
In the Supreme Court of the United States

Free Speech Coalition, et al.,

Petitioners,

v.

Ken Paxton, Attorney General of Texas,

Respondent.

On Writ of Certiorari
to the United States Court of Appeals
for the Fifth Circuit

**BRIEF OF *AMICUS CURIAE* YOTI LTD.
IN SUPPORT OF RESPONDENT**

JAMES R. MARSH
Counsel of Record
Marsh Law Firm PLLC
31 Hudson Yards, 11th Floor
New York, New York 10001
212–372–3030
jamesmarsh@marsh.law

Counsel for Amicus Curiae Yoti Ltd.

TABLE OF CONTENTS

TABLE OF AUTHORITIES	iii
INTEREST OF AMICUS CURIAE	1
SUMMARY OF THE ARGUMENT	3
ARGUMENT	4
I. Age Verification Technology is the Least Restrictive Alternative and is Narrowly Tailored to Achieve Texas’s Well-Founded Goal of Protecting Minors from Obscene Content on the Internet	4
A. Age Verification Technology Does Not Impose Categorically Different Burdens on Adults and is Akin to In-Person Age Verification.....	6
1. Estimating age from a user’s facial image using AI facial age estimation technology	8
2. Document-based verification using a physical identity document	12
3. Checking against databases, credit reference agencies, mobile phone operators, and other semi-public data sources	13
4. Reusable digital identity apps	14

5. Age assurance tokens in which
validation on one website is
transferrable to another website 14

B. Age Verification Technology Does Not
Unreasonably Implicate or Intrude
Upon the Privacy of Adults 16

C. Age Verification Technology is
Currently being Successfully Utilized
Worldwide 16

D. Texas Can Impose a Content-Based
Speech Restriction on Sexually Explicit
Materials on the Internet Even if it
Results in a Cost to the Websites and
Platforms Hosting the Content 18

CONCLUSION..... 21

TABLE OF AUTHORITIES

CASES

<i>Ashcroft v. ACLU (Ashcroft II),</i> 542 U.S. 656 (2004)	3
<i>Forsyth Cnty., Ga. v. Nationalist Movement,</i> 505 U.S. 123 (1992)	19
<i>iMatter Utah v. Njord,</i> 774 F.3d 1258 (10th Cir. 2014)	19
<i>Paris Adult Theatre I v. Slaton</i> 413 U.S. 49 (1973)	20, 21
<i>Rosenberger v. Rector & Visitors of Univ. of Virginia,</i> 515 U.S. 819 (1995)	19
<i>Surita v. Hyde,</i> 665 F.3d 860 (7th Cir. 2011)	19

INTEREST OF AMICUS CURIAE¹

Amicus curiae is Yoti Ltd., a London-based technology company founded in 2014, and the leading global provider of age assurance technology, undertaking tens of millions of age verification checks monthly for many of the largest global companies such as Meta, Tik Tok, Pinterest, Sony Playstation, Kids Web Services (part of Epic Games), XHamster, NCR, Diebold, PMI, BAT. Yoti's age assurance technologies include artificial intelligence facial age estimation which has been independently reviewed and approved by governmental regulatory bodies in Germany (KJM, FSM) and the United Kingdom (Age Check Certification Scheme). Yoti has undergone a facial age estimation algorithm benchmarking review by the United States National Institute of Standards and Technology (NIST)² and is listed in the highest-ranking participants.

Since 2016, Yoti has been an international leader in developing age verification standards, contributing significantly to the first publicly available standard for age checking, the PAS 1296:2018 Online Age Checking Provision and use of Online Age Check Service Code of Practice,³ and the subsequent international standards groups with the IEEE

¹ Pursuant to Rule 37.6, amicus affirms that no counsel for a party authored this brief in whole or in part, and no person other than amicus and their counsel made a monetary contribution to its preparation or submission.

² https://pages.nist.gov/frvt/html/frvt_age_estimation.html

³ <https://knowledge.bsigroup.com/products/online-age-checking-provision-and-use-of-online-age-check-services-code-of-practice?version=standard>

Standard for Online Age Verification⁴ and ISO Age Assurance System Framework.⁵ Yoti is a member of WeProtect Global Alliance,⁶ the Online Safety Tech Industry Association (OSTIA),⁷ the Age Verification Providers Association (AVPA)⁸ and has taken part in the euCONSENT Project⁹ to develop interoperable age verification and parental consent mechanisms to meet European laws on data protection, audio-visual media services, and the United Kingdom’s recently passed Online Safety Act¹⁰ and European Union Digital Services Act.¹¹

The issue before the Court is of interest to amicus Yoti Ltd. because of the profound impact this case will have on the development of age verification and child safety measures in the United States.

