segregate-and-suppress-approach-to-regulating-child-safety-online/.

Goodfellow, Ian J., Jonathan Shlens, and Christian Szegedy. “Explaining and Harnessing Adversarial Examples.” Preprint, submitted December 20, 2014. https://arxiv.org/abs/1412.6572.

Google. “Guide your child’s Gemini Apps experience.” Gemini Apps Help. Accessed January 23, 2026. https://support.google.com/gemini/answer/16109150?hl=en.

———. “Manage your child’s Google Play apps.” Google For Families Help. Accessed January 22, 2026. https://support.google.com/families/answer/7103028?hl=en.

———. “Play integrity and signing services.” Android Developers. Last modified January 14, 2025. https://developer.android.com/google/play/integrity.

———. “Providing a safe and trusted experience for everyone.” Developer Policy Center. Accessed January 23, 2026. https://play.google/developer-content-policy/.

Government Accountability Office. *Credit Cards: Pandemic Assistance Likely Helped Reduce Balances, and Credit Terms Varied Among Demographic Groups*. Government Accountability Office, September 29, 2023. https://www.gao.gov/products/gao-23-105269.

GPSPATRON. “GNSS Spoofing Scenarios with SDRs.” September 26, 2021. https://gpspatron.com/gnss-spoofing-scenarios-with-sdrs/.

Grosshans, Holly. “Comments in response to the New York State Office of the Attorney General’s (OAG) Advanced Notice of Proposed Rulemaking (ANPRM) for the Stop Addictive Feeds Exploitation (SAFE) for Kids Act (S7694A).” Common Sense Media. September 30, 2024. https://www.commonsensemedia.org/sites/default/files/featured-content/files/safe-for-kids-act-comments.pdf.

Hales, Daniel. “Re: Advanced Notice of Proposed Rulemaking pursuant to New York General Business Law section 1500 et seq.” Future of Privacy Forum. September 30, 2024. https://fpf.org/wp-content/uploads/2024/09/Future-of-Privacy-Forum-NY-SAFE-for-Kids-Act-Comments_09.30.24.pdf.

Hanacek, Natasha. *Face Analysis Technology Evaluation (FATE) Age Estimation & Verification*. National Institute of Standards and Technology, January 23, 2026. https://pages.nist.gov/frvt/html/frvt_age_estimation.html.

Hancock, Alexis, and Paige Collings. “Zero Knowledge Proofs Alone Are Not a Digital ID Solution to Protecting User Privacy.” Electronic Frontier Foundation. July 25, 2025. https://www.eff.org/deeplinks/2025/07/zero-knowledge-proofs-alone-are-not-digital-id-solution-protecting-user-privacy.

Hogg, Luke, and Evan Swarztrauber. *On the Internet, No One Knows You’re a Dog: Examining the Feasibility of Privacy-Preserving Age Verification Online*. Foundation for American Innovation, February 2025. https://cdn.sanity.io/files/d8lrla4f/staging/0287856bc4be1f8a80271e3e9048e48920f41f7b.pdf.

Howard, John J.. Richard O. Plesh, Yevgeniy B. Sirotin, Jerry L. Tipton, and Arun R. Vemury. *A Quantitative Framework for Evaluating Remote Identity Validation Systems: Technical Demonstration Analysis and Evaluation*. Department of Homeland Security, June 2025.

https://www.dhs.gov/sites/default/files/2025-07/25_0723_st_quantitative_framework_for_evaluating_remote_identity_validation_systems.pdf.

Hutchinson, Andrew. "Meta Calls for New Legislation That Would Force App Stores to Implement Age Restrictions." *Social Media Today*, November 15, 2023. https://www.socialmediatoday.com/news/meta-calls-new-legislation-would-force-app-stores-implement-age/699931/.

IEEE. "IEEE 2089.1-2024: IEEE Standard for Online Age Verification." IEEE Standards Association. September 23, 2021. https://standards.ieee.org/ieee/2089.1/10700/.

ilGur. "Change Geolocation (Location Guard)." Firefox Browser Add-Ons. Accessed January 22, 2026. https://addons.mozilla.org/en-US/firefox/addon/change-geolocation-locguard/.

Incode. "Privacy Policy." Last modified June 30, 2025. https://incode.com/privacy-policy/#transfers-of-personal-data.

Instagram. "Edit Instagram's in-app browser settings." Help Center. Accessed January 22, 2026. https://help.instagram.com/1255727395294269/.

———. "Instagram Teen Accounts: Inspired by 13+ Movie Ratings." Accessed January 21, 2026. https://about.instagram.com/community/teen-accounts.

———. "Instagram Teen Accounts Will Be Inspired by Movie Ratings for Ages 13+." Newsroom. Last modified December 19, 2025. https://about.fb.com/news/2025/10/instagram-teen-accounts-13-movie-ratings.

