Willful Disregard under the CPA

How Do Online Services Know Users are Minors?

May 5, 2025

Technology Impact Lab

Professors Paul Ohm, Meg Leta Jones and Jon Brescia

Abby Rochman, Susana Pretelt, Patrick Yurky, Niel Swanepoel, Kunjika Pathak

Table of Contents

Acknowledgements

We would like to thank Professors Paul Ohm, Meg Leta Jones and Jon Brescia for their support and guidance throughout this project. This would not have been possible without their constant help clarifying our doubts and steering us in the right direction. We would also like to extend our gratitude to Stevie DeGroff and Zoe Kirchoff from the Colorado AG's office who helped us clarify some key parts of the Colorado Privacy Act. And lastly, we thank our fellow classmates for their insights throughout the semester.

Executive Summary

This report examines the "willful disregard" standard under the Colorado Privacy Act (CPA) to assess whether a user is a minor on various online platforms. A 2024 amendment to the CPA introduced heightened duties for platforms that either know or willfully disregard that a user is under 18. Our project examines what this standard could look like in practice, and whether current platform practices already meet it. To answer this question, we analyzed three popular platforms used by minors: Roblox, TikTok, and Character.AI.

To do this, we identified five different age-related practices:
1) Self-declaration age at sign up; platforms are able to determine their users are minors using age assurance mechanisms such as self-declaration of age during sign up.
2) Age related data in network traffic; platforms are able to determine a user's age in the traffic once they have created an account.
3) Age segmentation in advertising systems; the advertising ecosystem takes into account users data that can be used to determine a user's age.
4) Age-based content and feature restrictions; limiting access to features or modifying content for its minor users shows that they know their user's age.
5) Linguistic inference of user age from behavior and comments; based on a proof of concept inference mockup tool to determine how platforms know their users are minors through daily interactions.

Using a combination of policy analysis and technical testing, we show that platforms already possess multiple ways to determine or infer whether a user is a minor. In many cases, they collect self-declared birthdates, transmit age fields in server traffic, restrict features by age, or allow advertisers to target users by age range.

However, this report has several limitations. Our analysis is limited to three platforms and a relatively small number of user interactions. The project is designed to illustrate methods related to age knowledge, rather than make legal determinations about willful disregard in specific cases. We also did not assess whether companies are complying with their own stated policies.

Additionally, we did not examine whether specific content on platforms is appropriate for its minor users.

Despite these limitations, our findings suggest that platforms have access to methods to help them determine whether some of their users are minors.

2. Background and Motivation

The Colorado Privacy Act (CPA) creates broad data privacy protections for consumers, including some heightened protections for minors.[1] Online service providers are obligated to obtain specific informed consent for certain activities whenever they "have actual knowledge or willfully disregard" that a particular user is a minor. The term "willful disregard" appears in similar state and federal privacy legislation but does not have a universally accepted definition.

This project investigates what it means for an online service provider to "willfully disregard" that a user is a minor. Answering this question can help provide clarity for online service providers and regulators alike.

Central to this regulatory challenge is the evolving field of age assurance. Policymakers and enforcement agencies are grappling with how to implement age-based protections in a privacy-preserving and equitable way—particularly for minors who may lack government-issued IDs. Broadly speaking, age assurance is a broad term for methods used to discern the age or age range of an individual.[2] The most common and traditional method is "age gating," which refers to when a user declares their birthdate without providing supporting evidence.[3] However, research shows that millions of users have lied about their age in order to gain access to platforms that require a minimum age.[4] Today, self-declaration is regarded as insufficient for protecting minors from harmful or inappropriate content by itself. Additionally, as more legislation stipulates age-based restrictions, there is an increasing need for platforms to know the ages of their users and act accordingly.

Furthermore, the volume and granularity of user data already collected for advertising, engagement, and content moderation purposes suggest that some controllers may already have the tools to identify minors, even if they do not explicitly verify age. This raises important

---

[1] Colo. Rev. Stat. § 6-1-1301 (2022).
[2] Future of Privacy Forum. *Unpacking Age Assurance: Technologies and Tradeoffs*. Infographic. June 2023. https://fpf.org/wp-content/uploads/2023/06/FPF_Age-Assurance_final_6.23.pdf.
[3] Ibid.
[4] Amanda Lenhart, Kristen Purcell, Aaron Smith, and Kathryn Zickuhr, *Social Media & Mobile Internet Use among Teens and Young Adults* (Washington, D.C.: Pew Internet & American Life Project, February 3, 2010), https://www.pewinternet.org/2010/02/03/social-media-and-young-adults/.; danah boyd et al., "Social Networking Sites as Networked Publics: Affordances, Dynamics, and Implications," *First Monday* 16, no. 6 (2011), https://firstmonday.org/ojs/index.php/fm/article/view/3850/3075.

questions about both the technical feasibility of age inference and the extent of controller accountability when they choose not to act on the data they collect.

The purpose of this project is to investigate five different practices currently used by platforms that indicate a user's age. These methods are processes already broadly employed, with the exception of an inference tool that we built based on a method analyzing user comments. The investigation reveals that platforms already have plenty of ways to determine their user is a minor and could be willfully disregarding that status. We investigate self-declaration, network traffic analysis, advertising, content and feature restrictions, and an inference tool for user content that we built.

Minors use many kinds of online platforms and services beyond just social media to connect with friends, play games and engage with content. For our project, we selected three distinct platforms with varied services to capture the range of offerings: Roblox, TikTok and Character.AI. All three platforms have a diverse user base consisting of adults as well as minors across the United States.

3. Legal Context

In 2024, Colorado enacted an amendment to its privacy framework through Senate Bill 24-041, which introduced new obligations to protect the personal data of minors online. This amendment will come into effect on October 1, 2025, meaning that enterprises must comply with these new standards. The amendment defines a "minor"[5] as anyone under the age of 18 and recognizes the concept of "heightened risk of harm to minors"[6] when their personal data is processed in ways that could lead to physical, financial, or reputational injury, or to unwanted intrusions into their privacy. Hence, through this amendment, new responsibilities arise for companies to avoid these heightened risks when handling minor's data.[7]

These obligations fall on the "controllers"[8] (i.e., companies offering online services), and they are triggered not only when a company "actually knows" that a user is a minor, but also when it "willfully disregards" that fact.[9] Hence, the standard of "willful disregard" means that companies may no longer shield themselves from liability simply by failing to ask for a user's age. If a

[5] Colorado Revised Statutes. *Senate Bill 24-041, § 6-1-1303(16.5)* (2024).

[6] *"(...) processing the personal data of minors in a manner that presents a reasonably Foreseeable risk that could cause: (a) unfair or deceptive treatment of, or unlawful disparate impact on, minors; (b) financial, physical, or reputational injury to minors; (c) unauthorized disclosure of the personal data of minors (...); or (d) physical or other intrusion upon the solitude or seclusion (...)." Colo. Rev. Stat. § 6-1-1303(14.5)*

[7] Particularly §6-1-1308.5 and §6-1-1309.5.

[8] *"'(...) person that, alone or jointly with others, determines the purposes for and means of processing personal data."* Colo. Rev. Stat. § 6-1-1303(7).

[9] *The Law states the following: "(...) a controller that offers any online service, product, or feature to a consumer whom the controller* ***actually knows or willfully disregards*** *is a minor shall use reasonable care to avoid any heightened risk of harm to minors caused by the online service, product, or feature."* Colo. Rev. Stat. § 6-1-1308.5(1)(a) *(emphasis added)*

controller has access to information that could inform them that a user is under 18 years-old and chooses not to act on it, they may still be held responsible. Therefore, controllers should obtain the minor's consent to process personal data for specific purposes, such as targeted advertising, the sale of this data, or profiling that has legal or similarly significant effects. The following table summarizes the rules established in the amendment[10] under which consent is required:

6-1-1308.5. Duties of controllers (2) Unless a controller has obtained consent, a controller that offers any online service, product, or feature to a consumer whom the controller actually knows or willfully disregards is a minor shall not:

(a)	**Process a minor's personal data**	**(I) For the purposes of:**
		• Targeted advertising
		• Sale of personal data
		• Profiling in furtherance of decisions that produce legal or similarly significant effects
		(II) For any other purpose than:
		• The one disclosed at the time of collection
		• Or a purpose reasonably necessary for and compatible with the disclosed purpose
		(III) For longer than reasonably necessary to provide the online service, product, or feature

(b)	**Use system design features**	That significantly increases, sustains, or extends a minor's use of the online service, product, or feature.

(c)	**Collect a minor's precise geolocation data**	**Unless all of the following are true:**
		(I) The data is reasonably necessary to provide the online service, product, or feature
		(II) The data is collected/retained only as long as needed to provide the service
		(III) The controller provides a visible signal for the entire duration of collection

What makes this reform important regarding minors is the knowledge standard it adopts. Across the U.S., privacy laws now use different formulations to express what kind of knowledge a company must have in order to trigger their obligations toward minors. In some cases, legislation uses an "actual knowledge"[11] standard, which implies direct awareness of the user's age. In

[10] For the complete quote, please refer to Appendix 9.3.
[11] Camille Altieri, *Now, on the Internet, Will Everyone Know If You're a Child?* Future of Privacy Forum, 2024, https://fpf.org/blog/now-on-the-internet-will-everyone-know-if-youre-a-child/

others, they use a "constructive knowledge"[12] standard, meaning that the enterprise should have known the age of the user.

