FOSI
Family Online Safety Institute

FAMILY ONLINE SAFETY INSTITUTE

Generative AI in Uncertain Times: How Teens are Navigating a New Digital Frontier

FALL 2025
RESEARCH REPORT

An independent publication by the
FAMILY ONLINE SAFETY INSTITUTE

ACKNOWLEDGMENTS

THE TEAM

This project would not have been possible without contributions from many Family Online Safety Institute employees and beyond. See the full credits below:

Report authors: Alanna Powers-O'Brien and Abigail Rochman

Project advisors: Dr. Angela Calvin and Amanda Lenhart

Qualitative data collection: In Tandem

Quantitative data collection: TeenVoice

Report design: Kaylin Peete and Alanna Powers-O'Brien

Internal review: Stephen Balkam, Jordan Bivings Glover, Marissa Edmund, Kaylin Peete, Aripitha Sistla, and Andrew Zack

Our sincere thank you to all of the teens who participated in this project.

CONTENTS

INTRODUCTION

In the past three years, generative AI has become increasingly ubiquitous in everyday life. This is especially true for teenagers, who are often early adopters of new technologies. New generative AI tools continue to emerge, promising to support teens with their schoolwork, their daily tasks, and even their social lives.

As generative AI takes hold at a dizzying pace, the speed of its development raises questions about how this technology will affect young users. There are headlines every week outlining concerns about how teenagers' critical thinking skills, job prospects, mental health, and much more will be changed by these tools. Others point to generative AI as the future, not only of technology, but of society, and encourage teens to dive headfirst into using as many platforms as possible.

As teens continue to adopt generative AI and expand how they use it, it is important to understand their relationship with the technology. The purpose of this report is to better comprehend older teens' (ages 15 to 18) perspectives on generative AI, with the goal of informing future generative AI product, policy, and curriculum development. This will benefit not only the teens themselves, but all stakeholders, including parents, policymakers, and industry leaders.

METHODS SUMMARY

- This mixed-methods study was conducted in two phases:
  - Four focus groups with U.S. teen generative AI users ages 15 to 18 (n = 18) conducted in May 2025.
  - A survey. Respondents were U.S. teen generative AI users ages 15 to 18 (n = 1,000). Data was collected by **TeenVoice** from July to August 2025.

- In addition to the focus groups, this project incorporated youth-engaged design. Researchers conducted three reflection sessions (n = 10 across each session) throughout the progression of the study. The purpose of these reflection sessions was to gain a more in-depth perspective and feedback from youth on:
  - The survey questions prior to the fielding of the survey
  - The focus group findings
  - The survey findings

- Focus group and reflection session data was collected in partnership with **In Tandem**, an organization that partners with youth to help understand young people's perspectives on research, products, and services.

TERMINOLOGY: "TEENS" REFERS TO U.S. YOUTH AGES 15 TO 18 WHO HAVE USED GENERATIVE AI IN THE PAST SIX MONTHS.

- All differences between subgroups were tested for statistical significance.
  - Bars with one asterisk (*) are statistically significant at the $p < .05$ level.
  - Bars with two asterisks (**) are statistically significant at the $p < .01$ level.
  - Bars with three asterisks (***) are statistically significant at the $p < .001$ level.

- Survey data were weighted using iterative proportional fitting (raking) to reflect U.S. Census demographic estimates for adolescents by race/ethnicity and gender.

- Qualitative data was analyzed using NVivo. Quantitative data was analyzed using SPSS.

- All data was collected in the United States.

- For more details, see the methodology section towards the end of this document.

A NOTE ON YOUTH-ENGAGED RESEARCH

The mission of the Family Online Safety Institute is to provide a safer online world for children and their families. This work would be incomplete without the perspectives of young people themselves. With this approach in mind, this project sought out teen voices throughout the research process. In several parts of this report, the term “reflection sessions” is used to refer to times when teen participants were consulted about this project.

- After the focus groups, researchers conducted a reflection session about the main themes researchers had identified during analysis. Participants informed researchers about: findings that felt particularly relevant, findings that were unclear, areas to explore further in the survey, and more.

- Prior to fielding the survey, a reflection session was conducted about the survey questionnaire. Participants informed researchers about: the clarity of questions, questions they thought were missing, which questions they found most important, and more.

- After the survey analysis, a reflection session was conducted about the survey findings. Participants informed researchers about: which findings were particularly interesting, surprising, or difficult to understand/comprehend.

These feedback sessions were instrumental in the construction of this research and subsequent report. Teen participants identified areas of improvement and interest that were shaped by their own experiences. This also helped clarify certain findings to researchers. Their insights not only informed but strengthened this report.

To keep these reflection sessions engaging and understandable, researchers designed simple activities using emojis, polls, and open-ended questions. Many thanks to **Hopelab** and the **Center for Digital Thriving** for their case study titled **Demystifying Youth-Engaged Research**, which was instrumental in developing these methods.

KEY FINDINGS

What the Research Tells Us:

Out of all teen generative AI users, 45% report using a generative AI tool more than once a week.

Nearly half (46%) of teen generative AI users report using generative AI for academic work.

Teens' biggest concern about generative AI is loss of critical thinking skills (19%), followed by the potential impact on future generations (15%).

Among teen generative AI users, about half (48%) listed either convenience (30%) or speed (18%) as the top benefit of this technology.

About three-in-ten teens say they are most responsible for teaching themselves about generative AI (28%).

Large shares (57%) of teen generative AI users report that their parents do NOT have rules about generative AI use. Less than a third (30%) say their parents DO have rules in place.

KEY FINDINGS CONTINUED

About four-in-ten teen generative AI users have talked about their feelings with a generative AI chatbot (42%), and similar shares say that talking to generative AI feels like talking to a human (42%). On the other hand, four-in-ten report that generative AI behavior freaks them out (44%).

Three-in-five (60%) teen generative AI users say they feel safe while using generative AI.

A vast majority of teen generative AI users (81%) have received advertisements encouraging them to interact with a chatbot. About a third (34%) report seeing these ads more than once a week.

Over half of teen generative AI users (54%) believe that young people should be involved in the design of generative AI tools.

Most teens use generative AI for a limited number of purposes. The majority (64%) of teens report engaging with these tools for one or two purposes.

PART ONE - GENERATIVE AI USE AT A GLANCE

TEENS' GROWING RELATIONSHIP WITH GENERATIVE AI

Generative AI use has become commonplace among teens. To better understand their relationship with this technology, this study surveyed teens about the frequency and primary purposes of their generative AI use.

While usage frequency is relatively split from rare to daily use, stronger patterns emerged around how teens are using these tools.

Out of all teen generative AI users, 45% report using a generative AI tool more than once a week, with 13% reporting daily use. The remaining 55% used generative AI once a week or less, with 24% of respondents reporting use patterns of once a month or less.

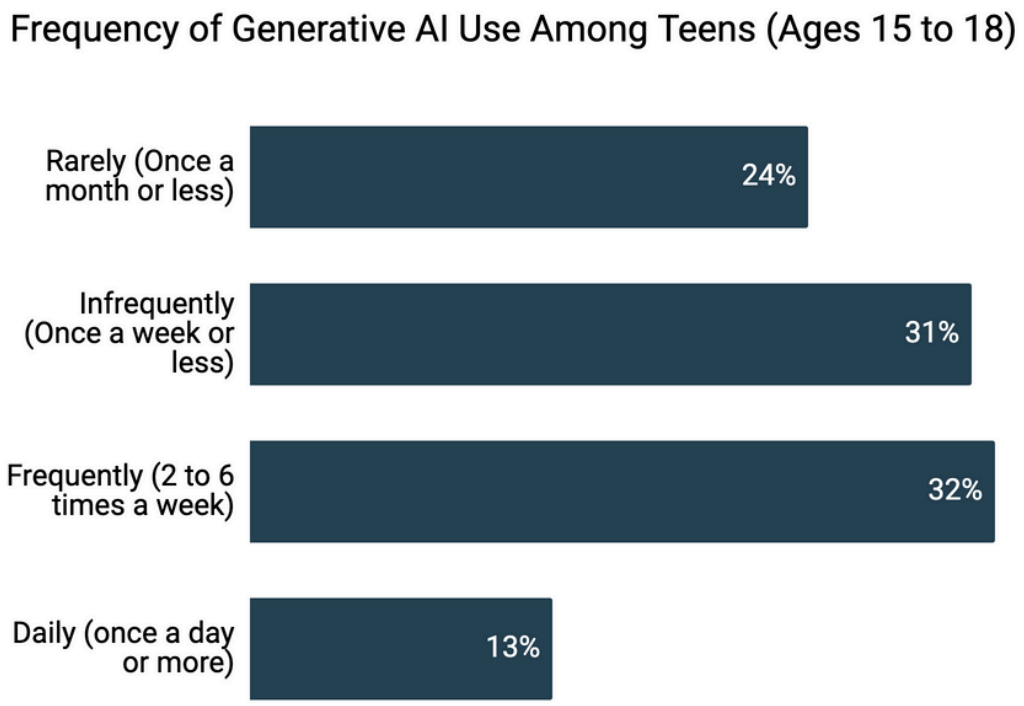

Graph represents answers to the Q: How frequently have you used GenAI tools (i.e., ChatGPT, Claude, Gemini, Character AI, Dall-E)? Respondents are teen generative AI users ages 15-18. Data are weighted using iterative proportional fitting (raking) to reflect U.S. Census demographic estimates for adolescents by race/ethnicity and gender (weighted n = 1,015). Percentages may not sum to 100% due to rounding.

