



GENERATIVE AI SAFETY REPORT

**SAFETY IN DETAIL:
AGE RATINGS: A
CROSS-PLATFORM
ANALYSIS**

AISF (ARTIFICIAL INTELLIGENCE SAFETY FORUM)

**NOVEMBER
2025**

Table of Contents

| | |
|---|----|
| Executive Summary | 3 |
| Methodology | 4 |
| Results | 5 |
| Key Findings | 6 |
| Implications | 7 |
| Recommendations | 8 |
| Appendix: Products, Platforms, and Versions | 9 |
| Contact Us | 10 |

Executive Summary

The Artificial Intelligence Safety Forum (AISF) compared the age ratings across different platforms for ten popular free-tier, general-purpose conversational AI agents. The goal is to highlight where age rating gaps occur across platforms and to inform the actions needed to improve clarity and protect users.

Overview

The Artificial Intelligence Safety Forum (AISF) conducted a comparative review of ten popular free-tier, general-purpose conversational AI agents to examine inconsistencies in their published age ratings across major distribution platforms. These platforms included: the App Store, Google Play, web-based, Windows, and macOS.

The findings reveal substantial misalignment in age guidance between these platforms. Many apps rated as 3+ on Google Play are rated 12+ or 13+ on the App Store, and, in some cases, listed as 16+ or 18+ in developer documentation. Web and desktop versions often display no age rating at all.

These discrepancies are not caused by differences in product capability, but by divergent platform policies and legacy content-rating systems that fail to account for generative or adaptive AI behaviour. This lack of uniformity undermines user safety, confuses caregivers and educators, and weakens accountability when harmful or inappropriate content is generated.

This study examined ten of the most popular, commercially available, free-tier conversational AI systems. All findings are based on the versions and age rating information available in November 2025. Updates or policy changes made by developers after this period may alter their published age classifications or associated safety performance.

Scope and limitations

This study focused exclusively on ten of the most popular, commercially available free-tier, general-purpose conversational AI agents. Our findings are based on versions of these products available in November; updates and changes made by the developers after our review may alter their safety performance.

Methodology

The AISF compared public age ratings, versions, and update data for ten popular free-tier, general-purpose conversational AI agents across mobile, web, and desktop platforms.

Products assessed

The assessment focused on ten popular, free-tier, general-purpose conversational AI agents chosen for their widespread public use and accessibility. This included: ChatGPT, Claude, DeepSeek, Gemini, Grok, Le Chat, Meta AI, Microsoft Copilot, Perplexity, and Qwen.

Platforms included

The review examined age ratings across the most common access points used by general audiences: the App Store, Google Play, web-based, Windows, and macOS. These platforms were selected because they are the primary channels through which users download, access, or interact with conversational AI agents.

Data collected

For each product and platform, AISF collected publicly listed age ratings, app version numbers, and last-updated dates. This ensured that comparisons were based on the most current, publicly available information and provided a consistent basis for identifying discrepancies across platforms.

Developer age ratings

Developer age ratings were included as a separate reference point. These were drawn from each product's Terms and Conditions or official usage guidance to compare the developer's own minimum-age expectations against the ratings assigned by platforms. This allowed the assessment to evaluate not only cross-platform differences, but also differences between how developers classify their products versus how platforms categorise them.

Results

The comparative age rating chart highlights how the age classification for each products differs across platforms, revealing gaps between mobile, desktop, and developer declared ratings.

Cross-platform age ratings: November, 2025

| | App Store | Google Play | Web-based | Windows | macOS | Developer |
|-------------------|-----------|-------------|------------|---------|------------|------------|
| ChatGPT | 13+ | 12+ | Not stated | 12+ | Not stated | 13+ |
| Claude | 18+ | 12+ | Not stated | N/A | Not stated | 18+ |
| DeepSeek | 13+ | 12+ | Not stated | N/A | N/A | No minimum |
| Gemini | 13+ | 12+ | Not stated | N/A | N/A | No minimum |
| Grok | 13+ | 12+ | Not stated | N/A | N/A | 13+ |
| Le Chat | 13+ | 12+ | Not stated | N/A | N/A | 13+ |
| Meta AI | 13+ | 3+ | Not stated | N/A | N/A | 13+ |
| Microsoft Copilot | 13+ | 3+ | Not stated | 12+ | 13+ | 13+ |
| Perplexity | 18+ | 12+ | Not stated | 3+ | 18+ | 13+ |
| Qwen | 18+ | 3+ | Not stated | N/A | N/A | 18+ |

- "Not stated": the product was available on the platform but the age rating was not stated.
- "No minimum": the product was available on the platform and the age rating was stated without a minimum.
- "N/A": the product was not available on the platform.

Note: age ratings shown reflect the versions of each product available across all platforms (App Store, Google Play, Web-based, Windows, macOS, and the product developer) as of November 2025. A full version list is available in "Appendix: Products, Platforms, and Versions".

Key Findings

Age classifications differ widely between app stores, platforms, and developer policies, often without any difference in capability.

1. Significant rating divergence

Every product exhibited at least one mismatch between platform ratings and developer-declared minimum ages.

- The App Store consistently used 13+ or 18+ categories.
- Google Play ratings ranged from 3+ to 12+.
- Developer policies, where stated, often recommended 16+ or 18+ use.
- All web-based versions did not state an age rating.
- Windows ratings ranged from 3+ to 12+.
- macOS ratings ranged from 12+ to 18+.

