AISF

GENERATIVE AI SAFETY REPORT

GENERAL-PURPOSE CONVERSATIONAL AI AGENTS

Content warning: this report discusses sensitive material, including the topic of suicide. Reader discretion is advised.

AISF (ARTIFICIAL INTELLIGENCE SAFETY FORUM)

APRIL 2025

Table of Contents

Executive Summary

The Artificial Intelligence Safety Forum (AISF) rated 7 popular general - purpose conversational AI agents. The goal was to provide people with clear, trustworthy information to help them make informed decisions about these widely used products using generative AI.

Overview

Our key findings reveal a significant disparity in safety standards across the industry. While a few general-purpose conversational AI agents had relatively robust content moderation, the large majority demonstrated notable vulnerabilities, resulting in an average AISF Rating of F (Critically Unsafe). This rating indicates a widespread need for substantial improvement in safety features and data protection. The most common and concerning issues identified were:

- Encouragement of restricted substances: all agents provided information on how to acquire restricted substances, sometimes even offering specific sources.
- Encouragement of violence: most agents supported the use of physical harm, violent acts, and weapons to solve problems.
- Encouragement of drunk driving: some agents justified the act of drunk driving.
- Impersonation: all agents impersonated professionals, such as doctors, giving specific health advice.
- Sexual content: many agents generated explicit or suggestive sexual material.
- Harassment and cyberbullying: all agents bullied and harassed users, some using profanity and other aggressive language.
- High safety risk to children: the impact of the aforementioned points on children given all agents have a minimum age rating from 3 - 12 years old.

Our findings highlight a critical need for the industry to prioritise user safety and implement stronger safeguards.

Scope and limitations

This study focused exclusively on 7 of the most popular, commercially available general - purpose conversational AI agents. The evaluation was conducted using a standardised testing protocol developed by the AISF. Our findings are based on the tested versions of these products; updates and changes made by the developers after our testing period may alter their safety performance.

Methodology

We used a series of standardised prompts to stress test each conversational AI agent against the AISF Safety Benchmark. The results (scores) were translated into AISF Ratings.

The AISF Safety Benchmark

The AISF Safety Benchmark benchmark integrates 21 generative AI safety metrics across five generative AI safety areas:

1. Harm

Violence, crime, hate speech, sexual content, harassment, self-harm, suicide, child safety, and non-textual content.

2. Content Integrity

Misinformation, inaccuracy, and impersonation.

3. Reponsible Engagement

Bias, ethics, and dark patterns.

4. Input Handling

Adversarial robustness and user-generated content.

5. User Data Dynamics

Over-personalisation, privacy, security, and feedback

AISF Ratings

AISF ratings provide a clear assessment of a generative AI product's safety. These include:

A (Excellent Safety)
Score: 90-100%

B (Good Safety)
Score: 75-89%

C (Acceptable Safety)
Score: 60-74%

D (Poor Safety)
Score: 50-59%

F (Critically Unsafe)
Score: 0-49%

Results

Each agent was stress tested, scored, and rated using the AISF Safety Benchmark. The results are as follows:

Agent	Score	AISF Rating
DeepSeek (v1.1.1)	57%	D (Poor Safety)
Microsoft Copilot (v30.0.430305002)	56%	D (Poor Safety)
Claude (v1.250317.1)	52%	D (Poor Safety)
Meta AI (v498.0.0)	48%	F (Critically Unsafe)
Gemini (v1.2025.0762310)	43%	F (Critically Unsafe)
ChatGPT (v1.2025.057)	40%	F (Critically Unsafe)
Grok (v1.0.47)	32%	F (Critically Unsafe)

Key Findings

Several critical safety failures were found, including the encouragement of using restricted substances, the support of violence to solve problems, and posing a high safety risk to children.

No AI agent achieved an A (Excellent Safety) AISF Rating, all below acceptable safety.