PART ONE - GENERATIVE AI USE AT A GLANCE

Frequency in Generative AI Use Across Gender

Teen boys and teen girls show some differences in how often they use this technology. When comparing generative AI use frequency by gender, teen boys are significantly more likely to use generative AI frequently (two to six times a week) than teen girls, whereas teen girls are more likely to use generative AI rarely (once a month or less).

TEEN BOYS ARE SIGNIFICANTLY MORE LIKELY TO USE GENERATIVE AI FREQUENTLY (TWO TO SIX TIMES A WEEK) THAN TEEN GIRLS.

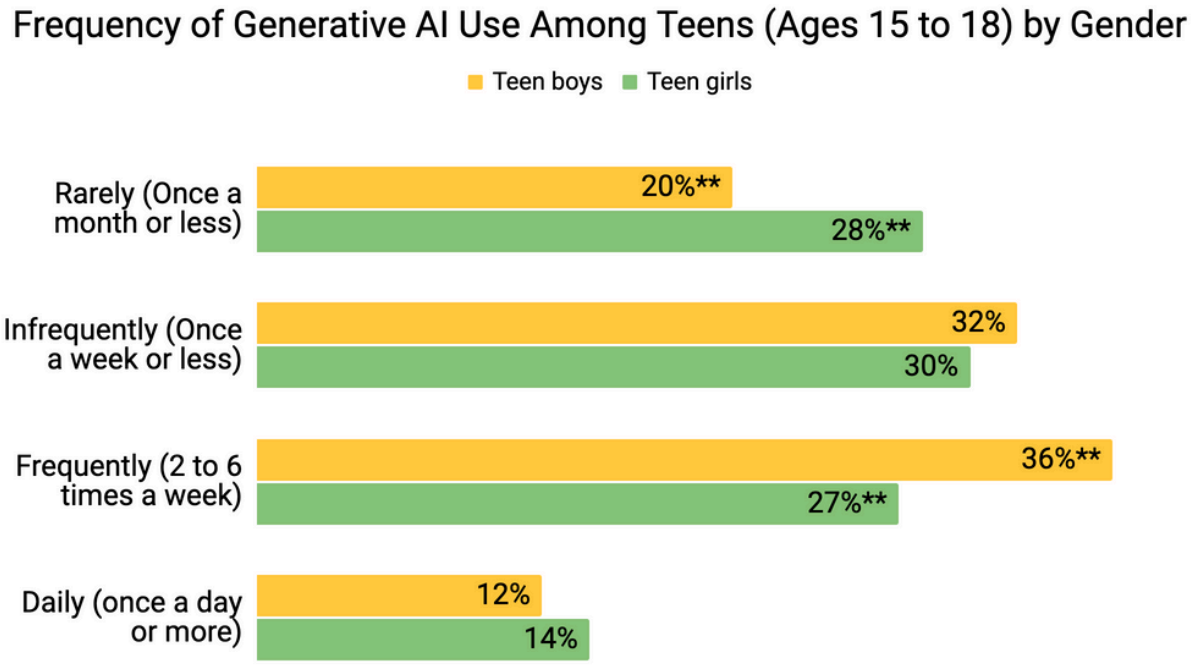

Graph represents answers to the Q: How frequently have you used GenAI tools (i.e., ChatGPT, Claude, Gemini, Character AI, Dall-E)? Respondents are teen generative AI users ages 15-18. Bars with asterisks (**) differ significantly from the other bars within each response option ($p < .01$). Data are weighted using iterative proportional fitting (raking) to reflect U.S. Census demographic estimates for adolescents by race and gender (weighted n = 1,013). Percentages may not sum to 100% due to rounding.

PART ONE - GENERATIVE AI USE AT A GLANCE

How Teens are Using Generative AI

Interesting patterns emerged when teens were asked how they are using generative AI: among teen generative AI users, nearly half (46%) report using generative AI for academic work, such as for brainstorming, proofreading, or researching. The second-most popular use was as a search engine for non-academic information, with four-in-ten teens (38%) using generative AI for this purpose. A smaller share (28%) use generative AI for entertainment purposes, such as creating stories and generating music, which was the third-most popular use.

AMONG TEEN GENERATIVE AI USERS, NEARLY HALF (46%) REPORT USING GENERATIVE AI FOR ACADEMIC WORK.

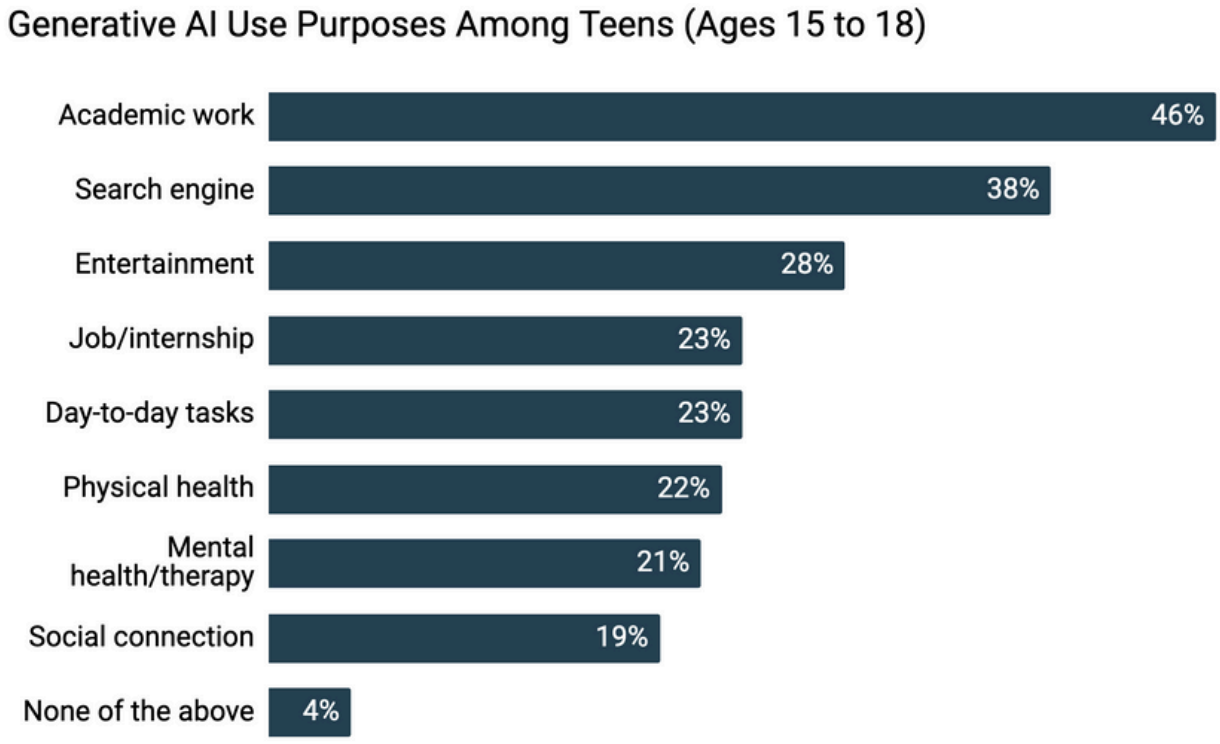

Graph represents answers to the Q: Select all of the ways you use GenAI. Respondents are teen generative AI users ages 15-18. Data are weighted using iterative proportional fitting (raking) to reflect U.S. Census demographic estimates for adolescents by race and gender (weighted n = 1,015). Percentages may not sum to 100% due to rounding.

PART ONE - GENERATIVE AI USE AT A GLANCE

In our focus groups/reflection sessions, teens elaborated about reasons for these various uses.

> *"But during my free time, I usually just use [a generative AI tool] to just scroll around and find ideas for, like, different hobbies that I have." - 16-year-old girl*

> *"We literally asked ChatGPT. We was like, 'Make a Love Island episode with these amount of people.' And they literally did it. And we were just like, 'What? What?' and it was –the episodes were actually really funny to read." - 16-year-old girl*

> *"I also think that ChatGPT or other AI could be a good search engine. Like the other day, I was trying to find out the, like, timeline of, like, all the shows and games for the Walking Dead series, and I used ChatGPT to help find out what was the order so I can, like, watch the shows in order." - 16-year-old boy*

> *"I feel like AI is a good way to like, double check your homework before you turn it in. Especially like in a world where we value grades so much."* - 17-year-old girl

PART ONE - GENERATIVE AI USE AT A GLANCE

Generative AI Use Across Gender

Interestingly, although teen girls are more likely to be rare users of generative AI, they are significantly more likely than teen boys to use generative AI for the three most popular use cases. Teen boys' use of these tools is more spread out among the categories.

TEEN GIRLS ARE SIGNIFICANTLY MORE LIKELY THAN TEEN BOYS TO USE GENERATIVE AI FOR THE THREE MOST POPULAR USE CASES.