Hence, across U.S. privacy law, different statutes adopt different standards of knowledge to trigger obligations toward minors. For example, the Federal Children's Online Privacy Protection Act (COPPA) uses an "actual knowledge" standard. In this case, the FTC clarified that "actual knowledge" under COPPA includes cases of "willful disregard" — situations where platforms "blind themselves" to clear evidence of a user's age.[13] However, COPPA itself does not explicitly define "willful disregard," and recent FTC rulemakings have declined to expand the standard so far as to include "constructive knowledge."[14]

At the state level, some recent laws have incorporated "willful disregard" language. Connecticut, Florida, and California's consumer privacy laws adopt this phrasing, possibly suggesting a somewhat broader knowledge trigger than COPPA's traditional "actual knowledge" standard.[15] Florida's regulatory rules provide some additional guidance, indicating that "willful disregard" could be found when a platform "should reasonably have been aroused to question" the user's age.[16]

Overall, while "willful disregard" does not have a universally accepted definition, it can increasingly be recognized as part of a slightly broader standard of actual knowledge, requiring companies to be alert to signs that users are minors rather than remaining willfully blind. Given the increasing complexity of user data and the growing potential for age inference through indirect signals, states could interpret "willful disregard" to impose a more proactive duty on platforms. These developments motivate our investigation: understanding what platforms know or could reasonably know about users' ages, and how "willful disregard" might operate in practice.

The CPA amendment uses the phrase "actually knows or willfully disregards," which situates it closer to an actual knowledge standard but could be interpreted more broadly by regulators. The CPA does not require companies to implement an age verification system,[17] but if they choose to

---

[12] Ibid.

[13] 89 Fed. Reg. 2034, 2037 (2024) ("The concept of actual knowledge includes willful disregard. See, e.g., Glob.-Tech Appliances, Inc. v. SEB S.A., 563 U.S. 754, 766 (2011) (noting that '[i]t is also said that persons who know enough to blind themselves to direct proof of critical facts in effect have actual knowledge of those facts'). Therefore, the Rule already applies to instances in which an operator of a general audience site or service willfully disregards the fact that a particular user is a child.").

[14] Ibid.

[15] Conn. Gen. Stat. § 53-420 (2024); Fla. Stat. § 501.702; Cal. Civ. Code § 1798.100 et seq. (2024).

[16] Florida Attorney General, Draft Rule Text for Section 501.72(5), Florida Statutes (effective July 18, 2024) ("A controller willfully disregards a consumer's age if it, based on the facts or circumstance readily available to the controller, should reasonably have been aroused to question whether a consumer was a child and thereafter failed to perform reasonable age verification.").

[17] Colo. Rev. Stat. § 6-1-1304(3)(f) (2024).

"conduct commercially reasonable age estimation" and still get it wrong, they are not liable. However, once a company has actual knowledge or has willfully disregarded that a user is a minor, the obligations are triggered.

Therefore, given the complexity of knowledge standards across states, we aimed to clarify how the "actual knowledge" or "willful disregard" standard should be understood under the CPA, and what that means in practice for online platforms. We hope our findings can help 1) illuminate how companies already possess and use their tools and information to identify minors on their platforms and 2) support compliance with new legal responsibilities that come with that knowledge.

4. Core Research Question

Given the legal background and motivation, we pose our core research question: How do companies know (or infer) that a user is a minor, and what does it mean to willfully disregard this information under the CPA? We aim to explore this using the data they already collect about their users.

5. Methods and Findings

The main purpose of our project is to provide five different practices that are used by platforms to determine their users' age. Using technical and legal analysis methods of how data is collected and inferred, our aim is to provide a useful framework to help assess whether a company's actions and policies meet the standard of willful disregard.

The five age-related platform practices we investigated:

1) Self-declaration: The age declared by users when prompted to create an account.
2) Network traffic analysis: The traffic captured in the apps for age-related signals.
3) Advertising analysis: The ad infrastructure, noting targeting advertising to and data collection of minors.
4) Content and feature restrictions analysis: The way apps present different content to minors and whether certain features are restricted or not.
5) User behavior inference analysis: Based on an inference mockup tool we built to predict whether a comment was made by a minor or an adult.

We analyzed Roblox, TikTok and Character.AI across these age-related platform practices. We chose to study apps that let minors make an account. In other words, these are not apps specific to minors but have safeguards in place for minors to use. Roblox is an online gaming platform that allows users to play as well as create their own games. TikTok is a social media platform for

creating, sharing and discovering short form video content. Character.AI is a chatbot service using artificial intelligence to create and interact with digital characters.

However, our research had several limitations. First, we are not making concrete determinations of whether a company is willfully disregarding a user's age. Additionally, we did not examine the content on apps to assess whether it was appropriate for minors or not. Our project also does not technically validate whether or not a company complies with their own stated policies.

Our analysis was limited to three apps and a small number of interactions. A more thorough investigation could expand this sample, examine additional app behaviors, or analyze traffic from non-iOS platforms. Additionally, future work could focus on analyzing traffic metadata at scale, or studying the role of shared data environments across apps run by a single parent company.

5.1. Method 1: Self-declaration

To better understand the information available to platforms to determine a user's age, we examined self-declaration: the most common and widely adopted form of age assurance. This method, where users must input their birth date upon account creation and before gaining access to the platform, is frequently implemented to determine eligibility for access and to apply age-based content restrictions. While users can lie about their age, at least some number of users do not.[18] Therefore, this is a valuable starting point to determine whether these three platforms have information that a portion of their registered users are minors.

We investigated the self-declaration mechanism by going to all three applications and attempting to use them. Each of the three applications required us to make an account to use it and specifically asked for a date of birth.

5.1.1. Roblox

Roblox prompts users to input their birth date when they first make an account. In terms of COPPA compliance, Roblox allows users under 13 to create accounts without immediately requiring parental consent.[19] Roblox requests in their Privacy Policy that under-13 users get parental consent before using Roblox, but there isn't anything in the sign-up flow that specifically mentions this or slows down the process for making an account as someone under the age of 13.[20]

---

[18] Lenhart et al., *Social Media & Mobile Internet Use*. boyd et al., "Social Networking Sites as Networked Publics.".
[19] Common Sense Media, *Roblox – Full Privacy Report*, Common Sense Privacy Program, n.d., https://privacy.commonsense.org/privacy-report/Roblox.
[20] Roblox Corporation, *Roblox Privacy and Cookie Policy*, March 14, 2025, https://en.help.roblox.com/hc/en-us/articles/115004630823-Roblox-Privacy-and-Cookie-Policy.

However, users who input an age under 13 at sign-up are automatically placed into a more restricted account experience, including: stricter chat filters, limited friend requests, locked-down profile visibility, and restricted content access to comply with COPPA.[21] Additionally, Roblox automatically limits their data collection on users who self-declare as under 13 years old. Roblox will not allow a user who declares to be under 5 years old to make an account.

Roblox self-declaration mechanism

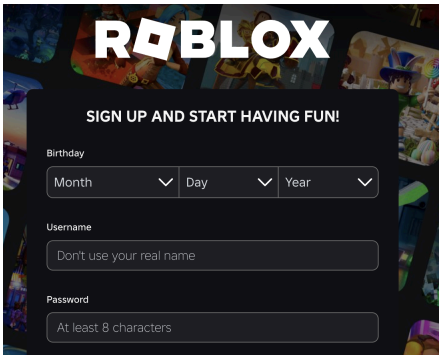

Roblox will not allow a user to make an account with a birthdate under 5 years of age

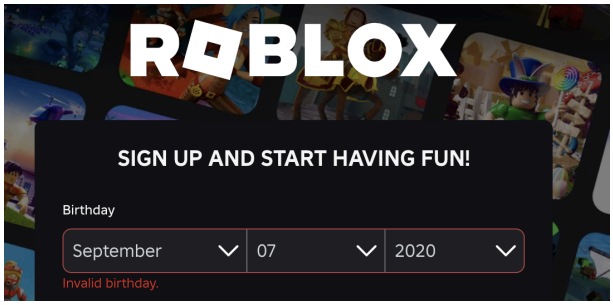

[21] Roblox Corporation, *Major Updates to Our Safety Systems and Parental Controls*, Roblox Newsroom, November 18, 2024, https://corp.roblox.com/newsroom/2024/11/major-updates-to-our-safety-systems-and-parental-controls; Jay Peters, "Roblox Is Making Changes for Pre-Teen Users after Reports That It Failed to Protect Children," *The Verge*, October 23, 2024, https://www.theverge.com/2024/10/23/24277992/roblox-pre-teen-children-parent-accounts-default-settings.

5.1.2. TikTok

TikTok requires all users to provide their date of birth upon registration.