———. "Introducing Instagram Teen Accounts: Built-In Protections for Teens, Peace of Mind for Parents." September 17, 2024. https://about.instagram.com/blog/announcements/instagram-teen-accounts/.

International Civil Aviation Organization. "Doc 9303 Machine Readable Travel Documents Part 1: Introduction." 2021. https://www.icao.int/sites/default/files/publications/DocSeries/9303_p1_cons_en.pdf.

International Organization for Standardization. "ISO/IEC 15438:2015 Information technology — Automatic identification and data capture techniques — PDF417 bar code symbology specification." 2015. https://www.iso.org/standard/65502.html.

———. "ISO/IEC 18013-5:2021 Personal identification — ISO-compliant driving licence Part 5: Mobile driving licence (mDL) application." 2021. https://www.iso.org/standard/69084.html.

———. "ISO/IEC 27566-1:2025(en) Information security, cybersecurity and privacy protection — Age assurance systems — Part 1: Framework." Online Browsing Platform. December 2025. https://www.iso.org/obp/ui/en/#iso:std:iso-iec:27566:-1:ed-1:v1:en.

———. "ISO/IEC CD 27566-3: Information security, cybersecurity and privacy protection – Age assurance systems - Part 3: Approaches to analysis or comparison." Accessed January 22, 2026. https://www.iso.org/standard/88147.html.

Iovine, Anna. "Do age-verification laws work? Not according to this study." *Mashable*, March 6, 2025. https://mashable.com/article/age-verification-laws-dont-work-nyu-study.

Jack Daniel's. "Jack's Birthdate Has Always Been a Mystery. Hopefully Yours Isn't." Accessed January 23, 2026. https://www.jackdaniels.com/en-us/ageGate.

Jackson, Dennis. "Who Bears the Burden? Technical Architectures for Age-Based Content Restriction." IAB/W3C Workshop on Age-Based Restrictions on Content Access. Datatracker.

August 25, 2025. https://datatracker.ietf.org/doc/slides-agews-paper-who-bears-the-burden-technical-architectures-for-age-based-content-restriction/.

Kaufman, Matt. “Revolutionizing Digital Connection: Roblox’s Vision for Age-Based Communication.” July 17, 2025. https://corp.roblox.com/newsroom/2025/07/advancing-safety-on-roblox-with-age-based-communication.

———. “Roblox Announces Ambitious Plan to Expand Age Estimation to All Users.” September 3, 2025. https://corp.roblox.com/newsroom/2025/09/roblox-to-expand-age-estimation-to-all-users.

Khan, Etienne, Anna Sperotto, Jeroen van der Ham, and Roland van Rijswijk-Deij. “Stranger VPNs: Investigating the Geo-Unblocking Capabilities of Commercial VPN Providers.” In *Passive and Active Measurement*, edited by Anna Brunstrom and Marco Fiore. Springer, 2023.

Kids Web Services. “Privacy Policy.” Last modified December 8, 2025. https://www.kidswebservices.com/en-US/privacy-policy#reasons-share.

Klapper, Leora, Dorothe Singer, Laura Starita, and Alexandra Norris. *The Global Findex Database 2025: Connectivity and Financial Inclusion in the Digital Economy*. The World Bank, 2025. https://openknowledge.worldbank.org/server/api/core/bitstreams/9288bdc5-7a9b-42de-a47c-3746fd68f22a/content.

KommAustria. “Guidelines for the promotion of self-regulatory bodies for the protection of minors adopted.” September 17, 2021. https://www.rtr.at/medien/aktuelles/neuigkeiten/2021/News09172021selbskontrolle_schutzminderjaehriger.de.html.

Kumar, Ratnesh. “How to Change or Fake Location in Chrome, Edge, or Firefox.” *GeekChamp*, December 27, 2025. https://geekchamp.com/how-to-change-or-fake-location-in-chrome-edge-and-firefox/.

Lai, Katherine. “Age assurance and online safety: What parents and children have to say.” Internet Matters. April 11, 2025. https://www.internetmatters.org/hub/research/age-assurance-online-safety-parents-children-opinions/.

Lang, David, Benjamin Listyg, Brennah V. Ross, Anna V. Musquera, and Zeve Sanderson. "Do Age-Verification Bills Change Search Behavior? A Pre-Registered Synthetic Control Multiverse." Center for Social Media and Politics, March 3, 2025. https://csmapnyu.org/research/academic-research/do-age-verification-bills-change-search-behavior-a-pre-registered-synthetic-control-multiverse

Lenhart, Amanda, Mary Madden, Aaron Smith, Kristen Purcell, and Kathryn Zickuhr. “Part 3: Privacy and Safety Issues.” Pew Research Center. November 9, 2011. https://www.pewresearch.org/internet/2011/11/09/part-3-privacy-and-safety-issues/.