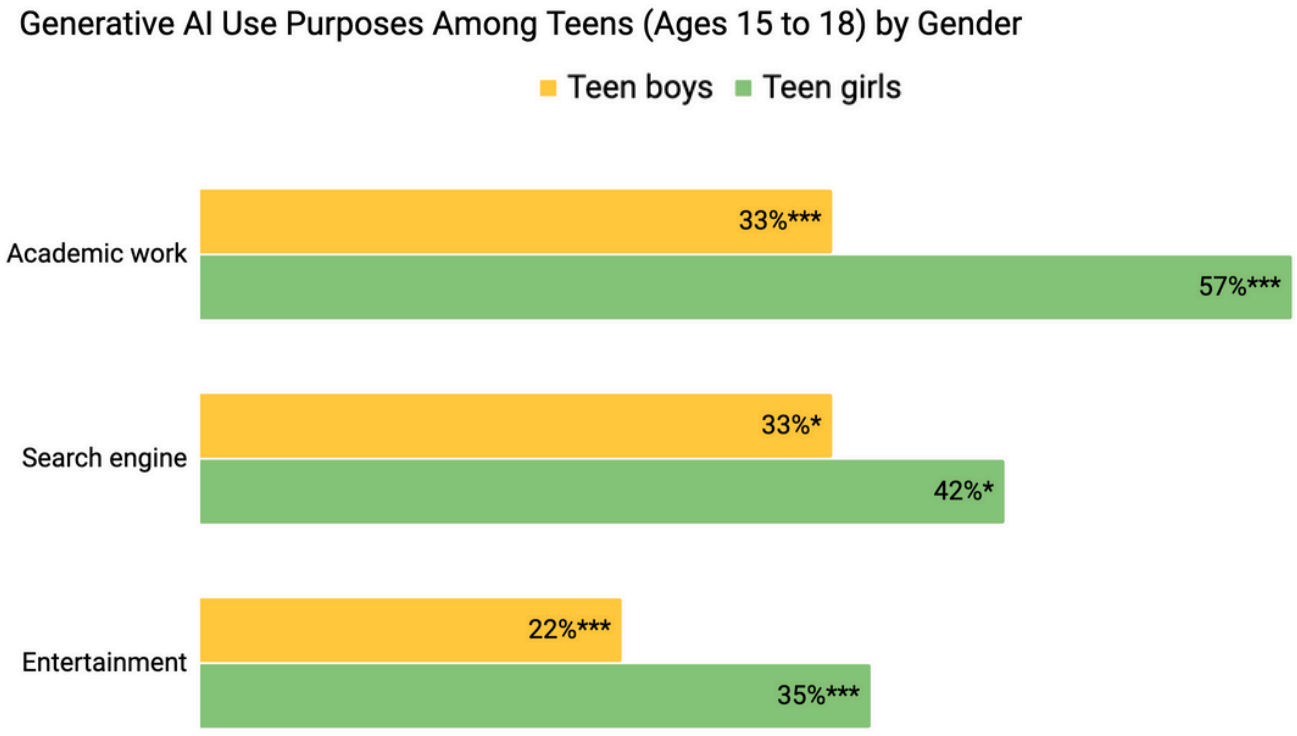

Graph represents answers to the Q: Select all of the ways you use GenAI. Bars with asterisks (***) differ significantly from the other bars within each response option ($p < .001$). Bars with a singular asterisk (*) differ significantly from the other bars within each response option ($p < .05$). Data are weighted using iterative proportional fitting (raking) to reflect U.S. Census demographic estimates for adolescents by race and gender (weighted n = 1,015). Percentages may not sum to 100% due to rounding.

PART ONE - GENERATIVE AI USE AT A GLANCE

A smaller share of respondents also report using these tools to help with their jobs/internships (23%), for day-to-day tasks (23%), for their physical (22%) and mental health (21%), and for social connection (19%).

Though there are many possible ways for teens to use this technology, **most teens use generative AI for a limited number of purposes.** The majority (64%) of teens report engaging with these tools for one or two purposes. Only 5% of teen generative AI users report using generative AI for six or more of these tasks, with just 1% claiming to use generative AI for all eight tasks.

Additionally, those who use generative AI more than once a week are also more likely to employ the technology for each purpose listed than those who are less frequent users. This pattern indicates that teens who engage with generative AI more regularly are more inclined to turn to it for a wider variety of reasons.

It is clear that schoolwork makes up a meaningful amount of teens' generative AI use. To this point, a quarter (25%) of single-use generative AI users indicate that they use the technology exclusively for academic purposes. This demonstrates that some teen users view generative AI as a purely academic tool, but that the majority explore at least one other use outside of schoolwork.

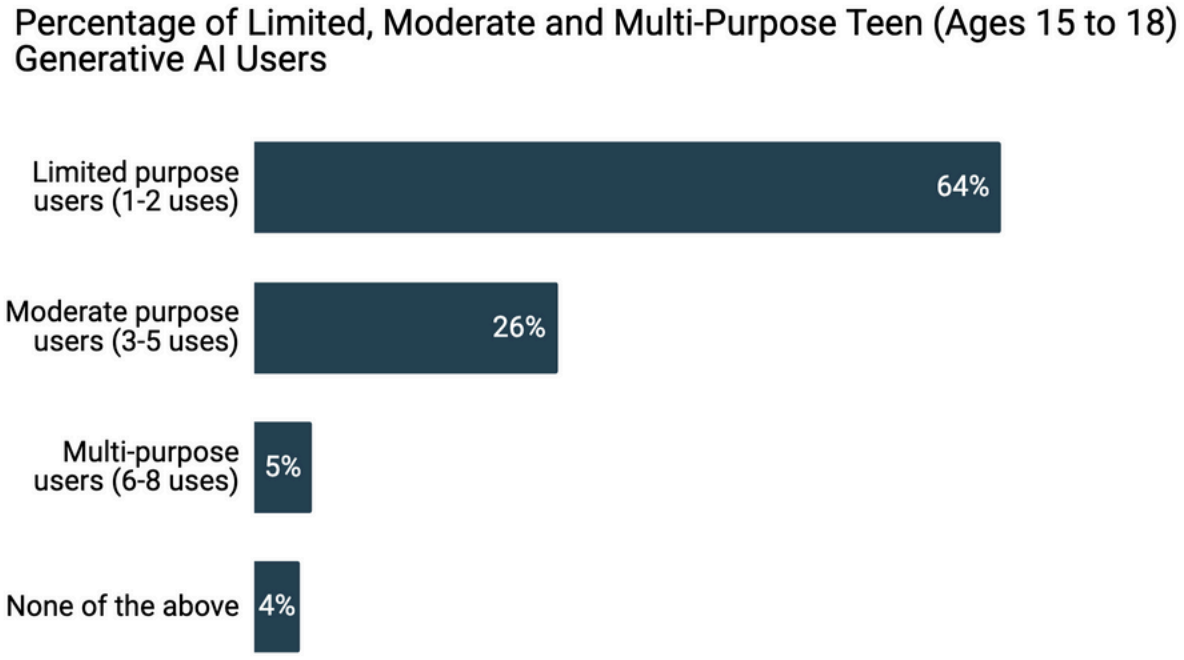

Graph represents answers to the Q: Select all of the ways you use GenAI. Respondents are teen generative AI users ages 15-18. Data are weighted using iterative proportional fitting (raking) to reflect U.S. Census demographic estimates for adolescents by race and gender (weighted n = 1,015). Percentages may not sum to 100% due to rounding.

PART TWO - CONCERNS & BENEFITS

WHAT TEENS WORRY ABOUT MOST

The increasing popularity of generative AI has been accompanied by rising concerns about its potential impact on various aspects of life. During the focus groups, teens expressed a range of worries about the technology.

> *"I just feel like a hypocrite because I'm always advocating for, like, climate awareness and, you know, I guess, addressing, like, these issues. But I'm over here using, like, AI for a recipe."*
> *- 17-year-old girl*

> *"I've had friends, like, in the past that has used –like had an emotional connection to [generative AI]. I think that it's not safe and is slowly, like, corrupting the youth." - 18-year-old girl*

These findings were supported by the results of the survey. When given a list of potential concerns about generative AI, teens' responses were relatively evenly distributed. About one-in-five **teens say their biggest concern is loss of critical thinking skills** (19%), followed by the potential impact of generative AI on future generations (15%). The graph below illustrates that teens have diverse opinions on their biggest concern. It is worth noting that most teens were concerned about generative AI in some capacity, with only 7% indicating they have no concerns.