TikTok self-declaration mechanism

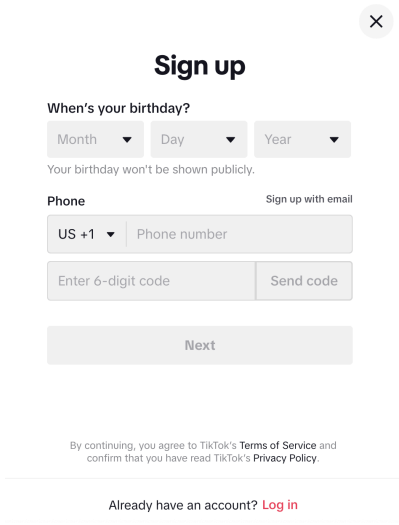

When signing up, if users self-report they are under 13, they do not access the standard version TikTok app. Instead, they access a version called Kids Mode.[22] In this version all content created is private, no messages are sent and the user feed is curated for age-appropriate content. Any data collected from the child user is collected for internal purposes and therefore does not require parental consent as required by COPPA.[23]

[22] TikTok, *TikTok Under 13 Experience*, TikTok Support, n.d., https://support.tiktok.com/en/safety-hc/account-and-user-safety/tiktok-under-13-experience

[23] Federal Trade Commission, *Complying with COPPA: Frequently Asked Questions*, n.d., accessed April 25, 2025, https://www.ftc.gov/business-guidance/resources/complying-coppa-frequently-asked-questions

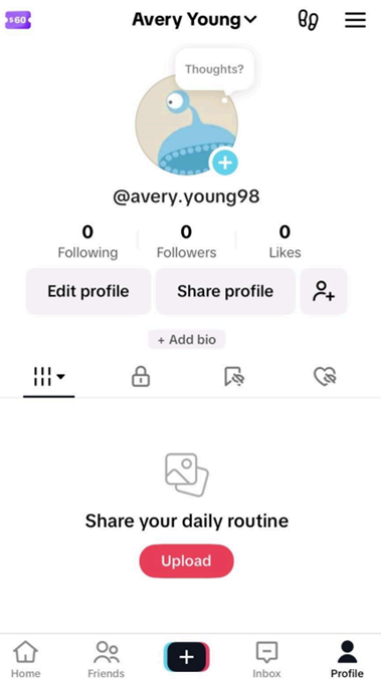	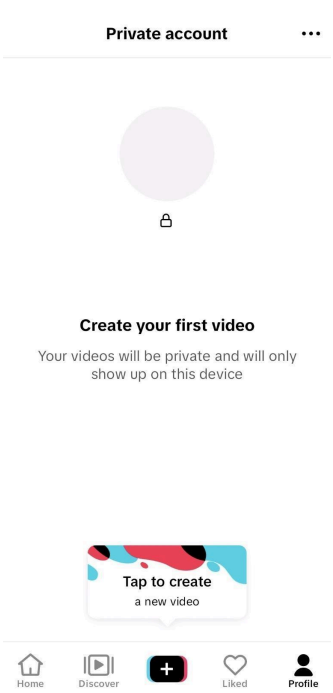
Standard TikTok	**Under 13 Kids Mode TikTok**

As shown in the screenshot above, the standard TikTok UI features a full navigation bar at the bottom. In contrast, the Under 13 Kids mode UI limits removes social tabs like Friends and Inbox, replacing them with more restricted icons like Liked and Discover, and does not allow public video uploads, maintaining a more private and controlled environment.

For users who self-identify as minors aged 13 to 17, no parental consent is required during the signup process. These users give consent solely through acceptance of the platform's terms of service.

TikTok additionally offers “family pairing” where parents can choose to link to their 13-17 kids account and impose certain restrictions.[24] Since this is not a mandatory feature, it is not a reliable indication of minors’ data being collected by the app and consequently whether the app is vulnerable to willful disregard. However, since this is a notable feature of TikTok's age related features, an illustration of the family pairing consent flows can be found in the appendix.[25]

[24] TikTok. (n.d.). Family pairing. TikTok Support.
[25] See Appendix 9.1.

5.1.3. Character.AI

Character.AI also has a self-declaration pop-up where users have to input their birth date. According to their policy, users must be 13 or older to use the app.[26] For users between the ages of 13-18, there are content and feature restrictions. This includes parental controls and a separate AI model for teens that reduces exposure to certain sensitive content.

Character.AI self-declaration mechanism

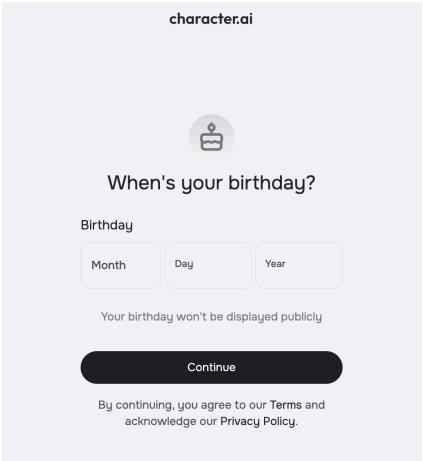

5.1.4. Implications for self-declaration

Across these platforms, millions of teenage users have declared their age to these companies. This suggests platforms are not unaware of their users' status as minors. In fact, they are knowingly collecting and storing this information. By failing to act on age data already in their possession, apps could raise questions about whether they are willfully disregarding the age of their users.

5.2. Method 2: Network Traffic

By directly inspecting the traffic sent between the user's device and the platform, we can look for additional evidence of age data being transferred between devices, platforms, and third parties. Some of the questions we intend to answer with this investigation:

- What are the different kinds of information that a platform receives as the result of a user's behavior and the device that they are using, and what does that information indicate about whether or not a user is a minor?

[26] Character.AI, "Privacy Policy," accessed May 5, 2025, https://character.ai/privacy.:contentReference[oaicite:2]{index=2}

- For a minor using an iPhone, with an AppleID registered for a minor, will the device communicate any information about their minor status? And, is there any difference between a minor's phone and an adult's phone in this regard?
- After creating an account, declaring age, or signing in — where a minor has explicitly communicated something about their age — can we demonstrate that the platform retains persistent knowledge of that information?
- What can a platform infer about a user's age as a result of data created through the user's behavior and usage of the app?

In this section, we will address all but the last question, detailing repeatable steps to document what data platforms receive regarding users' ages. The last question will be explored later, in "Method 5" on user inference behavior.

5.2.1. Technical setup

It is usually not possible to inspect the traffic coming to or from a phone, as the data is encrypted at the device/application level. To conduct our technical investigation, we used an open-source procedure to "jailbreak" the iPhone 10 devices, allowing the installation of software to intercept these traffic signals.[27] Using publicly available developer tools for testing and debugging HTTP traffic,[28] we configured a proxy server to intercept the traffic in a relatively human-readable format.

Overview of technical setup for network traffic interception and analysis.

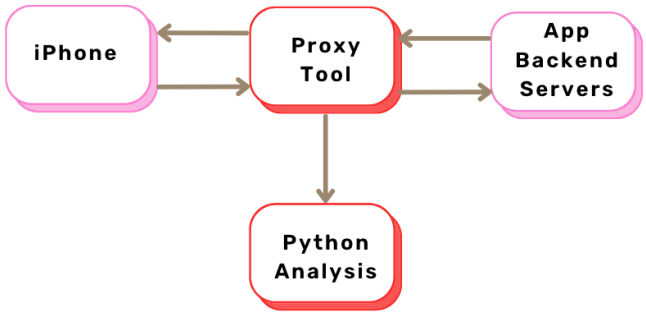

At this point, we were able to use our targeted apps on the phones and intercept all traffic being sent to and from the phone during that time. This allowed us to determine what the device communicates to the app, and what information the app stores and refuses to communicate with the user's device while running the application. While this data is relatively human-readable, we

[27] Jailbreak tool: Palera1n (https://github.com/palera1n/palera1n).
[28] Proxy and traffic interception: HTTP Toolkit (https://httptoolkit.com/).

wrote Python scripts to parse and analyze the recorded traffic in order to search for and target specific data fields — particularly those related to a user's identity or age.

For each of our sample applications, we implemented a few different behavior flows, targeted at different stages of the typical use process for an app.

5.2.2. Pre-declaration

We first investigated the data sent after the user downloads and opens an app, but before they take any further action.

Intercepted POST request from Roblox app upon initial launch, prior to user sign-up. Shows device metadata but no personal age-identifying data at this stage.

```
POST  dlsdk.appsflyer.com/v1.0/ios/id431946152
  ↳ Query Params:
     sdk_version: ['6.10']
     af_sig: ['0f2e1da610be389979e81ff48fb0482dca98f28573a72d40b89ebbe44ffe9202']
  ↳ Headers:
     host: dlsdk.appsflyer.com
     content-type: application/json
     connection: keep-alive
     accept: application/json
     user-agent: Roblox/2.667.668 CFNetwork/1408.0.4 Darwin/22.5.0
     content-length: 277
     accept-language: en-US,en;q=0.9
     accept-encoding: gzip, deflate, br
  ↳ Body:
{
    "idfa": {
        "value": "",
        "type": "unhashed"
    },
    "os": "16.5.1",
    "request_count": 1,
    "lang": "",
    "idfv": {
        "value": "C1E105B8-14AD-488A-BF3B-A9B434AB278C",
        "type": "unhashed"
    },
    "request_id": "1744554911947-1105814",
    "att_status": 1,
    "is_first": true,
    "timestamp": "2025-04-13T14:42:25.887",
    "type": "iPhone"
}
```

In these initial communications, most of the data was unremarkable. It was primarily related to session establishment and analytics. Specific information was limited to device type (iPhone 10),

OS version, and some session metadata. We found no evidence in this stage of data that would indicate the user's age or identity.[29]

5.2.3. Post-declaration

Next, we looked at traffic after the user declared their age by entering a birthdate. In each of the tested apps, this was requested early in the onboarding process. After submitting this data, we consistently observed fields such as “userAge” or “birthdate” appearing in traffic. These fields persisted across multiple requests and sessions, both sent from the device and returned by the server, suggesting that the app stores and reuses this data. This demonstrates that once age is declared, the platform has structured, persistent knowledge of that fact.