⁴ <https://standards.ieee.org/ieee/2089.1/10700/>

⁵ [ISO/IEC 27566 Information security, cybersecurity and privacy protection — Age assurance systems — Framework Part 1: Framework](#)

⁶ <https://www.weprotect.org/>

⁷ <https://ostia.org.uk/>

⁸ <https://avpassociation.com/>

⁹ <https://euconsent.eu/>

¹⁰ <https://www.gov.uk/government/publications/online-safety-act-explainer/online-safety-act-explainer>

¹¹ <https://digital-strategy.ec.europa.eu/en/faqs/digital-services-act-questions-and-answers>

SUMMARY OF THE ARGUMENT

In *Ashcroft v. ACLU (Ashcroft II)*, 542 U.S. 656 (2004), the Court considered a content-based speech restriction on sexually explicit materials on the Internet that Congress deemed harmful to minors. In considering this question, the Court assumed that certain protected speech may be regulated, with the key consideration being determining the least restrictive alternative that can be used to achieve that goal. In *Ashcroft II*, blocking and filtering software was found to be the least restrictive alternative. In this case, age verification technology is proposed as the least restrictive alternative to achieving Texas's uncontested goal of protecting minors from obscene content on the Internet.

The Court in *Ashcroft II* conducted an in-depth analysis of the viability and effectiveness of filtering software circa February 1999—over 25 years ago. The Court found that this technological solution was the least restrictive means of regulating Internet obscenity, with the crucial caveat that “it is reasonable to assume that other technological developments important to the First Amendment analysis have also occurred during that time.” *Id.* at 671.

This amicus brief is about “technological developments important to the First Amendment analysis,” acknowledging the Court’s finding that “the technology of the Internet evolves at a rapid pace.” *Id.*

In this case, the District Court soundly rejected age verification technology in favor of content filtering software, while the Fifth Circuit held that Texas’s

“age verification requirement likely passes constitutional muster under the rational-basis standard in *Ginsberg*.” Pet.App.27a.

Understanding age verification technologies is essential to deciding whether their application is the least restrictive alternative under the First Amendment.

ARGUMENT

I. Age Verification Technology is the Least Restrictive Alternative and is Narrowly Tailored to Achieve Texas’s Well-Founded Goal of Protecting Minors from Obscene Content on the Internet

This brief outlines practical examples of how current age verification technology (also known as age assurance technology) is working today, focusing on artificial intelligence (AI) driven facial age estimation. AI facial age estimation is the most popular age verification technology with individuals and the most prevalent approach by volume. It does not require using an identity document or other identity data, nor does it require using facial recognition technology. No database is created; no images are stored. The algorithm does not learn anything from each additional age check undertaken.

Yoti Ltd., a worldwide leader in age verification, offers a wide range of approaches, and has undertaken over 700 million age checks. Yoti is therefore uniquely qualified to explain the ease of use, accuracy, affordability, security, effectiveness, and accessibility, including details on how AI-driven facial age

estimation meets rigorous data protection requirements to preserve privacy. This brief details examples of how and where this low-friction technology is already operating globally at scale with leading web platforms.

Every month, AI-driven facial age estimation technology already enables tens of millions of adults to verify their age in about one second. The facial image is not stored and there is no reliance on any identity documentation. This technology enables platforms to easily distinguish an adult from a minor at a global scale and in the volumes required for global platforms such as Meta.

With the user's consent, Yoti can store the results of a previously completed facial age estimation on the user's device as a reusable digital ID or wallet. This digital ID or wallet can then be utilized to access any website containing adult content.

Taking this technology one step further, reusable approaches, termed 'tokens,' are now available so repeat visitors to an age-gated website, such as adult content websites, can prove their age once and then continue to use the generated age verification token indefinitely. The French government is the first government to issue guidelines authorizing the use of tokens as an appropriate method of age verification, beginning when age verification for adult content comes into force in France in January 2025.¹²

¹² <https://www.arcom.fr/nos-ressources/espace-juridique/textes-juridiques/referentiel-technique-sur-la-verification-de-lage-pour-la-protection-des-mineurs-contre-la-pornographie-en-ligne>

There is a range of fallback options and exciting new possibilities for how visitors to age-gated websites can prove their age and only share their adult status, while ensuring privacy and security across platforms and websites.