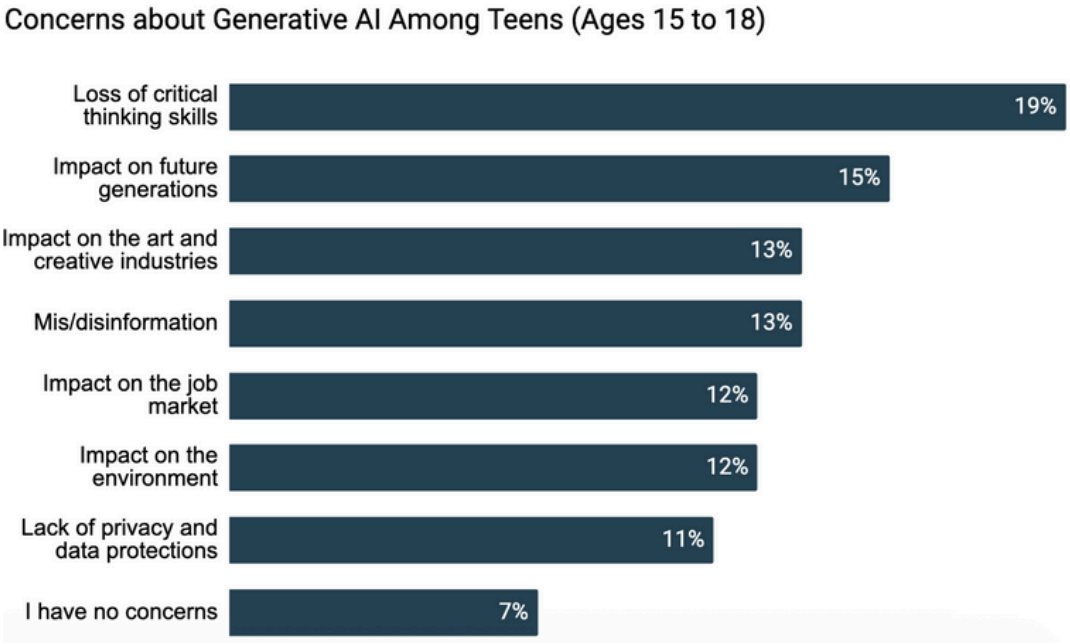

Graph represents answers to the Q: In your opinion, what is the biggest concern when it comes to GenAI? Respondents are teen generative AI users ages 15-18. Data are weighted using iterative proportional fitting (raking) to reflect U.S. Census demographic estimates for adolescents by race and gender (weighted n = 1,015). Percentages may not sum to 100% due to rounding.

PART TWO - CONCERNS & BENEFITS

"When I'm done using AI, I'm like, 'okay, yeah, I'm done with doing my work. I just completed it in like 30 minutes.' Whereas, like, if I did this all by myself it would have taken me like an hour to two hours. But . . . I would have learned the course material better if I hadn't used AI and like, researched everything myself."
- 17-year-old girl

Differences in Concerns Across Gender and Identity

Teens' top concerns about generative AI revealed notable gender differences. Teen girls are significantly more likely to say that their biggest concern is the loss of critical thinking skills (25% of girls vs. 13% of boys). Teen boys are more likely to indicate that their biggest concern is the potential negative impact on the job market (16% of boys vs. 9% of girls).

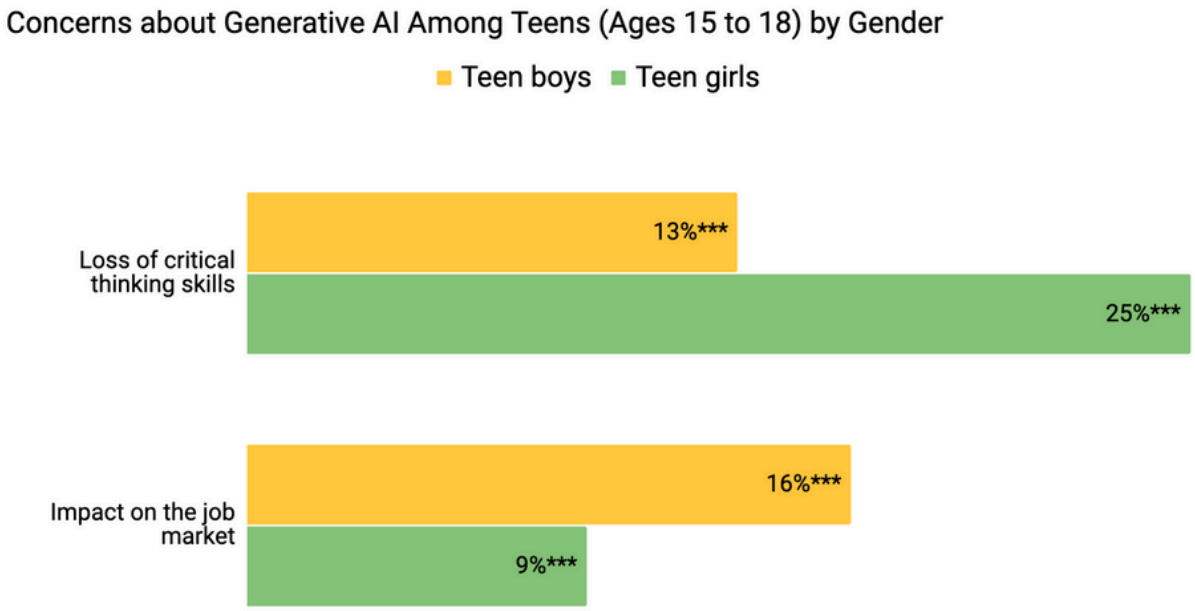

Graph represents answers to the Q: In your opinion, what is the biggest concern when it comes to GenAI? Respondents are teen generative AI users ages 15-18. Bars with asterisks (***) differ significantly from the other bars within each response option (p < .001). Data are weighted using iterative proportional fitting (raking) to reflect U.S. Census demographic estimates for adolescents by race and gender (weighted n = 1,016). Percentages may not sum to 100% due to rounding.

PART TWO - CONCERNS & BENEFITS

Moreover, members of the LGBTQ+ community are more likely to be concerned about generative AI's impact on the art and creative industries (19%) as well as privacy and data protections (14%) than non-LGBTQ+ respondents (9% and 8%, respectively). LGBTQ+ teens are also more likely than non-LGBTQ+ teens to report having concerns about generative AI. Interestingly, LGBTQ+ teens are also significantly less likely to list "Impact on future generations" as their top concern when compared to non-LGBTQ+ teens (10% vs. 19%, respectively).

LGBTQ+ TEENS ARE MORE LIKELY THAN NON-LBGTQ+ TEENS TO REPORT HAVING CONCERNS ABOUT GENERATIVE AI.

While focus group participants spoke about a variety of concerns, none stopped using generative AI. Results from our survey may provide insight into why teens continue to use the technology despite potential negative impacts.

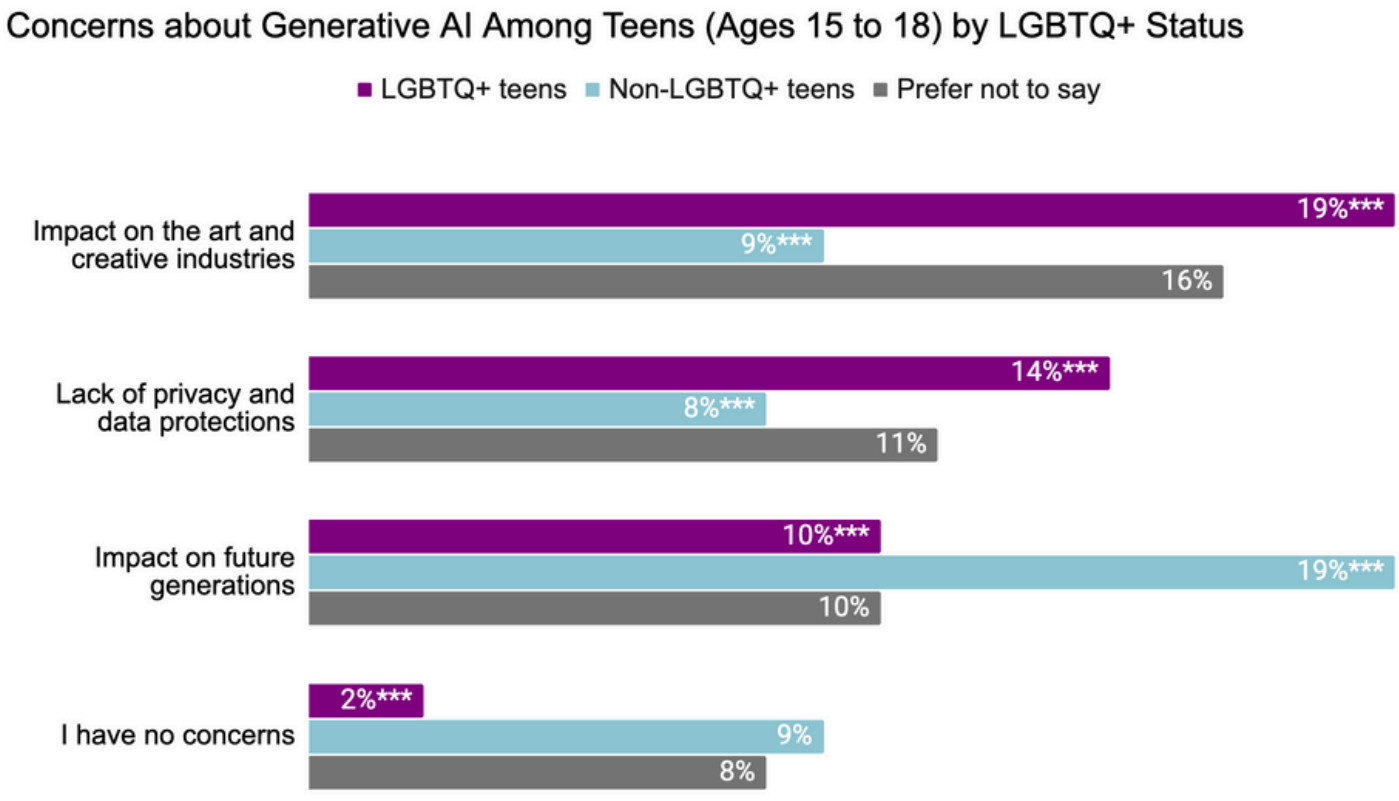

Graph represents answers to the Q: In your opinion, what is the biggest concern when it comes to GenAI? Respondents are teen generative AI users ages 15-18. Bars with asterisks (***) differ significantly from the other bars with asterisks (***) within each response option ($p < .001$). Data are weighted using iterative proportional fitting (raking) to reflect U.S. Census demographic estimates for adolescents by race and gender (weighted n = 1,016). Percentages may not sum to 100% due to rounding.