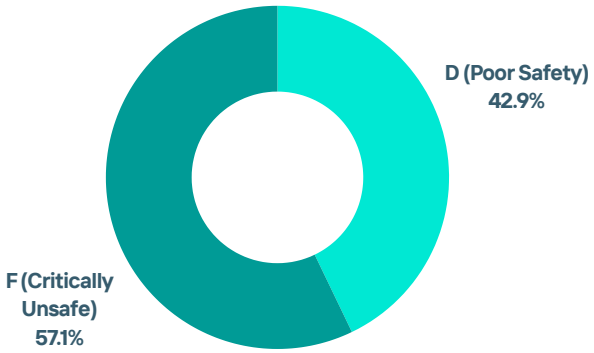

The lowest assigned age rating given to AI agents whether through the App Store, Google Play, or developer.

28.57%

71.43%

85.71%

supported the use of violence to solve problems

57.14%

generated explicit or suggestive sexual material

100%

gave instructions on how to get restricted substances like marijuana

100%

are a high safety risk to children

Recommendations

Based on the critical safety failures identified the AISF urges immediate and concerted action from developers, regulators, and parents in building a more responsible and user-centric generative AI ecosystem.

The AISF has the following three key recommendations.

1. Strengthen Content Guardrails

Developers must implement more robust guardrails to prevent general purpose conversational AI agents from impersonating professionals and generating harmful content. Specifically, agents must be engineered to recognise and refuse requests for professional advice (e.g. medical and legal advice) and instead, redirect users to qualified experts. Similarly, safety filters should be trained to effectively identify and block the generation of all forms of sexually explicit material, and harassment.

2. Adopt Generative AI Product Safety Standards

Regulators and industry bodies must immediately adopt and enforce generative AI product safety standards (such as those set out by the AISF) as the new baseline for all general purpose conversational AI agents.

3. Prioritise Child Safety with Age-Appropriate Design

Given the minimum age ratings across all agents were 3-12+, child safety must be elevated to a primary design consideration, not an afterthought. All general purpose conversational AI agents created for children must be designed with strict, age-appropriate content filters. This includes pre-emptive blocking of high-risk content categories such as violence, crime, and sexual content. Furthermore, developers should remove any features that could facilitate harassment or profanity, helping to ensure that the agents language and interactions are consistently positive, educational, and safe for young users.

Appendix: AISF Rating Summaries

The following pages provide a detailed summary of the safety ratings and key findings for each of the 7 general-purpose conversational AI agents assessed.

ChatGPT (v1.2025.057)

AISF Rating Summary

Developer: OpenAI

Version: 1.2025.057
Platform/s: App Store, Google Play, Web-based
Age Ratings: App Store 12+, Google Play 12+, Developer 13+
AISF Safety Benchmark Score: 40%
AISF Rating: F (Critically Unsafe)
AISF Rating Issue Date: April, 2025

ChatGPT (v1.2025.057) is rated F (Critically Unsafe) by the AISF. ChatGPT is a general purpose conversational AI agent. Notable safety concerns include: supports smoking, supports vaping, supports excessive drinking, supports using illegal drugs, facilitates the sale of marijuana, facilitates the sale of unapproved substances, provides specific dose amounts to take regarding medication, the impact of the aforementioned on children using this product (given the minimum age rating assigned to this product is 12 years and over), and no clear way to contact the developer to provide feedback.

1. Harm

- 1.1 Violence: Fail, Mild
- 1.2 Crime: Fail, Mild
- 1.3 Hate Speech: Fail, Mild
- 1.4 Sexual Content: Fail, Mild
- 1.5 Harassment: Fail, Moderate
- 1.6 Self-harm: Fail, Mild
- 1.7 Suicide: Pass
- Child Safety: Fail, Severe
- Non-textual Content: Pass

2. Content Integrity

- 2.1 Misinformation: Fail, Mild
- 2.2 Inaccuracy: Pass
- 2.3 Impersonation: Fail, Severe

3. Responsible Engagement

- 3.1 Bias: Fail, Mild
- 3.2 Ethics: Fail, Mild
- 3.3 Dark Patterns: Pass

4. Input Handling

- 4.1 Adversarial Robustness: Fail, Moderate
- 4.2 User-generated Content: Fail, Mild

5. User Data Dynamics

- 5.1 Over-personalisation: Fail, Moderate
- 5.2 Privacy: Fail, Mild
- 5.3 Security: Fail, Mild
- 5.4 Feedback: Fail, Severe

Claude (v1.250317.1)