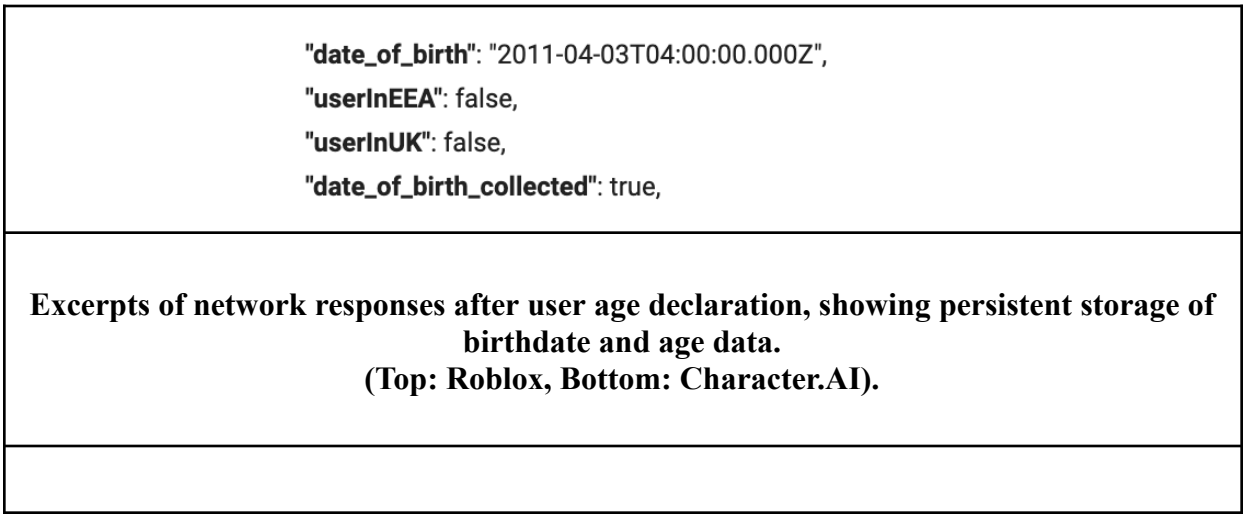

Excerpts of network responses after user age declaration, showing persistent storage of birthdate and age data.
(Top: Roblox, Bottom: Character.AI).

[29] The “idfa” field refers to Apple's Identifier for Advertisers, a device-level identifier intended to facilitate user tracking across apps and services. While IDFA values could, in theory, allow an app controller to link a user's behavior across multiple apps they operate, we did not analyze cross-app tracking behaviors in this study.

```
"user": {
  "userID": "771772751",
  "email": "offonagainc@gmail.com",
  "customIDs": {
    "deviceID": "55304c0c-c494-4d2b-8b3f-3aae38fab461"
  },
  "custom": {
    "subscription": "NONE",
    "entitlements": [],
    "isEarlyAccess": false,
    "selectedLanguage": "en",
    "isAnonymous": false,
    "username": "TenderKingfisher5621",
    "userAgeInYears": 14,
    "dateJoinedMs": 1744563792264
  },
```


5.2.4. General usage and broader implications

Finally, we captured traffic during regular app usage, including navigating screens, clicking buttons, and interacting with various app features. These activities generated a large volume of varied network traffic. Our review of this traffic remained preliminary, with more detailed analysis reserved for future work. Broader behavioral inferences are discussed further in Method 5.

Despite the limited exploration of general usage flows, the network traffic captures consistently demonstrate that once age-related information is collected during account creation, it persists across multiple sessions and interactions. This finding highlights a concrete and reproducible method for documenting what information platforms have at specific stages of user interaction. In particular, this approach could be valuable for enforcement, compliance audits, or documenting baseline knowledge conditions in specific cases.

While there were minor differences in how the apps structured their network requests, all consistently transmitted age-related information once it had been collected.

5.3. Method 3: Advertising

This section outlines how the observed platforms structure their advertising systems and the types of data they use to inform ad targeting, particularly, in relation to their user's age. The aim is to recognize specific design choices and practices that may suggest a platform has the ability

to know whether a user is a minor. In doing so, we identify common examples of how platforms signal or infer age internally. We begin with Roblox, whose ad infrastructure is closely linked to user behavior and age filters. Then we turn to TikTok, which openly outlines age-based targeting in its ad policies. Finally, we consider Character.AI, which does not currently serve ads but articulates future possibilities that involve user profiling.

5.3.1. Roblox

Roblox's advertising infrastructure is unique. It doesn't use traditional banner or interstitial ads. Instead, Roblox offers immersive ads and sponsored content, like branded experiences and in-game billboards. Since users can create and sell virtual items on Roblox, users can advertise for the items that they sell in the Roblox marketplace. Brands can also place immersive ads on Roblox. For example, Nike could have a billboard in a skateboarding experience.

Roblox in-game immersive advertisement example

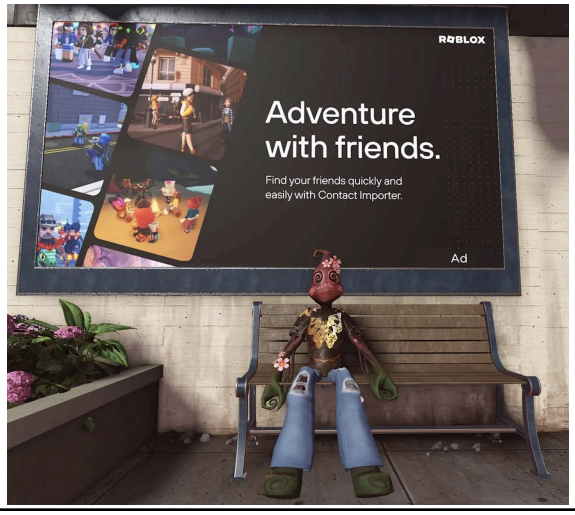

Roblox developers can place ads in Roblox in the home and search functions, as well as in-game via virtual billboard.[30] In addition to age, people can target their ads based on "experience genre targeting" or "audience genre targeting."[31] "Experience genre targeting" is for immersive ads, and means targeting based on the genre of games where the ad is served.[32] For example, if a developer were to only select "action," their ad would only appear in experiences classified as "action."[33] "Audience genre targeting" is for sponsored experiences, and means that the

[30] Roblox Corporation, *Ad Campaigns*, Roblox Creator Hub, 2025, https://create.roblox.com/docs/production/promotion/ads-manager
[31] Ibid.
[32] Ibid.
[33] Ibid.

experiences are targeted based on a user's preferences, like users who like to play experiences in selected genres.[34]

Both brands and developers can specify the age range of the players they want to see their ad.[35] Roblox says that ads are not shown to users under the age of 13.[36] Roblox is saying that it is actively tracking users' ages and allowing developers and brands to target them for ads. This implies that they wouldn't be able to claim ignorance about whether a user is a minor between 13-17 year-olds. Currently, their approach raises compliance concerns with the Colorado law, which requires minors to consent to process their personal data for the purposes of targeted advertising.

In sum, Roblox confirms it allows age-based targeting in its ad infrastructure. This means it either collects or infers users' ages, including for minors. Therefore, continued ad targeting to 13-17 year-olds could possibly suggest "willful disregard" of minor status under the CPA.

5.3.2. TikTok

TikTok states in its terms of service how they let advertisers target ads and to which audience. This language is a first step to evidence how TikTok uses age for their advertising strategy:

- According to TikTok's Ad Review Policy, part of the review process for any ad includes assessing the "targeted market and age group."[37] This means that people who want to advertise on the site must limit the age of their audience as a required input when designing ad campaigns. Moreover, TikTok places restrictions on certain types of ads (such as those for pharmaceuticals, alcohol, gambling, or even nutritional supplements), stating explicitly that these can only be shown to adults.[38] The Terms of Service say that TikTok uses user activity to deliver targeted advertising. This includes both on-platform and off-platform behavior. They specify that personalized ads are turned on by default, and are informed by data such as location, inferred interests, age range, and behavioral patterns.[39] Also, the Ad Settings allow users to view and adjust the inferred targeting categories for ad delivery.[40]

---

[34] Ibid.
[35] Ibid.
[36] Ibid.
[37] TikTok. *Advertising on TikTok: First Things to Note.* TikTok Ads Help Center, 2025. https://ads.tiktok.com/help/article/advertising-on-tiktok-first-things-to-note
[38] TikTok. *About Advertising to People Under the Age of 18.* TikTok Business Help Center, 2025. https://ads.tiktok.com/help/article/about-advertising-to-people-under-the-age-of-18; TikTok. *Protecting Minors on TikTok: Advertising Initiatives.* TikTok Business Help Center, 2025. https://ads.tiktok.com/help/article/protecting-minors-on-tiktok-advertising-initiatives?lang=en.
[39] TikTok. *Privacy Policy.* 2024. https://www.tiktok.com/legal/page/row/privacy-policy/en.
[40] TikTok. *Ads and Your Data.* 2025. https://www.tiktok.com/privacy/ads-and-your-data/en.