Livingstone, Sonia, Abhilash Nair, Mariya Stoilova, Simone van der Hof, and Cansu Cagla. “Children’s Rights and Online Age Assurance Systems.” *The International Journal of Children’s Rights 32* (2024): 721-747, https://brill.com/view/journals/chil/32/3/article-p721_009.xml.

Louisiana. “Act No. 481.” August 1, 2025. https://www.legis.la.gov/Legis/ViewDocument.aspx?d=1427667.

Lucchesi, Marisa. “A Family Affair: The Reality of Household Device Sharing.” Upwave. October 18, 2020. https://www.upwave.com/reality-household-device-sharing/.

Maryland Department of Transportation. “How to add your Maryland Mobile ID to Apple Wallet.” Video, January 2023. Posted by MDOT MVA. YouTube. https://www.youtube.com/watch?v=Xdns9qbzWmo.

Mathewson, Tara Garcia. “Schools Were Just Supposed to Block Porn. Instead They Sabotaged Homework and Censored Suicide Prevention Sites.” *The Markup,* April 13, 2024. https://themarkup.org/digital-book-banning/2024/04/13/schools-were-just-supposed-to-block-porn-instead-they-sabotaged-homework-and-censored-suicide-prevention-sites.

MaxMind. “Geolocation Accuracy.” Accessed January 22, 2026. https://support.maxmind.com/hc/en-us/articles/4407630607131-Geolocation-Accuracy.

McConvey, Joel R. “Email address age assurance is private, complaint, and simple: VerifyMy, Yoti.” *Biometric Update*, February 19, 2025. https://www.biometricupdate.com/202502/email-address-age-assurance-is-private-compliant-and-simple-verifymy-yoti.

McMahon, Liv. “VPNs top download charts as age verification law kicks in.” *BBC*, July 28, 2025. https://www.bbc.com/news/articles/cn72ydj70g5o.

MDN. “Web Authentication API.” Mozilla. Last modified October 9, 2025. https://developer.mozilla.org/en-US/docs/Web/API/Web_Authentication_API.

Merry, Krista, and Pete Bettinger. “Smartphone GPS accuracy study in an urban environment.” *PLoS One* 14, no. 7 (2018), https://journals.plos.org/plosone/article/file?id=10.1371/journal.pone.0219890&type=printable.

Meta. “About the in-app browser for Facebook and Instagram.” Business Help Center. Accessed January 22, 2026. https://www.facebook.com/business/help/206578174518231.

———. “Introducing New Ways to Verify Age on Instagram.” Newsroom. Accessed January 23, 2026. https://about.fb.com/news/2022/06/new-ways-to-verify-age-on-instagram/.

———. “Terms of Service.” Last modified January 1, 2025. https://www.facebook.com/legal/terms/update/.

———. “We’re Introducing New Built-In Restrictions for Instagram Teen Accounts, and Expanding to Facebook and Messenger.” Newsroom. April 8, 2025. https://about.fb.com/news/2025/04/introducing-new-built-in-restrictions-instagram-teen-accounts-expanding-facebook-messenger/.

———. “Working With Parents and New Technology to Enroll More Teens Into Teen Accounts.” Newsroom. Last modified September 21, 2025. https://about.fb.com/news/2025/04/meta-parents-new-technology-enroll-teens-teen-accounts/.

Microsoft. “List of Participants - Microsoft Trusted Root Program.” Microsoft Security. Last modified August 6, 2025. https://learn.microsoft.com/en-us/security/trusted-root/participants-list.

———. “Secure boot.” Windows Hardware Developer. Last modified February 8, 2023. https://learn.microsoft.com/en-us/windows-hardware/design/device-experiences/oem-secure-boot.

———. “What is Artifact Signing?” Learn. Last modified January 8, 2026. https://learn.microsoft.com/en-us/azure/trusted-signing/overview.

Minnesota. “SF 2105.” Office of Revisor of Statutes. March 4, 2025. https://www.revisor.mn.gov/bills/94/2025/0/SF/2105/versions/latest/.

Minocha, Shreyas, Isaac Sheridan, Harry Oppenheimer, Paul Pearce, and Michael A. Specter. “Papers, Please: A First Look at Age Verification on the Web.” Georgia Institute of Technology. In submission.

Mozilla. “Codecs used by WebRTC.” MDN. May 23, 2025. https://developer.mozilla.org/en-US/docs/Web/Media/Guides/Formats/WebRTC_codecs.

National Institute of Standards and Technology. “FIPS 186-5 Digital Signature Standard.” February 3, 2025. https://csrc.nist.gov/pubs/fips/186-5/final.

New York. “Stop Addictive Feeds Exploitation (SAFE) for Kids Act.” The New York State Senate. June 20, 2024. https://www.nysenate.gov/legislation/bills/2023/S7694/amendment/A.