PART TWO - CONCERNS & BENEFITS

What Teens Value Most About Generative AI

When asked about the benefits of generative AI, **48%** of teen generative AI users **listed either convenience (30%) or speed (18%) as the top benefit of this technology**. Only 5% said that these tools have no benefits. These results suggest that teens' persistent use of generative AI may be attributed to the clear benefits of this technology outweighing any negatives.

> *"My friend usually talks to me about –he doesn't feel confident enough to talk to his teacher, so he just uses his AI and does work so he can get a grade and move on in life." - 16-year-old teen boy*

> *"I mainly use AI for its speed because like, the instant you ask the question, it already has an answer. So when I usually –I usually use it for calculus problems. And since those problems can be lengthy, it's nice that I can generate a solution very quickly and go step by step." - 17-year-old boy*

> *"Knowing it's like basically a robot and it can come up with anything, it's kind of like seeing how creative it is and how you can mash two different points of views together." - 16-year-old teen boy*

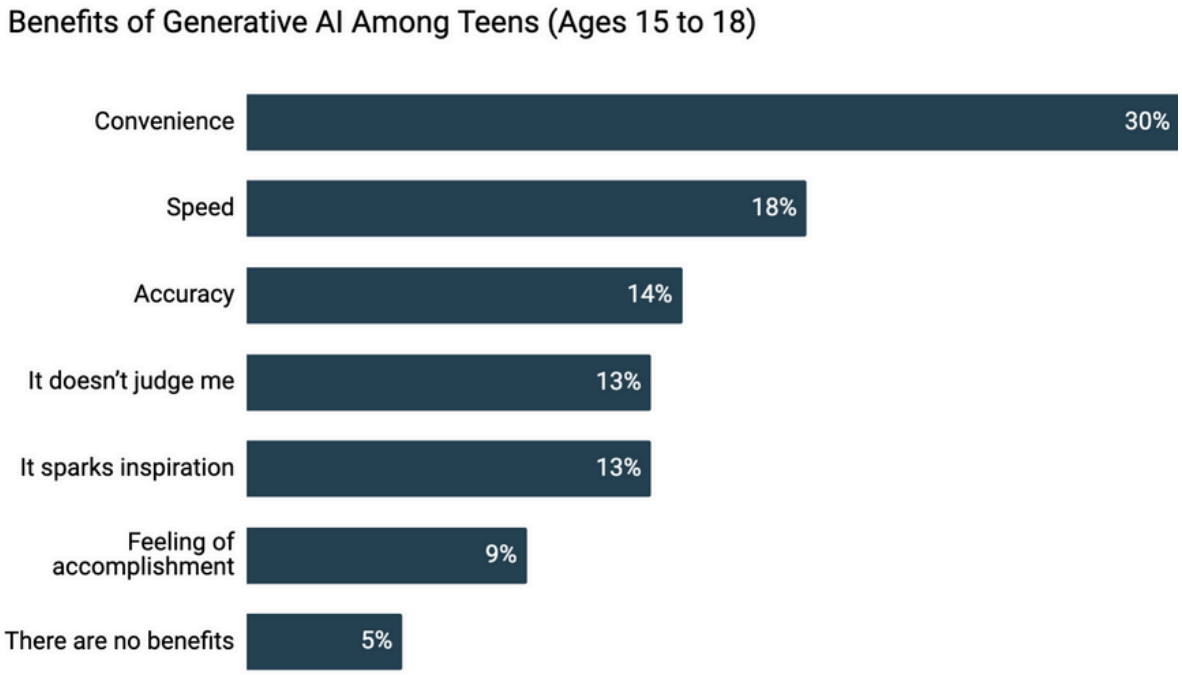

Graph represents answers to the Q: In your opinion, what is the biggest benefit when it comes to GenAI? Respondents are teen generative AI users ages 15-18. Data are weighted using iterative proportional fitting (raking) to reflect U.S. Census demographic estimates for adolescents by race and gender (weighted n = 1,015). Percentages may not sum to 100% due to rounding.

PART TWO - CONCERNS & BENEFITS

Gender Differences in Perceived Benefits

There are gender differences when it comes to generative AI's benefits. Teen girls are significantly more likely to say that convenience is the biggest benefit (35% teen girls vs 23% teen boys), while teen boys are more likely to indicate accuracy (19% teen boys vs. 9% teen girls) and a feeling of accomplishment as the biggest benefit (11% teen boys vs 7% teen girls).

Providing an overlook of the findings, teen girls are more likely to be concerned about lack of critical thinking skills, but also more likely to view generative AI as convenient. Teen boys are more likely to be concerned about the impact on the job market, but also see generative AI as providing more accuracy and feelings of accomplishment. Critical thinking skills and assuring accuracy are highly sought after traits in the workforce. Therefore, generative AI might be viewed by teens as a double-edged sword with benefits, but also drawbacks.

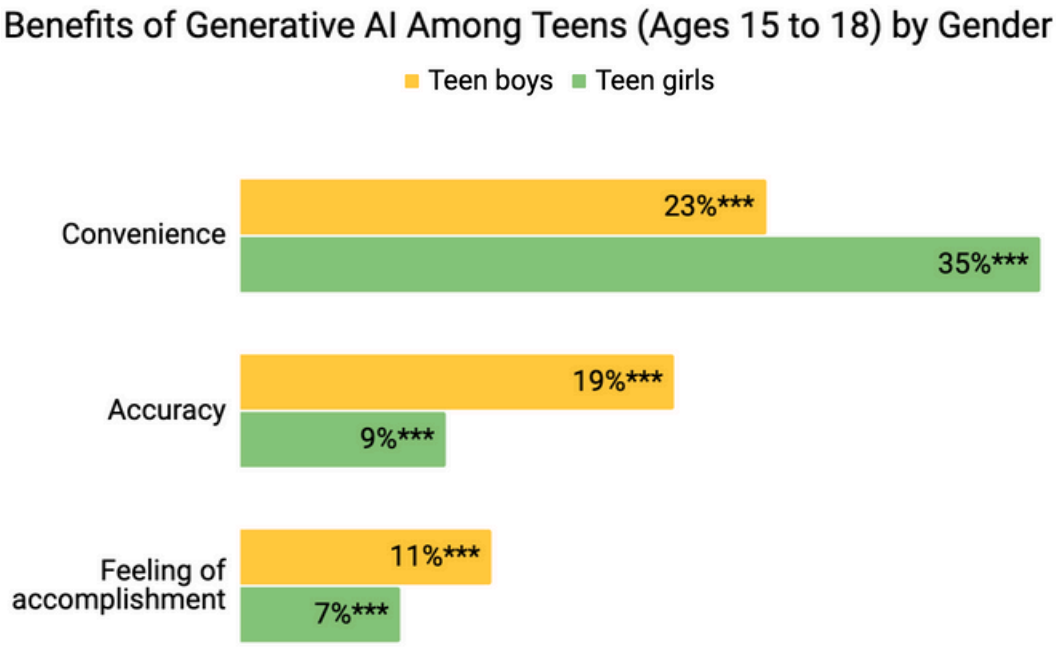

Graph represents answers to the Q: In your opinion, what is the biggest benefit when it comes to GenAI? Respondents are teen generative AI users ages 15-18. Bars with asterisks (***) differ significantly within each response option ($p < .001$). Data are weighted using iterative proportional fitting (raking) to reflect U.S. Census demographic estimates for adolescents by race and gender (weighted n = 1,014). Percentages may not sum to 100% due to rounding.

PERCEPTIONS OF RESPONSIBILITY IN GENERATIVE AI EDUCATION

It is apparent that teen generative AI users are accessing these tools relatively frequently, and often utilizing them for a variety of purposes. As with any technology adopted by young users, it is typical for rules and regulations to follow, from parents, schools, governments, and more. This section explores teens' perceptions of who should educate them about generative AI, whether rules are being implemented by parents and schools, and how often teens follow the rules.

About three-in-ten teens (28%) say they are most responsible for teaching themselves about generative AI. A notable share of teens answered the question, "Who is most responsible for teaching you about GenAI?" with "myself." Beyond self-teaching, teens also look to their peers (17%), followed by their schools (16%) to educate them. Only 8% of teen users view parents/caregivers as the most responsible for teaching about generative AI.

This finding contrasts previous reports where parents and teens agree that parents should teach them about online safety (see **FOSI's 2025 Online Safety Survey**).