- In a policy page dedicated to advertising restrictions for people under 18, TikTok states that it does not deliver personalized ads to users aged 13–17 in the European Economic Area, the UK, Switzerland, Brazil, Mexico, and Colombia.[41]

Thus, TikTok collects data to estimate a user's age, and uses it to restrict ad content based on age thresholds. And while it may appear that TikTok implements these policies on the basis of the self-declared age, the platform's Terms of Services state that they use the information they collect "to infer additional information about you, such as your age, gender, and interests," and that the user can "find more information about how an ad was personalized using various targeting information selected by the advertiser, such as location information, age range, inferred gender information, inferred interests, and your behavioral information based both on your on- and off-TikTok activity."[42] As such, in addition to self-declared age data, TikTok also uses other types of data to classify users into age thresholds and deliver them targeted ads.

TikTok advertisement examples

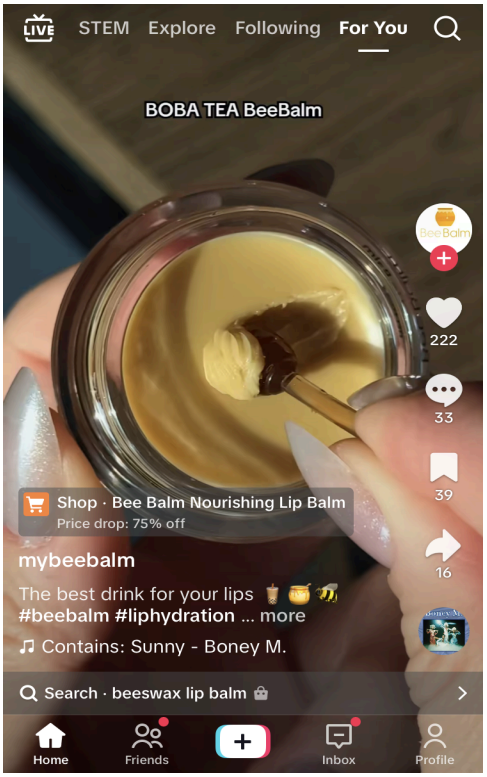

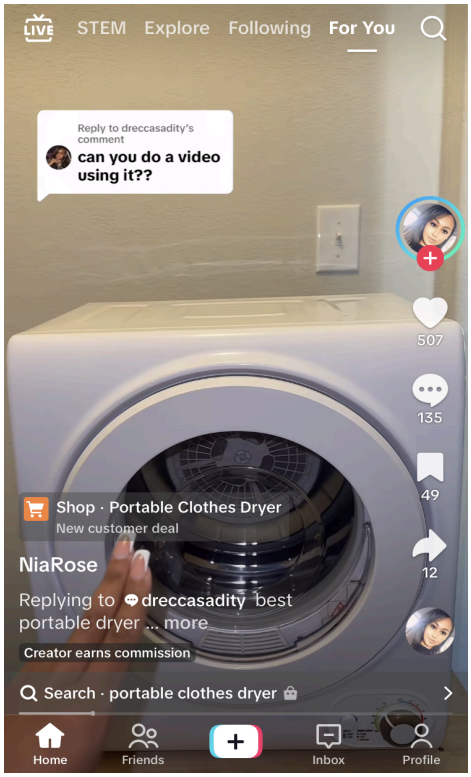

[41] TikTok. *About Advertising to People Under the Age of 18*. TikTok Ads Help Center. Accessed May 5, 2025. https://ads.tiktok.com/help/article/about-advertising-to-people-under-the-age-of-18.
[42] TikTok. *Privacy Policy.* 2024.

5.3.3. Character.AI

Character.AI's advertising framework is much less developed, because they clarify that for the moment they do not deliver ads on the app. However, the policy does say that in the future, data may be shared with advertising and analytics partners who may use cookies and tracking technologies to "deliver advertising and content targeted to your interests."[43] Hence, although this is framed as a possibility rather than a current practice, it still implies that platforms are able to collect enough data to enable interest-based ad targeting. Therefore, even if Character.AI is not yet serving ads, it appears to have access to the type of data that would allow it to profile users by age in the future.

5.3.4. Implications for advertising analysis

Overall, what Roblox, TikTok and Character.AI's advertising policies reveal is that these platforms are collecting user's data, giving them the ability to use it to build detailed user profiles, including age segmentation, for the purpose of targeted advertising. These actions are explained in their own terms of service and ad policies. Platforms acknowledge that they receive and process data that can be used to determine a user's age.

5.4. Method 4: Content and feature restrictions

In this section, we examine content and feature restrictions employed by apps for its minor users. This could be restricting access to features such as voice chats or content moderation. The aim is to examine how platforms know that the user is a minor through the content and features they let minors access or restrict.

Roblox has content restrictions automatically applied for users who self-declare as under-13, and additional features that are age-gated until users confirm using ID that they are over 13 or over 17 years of age. TikTok has hard and soft restrictions for its minor users while Character.AI has different input and output models for minor only content.

5.4.1. Roblox

Beyond just asking users to self-declare their ages at sign-up, Roblox has additional age verification processes in place that users must comply with in order to access certain features, like voice chat and experiences for 17+ users. In order to access voice chat, users must prove that they are at least 13 years of age by uploading a government-issued photo ID (this can be a

[43] Character.AI. *Privacy Policy.* 2023. https://character.ai/privacy.

driver's license, passport, residency card, or any other government-issued identification document with the user's picture on it).[44]

How the age verification process works:

1. Document submission
2. Selfie capture
3. Analysis: 3rd party actor Persona analyzes the user's age and identity.[45]

While Roblox does not currently require ID-based age verification for all users, data from their verification practices could give them relevant information to trigger CPA protections when a user does demonstrate that they are a minor in this way.

Roblox age verification process
Step 1: Click "Verify my Age"
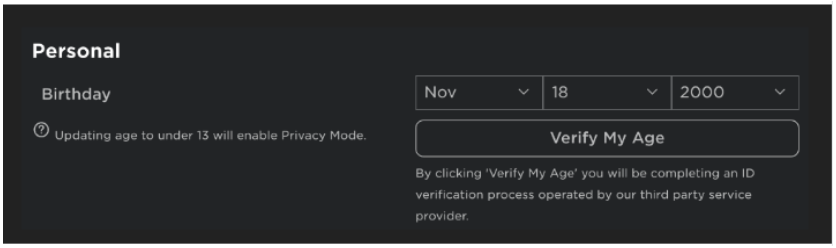
Step 2: Scan QR code to begin age verification flow.
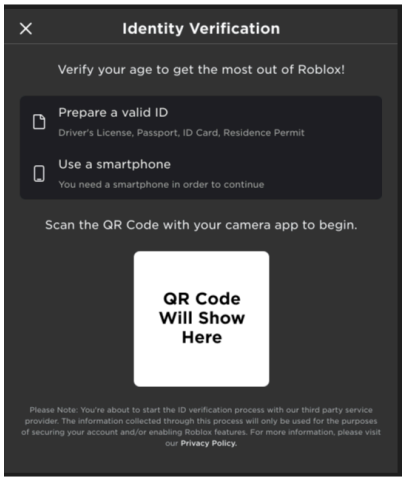
Step 3: Begin age verification flow on Persona.

[44] Roblox Corporation, *Age ID Verification*, Roblox Support, n.d., https://en.help.roblox.com/hc/en-us/articles/4407282410644-Age-ID-Verification.
[45] Ibid; Roblox Corporation, *Roblox Biometric Privacy Notice*, Roblox Support, August 29, 2024, https://en.help.roblox.com/hc/en-us/articles/4412863575316-Roblox-Biometric-Privacy-Notice.

Roblox age verification process

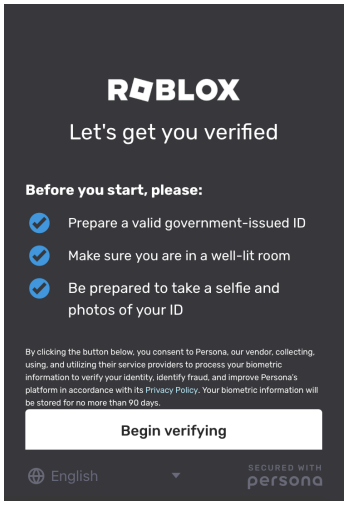

Step 4: Choose ID type, upload ID, and provide a selfie to Persona for validation.

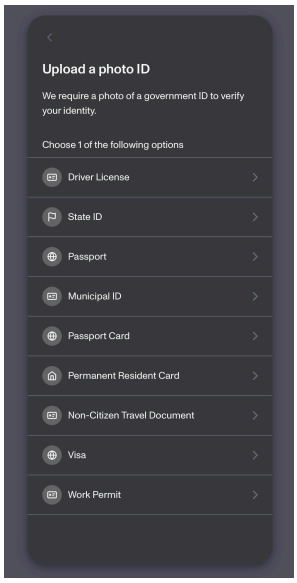

5.4.2. TikTok

TikTok has several content restriction mechanisms based on the age of users and which are relevant to minors.[46]

Feature	Feature Category	Restriction Type	Applies to Age	Age Verification Process	Notes

[46] TikTok, *Privacy and Safety Settings for Users Under Age 18*, TikTok Help Center, n.d., accessed April 25, 2025, https://support.tiktok.com/en/account-and-privacy/account-privacy-settings/privacy-and-safety-settings-for-users-under-age-18.