Ngan, Mei, Patrick Grother, and Austin Hom. *Face Analysis Technology Evaluation (FATE) Part 10: Performance of Passive, Software-Based Presentation Attack Detection (PAD) Algorithms*. National Institute of Standards and Technology, September 2023. https://nvlpubs.nist.gov/nistpubs/ir/2023/NIST.IR.8491.pdf.

NordVPN. “Thousands of ultra-fast VPN servers across 178 locations.” Accessed January 22, 2026. https://nordvpn.com/offer/servers/.

Norman, Tré. “What Age Can You Get a Credit Card.” Edvisors. Last modified February 17, 2025. https://www.edvisors.com/credit-cards/credit-card-faqs/what-age-can-you-get-a-credit-card/.

North Carolina. “HB 301.” North Carolina General Assembly. March 5, 2025. https://www.ncleg.gov/BillLookup/2025/H301.

OBS Project. “OBS Studio.” Accessed January 23, 2026. https://obsproject.com/.

OECD. *Age Assurance Practices of 50 Online Services Used by Children*. OECD Publishing, June 2025. https://www.oecd.org/content/dam/oecd/en/publications/reports/2025/06/age-assurance-practices-of-50-online-services-used-by-children_3fbe5e92/a19853ab-en.pdf.

———. *How’s Life for Children in the Digital Age?* OECD Publishing, May 15, 2025. https://www.oecd.org/en/publications/how-s-life-for-children-in-the-digital-age_0854b900-en.html.

———. *PISA 2022 Results (Volume IV): How Financially Smart Are Students?* OECD Publishing, 2024. https://www.oecd.org/content/dam/oecd/en/publications/reports/2024/06/pisa-2022-results-volume-iv_125a58b3/5a849c2a-en.pdf.

Ofcom. “Age checks for online safety - what you need to know as a user.” June 26, 2025. https://www.ofcom.org.uk/online-safety/protecting-children/age-checks-for-online-safety--what-you-need-to-know-as-a-user.

———. “Children’s Online User Ages Quantitative Research Study.” Accessed January 22, 2026. https://www.ofcom.org.uk/siteassets/resources/documents/research-and-data/online-research/keeping-children-safe-online/childrens-online-user-ages/children-user-ages-chart-pack.pdf?v=328540.

———. “Investigation into 4chan and its compliance with duties to protect its users from illegal content.” Last modified December 4, 2025.

https://www.ofcom.org.uk/online-safety/illegal-and-harmful-content/investigation-into-4chan-and-its-compliance-with-duties-to-protect-its-users-from-illegal-content.

———. "Investigation into AVS Group Ltd's compliance with the duty to prevent children from encountering pornographic content through the use of age assurance." Last modified January 15, 2026. https://www.ofcom.org.uk/online-safety/protecting-children/investigation-into-avs-group-ltds-compliance-with-the-duty-to-prevent-children-from-encountering-pornographic-content-through-the-use-of-age-assurance.

———. "Investigation into the provider of xgroovy.com's compliance with the duty to prevent children from encountering pornographic content through the use of age assurance." Last modified November 20, 2025. https://www.ofcom.org.uk/online-safety/protecting-children/investigation-into-the-provider-of-xgroovy.coms-compliance-with-the-duty-to-prevent-children-from-encountering-pornographic-content-through-the-use-of-age-assurance.

———. "Protecting Children from Harms Online: Guidance on Content Harmful to Children." April 24, 2025. https://www.ofcom.org.uk/siteassets/resources/documents/consultations/category-1-10-weeks/statement-protecting-children-from-harms-online/main-document/guidance-on-content-harmful-to-children.pdf?v=395445.

———. "Quick guide to implementing highly effective age assurance." Last modified January 16, 2025. https://www.ofcom.org.uk/online-safety/illegal-and-harmful-content/age-assurance?url=https://www.ofcom.org.uk/online-safety/illegal-and-harmful-content/online-pornography&data=05%7c02%7cJohn.Eccleston@ofcom.org.uk%7c4682c653674e47cace2b08ddacdcf2da%7c0af648de310c40688ae4f9418bae24cc%7c0%7c0%7c638856787157089173%7cUnknown%7cTWFpbGZsb3d8eyJFbXB0eU1hcGkiOnRydWUsIlYiOiIwLjAuMDAwMCIsIlAiOiJXaW4zMiIsIkFOIjoiTWFpbCIsIldUIjoyfQ%3d%3d%7c0%7c%7c%7c&sdata=oTw%2bPsPihVYlZdR4lNwQ311vnlF9fAoonRN6WEJ1qkA%3d&reserved=0.

New York State Office of the Attorney General. "Notice of Proposed Rulemaking." September 15, 2025. https://ag.ny.gov/sites/default/files/regulatory-documents/safe-for-kids-act-nprm.pdf.