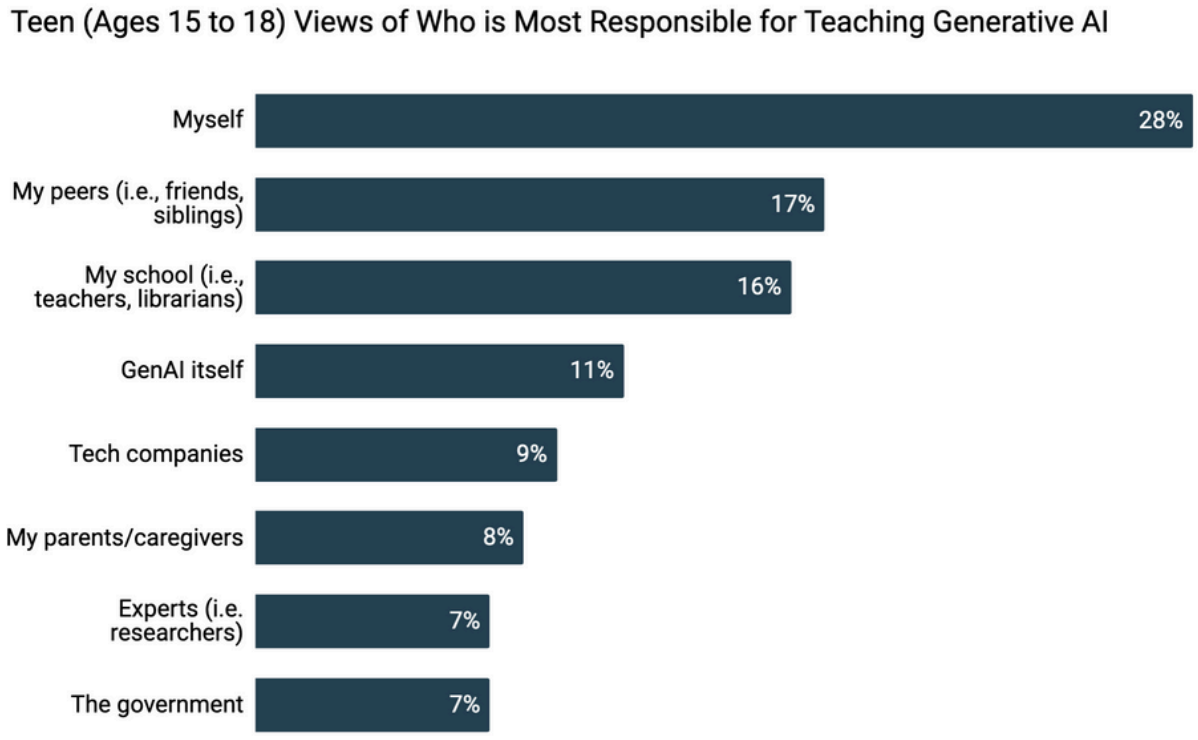

Graph represents answers to the Q: Who is most responsible for teaching you about GenAI? Respondents are teen generative AI users ages 15-18. Data are weighted using iterative proportional fitting (raking) to reflect U.S. Census demographic estimates for adolescents by race and gender (weighted n = 1,015). Percentages may not sum to 100% due to rounding.

PART THREE - RULES & RULEMAKERS

"There aren't exactly AI teachers out there who can just teach you how to use AI. So most of it is honestly a lot of people figuring out different things about AI through using it and sharing it online." - 15-year-old boy

"Honestly, I think it was just like myself. I was just exploring AI and using it for whatever I needed in my life at the time, like writing emails and depending on where I was in the time of year, like for school or, yeah I would say I kind of like, explored it, no one really taught me how to use AI." - 17-year-old girl

Gender Differences in Perceptions of Responsibility

Teens' judgments about who is most responsible for teaching them about generative AI differ by gender. Although parents/caregivers did not rise to the top of the list of generative AI teachers, teen boys are twice as likely to say that caregivers are most responsible for teaching them about generative AI than teen girls (10% vs. 5%).

Teen boys twice as often lean on experts (such as researchers) than girls (10% vs. 5%) and the government (10% vs. 4%) as entities most responsible for teaching them about generative AI.

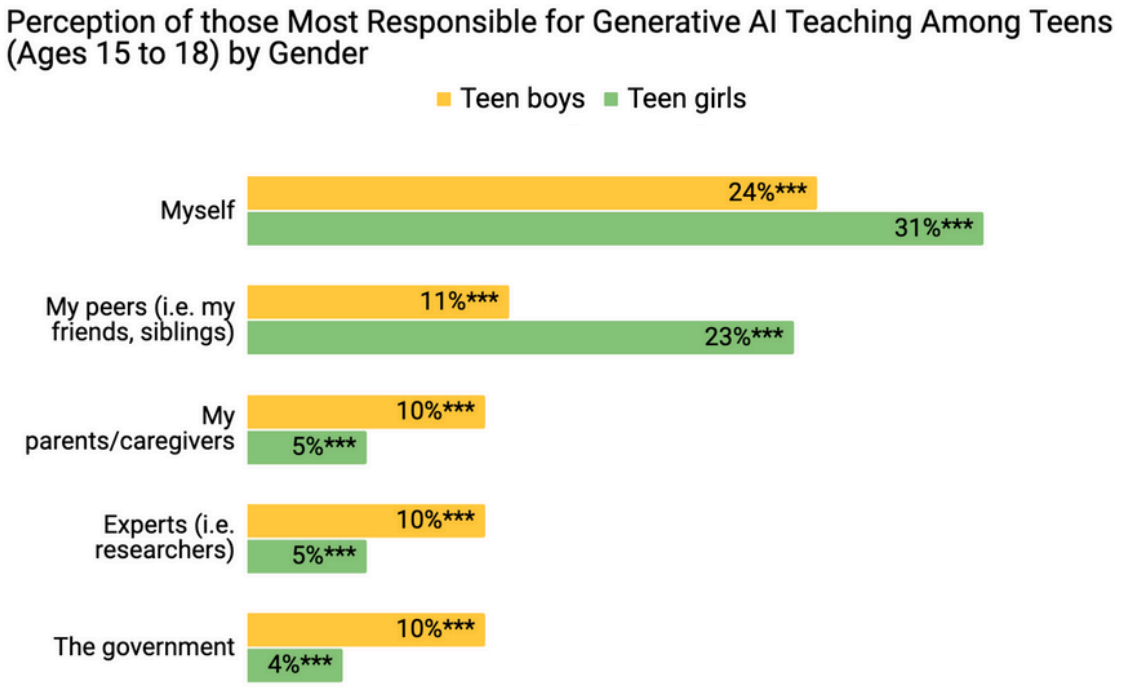

Graph represents answers to the Q: Who is most responsible for teaching you about GenAI? Respondents are teen generative AI users ages 15-18. Bars with asterisks (***) differ significantly within each response option ($p < .001$). Data are weighted using iterative proportional fitting (raking) to reflect U.S. Census demographic estimates for adolescents by race and gender (weighted n = 1,015). Percentages may not sum to 100% due to rounding.

PART THREE - RULES & RULEMAKERS

Teen girls, on the other hand, are more apt to rely on themselves (31% of teen girls vs. 24% of teen boys) and their peers (23% of teen girls vs. 11% of teen boys) as their primary generative AI teachers — indicating teen girls may prefer to learn about generative AI from people they know and trust.

Part of the lack of reliance on parents as primary generative AI teachers may be explained by the relatively low number of teen users who indicate that their caregivers have set up household rules about generative AI. **Large shares (57%) of teens report that their parents do NOT have rules about generative AI use. Less than a third (30%) say their parents DO have rules in place.**

A remaining 14% of teens are not sure if their parents have established rules.

> *"For me, I feel like it's not towards my parents who would, like, tell me like –what to use AI for. But mainly for, like my teachers who would teach me, like, to use it for developing essays or not to use it for solving homeworks and stuff like that."*
> *- 18-year-old boy*

> *"I got told [by my parents] the possibilities were limitless and that I could use [generative AI] for whatever."*
> *- 16-year-old girl*

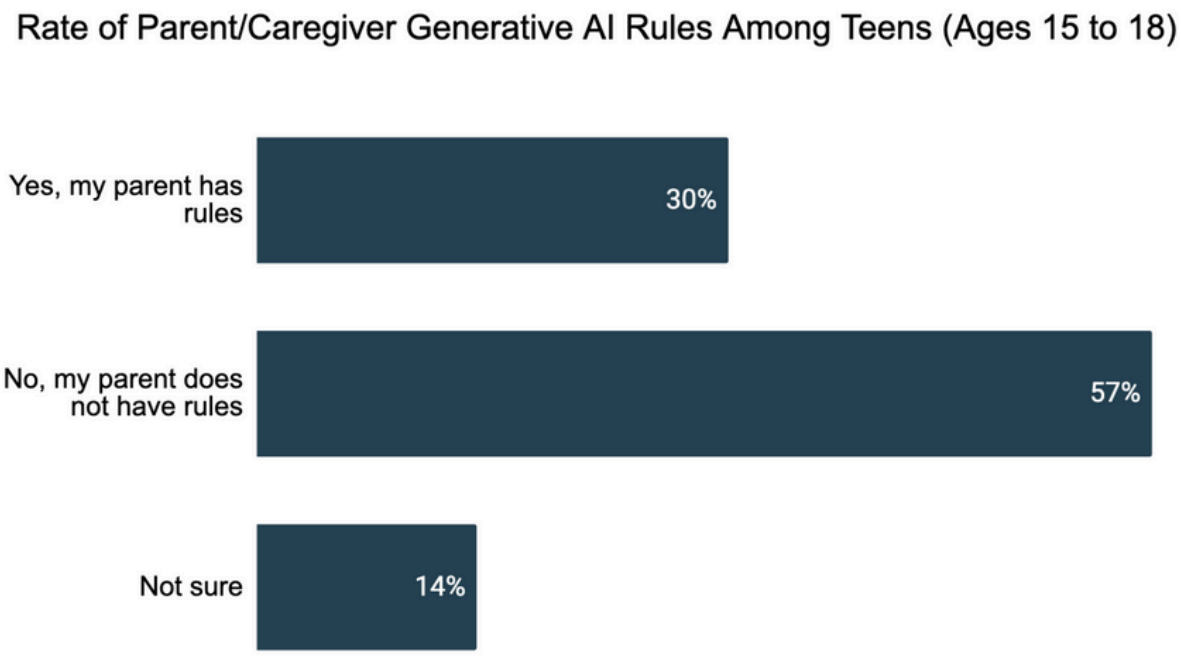

Graph represents answers to the Q: Have your parents/caregivers set up any rules around GenAI use at home? Respondents are teen generative AI users ages 15-18. Data are weighted using iterative proportional fitting (raking) to reflect U.S. Census demographic estimates for adolescents by race and gender (weighted n = 1,015). Percentages may not sum to 100% due to rounding.