Privacy set to private by default	Privacy Settings	Soft (default setting, user can change)	Under 18	Self-declared age	Users under 18 start with private accounts but can opt out
Videos by under 16 users excluded from 'For You' recommendations	Privacy settings	Hard (algorithmic exclusion)	Under 16	Self-declared age	Reduces visibility and discovery of under 16 content
DMs off by default	Direct Messaging	Soft (default setting, user can change)	16-17	Self-declared age	Messaging turned off by default but can be enabled
DMs only to friends/suggested friends	Direct Messaging	Soft (limited to known contacts)	16-17	Self-declared age	Even when enabled, chatting is limited to friends
DMs not available at all	Direct Messaging	Hard (blanket ban)	Under 16	Self-declared age	Under 16s cannot access messaging
Videos cannot be downloaded	Video Downloads	Hard (blanket ban)	Under 16	Self-declared age	Prevents saving and re-sharing of content
Comments only from friends	Comments & Interactions	Hard (feature limitation)	13-15	Self-declared age	Only friends can comment on videos
Duet and Stitch disabled	Comments & Interactions	Hard (feature limitation)	Under 16	Self-declared age	Duet and Stitch unavailable to under 16s
Mature content filtered	Content Exposure	Hard (content filtering)	Under 18	Algorithmic tagging and moderation	Age-tagged mature content is filtered out
TikTok Shop restricted	Platform Features	Hard (access restriction)	Under 18	Age-gating verification	Monetization features for adults only
TikTok Live restricted	Platform Features	Hard (access restriction)	Under 18	Age-gating verification	Live posting features for adults only

What this shows us is again that TikTok curates user features and restrictions explicitly on age categories, which necessarily involves them having the data needed to make these feature segments.

5.4.3. Character.AI

Once users have declared that they are minors, Character.AI has certain content and feature restrictions. They filter content for teens based on outputs and inputs. For outputs, they have a

separate model for their teen users which is supposed to reduce the likelihood of them viewing sensitive content. They employ classifiers to identify certain inappropriate content and enforce their content policies. For inputs, they rely on the user's conversation with the characters.[47]

Character AI inputs for user's conversation with characters

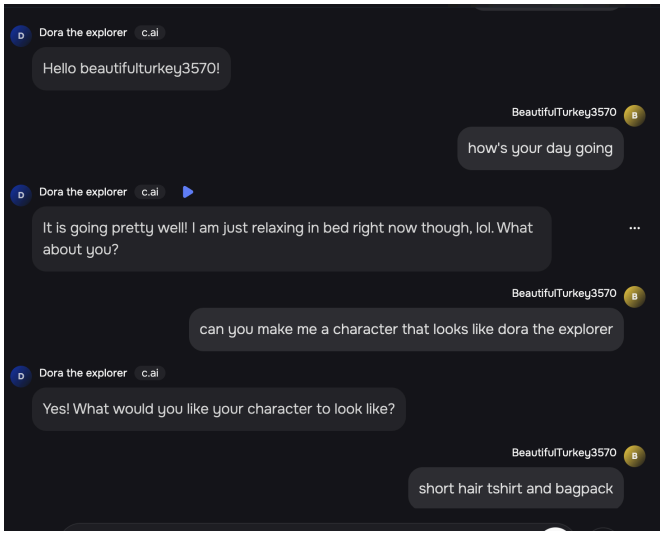

Character.AI has also recently rolled out a parental insights feature which is a toggle feature on the app. The minors have to add their guardians or parents in their app and invite them. After they accept their invite, the parental insights feature is turned on and parents can access a summary of the teen's activity on the platform.[48] The parental insights feature gives them access to three things on the platform; daily time spent on the platform, characters most interacted with and time spent with each character.[49]

Character AI Parental Insights Feature

[47] "How Character.AI Prioritizes Teen Safety," *Character.AI Blog*, December 12, 2024, https://blog.character.ai/how-character-ai-prioritizes-teen-safety/
[48] "Parental Insights," *C.AI Help Center*, March 7, 2025, https://support.character.ai/hc/en-us/articles/34659826266267-Parental-Insights
[49] "Introducing Parental Insights: Enhanced Safety for Teens," *Character.AI Blog*, March 25, 2025, https://blog.character.ai/introducing-parental-insights-enhanced-safety-for-teens/

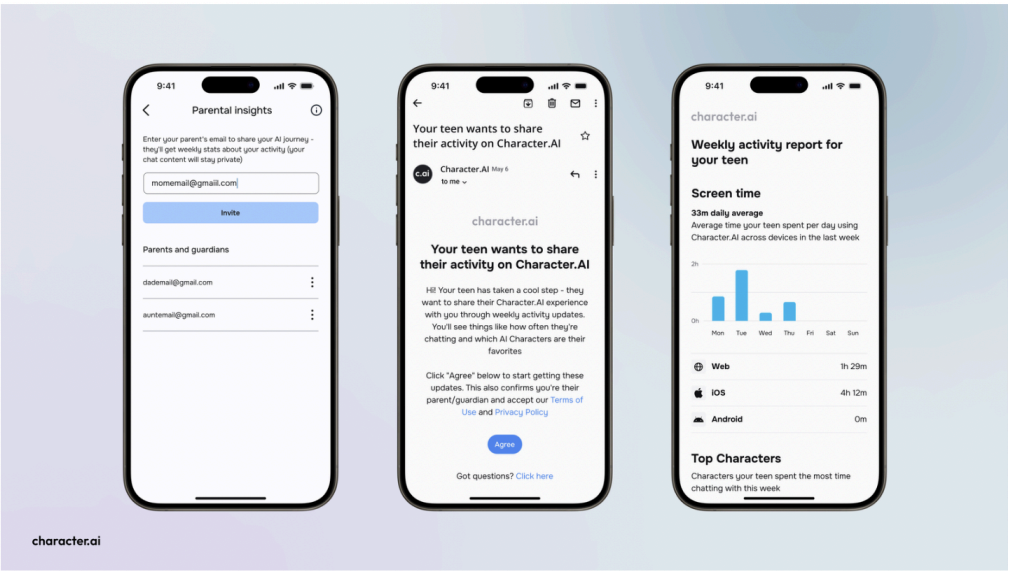

If they see that the content violates their Terms of Service or Community guidelines, they block that content or in some cases suspend accounts if they see them repeatedly inputting similar content. The content that is not allowed as per the Terms of Service is a long list including but not limited to sexual harassment, self-harm, pornographic content and criminal activities. In an ongoing lawsuit, plaintiffs allege that two minors from Texas were led to suffer severe physical and mental harms by the platform.[50]

However, it is important to note that the content restricted per the community guidelines is restricted for all users such as obscenities and they make no distinction between that content and minor specific content.

5.4.4. Implications for content restrictions

Across all platforms, age-based content restrictions lead to a clear implication: if apps tailor features based on a user's indicated or verified age, they must necessarily collect and process age-related data from minors. Therefore, if it can be shown that an app adjusts its features according to age, it follows that the platform has information about user age—and may be held liable under the CPA for willfully disregarding rules on the protection of minors.

5.5. Method 5: User behavior inference

[50] Meetali Jain, "Center for Humane Technology: New Federal Lawsuit Reveals How Character.AI's Inherently Dangerous Product Designs Harm Children," *Tech Justice Law Project*, December 10, 2024, https://techjusticelaw.org/2024/12/10/center-for-humane-technology-new-federal-lawsuit-reveals-how-character-ais-inherently-dangerous-product-designs-harm-children/.

We developed a proof-of-concept mockup of an inference tool to demonstrate how a platform might determine whether a given message originated from a minor (13-17) or not based on the linguistic characteristics of their text. Our approach was inspired by the 2016 study "Your Age Is No Secret: Inferring Microbloggers' Ages via Content and Interaction Analysis."[51] In their study, they conducted a content analysis of tweets using unigram and bigram features extracted from users' tweets, and weighted them to highlight age-indicative phrases (like "homework" vs. "job interview").[52] Rather than predicting exact ages, they focused on classifying users into specific age ranges, like we were trying to do with minors vs. non-minors. While their study included social interaction modeling as well as content analysis, we chose to focus only on what we could glean from the content of user comments on the platforms.

For each platform, we generated platform-specific training datasets reflecting the language used by users across three age groups. Because real labeled age data is not typically available at scale, all datasets were synthetically constructed with ChatGPT, with messages written to reflect plausible, platform-specific behavior and linguistic tone for different age categories. The datasets for each platform were developed over several sessions, about 100 phrases at a time. This is because, when asked to deliver more than 100 minor or adult specific phrases, ChatGPT'squality would noticeably decline. It would especially lead to several repeated messages, which would negatively impact the tool. We used prompts like "Give me 100 messages that a minor user (13-17 years old) might send on Roblox," followed by "give me 100 messages that an adult (18+) might send on Roblox," and used the same prompts for the other platforms. We also included more specific instructions like avoiding overly structured patterns and encouraged the inclusion of slang and casual language. To ensure data quality, scripts were developed to detect duplicates and formatting inconsistencies (e.g., stray commas, repeated messages), in addition to manual data cleaning.