Oladipo, Oluwasegun, Elijah Olusayo Omidiora, and Victor Chukwudi Osamor. "Face Age Estimation and the Other-race Effect." *International Journal of Advanced Computer Science and Applications* 12, no. 11 (2021), https://thesai.org/Publications/ViewPaper?Volume=12&Issue=11&Code=IJACSA&SerialNo=24.

OneID. "How OneID's open banking-powered identity verification services help boost productivity for small businesses." June 4, 2025. https://www.openbanking.org.uk/insights/how-oneids-open-banking-powered-identity-verification-services-help-boost-productivity-for-small-businesses/.

Opanasets, Alex. "Almost a Quarter of Married Couples Didn't Have Joint Accounts in 2023, Up From 15% in 1996." United States Census Bureau. September 24, 2025. https://www.census.gov/library/stories/2025/09/married-but-separate.html.

OpenAge Initiative. "OpenAge." Accessed January 22, 2026. https://www.openageinitiative.org/.

OpenAI. “Our approach to age predictions.” January 20, 2026. https://openai.com/index/our-approach-to-age-prediction/.

———. “Updating our Model Spec with teen protections.” December 18, 2025. https://openai.com/index/updating-model-spec-with-teen-protections/.

Open Rights Group. “Content filtering by UK ISPs.” Accessed January 23, 2026. https://wiki.openrightsgroup.org/wiki/Content_filtering_by_UK_ISPs#Mobile_Broadband_Group_code_of_practice.

Panić, Nenad, Marina Marjanovic, and Timea Bezdan. “Addressing Demographic Bias in Age Estimation Models through Optimized Dataset Composition.” *Mathematics* 12, no. 15 (2024), https://www.mdpi.com/2227-7390/12/15/2358.

Papandrea, Dawn, and Julie Sherrier. “46% of Parents Say Their Child Used Their Credit or Debit Card Without Permission, Racking Up $500+.” LendingTree. March 1, 2022. https://www.lendingtree.com/credit-cards/study/kids-and-credit-cards-survey/.

Pedley, Kieran. “Britons back Online Safety Act’s age checks, but are sceptical of effectiveness and unwilling to share ID.” Ipsos. August 17, 2025. https://www.ipsos.com/en-uk/britons-back-online-safety-acts-age-checks-are-sceptical-effectiveness-and-unwilling-share-id.

Peters, Jay. “Discord customer service data breach leaks user info and scanned photo IDs.” *The Verge*, October 3, 2025. https://www.theverge.com/news/792032/discord-customer-service-data-breach-hack.

Pirogov, Viacheslav, and Maksim Artemev .“Evaluating Deepfake Detectors in the Wild.” Preprint, submitted July 29, 2025. https://arxiv.org/abs/2507.21905.

Privacy & Scaling Exploration. “Rate-Limiting Nullifier.” Accessed January 23, 2026. https://rate-limiting-nullifier.github.io/rln-docs/.

Promly. “Promly.” Accessed January 23, 2026. https://www.promly.org/.

Puc, Andrez, Vitomir Struc, and Klemen Grm. “Analysis of Race and Gender Bias in Deep Age Estimation Models.” *Proceedings of the 28th European Signal Processing Conference* (2021), https://ieeexplore.ieee.org/stamp/stamp.jsp?tp=&arnumber=9287219.

Reddit. “Why is Reddit asking for my age?” Last modified January 15, 2026. https://support.reddithelp.com/hc/en-us/articles/36429514849428-Why-is-Reddit-asking-for-my-age.

Ridley, Jacob. “Brits can get around Discord's age verification thanks to Death Stranding's photo mode, bypassing the measure introduced with the UK's Online Safety Act. We tried it and it works—thanks, Kojima.” *PC Gamer*, July 25, 2025. https://www.pcgamer.com/hardware/brits-can-get-around-discords-age-verification-thanks-to-death-strandings-photo-mode-bypassing-the-measure-introduced-with-the-uks-online-safety-act-we-tried-it-and-it-works-thanks-kojima/.

Roblox. “Roblox Requires Users Worldwide to Age-Check to Access Chat.” News Details. January 7, 2026. https://ir.roblox.com/news/news-details/2026/Roblox-Requires-Users-Worldwide-to-Age-Check-to-Access-Chat/default.aspx.

Rothschild, Jillian Andres, Samuel B. Novey, and Michael J. Hanmer. *Who Lacks ID in America Today? An Exploration of Voter ID Access, Barriers, and Knowledge.* Center for Democracy and Engagement, January 2024. https://cdce.umd.edu/sites/cdce.umd.edu/files/pubs/Voter%20ID%202023%20survey%20Key%20Results%20Jan%202024%20%281%29.pdf.

Saxon, James, and Nick Feamster. "GPS-Based Geolocation of Consumer IP Addresses." In *Passive and Active Measurement*, edited by Oliver Hohfeld, Giovane Moura, and Cristel Pelsser. Springer, 2022.