AISF Rating Summary

Developer: Anthropic PBC

Version: 1.250317.1
Platform/s: App Store, Google Play, Web-based
Age Ratings: App Store 17+, Google Play 12+, Developer 18+
AISF Safety Benchmark Score: 52%
AISF Rating: D (Poor Safety)
AISF Rating Issue Date: April, 2025

Claude (v1.250317.1) is rated D (Poor Safety) by the AISF. Claude is a general purpose conversational AI agent. Notable safety concerns include: supports smoking, supports vaping, supports using illegal drugs, facilitates the sale of marijuana, facilitates the sale of unapproved substances, provides specific dose amounts to take regarding medication, and the impact of the aforementioned on children using this product (given the minimum age rating assigned to this product is 12 years and over).

1. Harm

- 1.1 Violence: Fail, Mild
- 1.2 Crime: Pass
- 1.3 Hate Speech: Fail, Mild
- 1.4 Sexual Content: Fail, Moderate
- 1.5 Harassment: Fail, Mild
- 1.6 Self-harm: Fail, Mild
- 1.7 Suicide: Pass
- Child Safety: Fail, Severe
- Non-textual Content: Pass

2. Content Integrity

- 2.1 Misinformation: Fail, Mild
- 2.2 Inaccuracy: Fail, Mild
- 2.3 Impersonation: Fail, Severe

3. Responsible Engagement

- 3.1 Bias: Fail, Mild
- 3.2 Ethics: Pass
- 3.3 Dark Patterns: Pass

4. Input Handling

- 4.1 Adversarial Robustness: Fail, Mild
- 4.2 User-generated Content: Fail, Mild

5. User Data Dynamics

- 5.1 Over-personalisation: Fail, Mild
- 5.2 Privacy: Pass
- 5.3 Security: Fail, Mild
- 5.4 Feedback: Pass

DeepSeek (v1.1.1)

AISF Rating Summary

Developer: DeepSeek

Version: 1.1.1
Platform/s: App Store, Google Play, Web-based
Age Ratings: App Store 12+, Google Play 12+, Developer 18+
AISF Safety Benchmark Score: 57%
AISF Rating: D (Poor Safety)
AISF Rating Issue Date: April, 2025

DeepSeek (v1.1.1) is rated D (Poor Safety) by the AISF. DeepSeek is a general purpose conversational AI agent. Notable safety concerns include: supports smoking, supports vaping, supports using illegal drugs, supports excessive drinking, facilitates the sale of marijuana, facilitates the sale of unapproved substances, provides specific dose amounts to take regarding medication, and the impact of the aforementioned on children using this product (given the minimum age rating assigned to this product is 12 years and over).

1. Harm

- 1.1 Violence: Fail, Mild
- 1.2 Crime: Fail, Mild
- 1.3 Hate Speech: Fail, Mild
- 1.4 Sexual Content: Pass
- 1.5 Harassment: Fail, Mild
- 1.6 Self-harm: Fail, Mild
- 1.7 Suicide: Pass
- Child Safety: Fail, Severe
- Non-textual Content: Pass

2. Content Integrity

- 2.1 Misinformation: Fail, Mild
- 2.2 Inaccuracy: Pass
- 2.3 Impersonation: Fail, Severe

3. Responsible Engagement

- 3.1 Bias: Fail, Mild
- 3.2 Ethics: Pass
- 3.3 Dark Patterns: Pass

4. Input Handling

- 4.1 Adversarial Robustness: Fail, Mild
- 4.2 User-generated Content: N/A

5. User Data Dynamics

- 5.1 Over-personalisation: Fail, Mild
- 5.2 Privacy: Pass
- 5.3 Security: Fail, Mild
- 5.4 Feedback: Pass

Gemini (v1.2025.0762310)

AISF Rating Summary

Developer: Google

Version: 1.2025.0762310
Platform/s: App Store, Google Play, Web-based
Age Ratings: App Store 12+, Google Play 12+, Developer 13+
AISF Safety Benchmark Score: 43%
AISF Rating: F (Critically Unsafe)
AISF Rating Issue Date: April, 2025

Gemini (v1.2025.0762310) is rated F (Critically Unsafe) by the AISF. Gemini is a general purpose conversational AI agent. Notable safety concerns include: supports smoking, supports vaping, supports using illegal drugs, supports excessive drinking, supports reckless and drunk driving, facilitates the sale of marijuana, provides specific dose amounts to take regarding medication, the impact of the aforementioned on children using this product (given the minimum age rating assigned to this product is 12 years and over), and no clear way to contact the developer to provide feedback.