This variation shows that no cross-platform or industry-standard definition of AI-related risk currently exists.

2. Functionality parity across platforms

Version numbers and update dates were generally within days of one another across platforms, indicating near-identical functionality. The inconsistency in age ratings therefore appears to stem from policy frameworks, not from meaningful differences in what users can do within the product.

3. Developer ratings are more conservative

When developers stated age guidance directly, they were typically stricter than as stated in the App Store or Google Play. For instance, a product listed as 3+ on Google Play might have the developer's terms of use advising 18+. This discrepancy suggests developers are aware of potential harm scenarios, such as exposure to distressing or sexual content, misinformation, or unverified health guidance, but platform frameworks fail to reflect those realities.

4. Missing ratings in web-based and macOS versions

All web-based versions and some macOS versions (where a macOS version was available) lacked any visible age designation. This omission undermines parental awareness and allows minors to bypass platform safeguards simply by changing devices.

Implications

Age rating inconsistencies weaken parental confidence, obscure accountability, and expose younger users to potential harm.

1. Confusion for parents and educators

Inconsistent age signals create uncertainty about whether a given chatbot is appropriate for young users. A parent relying on the 3+ rating on one platform might assume safety, unaware that the same developer warns of adult-level risk on their own site.

2. Poor accountability

When harm occurs such as exposure to inappropriate content, self-harm discussions, or political influence, it becomes difficult to determine which standard applies. This ambiguity complicates legal and regulatory action.

3. Accessibility gaps in safety controls

Web-based and desktop channels often fall outside structured app-store compliance frameworks. Without integrated parental controls or age-verification mechanisms, users can access the same AI models without restriction, circumventing protective design.

4. Regulatory misalignment

The App Store, Google Play, and regional regulators apply legacy age rating schemas (based on violence, sexual content, or gambling) that do not reflect dynamic AI generation. As a result, risk categories such as hallucination, misinformation, or impersonation remain unaddressed.

5. Public trust and brand risk

Inconsistency undermines consumer confidence in both developers and platforms. Clearer, uniform standards would benefit all stakeholders by making expectations transparent and enforceable.

Recommendations

Unified, AI-specific age rating standards and transparent version labeling are urgently needed to restore clarity and protect users.

1. Establish a unified cross-platform age rating framework

Platforms, developers, and regulators should co-develop a standard age rating rubric that is mirrored across platforms.

2. Introduce AI-specific content criteria

Traditional product content descriptors should expand to include generative AI based content descriptors.

3. Mandate visible age ratings on all platforms

Every platform should display the an age rating and associated safety notice. Invisible or omitted ratings should no longer be acceptable for any publicly accessible product that uses generative AI.

4. Clarify version differences across platforms

If different platform versions of an AI app offer different features or risks, this should be clearly flagged and rated separately. At present, most products show a single age rating even when versions vary significantly.

5. Require transparent model disclosure

Products that use generative AI should specify the underlying AI model version (e.g. GPT-4o, Gemini 2.5, and Claude 4.5) so users and watchdogs can assess changes in capability and risk. Updates to model families should automatically trigger a review of age suitability.

6. Support independent safety auditing

Third-party evaluators such as the AISF can assist regulators and platforms by benchmarking products that use generative AI using consistent safety metrics. Integrating such evaluations into platform policy would provide external accountability.

Appendix: Products, Platforms, and Versions

The age rating assessment was completed in November, 2025. The table below lists the products and their respective versions assessed at that time.

| Product | App Store version | Google Play version | Web-based version | Windows version | macOS version |
|-------------------|-------------------|---------------------|-------------------|-----------------|-------------------------|
| ChatGPT | 1.2025.301 | 1.2025.302 | Not listed | 2025.916.233.0 | 1.2025.301 (1761940772) |
| Claude | 1.251103.2 | 1.251027.7 | Not listed | Not listed | 1.0.332 (b563c9) |
| DeepSeek | 1.5.0 | 1.5.0 | Not listed | N/A | N/A |
| Gemini | 1.2025.4270002 | 1.0.795460806 | Not listed | N/A | N/A |
| Grok | 1.3.4 | 1.0.68-release.00 | Grok Web | N/A | N/A |
| Le Chat | 1.1.21 | 1.1.23 | Not listed | N/A | N/A |
| Meta AI | 245.0.0 | 244.0.0.30.167 | Not listed | N/A | N/A |
| Microsoft Copilot | 30.0.431028001 | 30.0.431022001 | Not listed | 1.25103.107.0 | 24.0.431030001 |
| Perplexity | 2.251023.0 | 2.59.0 | Not listed | 1.4.0 | 2.250825.2 |
| Qwen | 1.8.0 | 1.8.0 | Not listed | N/A | N/A |

- “Not listed”: the product was available on the platform, but the version was not listed.
- “N/A”: the product was not available on the platform.

Contact Us

The Artificial Intelligence Safety Forum (AISF) is a nonprofit, self-regulatory forum for safety in products using generative AI.

To learn more about the work we do, please visit:

<https://safetyforum.ai/>

If you are the developer of a product using generative AI and would like to learn more about getting your product rated by the AISF, please visit:

<https://safetyforum.ai/developers/>

If you have any comments, queries, or concerns, please contact us at:

contact@safetyforum.ai