Silberling, Amanda, and Zack Whittaker. "TeaOnHer, a rival Tea app for men, is leaking users' personal data and driver's licenses." *TechCrunch*, August 6, 2025. https://techcrunch.com/2025/08/06/a-rival-tea-app-for-men-is-leaking-its-users-personal-data-and-drivers-licenses/.

Singh, Nitish. "Top iOS Location Changer Apps in 2025." *Geekflare*, January 8, 2025. https://geekflare.com/consumer-tech/best-ios-gps-location-changer-apps/.

Snap. "Implementing Australia's Social Media Minimum Age Law." Newsroom. November 22, 2025. https://newsroom.snap.com/australia-social-media-minimum-age-law.

Sommese, Raffaele, Anna Sperotto, Antonio Prado, Jeroen van der Ham, and Antonia Affinito. "Disrupting the Internet in the name of copyright: An Italian Story." Internet Engineering Taskforce. November 6. 2025. https://datatracker.ietf.org/meeting/124/materials/slides-124-iabopen-piracy-shield-talk-00.

Spur. "Advanced detection of anonymization and threats." Accessed January 22, 2026. https://spur.us/.

Stapelberg, Alan. "It's now easier to prove age and identity with Google Wallet." Google. *The Keyword,* April 29. 2025. https://blog.google/products/google-pay/google-wallet-age-identity-verifications/.

Starmer, Keir. "A bold new mission." Substack. December 18, 2025. https://keirstarmer.substack.com/p/a-bold-new-mission.

Statcounter. "Android Version Market Share Worldwide." December 2025. https://gs.statcounter.com/android-version-market-share.

———. "Desktop Browser Market Share Worldwide." Accessed January 23, 2026. https://gs.statcounter.com/browser-market-share/desktop/worldwide/#monthly-202407-202507.

———. "Mobile & Tablet iOS Version Market Share Worldwide." Last modified January 19, 2026. https://gs.statcounter.com/os-version-market-share/ios/mobile-tablet/worldwide.

Stockwell, Sam, and Rosamund Powell. *Age Assurance Technologies and Online Safety*. Centre for Emerging Technology and Security, October 8, 2025. https://cetas.turing.ac.uk/publications/age-assurance-technologies-and-online-safety.

Supreme Court of the United States. "Ashcroft v. ACLU." June 29, 2004. https://www.courtlistener.com/opinion/137005/ashcroft-v-american-civil-liberties-union/.

———. "Free Speech Coalition v. Paxton." June 27, 2025. https://www.supremecourt.gov/opinions/24pdf/23-1122_3e04.pdf.

———. "Reno v. ACLU." Courtlistener. May 19, 1997. https://www.courtlistener.com/opinion/118147/reno-v-american-civil-liberties-union/.

Swapface. “Easy-to-use Faceswap AI tool in the world!” Accessed January 23, 2026. https://www.swapface.org/.

Sweeney, Latanya. “k-anonymity: a model for protecting privacy.” *International Journal of Uncertainty, Fuzziness, and Knowledge-Based Systems* 10, no. 5 (2002): 557-570, https://dl.acm.org/doi/10.1142/S0218488502001648.

Tenbarge, Kat. “Fewer than 1% of parents use social media tools to monitor their children’s accounts, tech companies say.” *NBC News*, March 29, 2024. https://www.nbcnews.com/tech/social-media/fewer-1-parents-use-social-media-tools-monitor-childrens-accounts-tech-rcna145592.

Texas. “App Store Accountability Act.” May 27, 2025. https://capitol.texas.gov/tlodocs/89R/billtext/pdf/SB02420F.pdf.

Thaler, Justin. “Proofs, Arguments, and Zero-Knowledge.” *Foundations and Trends in Privacy and Security 4*, no. 2-4 (2022): 117-660, https://par.nsf.gov/servlets/purl/10417638.

Thurman, Neil, and Fabian Obster. “The regulation of internet pornography: What a survey of under-18s tells us about the necessity for and potential efficacy of emerging legislative approaches.” *Policy & Internet* 13, no. 3 (2021): 415-432, https://onlinelibrary.wiley.com/doi/full/10.1002/poi3.250.

TikTok. “An update on our work to provide teens with age appropriate experiences.” News. October 8, 2025. https://newsroom.tiktok.com/an-update-on-our-work-to-provide-teens-with-age-appropriate-experiences?lang=en-150.

Tinder. “FAQ Mandatory Liveness Check.” Accessed January 21, 2026. https://policies.tinder.com/faq-mandatory-liveness-check/intl/en/.

———. “ID + Photo Verification.” Accessed January 21, 2026. https://www.help.tinder.com/hc/en-us/articles/19868368795917-ID-Photo-Verification.

Tor. “Browse Privately. Explore freely.” Accessed January 22, 2026. https://www.torproject.org/.