A. Age Verification Technology Does Not Impose Categorically Different Burdens on Adults and is Akin to In-Person Age Verification

At a bar, bartenders routinely inspect an individual and, based on an initial estimation of their age, decide whether to request a physical identification document such as a driver's license. To complete manual age verification, a government-issued identity document needs to be reviewed to determine whether it is a bona fide document, whether the image of the person in the document matches the person presenting it, and check the date of birth to ascertain if the person is old enough to consume alcohol.

In today's online world, platforms are already harnessing a range of technologies to assess age, frequently working with independent age verification providers. Platforms can select a range of age verification options to offer customers a choice and can make their own determination about what range of options suits their customers, accounting for different markets, regulatory requirements, and demographics.

One technological method emulates the human ability to estimate age, as demonstrated by the bartender example of someone with years of experience discerning adults from minors – this is facial age estimate.

Some age verification methods build on the physical review of the document model by technologically checking the document's authenticity and then checking that the person presenting it is, in fact, the rightful owner.

This technology typically utilizes ID document databases to determine if the format of the document being presented is authentic, liveness detection to ensure that the person presenting the document is alive, and face matching to compare the person presenting the document to the face on the document. In addition, depending on the circumstances, it may be necessary to verify individuals using data points from authoritative sources such as department of motor vehicle databases, voter registration records, or credit reference databases.

It is common for visitors of adult content websites to visit numerous sites over several hours. Since it is inconvenient and excessive for such users to undergo age verification every time they access adult websites, the age verification technology industry has developed reusable tokens that allow users to repeatedly visit adult content sites without undergoing repeated age verification checks.

The typical adult content site consists of an initial 'landing page' that can offer several types of age verification options, enabling the consumer to pick the one they prefer. Some methods may involve just one or two clicks and take just one or two seconds, equivalent to or less than the time it takes a bartender to do a visual estimation and ask an individual to dig out their ID from their wallet, hand it over, and present it for inspection.

Current age verification methods include:

- 1. Estimating age from a user's facial image using AI facial age estimation technology**

In everyday life, we are all human age estimators; we subconsciously assess the ages of those we encounter, drawing on our lived experience. Consider a long-tenured high school principal. His experience dealing with teenagers would make him more accurate than the average adult at differentiating those under 18 from those over 18.

The most developed age estimation approach, which has been around for over a decade, utilizes a facial image. This image-based technology was initially reviewed by NIST over a decade ago and is now part of an ongoing NIST benchmarking assessment¹³ and is currently under review by nationally accredited audit bodies such as the Age Check Certification Scheme (ACCS)¹⁴ with well-established companies such as Google, Fujitsu, Idemia providing this service. Based on NIST data, these models are likely to be accurate within one to three years of the individual's age, depending on the age range. Specifically, the accuracy for younger individuals is approximately 1 to 1.5 years, whereas the accuracy for older individuals is around 2-4 years. There is higher accuracy for younger versus older people given that physical appearance is impacted over time in relation to the

¹³ https://pages.nist.gov/frvt/html/frvt_age_estimation.html

¹⁴ See <https://accscheme.com/> and <https://accscheme.com/registry/>

individual's lived experience. The error rate increases with age, as well.

When a visitor arrives on an adult website landing page, the request to prove age may appear in a new window or a pop-up frame. Most adult content users access websites using mobile devices and are asked to take a photo using their phone's camera. Desktop users can do the same via their laptop webcam.

Facial age estimation technology determines someone's age using a facial image. The user simply looks at the camera on their device, positions their face in the oval shape provided, and takes their photo. The image is then analyzed by an algorithm trained to determine age by analyzing facial features from a diverse training set of images where the actual age was known in months and years.

To the software model, the new image presented is simply a pattern of pixels, and the pixels are numbers. Facial age estimation technology has been trained to spot patterns in numbers and pinpoint how people of a particular age appear in those number patterns. The facial age estimation process usually takes less than one second. It can produce either a yes/no result on whether the individual in the image meets a designated age threshold or an estimated age. The image is assessed directly and there is no need to store it, so nothing is stored, ensuring the visitor's privacy. If the visitor meets the age required (e.g., 18+), the age estimation technology confirms to the adult website that the user can proceed, and the user can then access the adult content on the website.