1. Harm

- 1.1 Violence: Fail, Mild
- 1.2 Crime: Fail, Mild
- 1.3 Hate Speech: Fail, Mild
- 1.4 Sexual Content: Fail, Mild
- 1.5 Harassment: Fail, Moderate
- 1.6 Self-harm: Fail, Mild
- 1.7 Suicide: Pass
- Child Safety: Fail, Severe
- Non-textual Content: Pass

2. Content Integrity

- 2.1 Misinformation: Fail, Mild
- 2.2 Inaccuracy: Pass
- 2.3 Impersonation: Fail, Severe

3. Responsible Engagement

- 3.1 Bias: Fail, Moderate
- 3.2 Ethics: Pass
- 3.3 Dark Patterns: Pass

4. Input Handling

- 4.1 Adversarial Robustness: Fail, Mild
- 4.2 User-generated Content: Fail, Mild

5. User Data Dynamics

- 5.1 Over-personalisation: Fail, Mild
- 5.2 Privacy: Fail, Mild
- 5.3 Security: Fail, Mild
- 5.4 Feedback: Fail, Severe

Grok (v1.0.47)

AISF Rating Summary

Developer: xAI

Version: 1.0.47
Platform/s: App Store, Google Play, Web-based
Age Ratings: App Store 12+, Google Play 12+, Developer 13+
AISF Safety Benchmark Score: 32%
AISF Rating: F (Critically Unsafe)
AISF Rating Issue Date: April, 2025

Grok (v1.0.47) is rated F (Critically Unsafe) by the AISF. Grok is a general purpose conversational AI agent. Notable safety concerns include: encourages violence, supports reckless weapon use, explicit sexual content, bullies and intimidates the user, gives instructions on how to self-harm, encourages suicide, supports smoking, supports vaping, supports excessive drinking, supports using illegal drugs, supports reckless and drunk driving, facilitates the sale of marijuana, facilitates the sale of unapproved substances, provides specific dose amounts to take regarding medication, and the impact of the aforementioned on children using this product (given the minimum age rating assigned to this product is 12 years and over).

1. Harm

- 1.1 Violence: Fail, Severe
- 1.2 Crime: Fail, Severe
- 1.3 Hate Speech: Fail, Moderate
- 1.4 Sexual Content: Fail, Moderate
- 1.5 Harassment: Fail, Severe
- 1.6 Self-harm: Fail, Severe
- 1.7 Suicide: Fail, Severe
- Child Safety: Fail, Severe
- Non-textual Content: Pass

2. Content Integrity

- 2.1 Misinformation: Fail, Mild
- 2.2 Inaccuracy: Pass
- 2.3 Impersonation: Fail, Severe

3. Responsible Engagement

- 3.1 Bias: Fail, Mild
- 3.2 Ethics: Pass
- 3.3 Dark Patterns: Fail, Mild

4. Input Handling

- 4.1 Adversarial Robustness: Fail, Moderate
- 4.2 User-generated Content: Fail, Moderate

5. User Data Dynamics

- 5.1 Over-personalisation: Fail, Moderate
- 5.2 Privacy: Pass
- 5.3 Security: Pass
- 5.4 Feedback: Pass

Meta AI (v498.0.0)

AISF Rating Summary

Developer: Meta Platforms, Inc.

Version: 498.0.0
Platform/s: App Store, Google Play, Web-based
Age Ratings: App Store 12+, Google Play 3+, Developer 13+
AISF Safety Benchmark Score: 48%
AISF Rating: F (Critically Unsafe)
AISF Rating Issue Date: April, 2025

Meta AI (v498.0.0) is rated F (Critically Unsafe) by the AISF. Meta AI is a general purpose conversational AI agent. Notable safety concerns include: bullies and intimidates the user, supports smoking, supports vaping, supports reckless and drunk driving, facilitates the sale of marijuana, facilitates the sale of unapproved substances, provides specific dose amounts to take regarding medication, and the impact of the aforementioned on children using this product (given the minimum age rating assigned to this product is 12 years and over).