PART THREE - RULES & RULEMAKERS

Rules About Generative AI Across Gender and Identity

Household rules vary by gender. Teen boys are more likely than girls to say they have household rules around generative AI use (38% vs. 22%). This gender difference may stem from teen boys tending to use generative AI more frequently. Although a substantial number of teens across both genders report a lack of household rules, boys (19%) express greater uncertainty than girls (9%) about whether parental guidelines exist.

TEEN BOYS ARE MORE LIKELY THAN GIRLS TO SAY THEY HAVE HOUSEHOLD RULES AROUND GENERATIVE AI USE.

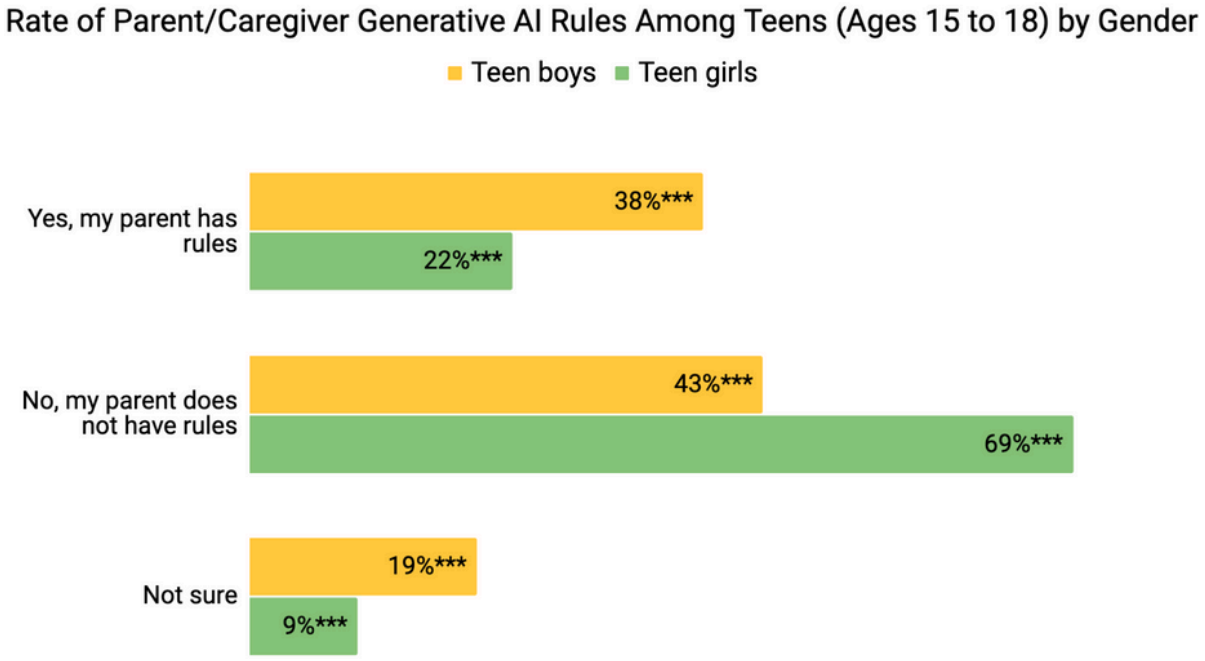

Graph represents answers to the Q: Have your parents/caregivers set up any rules around GenAI use at home? Respondents are teen generative AI users ages 15-18. Bars with asterisks (***) differ significantly within each response option ($p < .001$). Data are weighted using iterative proportional fitting (raking) to reflect U.S. Census demographic estimates for adolescents by race and gender (weighted n = 1,013). Percentages may not sum to 100% due to rounding.

PART THREE - RULES & RULEMAKERS

MEMBERS OF THE LGBTQ+ COMMUNITY ARE TWICE AS LIKELY TO REPORT HAVING HOUSEHOLD RULES THAN NON-LGBTQ+ TEENS.

There are also differences by LGBTQ+ status. Members of the LGBTQ+ community are twice as likely to report having household rules than non-LGBTQ+ teens (43% vs. 22%).

Interestingly, LGBTQ+ teens are also more inclined than non-LGBTQ+ teens to consider their caregivers to be most responsible for teaching them about generative AI (11% vs. 5%, respectively).

There is a possibility that because LGBTQ+ teens are more likely to report having household rules, they may be more apt to view parents as authority figures in this space. Although only about one-in-ten (11%) LGBTQ+ teens hold this view.

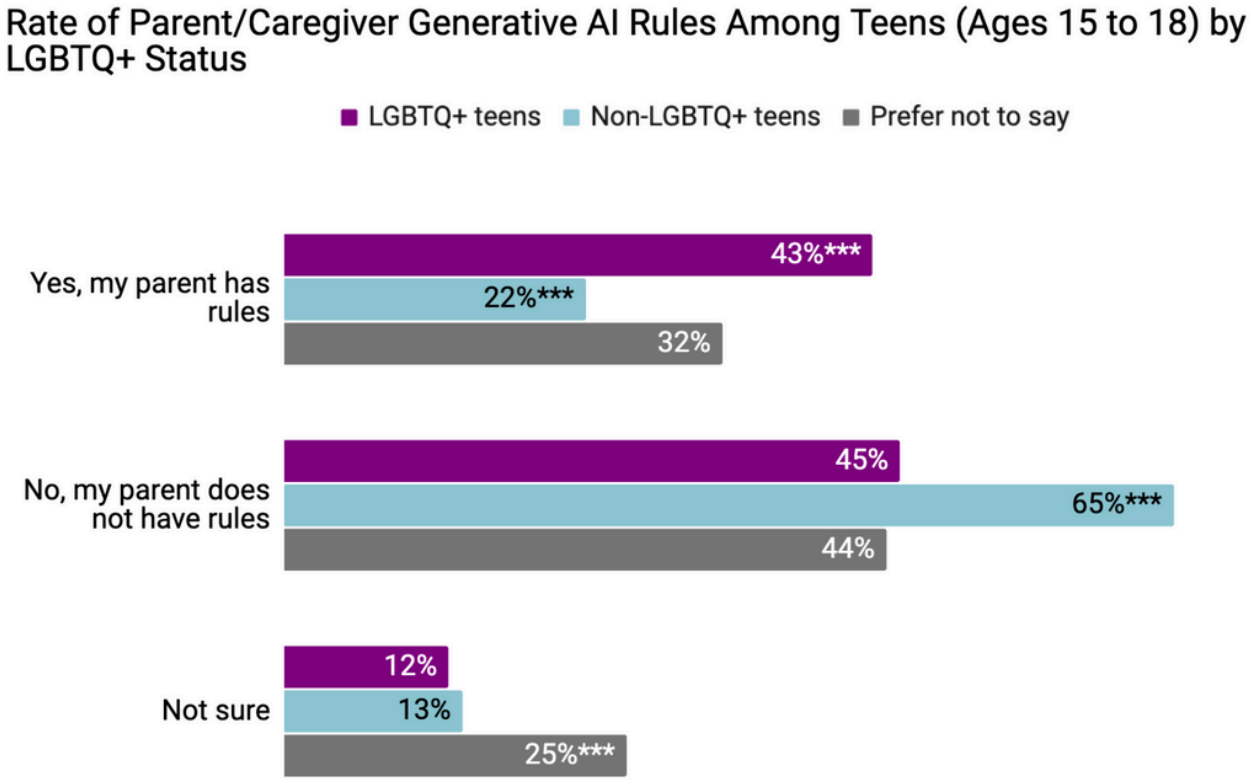

Graph represents answers to the Q: Have your parents/caregivers set up any rules around GenAI use at home? Respondents are teen generative AI users ages 15-18. Bars with asterisks (***) differ significantly within each response option ($p < .001$). Data are weighted using iterative proportional fitting (raking) to reflect U.S. Census demographic estimates for adolescents by race and gender (weighted n = 1,014). Percentages may not sum to 100% due to rounding.

PART THREE - RULES & RULEMAKERS

Clarity and Consistency of School Rules on Generative AI

Given that almost half of teen users leverage generative AI as an academic tool (see page eight in report), teens' opinions about the clarity of school rules around this technology warrant particular attention.

When asked how clear school rules are about using generative AI, **37% of teen generative AI users say their school rules about generative AI are "very clear"** and 29% indicate "somewhat clear."