A website operator can require an ‘age buffer’ in years above the age of access. For example, the website operator could decide that if the age of access is 18 years and a method is accurate to within 1 to 1.5 years, then this method can be used by all visitors who are, for instance, two, three, or five years older than the necessary age, translating to 20, 21 or 23 years old. Visitors within this buffer—below 20 in this example—will need to use an additional age verification check to help pinpoint their age.

In Yoti’s publicly accessible white papers, we regularly publish the mean absolute errors, false positives, and false negatives for each age group across gender and skin tone. In our most recent white paper, published in September 2024, the Yoti algorithm produces a true positive rate for thirteen- to seventeen-year-old minors as correctly estimated as under age twenty-one 99.2% of the time. The false positives rate reveals that under 1% of fourteen to seventeen-year-olds incorrectly judged to be over age twenty-one as 0.03% for fourteen-year-olds, 0.34% for fifteen-year-olds, 0.73% for sixteen-year-olds and 2.43% of seventeen-year-olds.¹⁵

Using modern age verification technology, all adults above age twenty or twenty-one can easily and securely assert their adult status in a privacy-ensured way in around one second. This method is highly accurate against different ages and skin tones, where the technology is consistently trained at scale and

¹⁵ <https://www.yoti.com/wp-content/uploads/2024/11/Yoti-Age-Estimation-White-Paper-September-2024-PUBLIC.pdf> page 31.

using artificial intelligence is constantly improving over time.

To achieve a high level of assurance, facial age estimation should always be used in combination with liveness detection— including both presentation and injection attack detection— and be subject to continuous independent testing and monitoring. Liveness checks prevent people from using a photo, video, or mask instead of taking a live photo. Facial injection attack detection ensures that the image being analyzed originates from a live source, such as the device camera, and not from a fraudulent or pre-recorded input. This technique verifies the authenticity of the face by distinguishing between real-time facial data and injected, manipulated, or replayed images, helping to prevent spoofing or unauthorized access.

With the rise of generative AI technology and deepfakes to circumvent age verification checks, the ability to prevent direct and indirect injection attacks is also required.¹⁶ Website operators must identify and prevent sophisticated injection attacks, ensuring that the images captured during a verification process are genuine and remain untampered with.¹⁷

The ongoing large scale NIST evaluation uses over 11 million data points and provides scientific certainty for businesses and governments that facial age estimation is an accurate, fair, and privacy-preserving

¹⁶ <https://www.yoti.com/wp-content/uploads/2024/03/Yoti-How-to-combat-Generative-AI-and-deepfakes-white-paper.pdf>

¹⁷ <https://www.yoti.com/blog/yoti-releases-white-paper-detailing-approach-to-combating-generative-ai-and-deepfakes/>

age assurance solution.¹⁸ Audit bodies such as ACCS have been auditing facial age estimation since 2020. They also provide testing, using high-quality test images captured from mobile phones, to provide testing conditions closer to those of adult content and social media companies.

Facial age estimation has been approved by a range of regulators as a method for age verification to access adult content: the German regulators (KJM and FSM) have been using it since 2021 with a 5-year buffer. The French content and data protection regulators (Arcom and CNIL) and the UK data protection and content regulators (ICO and OFCOM) clearly state it as an appropriate method.

2. Document-based verification using a physical identity document

For individuals within a defined buffer zone—for example between the ages of eighteen and twenty-one—other forms of verification, such as document-based verification, can be utilized as a fallback or supplemental method of establishing age.

Document-based verification requires assurances that the actual owner of the document is the person providing the document. So, in the same way a thorough bartender will carefully compare the document to the person presenting it, effective online document-based age verification requires document authenticity checks, face matching, and liveness verification. Using these three steps, website

¹⁸ https://pages.nist.gov/frvt/reports/aev/fate_aev_report.pdf

operators have a high level of assurance that the document is authentic, belongs to the right person, and that the document owner is the person presenting the document at that discrete moment in time. Document-based age verification technology and human analysts in Yoti's 24-hour data security centers review documents to detect a range of fraud vectors, including fake, tampered-with, borrowed, or counterfeit documents.

Sharing an identity document might be perceived as disclosing more individually identifying details, such as name, nationality, and address, alongside a live facial scan; however, in actuality, just the anonymized age attribute or "age over" such as "18" is shared with the adult content site.

The Yoti terms of service and industry best practices, along with legal and regulatory requirements in the UK and EU, require that only the age result is transferred and that any personal data captured as part of a one-time upload of an identity document verification check is promptly deleted after the age verification check is completed.

