Limitations and Pitfalls of Integrating PETs in Online Age Verification

Sylvain Chatel, Christian Knabenhans, Wouter Lueks, Mathilde Raynal, **Carmela Troncoso**, and Ádám Vécsi

MPI-SP & CISPA & EPFL

EPFL

CISPA HELMHOLTZ CENTER FOR INFORMATION SECURITY

MAX PLANCK INSTITUTE FOR SECURITY AND PRIVACY

Age verification: a privacy nightmare!

Measuring User Responses to Age Verification Architectures: Evidence from a Deceptive Online Experiment

Yanzi Lin, Vivianna Lieu, Cheng Zhang, Weiqian Zhang, Lorrie Faith Cranor, Sarah Scheffler

Carnegie Mellon University

Abstract

Following the U.S. Supreme Court's decision in Free Speech Coalition v. Paxton (2025), which established that age verification systems must be "adequately tailored," understanding user behavior has become legally relevant for system design. This preliminary study empirically examines how different age verification methods affect user behavior through a deceptive online experiment framed as usability testing for a mock gambling website. Participants (n=99 U.S. residents) were randomly assigned to six verification conditions, including simple checkbox self-declaration, government-issued ID upload, and AI-based facial age estimation. Results show stark differences in user responses: checkbox verification achieved 95.2% completion rates, while government ID methods drove up to 60.5% of users to return their study without finishing. We also tested the effects of privacy disclosures on completion rates. These had mixed effects, with detailed data handling information both increasing completion rates and polarizing user comfort levels. In a survey accompanying the empirical study, participants expressed significant privacy concerns about document-based methods, citing fears of identity theft and data misuse. These findings provide empirical evidence that can be applied to the U.S. Constitutional requirement for "adequate tailoring" of age verification systems, as well as policy analysis and technical design of age verification more broadly. We outline plans for expanded research using R-rated movie content to examine these effects at larger scale.

The EU approach to age verification

The European Commission is working towards an EU-harmonised approach to age verification.

To help online platforms implement a user-friendly and privacy-preserving age verification method, the Commission is developing a harmonised approach across the EU in close collaboration with the Member States.

Share

Quick Links

PETs solve privacy problems

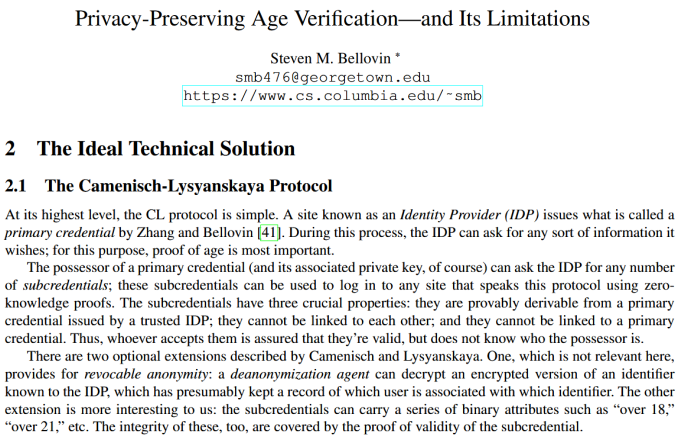

Privacy-Preserving Age Verification—and Its Limitations

Steven M. Bellovin *
smb476@georgetown.edu
https://www.cs.columbia.edu/~smb

2 The Ideal Technical Solution

2.1 The Camenisch-Lysyanskaya Protocol

At its highest level, the CL protocol is simple. A site known as an *Identity Provider (IDP)* issues what is called a *primary credential* by Zhang and Bellovin [41]. During this process, the IDP can ask for any sort of information it wishes; for this purpose, proof of age is most important.

The possessor of a primary credential (and its associated private key, of course) can ask the IDP for any number of *subcredentials*; these subcredentials can be used to log in to any site that speaks this protocol using zero-knowledge proofs. The subcredentials have three crucial properties: they are provably derivable from a primary credential issued by a trusted IDP; they cannot be linked to each other; and they cannot be linked to a primary credential. Thus, whoever accepts them is assured that they're valid, but does not know who the possessor is.

There are two optional extensions described by Camenisch and Lysyanskaya. One, which is not relevant here, provides for *revocable anonymity*: a *deanonymization agent* can decrypt an encrypted version of an identifier known to the IDP, which has presumably kept a record of which user is associated with which identifier. The other extension is more interesting to us: the subcredentials can carry a series of binary attributes such as "over 18," "over 21," etc. The integrity of these, too, are covered by the proof of validity of the subcredential.

Proving attributes (like age) in a privacy-preserving way is a well-studied problem with well-understood solution

Anonymous credentials
Attribute-based credentials
Verifiable credentials

Libraries for privacy-preserving proofs start popping (e.g., Google age verification)

Problem solved! or problem narrowed?

PETs ONLY solve privacy problems

A Study of China's Censorship and Its Evasion Through the Lens of Online Gaming

Yuzhou Feng, *Florida International University;* Ruyu Zhai, *Hangzhou Dianzi University;* Radu Sion, *Stony Brook University;* Bogdan Carbunar, *Florida International University*

https://www.usenix.org/conference/usenixsecurity23/presentation/feng

PETs cannot prevent circumvention

Trivial circumvention orthogonal to how age is proved/verified

- VPNs and off-band access to content
- Take mum's ID, buy/rent accounts online

(and might worsen privacy: delegitimize VPNs to avoid circumvention)

Focusing the discussion on privacy, prevents discussion on effectiveness

- can we provide credentials to everyone?
- can we label content?
- can we implement widely? At what cost?

Is there a sweet win situation? Or is all lose-lose?

PETs ONLY solve privacy problems

PETs cannot prevent censorship

Repurpose age verification: selecting values to target subpopulations

Censor by availability: PETs are not universal and not always inclusive
May force adoption of particular vendors/software (centralization)

Induced censorship possibilities cannot be eliminated
PETs (in general security) can make it worse

PETs ONLY solve SOME privacy problems

PETs cannot prevent privacy leaks outside of age verification

What are the privacy properties PETs provide in this context?

Avoid collection of ID & biometrics

but cannot prevent current tracking practices based on meta-data

and it is one more attribute revealed (contribution to quasi identifier)

e.g., EUDI Wallet

The discussion on privacy must address the big picture and point limits of protection

PETs centralize power

PETs reduce freedom of application developers

PETs are advanced and complex – libraries become black boxes

Functionality: the library determines what can be proven and how

Formats: the library determines data formats for the application

Creation of dependencies on software provider (closed, too complex) (or in the OS/Browser…)

Centralization is not only architectural. Libraries can become prisons. PETs are especially dangerous because by nature are restrictive

What then?

Acknowledge that application-layer privacy is only one of the many problems around age verification (and we only talk here about technical issues)

Acknowledge the limitations of PETs: is a win-win possible?

And also the risks: privacy-washing undesirable functionalities

Be careful with what you wish: privacy can turn an open market into a closed shop instead of bringing freedom

Standardization efforts must focus solutions that can truly be open