Trotman, Rachael. “Introducing Yoti Keys: privacy-focused, seamless and anonymous age verification.” Yoti. April 8, 2025. https://www.yoti.com/blog/yoti-keys-private-seamless-anonymous-age-verification/.

TrustCloud. “The AU10TIX case: millions of records exposed in a security breach affecting major apps.” July 11, 2024. https://trustcloud.tech/blog/au10tix-case-records-exposed-security-breach-major-apps/.

Tutor, Annie Chestnut. *Age Verification: What It Is, Why It Is Necessary, and How to Achieve It*. Center for Technology and the Human Person at the Heritage Foundation, March 6, 2025. https://www.heritage.org/sites/default/files/2025-03/BG3895.pdf.

———. *Parents’ Survey: Online Filters and Blocking Software Still Only Work Sometimes*. Center for Technology and the Human Person at the Heritage Foundation, October 29, 2025. https://www.heritage.org/sites/default/files/2025-10/BG3940.pdf.

Parliament of the United Kingdom. “Children’s Wellbeing and Schools Bill: Running List of All Amendments on Report.” December 9, 2025. https://bills.parliament.uk/publications/63901/documents/7465.

———. “Online Safety Act 2023: Virtual Private Networks.” September 15, 2025. https://hansard.parliament.uk/lords/2025-09-15/debates/57714CE6-0CE4-49F6-B028-E271D5100F7F/OnlineSafetyAct2023VirtualPrivateNetworks.

United States. “18 U.S. Code § 2721 - Prohibition on release and use of certain personal information from State motor vehicle records.” 1994. https://www.law.cornell.edu/uscode/text/18/2721.

———. “Child Online Protection Act.” October 21, 1998. https://archive.epic.org/free_speech/censorship/copa.html.

———. “Children’s Online Privacy Protection Act.” October 21, 1998. https://uscode.house.gov/view.xhtml?req=granuleid%3AUSC-prelim-title15-section6501&edition=prelim.

———. “Communications Decency Act.” February 1, 1996. https://archive.epic.org/free_speech/cda/cda.html.

———. “S.737 - SCREEN Act.” February 26, 2025. https://www.congress.gov/bill/119th-congress/senate-bill/737/text.

———. “S.2714 - CHAT Act.” September 4, 2025. https://www.congress.gov/bill/119th-congress/senate-bill/2714.

———. “S.3062 - GUARD Act.” October 28, 2025. https://www.congress.gov/bill/119th-congress/senate-bill/3062/text,

United States Department of Homeland Security. “REAL ID Frequently Asked Questions.” Accessed January 23, 2026. https://www.tsa.gov/real-id/real-id-faqs.

United States District Court for the Northern District of California. “Order Granting in Part and Denying in Part Motion for Preliminary Injunction.” December 31, 2024. https://oag.ca.gov/system/files/attachments/press-docs/Order%20on%20Preliminary%20Injunction.pdf.

———. “Plaintiffs’ Second Amended Master Complaint.” November 14, 2023. https://www.motleyrice.com/sites/default/files/documents/social_media_addiction-redacted_master_complaint.pdf.

United States District for the Western District of Arkansas. “Memorandum Opinion and Order.” December 15, 2025. https://netchoice.org/wp-content/uploads/2025/12/47-Order-Granting-Motion-for-Preliminary-Injunction.pdf.

Utah. “App Store Accountability Act.” Utah State Legislature. March 26, 2025. https://le.utah.gov/~2025/bills/static/SB0142.html.

Valev, Neven, Franziska Bieri, and Menna Bizuneh. “Percent people with credit cards - Country rankings.” The Global Economy. Accessed January 23, 2026. https://www.theglobaleconomy.com/rankings/people_with_credit_cards/.

Vallance, Chris. “4chan launches legal action against Ofcom in US.” *BBC*, August 27, 2025. https://www.bbc.com/news/articles/clyjq40vjl7o.

Vekaria, Yash, Yohan Beugin, Shaoor Munir, et al. “SoK: Advances and Open Problems in Web Tracking.” Preprint, submitted June 16, 2025. https://arxiv.org/pdf/2506.14057.

VerifyMy. *Innovative age assurance: Email address as the new benchmark for fricitonless age estimation*. VerifyMy, June 2024.

https://verifymy.io/wp-content/uploads/2024/11/Verifymy-White-Paper-Innovative-age-assurance-Email-address-as-the-new-benchmark-for-frictionless-age-estimation.pdf.

———. "Privacy Policy." Last modified November 2025. https://verifymy.io/age-verification-and-estimation/age-assurance-privacy-policy/#elementor-toc__heading-anchor-7.

Vrachnos, Athanasios, Evangelia Papadaki, Panagiota Lagou, Nikolaos Soumelidis, and Eirini Papamichail. *Remote ID Proofing Good Practices*. ENISA, March 2024. https://www.enisa.europa.eu/sites/default/files/2024-11/Remote%20ID%20Proofing%20Good%20Practices_en_0.pdf.