3. Checking against databases, credit reference agencies, mobile phone operators, and other semi-public data sources

There are a number of techniques that have long been used for Know Your Customer checks in financial services and gambling, and as in the example above, they can also be utilized to determine whether a user is an adult or a minor.

4. Reusable digital identity apps

Repeatedly undergoing age verification checks is not only tiresome, but it can limit users' engagement and access to legally accessible content. Especially for users close in age to legal adulthood, a more permanent method of assuring age is possible by using a digital identity app that allows them to establish their age once and then utilize this proof of age in many different contexts, on many different websites and platforms.

In addition to storing document-based verification, a digital ID app can utilize facial age estimation to validate age once and then share that verification when required across the Internet. In the physical world, a location might allow visitors seeking access to an adult venue to scan a QR code and share a verified age estimation via a reusable digital identity app.

In the online world, a digital identity app simply certifies that the visitor is an adult through a simple "yes" or "no" check. Nothing further is shared with the website operator, and no personal information is revealed, ensuring complete privacy for the user.

5. Age assurance tokens in which validation on one website is transferrable to another website

All the above age assurance options can be independently assessed, accredited, and tokenized. Websites can determine which methods they accept, whether to tokenize, and the duration of the tokens.

Using tokens decreases the costs for relying parties and can further reduce any burden by the visitor to keep verifying their age at each adult site they visit.

Age assurance tokens allow a visitor to determine their age once using any of the methods above and then use that same verification repeatedly across websites. An anonymous cookie is added to their browser which is then accessed by subsequent websites. Tokens don't contain any identifiable information, just the age verification result and information about how the check was performed. Tokens can be stored in a confidential age verification account, cached, and utilized over and over again indefinitely.

Users can pass their age tokens to another browser or, after clearing their cache, by logging on to their age verification account. Once cached or re-cached, users can freely visit other sites that accept the token criteria. If a user visits a site with different criteria, such as a higher age threshold or one requiring a more recently issued token, the user will be asked to undertake a new age verification.

The information shared with the relying party is the sum of the token's information, as detailed above. This comprises the method of age assurance used, the type of age recorded—such as age over, age under, or age range—the issuer of the age token, the type of liveness check performed, the type of authenticity check performed, and the time and date the check was performed.


B. Age Verification Technology Does Not Unreasonably Implicate or Intrude Upon the Privacy of Adults

Using modern age verification technology, as soon as a visitor's age is estimated, their facial image is deleted, ensuring their privacy at all times. Visitors do not need to provide personal details like their name or date of birth or share any identity documents. The technology is specifically designed not to identify anyone. Unlike facial-recognition technology, age verification technology estimates someone's age without identity matching that person or acquiring any details from them other than the facial image which is instantly assessed and then is no longer needed, and therefore does not need to be stored.¹⁹

C. Age Verification Technology is Currently being Successfully Utilized Worldwide

Many of the largest and most successful social media, adult content websites, and gaming platforms already use age verification technology, with millions of checks being conducted monthly.

¹⁹ https://ico.org.uk/media/for-organisations/documents/4020427/yoti-sandbox-exit-report_20220522.pdf



Examples include:

Meta, including Facebook Dating, Instagram, and, more recently, Instagram Teen Accounts, utilizes age verification technology to support age-appropriate access to these services. When Facebook Dating and Instagram detect a user attempting to change their date of birth from under 18 to over 18, they deploy Yoti's facial age estimation as one of the ways the person can establish that they are an adult. This tool has been deployed by Meta since 2022. Meta recently introduced Instagram Teen Accounts with built-in default protections limiting who can contact and the content teen users can access. Teens under sixteen need an adult's permission to use less protective settings. For teen users who want to switch to an adult account, Instagram uses Yoti facial age estimation, among other tools, to establish eligibility for such an account.

OnlyFans, a subscription-based social media platform for adults where both creators and subscribers must be over eighteen, deploys age verification technology in many countries to ensure that those engaging with the platform are adults.

Yubo, a live-streaming social media platform with almost 60 million users, offers a service for adults and a separate service for thirteen to seventeen-year-olds. Since 2022, Yubo has been using facial age estimation as its first point of contact to ensure that users are old enough and provided access to the proper content. Yubo employs a fallback option for users close to age 18, such as using the Yoti reusable digital ID app or a one-time document submission.