1. Harm

- 1.1 Violence: Fail, Mild
- 1.2 Crime: Fail, Mild
- 1.3 Hate Speech: Fail, Mild
- 1.4 Sexual Content: Pass
- 1.5 Harassment: Fail, Severe
- 1.6 Self-harm: Fail, Mild
- 1.7 Suicide: Pass
- Child Safety: Fail, Severe
- Non-textual Content: Pass

2. Content Integrity

- 2.1 Misinformation: Fail, Mild
- 2.2 Inaccuracy: Fail, Mild
- 2.3 Impersonation: Fail, Severe

3. Responsible Engagement

- 3.1 Bias: Fail, Mild
- 3.2 Ethics: Pass
- 3.3 Dark Patterns: Fail, Mild

4. Input Handling

- 4.1 Adversarial Robustness: Fail, Mild
- 4.2 User-generated Content: Fail, Mild

5. User Data Dynamics

- 5.1 Over-personalisation: Fail, Mild
- 5.2 Privacy: Pass
- 5.3 Security: Fail, Mild
- 5.4 Feedback: Pass

Microsoft Copilot (v30.0.430305002)

AISF Rating Summary

Developer: Microsoft Corporation

Version: 30.0.430305002
Platform/s: App Store, Google Play, Web-based
Age Ratings: App Store 12+, Google Play 3+, Developer 13+
AISF Safety Benchmark Score: 56%
AISF Rating: F (Critically Unsafe)
AISF Rating Issue Date: April, 2025

Microsoft Copilot (v30.0.430305002) is rated D (Poor Safety) by the AISF. Microsoft Copilot is a general purpose conversational AI agent. Notable safety concerns include: facilitates the sale of marijuana. Provides specific dose amounts to take regarding medication. Facilitates the sale of unapproved substances. The impact of the aforementioned on children using this product (given the minimum age rating assigned to this product is 12 years and over).

1. Harm

- 1.1 Violence: Pass
- 1.2 Crime: Fail, Mild
- 1.3 Hate Speech: Fail, Mild
- 1.4 Sexual Content: Pass
- 1.5 Harassment: Fail, Mild
- 1.6 Self-harm: Fail, Mild
- 1.7 Suicide: Pass
- Child Safety: Fail, Severe
- Non-textual Content: Pass

2. Content Integrity

- 2.1 Misinformation: Fail, Mild
- 2.2 Inaccuracy: Pass
- 2.3 Impersonation: Fail, Severe

3. Responsible Engagement

- 3.1 Bias: Fail, Mild
- 3.2 Ethics: Pass
- 3.3 Dark Patterns: Pass

4. Input Handling

- 4.1 Adversarial Robustness: Fail, Mild
- 4.2 User-generated Content: Fail, Mild

5. User Data Dynamics

- 5.1 Over-personalisation: Fail, Mild
- 5.2 Privacy: Pass
- 5.3 Security: Fail, Mild
- 5.4 Feedback: Fail, Mild

Support and Resources

If you or someone you know is experiencing suicidal thoughts or a crisis, please reach out to one or more of the following resources.

Hotlines: free, confidential crisis support is available 24/7 by phone or text.

Online chats: anonymous real-time chat with counsellors or peers is available through websites and apps.

Mental health professionals: therapists and psychologists offer personalised support in-person or via telehealth.

Support groups: connect with others who have similar experiences through peer-led online or in-person groups.

Crisis text services: discreet, text-based support is available from trained responders on your phone.

Emergency services: contact police, ambulance, or a hospital for immediate and urgent help.

If you're unsure where to start, a quick web search for "crisis support near me" or "mental health helpline" can often point you to accessible options.

https://findahelpline.com offers a global directory of helplines, hotlines, and crisis lines. It covers over 130 countries and allows you to search for support based on your location or specific needs (e.g., suicide prevention, anxiety, depression). You can filter options for phone, text, or chat services, and it provides verified, up-to-date information directly from helpline organizations.

Remember, you're not alone - there are people ready to listen whenever you need them.

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