Wilson, Cam. "1 in 3 parents will help kids get around teen social media ban, government privately warned." *Crikey*, October 20, 2025. https://www.crikey.com.au/2025/10/20/1-in-3-parents-will-help-kids-get-around-teen-social-media-ban/.

———. "How Australian teens are already planning to dodge the social media ban." *Crikey*, October 10, 2025. https://www.crikey.com.au/2025/10/10/teen-social-media-ban-workarounds/.

Wise, Alana. "Tea encouraged its users to spill. Then the app's data got leaked." *NPR*, August 2, 2025. https://www.npr.org/2025/08/02/nx-s1-5483886/tea-app-breach-hacked-whisper-networks.

Wiser, Ben, Borbala Benko, Philipp Pfeiffenberger, and Sergey Kataev. "Web Environment Integrity Explainer." Google. Last modified December 3, 2024. https://github.com/explainers-by-googlers/Web-Environment-Integrity/blob/main/explainer.md.

Wrinn, Corey, and Bob Savvy. "Youth Accounts Map a Promising Path Forward for Banking Providers." *The Financial Brand*, March 11, 2025. https://thefinancialbrand.com/news/financial-education/the-kids-are-alright-youth-accounts-map-a-promising-path-forward-187314.

Wu, Mingshi, Jackson Sippe, Danesh Sivakumar, et al. "How the Great Firewall of China Detects and Blocks Fully Encrypted Traffic." *Proceedings of the 32nd USENIX Security Symposium* (2023): 2653-2670, https://www.usenix.org/conference/usenixsecurity23/presentation/wu-mingshi.

Wukovits, Nora. *Transposition of the 2018 Audiovisual Media Services Directive: Implementation in Action*. European Parliamentary Research Service, October 2022. https://www.europarl.europa.eu/RegData/etudes/IDAN/2022/730354/EPRS_IDA(2022)730354_EN.pdf.

Yimeng, Zhao. "Parents help children dodge time limits on online games." *China Daily*, August 20, 2024. https://global.chinadaily.com.cn/a/202408/20/WS66c3e9a1a31060630b923e74.html.

Yoti. "Adult Content Age Verification." Accessed January 22, 2026. https://www.yoti.com/adult-content-age-verification/.

———. "Age Verification." Accessed January 23, 2026. https://www.yoti.com/business/age-verification/.

———. "Secure Image Capture." Accessed January 23, 2026. https://developers.yoti.com/age-verification/secure-image-capture.

———. "Tokens (Yoti Key)." Accessed January 23, 2026. https://developers.yoti.com/age-verification/age-token?redirect_from=%2Fage-verification%2Fage-tokens.

———. “Yoti Age Verification Service - Privacy Notice.” Last modified September 15, 2025. https://www.yoti.com/privacy/age-verification/.

———. “Yoti Developer Documentation.” Accessed January 23, 2026. https://developers.yoti.com/.

———. “Yoti Facial Age Estimation.” July 2025. https://cdn.aws.yoti.com/wp-content/uploads/2026/01/Yoti-Age-Estimation-White-Paper-July-2025-PUBLIC-v1.pdf.

———. “Yoti ID is your secure Digital ID.” Accessed January 24, 2026. https://www.yoti.com/personal/.

———. ‘You’re in safe hands.” Accessed January 23, 2026. https://www.yoti.com/security/.

YouTube. “Building content recommendations to meet the unique needs of teens and pre-teens.” Accessed January 23, 2026. https://services.google.com/fh/files/misc/vibe-expansion-gtm-casestudy.pdf.

Yubo. “Staying Safe on Yubo: A Guide for Parents, Carers, & Educators.” Accessed January 21, 2026. https://yubo.cdn.prismic.io/yubo/45c05356-1a36-439b-905a-e2c112d93673_Yubo-SafetyGuide-a-guide-for-parents-carers-educators.pdf.

Yubo Team. “How Yubo Pioneered 100% Age Verification to Set a New Standard for Trust & Safety on Social Media.” Accessed January 23, 2026. https://www.yubo.live/blog/how-yubo-pioneered-100-percent-age-verification.

Yürüten, Onur. “VoiceAssure: A ROBUST, PRIVACY-FIRST, VOICE-BASED AGE ESTIMATION TECHNOLOGY.” Medium. April 3, 2023. https://medium.com/@onuryrten/voiceassure-a-robust-privacy-first-voice-based-age-estimation-technology-f9c9d6c8340d.

Zendle, David, Catherine Flick, Elena Gordon-Petrovskaya, Nick Ballou, Leon Y. Xiao, and Anders Drachen. “No evidence that Chinese playtime mandates reduced heavy gaming in one segment of the video games industry.” *Nature Human Behavior* 7 (October 2023): 1753-1766, https://www.nature.com/articles/s41562-023-01669-8.pdf.