- Roblox: 1540 data points
- Character.AI: 673 data points
- TikTok: 1781 data points

We built a custom language model for each platform (Roblox, TikTok, and Character.AI) to help our system guess a user's age based on comments they might post. For each platform, we trained a separate version of the model so it could learn the communication style used on that platform, since people communicate differently depending on the app. We did not hard-code any age-indicative phrases or rules into the model, meaning we did not manually program the system to look for specific words like "taxes" or "homework" to guess someone's age. Instead, we used machine learning techniques to train the model to recognize language patterns based on the

[51]Jinxue Zhang, Yanchao Zhang, and Huan Liu, "Your Age Is No Secret: Inferring Microbloggers' Ages via Content and Interaction Analysis," in *Proceedings of the Tenth International AAAI Conference on Web and Social Media (ICWSM 2016)*, 476–485, 2016, https://ojs.aaai.org/index.php/ICWSM/article/view/14731.
[52] Ibid.

synthetic datasets. The datasets were labeled by age (0 for adult, 1 for minor, like: "adulting is hard,0" or "almost fell asleep in history class,1").

We used a technique that looks at the words people might use, including both single words (like "like"), and two-word phrases (like "so fun"). These combinations help the model pick up on patterns tied to different age groups. We kept these small common words because they can be especially useful clues for someone's age in casual online chats. We kept the number of features (or signals) the model looks at fairly small–limited to the top 1,000 words or phrases–so the model stays focused. We also told the model to pay equal attention to each age group, to try to prevent it from favoring one group just because there's more training data for them.

The Roblox model was trained as a three-class classifier to distinguish between Children (0-12), Teens (13-17), and Adults (18+), because Roblox is the only platform of the three that explicitly allows children under 13 to be users. The TikTok and Character.AI models were trained as binary classifiers, with classes of Minor (13-17) or Adult (18+).

5.5.1. Results of user behavior-based age inference tool

Each model was trained using an 80/20 train-test split. This test set was not seen by the model during training, so it provided an unbiased estimate of how well the model generalizes to new, unseen data. Final accuracy scores were:

- Roblox: 80%
- TikTok: 73%
- Character.AI: 76%

These accuracy scores are similar to those found in the 2016 study using tweets, which had an accuracy score of 81.3%.

5.5.2. Implications for user behavior-based age inference tool

While this approach has potential, these final accuracy scores may not meet the legal standard of "willfully disregarding" age. Additionally, they were trained on synthetic data that might lack nuance of real-life interactions. Still, this version is only a proof of concept, and it shows that platforms could train similar models using their own internal data. Many already have vast stores of human-labeled data from content moderation workflows.[53] For example, Meta announced on April 21, 2025, that it would start using its AI age verification system based on user behavior to

[53] Spandana Singh, "Case Study: Facebook," in *Everything in Moderation: An Analysis of How Internet Platforms Are Using Artificial Intelligence to Moderate User-Generated Content* (Washington, D.C.: New America, 2019), https://www.newamerica.org/oti/reports/everything-moderation-analysis-how-internet-platforms-are-using-artificial-intelligence-moderate-user-generated-content/case-study-facebook/.

enforce age-informed account settings.[54] Meta says that it will use its AI tools to proactively find suspected teen accounts with adult birthdays and automatically place the account under more restrictive teen settings.[55]

6. Conclusions

Based on this analysis, we identify five key indicators that suggest a platform may have knowledge of a user's minor status and could be willfully disregarding this information in its data practices and commercial operations.

First, platforms that request a user's date of birth during account registration or through other interactions collect direct self-declared age data. Even when this information is self-reported, the platform is aware that at least a subset of its users are minors. This creates a clear and high-certainty basis for concluding that the platform possesses knowledge of underage users.

Second, network traffic analysis provides a technical means of confirming age-related data collection. If packet captures reveal fields containing age, date of birth, or metadata linked to age verification attempts, then the platform can reasonably be said to have access to information sufficient to identify a user's minor status. As with self-declaration, this indicator carries a high level of certainty.

Third, advertising practices offer indirect yet telling evidence of age awareness. When platforms apply age-based targeting restrictions—such as disabling detailed targeting based on gender, income, or behavior for users under 18—they implicitly acknowledge the user's age. While this does not always constitute direct knowledge, the use of such segmentation provides a medium level of certainty that age is being inferred and operationalized.

Fourth, age-specific content and feature restrictions provide further evidence of age awareness. Platforms often limit functionality (such as live streaming or messaging) or alter recommendation algorithms for younger users. These differentiated design choices, made based on age thresholds, indicate that the platform is actively using age information—even if indirectly inferred—to govern user experience. This, too, offers a medium level of certainty.

Fifth, behavioral inference techniques may allow platforms to estimate user age through patterns in engagement, preferences, or device use. For example, high engagement with teen-centric content or certain temporal usage patterns could serve as proxies for youth. Although potentially powerful, this method is probabilistic and less precise than others, making it a lower-certainty indicator of platform knowledge.

---

[54]Mia Sato, "Meta Is Ramping Up Its AI-Driven Age Detection," *The Verge*, April 21, 2025, https://www.theverge.com/news/651826/meta-instagram-age-detection-ai-settings.
[55] Ibid.

The findings suggest that many platforms are already collecting a wide array of user data that can reveal or strongly suggest a user's age. In fact, some platforms likely possess sufficient technical capacity to detect when a user is a minor—even if they do not formally verify age at registration. If the same data is used to personalize services, curate content, or target advertising based on inferred age groups, then it can—and should—also be used to trigger the corresponding legal obligations under the CPA.

8. Bibliography

Altieri, Camille. *Now, on the Internet, Will Everyone Know If You're a Child?* Future of Privacy Forum, 2024. https://fpf.org/blog/now-on-the-internet-will-everyone-know-if-youre-a-child/.

boyd, danah, Eszter Hargittai, Jason Schultz, and John Palfrey. "Social Networking Sites as Networked Publics: Affordances, Dynamics, and Implications." *First Monday* 16, no. 6 (2011). https://firstmonday.org/ojs/index.php/fm/article/view/3850/3075.

Cal. Civ. Code § 1798.100 et seq. (2024).

Character.AI. "Privacy Policy." Accessed May 5, 2025. https://character.ai/privacy.

Colorado Revised Statutes. *Senate Bill 24-041*. 2024.

Colorado Revised Statutes § 6-1-1301 (2022).

Conn. Gen. Stat. § 53-420 (2024).

Common Sense Media. *Roblox – Full Privacy Report*. Common Sense Privacy Program, n.d. https://privacy.commonsense.org/privacy-report/Roblox.

Federal Trade Commission. *Complying with COPPA: Frequently Asked Questions*. n.d. Accessed April 25, 2025. https://www.ftc.gov/business-guidance/resources/complying-coppa-frequently-asked-questions.

Federal Trade Commission. *Children's Online Privacy Protection Rule: Notice of Proposed Rulemaking*. 89 Fed. Reg. 2034 (January 11, 2024)

Fla. Stat. § 501.702 (2024).

Future of Privacy Forum. *Unpacking Age Assurance: Technologies and Tradeoffs*. Infographic. June 2023. https://fpf.org/wp-content/uploads/2023/06/FPF_Age-Assurance_final_6.23.pdf.

"How Character.AI Prioritizes Teen Safety." *Character.AI Blog*. December 12, 2024. https://blog.character.ai/how-character-ai-prioritizes-teen-safety/.

Jain, Meetali. "Center for Humane Technology: New Federal Lawsuit Reveals How Character.AI's Inherently Dangerous Product Designs Harm Children." *Tech Justice Law Project*, December 10, 2024. https://techjusticelaw.org/2024/12/10/center-for-humane-technology-new-federal-lawsuit-reveals-how-character-ais-inherently-dangerous-product-designs-harm-children/.

"Introducing Parental Insights: Enhanced Safety for Teens." *Character.AI Blog*. March 25, 2025. https://blog.character.ai/introducing-parental-insights-enhanced-safety-for-teens/

Lenhart, Amanda, Kristen Purcell, Aaron Smith, and Kathryn Zickuhr. *Social Media & Mobile Internet Use among Teens and Young Adults*. Pew Internet & American Life Project, February 3, 2010. https://www.pewinternet.org/2010/02/03/social-media-and-young-adults/.

"Parental Insights." *C.AI Help Center*. March 7, 2025. https://support.character.ai/hc/en-us/articles/34659826266267-Parental-Insights.

Peters, Jay. "Roblox Is Making Changes for Pre-Teen Users after Reports That It Failed to Protect Children." *The Verge*, October 23, 2024. https://www.theverge.com/2024/10/23/24277992/roblox-pre-teen-children-parent-accounts-default-settings.

Roblox Corporation. *Ad Campaigns*. Roblox Creator Hub, 2025. https://create.roblox.com/docs/production/promotion/ads-manager.

Roblox Corporation. *Age ID Verification*. Roblox Support, n.d. https://en.help.roblox.com/hc/en-us/articles/4407282410644-Age-ID-Verification.