D. Texas Can Impose a Content-Based Speech Restriction on Sexually Explicit Materials on the Internet Even if it Results in a Cost to the Websites and Platforms Hosting the Content

Cost is an understated but important issue in deciding the least restrictive means of regulating Internet obscenity. The end user almost always bears the cost of blocking and filtering software, while the website profiting from obscenity and adult content almost always bears the cost of age verification technology. Cost-shifting refers to allocating a policy or practice's

expenses or financial burdens to one party while allowing others to shirk or reduce these costs.

Under the First Amendment, the government cannot impose a financial burden on speakers based on the content of their speech. *Rosenberger v. Rector & Visitors of Univ. of Virginia*, 515 U.S. 819, 828 (1995). However, imposing a cost is permissible if it is “designed to meet the expenses incident to the administration of the law.” *Forsyth Cnty., Ga. v. Nationalist Movement*, 505 U.S. 123, 129 n. 8 (1992).

One example is requiring fees and paid permits for groups and individuals exercising their First Amendment right to protest. Concerning permit fees, the Supreme Court has ruled that costs must be reasonable and not prohibitively expensive, that they cannot be based on the content of the speech or the viewpoint of the demonstrators, and that they typically cover administrative expenses such as processing the permit application, coordinating with law enforcement, and ensuring public safety during the event. See e.g. *iMatter Utah v. Njord*, 774 F.3d 1258 (10th Cir. 2014); *Surita v. Hyde*, 665 F.3d 860 (7th Cir. 2011).

The same principles can be applied when online adult content providers require age verification. In the case of age verification, cost-shifting legitimately places the cost of age verification technology on the providers of adult content rather than on the consumers who would otherwise incur costs by having to install and maintain monitoring and filtering software on their own devices. Since the costs of operating and maintaining modern age verification technology are no longer prohibitively expensive or unreasonable,

requiring adult websites to incur these costs does not violate the websites' First Amendment rights.²⁰

Furthermore, this cost evaluation does not consider the actual negative societal costs when minors can access explicit adult content, nor does it consider the undue burden that parents incur for the entire responsibility of filtering, monitoring, and installing monitoring and filtering software on their children's and shared family devices.

When it comes to obscenity, the government has a much wider latitude to determine and impose costs, both tangible and intangible. In *Paris Adult Theatre I v. Slaton*, 413 U.S. 49, 57–58 (1973), the Court held that “there are legitimate state interests at stake in stemming the tide of commercialized obscenity, even assuming it is feasible to enforce effective safeguards against exposure to juveniles and to passersby. Rights and interests ‘other than those of the advocates are involved.’ These include the interest of the public in the quality of life and the total community environment, the tone of commerce in the great city centers, and, possibly, the public safety itself.” Internal citations removed.

²⁰ See Declaration of Tony Allen: “Yoti state that their Age Verification Service (AVS) pricing ranges between \$0.03 (for large volumes e.g. circa 100 M, \$0.10 for circa 5M checks and \$0.31 for lower volumes, one time account based checks e.g. under 100,000). They also offer free, \$0.0 shares of 18 plus attributes from the reusable Yoti digital identity app....” J.A.201.

CONCLUSION

“Understandably those who entertain an absolutist view of the First Amendment find it uncomfortable to explain why rights of association, speech, and press should be severely restrained in the marketplace of goods and money, but not in the marketplace of pornography.” *Paris Adult Theatre I v. Slaton*, 413 U.S. 49, 60–62 (1973).

Considering the arguments presented, it is evident that current age verification technologies, particularly AI facial age estimation, offer practical, privacy-preserving, and scalable solutions for online age verification. This technology provides a secure and efficient method for websites and platforms to comply with legal requirements while respecting user privacy. The methods outlined, including the reuse of age tokens and digital identity apps, further reduce user friction while maintaining high accuracy and security at a low cost.

Given the demonstrated effectiveness of these technologies in protecting minors and enabling adults to assert their age without unnecessary intrusion, the Court should recognize that AI-driven facial age estimation solutions are the least restrictive alternative under the First Amendment when restricting sexually explicit materials on the Internet.

For these reasons, Yoti Ltd. respectfully submits that the Court should support the implementation of these advanced age assurance technologies as a reliable method to comply with age verification requirements while also safeguarding privacy and promoting accessibility for all users.

RESPECTFULLY SUBMITTED

JAMES R. MARSH
Counsel of Record
Marsh Law Firm PLLC
31 Hudson Yards, 11th Floor
New York, New York 10001
212-372-3030
jamesmarsh@marsh.law

November 2024