Roblox Corporation. *Major Updates to Our Safety Systems and Parental Controls*. Roblox Newsroom, August 29, 2024. https://corp.roblox.com/newsroom/2024/11/major-updates-to-our-safety-systems-and-parental-controls.

Roblox Corporation. *Roblox Biometric Privacy Notice*. Roblox Support, August 29, 2024. https://en.help.roblox.com/hc/en-us/articles/4412863575316-Roblox-Biometric-Privacy-Notice.

Roblox Corporation. *Roblox Privacy and Cookie Policy*. March 14, 2025. https://en.help.roblox.com/hc/en-us/articles/115004630823-Roblox-Privacy-and-Cookie-Policy.

Sato, Mia. "Meta Is Ramping Up Its AI-Driven Age Detection." *The Verge*, April 21, 2025. https://www.theverge.com/news/651826/meta-instagram-age-detection-ai-settings.

Singh, Spandana. "Case Study: Facebook." In *Everything in Moderation: An Analysis of How Internet Platforms Are Using Artificial Intelligence to Moderate User-Generated Content*. Washington, D.C.: New America, 2019. https://www.newamerica.org/oti/reports/everything-moderation-analysis-how-internet-platforms-are-using-artificial-intelligence-moderate-user-generated-content/case-study-facebook/.

TikTok. *About Advertising to People Under the Age of 18*. TikTok Ads Help Center. Accessed May 5, 2025. https://ads.tiktok.com/help/article/about-advertising-to-people-under-the-age-of-18.

TikTok. *Privacy and Safety Settings for Users Under Age 18*. TikTok Help Center, n.d. Accessed April 25, 2025. https://support.tiktok.com/en/account-and-privacy/account-privacy-settings/privacy-and-safety-settings-for-users-under-age-18.

TikTok. *TikTok Under 13 Experience*. TikTok Support, n.d. https://support.tiktok.com/en/safety-hc/account-and-user-safety/tiktok-under-13-experience.

Zhang, Jinxue, Yanchao Zhang, and Huan Liu. "Your Age Is No Secret: Inferring Microbloggers' Ages via Content and Interaction Analysis." In *Proceedings of the Tenth International AAAI Conference on Web and Social Media (ICWSM 2016)*, 476–485. 2016. https://ojs.aaai.org/index.php/ICWSM/article/view/14731.

9. Appendix

9.1. TikTok "family pairing" mode consent flows process

Below is a breakdown of the explicit consent points in TikTok's Family Pairing process:

TikTok "Family Pairing" mode consent flow
Step 1: The teen selects "Teen" in Family Pairing and scans the QR code shown on the parent's device.
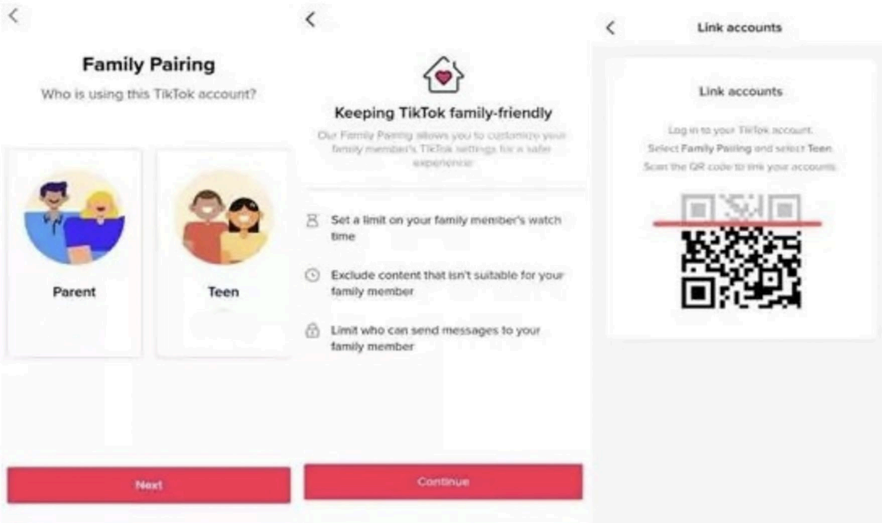
Step 2: Consent Action and account linkage. Next, TikTok displays a prompt asking the teen to confirm the link with the specified parent account. This is the first and primary point of consent in the Family Pairing process.
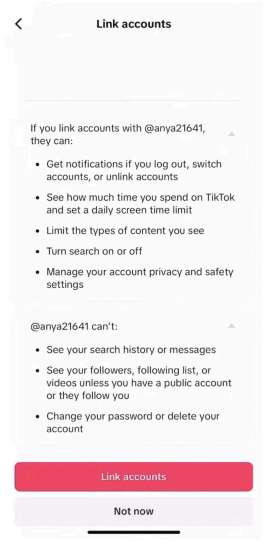

Outcome: Once the teen taps “Link Accounts,” TikTok: Links the two accounts, activates parental controls set by the parent (i.e. no messaging feature) and logs the consent action on its backend systems.
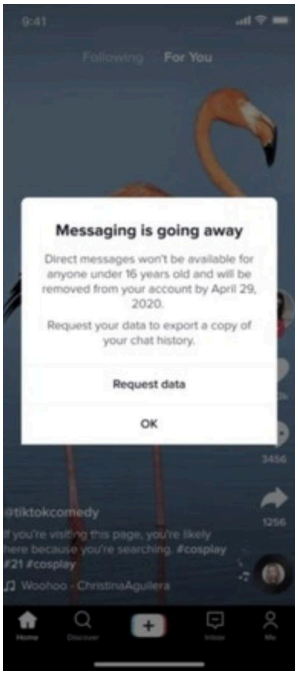

9.2. User behavior inference tool

- Data set
- Python code
- Example output (Roblox classifier)

9.3. Textual quotations and definitions

Footnote Number	***Textual Quote***
10	*“6-1-1308.5. Duties of controllers - duty of care - rebuttable presumption. (1) (a) a controller that offers any online service, product, or feature to a consumer whom the controller actually knows or willfully disregards is a minor shall use reasonable care to avoid any heightened risk of harm to minors caused by the online service, product, or feature. (b) in any enforcement action brought by the attorney general or a district attorney pursuant to section 6-1-1311, there is a rebuttable presumption that a controller used reasonable care as required under this section if the controller complied with this section. (2) unless a controller has obtained consent in accordance with subsection (3) of this section, a controller that offers any online service, product, or feature to a consumer whom the controller actually knows or willfully disregards is a minor shall not: (i) for the purposes of: (a) targeted advertising; (B) the sale of personal data; or (c) profiling in furtherance of decisions that produce legal or similarly significant effects concerning a consumer; (ii) for any processing purpose other than the processing purpose that the controller disclosed at the time the controller collected the minor's personal data or that is reasonably necessary for, and compatible with, the processing*

	purpose that the controller disclosed at the time the controller collected the minor's personal data; or (iii) for longer than is reasonably necessary to provide the online service, product, or feature; (b) use any system design feature to significantly increase, sustain, or extend a minor's use of the online service, product, or feature; or (c) collect a minors precise geolocation data unless: (i) the minor's precise geolocation data is reasonably necessary for the controller to provide the online service, product, or feature; (ii) the controller only collects and retains the minor's precise geolocation data for the time necessary to provide the online service, product, or feature; and (iii) the controller provides to the minor a signal indicating that the controller is collecting the minors precise geolocation data and makes the signal available to the minor for the entire duration of the collection of the minor's precise geolocation data; except that this subsection (2)(c)(iii) does not apply to any service or application that is used by and under the direction of a ski area operator, as defined in section 33-44-103 (7). (3) (a) a controller shall not engage in the activities described in subsection (2) of this section unless the controller obtains: (i) the minor's consent; or (ii)(a) if the minor is a child, the consent of the minor's parent or legal guardian. (b) a controller that complies with the verifiable parental consent requirements established in the 'children's online privacy protection act of 1998', 15 u.s.c. sec. 6501 et seq., as amended, and the regulations, rules, guidance, and exemptions adopted pursuant to said act, as amended, is deemed to have satisfied any requirement to obtain parental consent under this subsection (3)(a)(ii). (b)(1) a controller that offers any online service, product, or feature to a consumer whom that controller actually knows or willfully disregards is a minor shall not: (a) provide any consent mechanism that is designed to substantially subvert or impair, or is manipulated with the effect of substantially subverting or impairing, user autonomy, decision-making, or choice; or (b) except as provided in subsection (3)(b)(ii) of this section, offer any direct messaging apparatus for use by a minor without providing readily accessible and easy-to-use safeguards to limit the ability of an adult to send unsolicited communications to the minor with whom the adult is not connected. (ii) subsection (3)(b)(1)(b) of this section does not apply to an online service, product, or feature of which the predominant or exclusive function is: (a) electronic mail; or (b) direct messaging consisting of text, photos, or videos that are sent between devices by electronic means, where messages are: shared between the sender and the recipient; only visible to the sender and the recipient; and not posted publicly. (4) subsections (2)(a) and (2)(b) of this section do not apply to any service or application that is used by and under the direction of an educational entity, including a learning management system or a student engagement program." Colorado Revised Statutes § 6-1-1308.5