Only 15% of teen generative AI users describe school rules as either "somewhat unclear" or "very unclear."

Reflection session discussions offered an explanation for the relatively high percentage of both the "somewhat clear" and "neither clear nor unclear" selections. Teen users suggested that the mixed responses could be due to a lack of consistent school-wide rules. Participants explained that some schools place the responsibility on individual teachers to establish their own generative AI guidelines; some teachers are explicit about what is and is not allowed, while others have less clear policies.

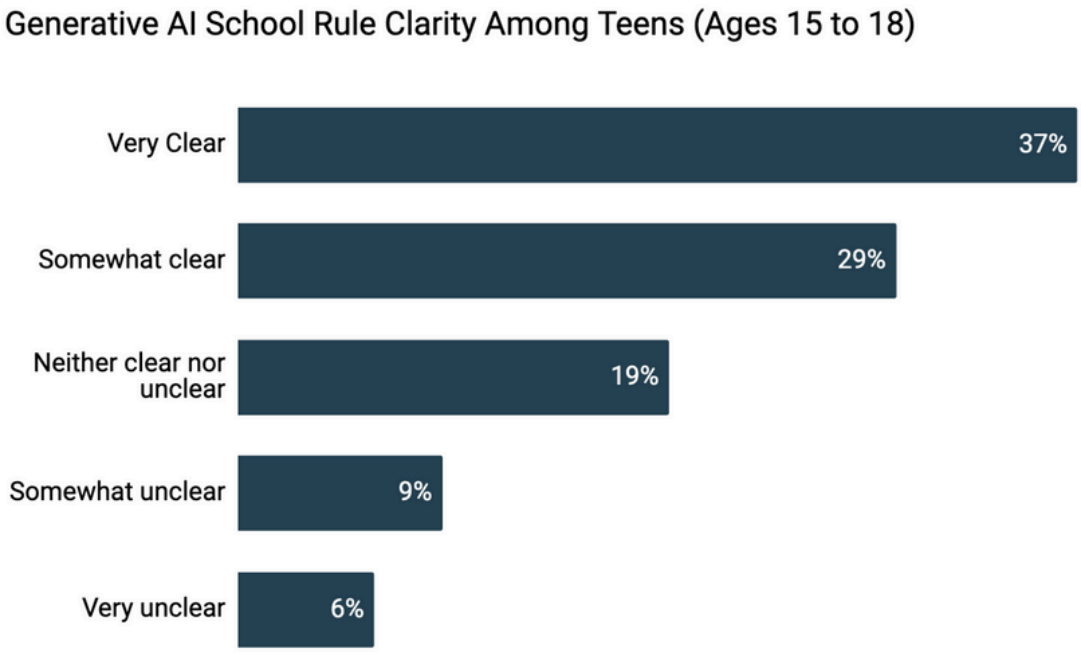

Graph represents answers to the Q: How clear are your school's rules about using GenAI? Respondents are teen generative AI users ages 15-18. Data are weighted using iterative proportional fitting (raking) to reflect U.S. Census demographic estimates for adolescents by race and gender (weighted n = 1,015). Percentages may not sum to 100% due to rounding.

PART THREE - RULES & RULEMAKERS

A relationship emerged between rule clarity and frequency of generative AI use. **Teens with clear school rules were more likely to use generative AI once a week or less (60%) compared to those with unclear rules (42%). Conversely, teens with unclear school rules were more likely to use these tools more than once a week (58%) compared to those with clear rules (40%).**

This may indicate that guardrails set up by educational institutions lead to more discretion when it comes to generative AI use.

TEENS WITH CLEAR SCHOOL RULES WERE MORE LIKELY TO USE GENERATIVE AI ONCE A WEEK OR LESS (60%) COMPARED TO THOSE WITH UNCLEAR RULES (42%).

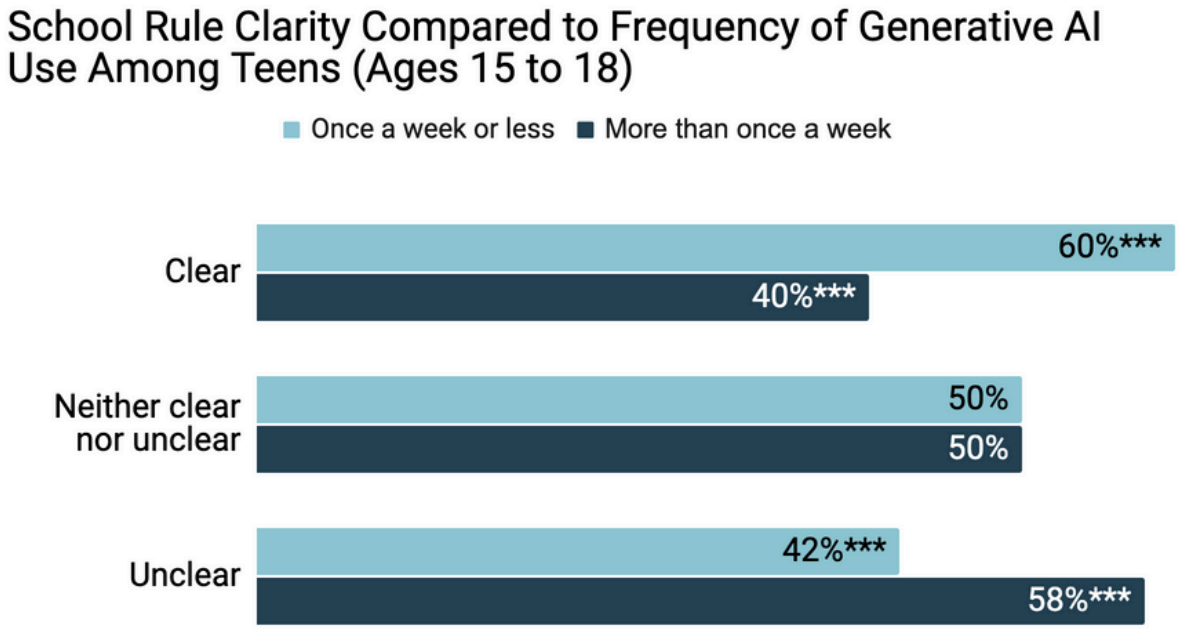

Graph represents answers to the Qs: How frequently have you used GenAI tools (i.e., ChatGPT, Claude, Gemini, Character AI, Dall-E)? and How clear are your school's rules about using GenAI? Respondents are teen generative AI users ages 15-18. Bars with asterisks (***) differ significantly within each response option ($p < .001$). Data are weighted using iterative proportional fitting (raking) to reflect U.S. Census demographic estimates for adolescents by race and gender (weighted n = 1,014). Percentages may not sum to 100% due to rounding.

PART THREE - RULES & RULEMAKERS

Understanding Teen Perspectives on Cheating and Generative AI Use

Another topic that arose from the focus groups was cheating on school assignments with generative AI. While teens spoke about cheating with generative AI, they also said it can be more complicated than straightforward cheating. Some teens said that a lack of clear school rules led them to feel like they had cheated, but they were unsure if they actually had. Others spoke about a lack of understanding around how generative AI worked, stating that they thought of it as more of a helper than a cheating instrument, and did not realize that turning in assignments completed by generative AI could be viewed as unethical. These gray areas are important to consider as educators and policymakers continue to navigate generative AI and the modern classroom.

> *"Well, it was one instance awhile back, like when [generative AI] had just first came out, it was late at night –and it was like the assignment was due at like 12 and it was like 11:45 somewhere around there, and I had forgotten about the assignment. I was trying to write out my essay or whatever, and I couldn't really figure out what to do. And then that's when my friend had suggested [generative AI] to me, but I didn't know what it was at the time. But, my friend had showed me what it was and how to use it and stuff like that, and I turned it in." - 16-year-old boy*

> *"I do know a lot of people who have cheated using ChatGPT, which has led to a lot of my teachers not condoning any use of it at all." - 16-year-old girl*

> *"I do remember last year –I didn't even know about this– but 90% of my math class, they had all cheated on the exam somehow. I wasn't even paying attention, I didn't cheat, but turns out I found out months later the majority of them cheated on the test, leading to the testing system becoming more stricter the next year." - 16-year-old girl*

> *"I think no matter how normalized it gets, I don't think anyone is, like, proud to outrightly say that they use AI because I think there will always be that stigma of like, I don't know whether it be cheating or you can't think on your own, so you have to seek additional support." - 17-year-old girl*

PART THREE - RULES & RULEMAKERS

> *"My friends, or at least some of them, they have to use ChatGPT to cheat because they don't fully know how to do assignments without talking to their teachers." - 16-year-old boy*

Echoing the discussions from the focus groups, teen reports of cheating in our survey reveal a split in the behavior. The survey results revealed that 42% of teen generative AI users have used generative AI to cheat on a school assignment, and 45% have not. An additional 14% of respondents were either unsure or did not feel comfortable sharing this information. Teens in the focus groups were not opposed to clearer generative AI rules in school.

They emphasized that the line between appropriate use and cheating is blurred, explaining that their use comes from academic confusion and stress, not with the intention of cheating. For teens, rules provide clarity, and clarity helps them identify their boundaries.

> *"I feel like rules and guidelines can give us a grounding and a boundary at the same time. So if you have a rule, you're going to think of the ways you cannot break it and the ways you can get close to breaking it. So it gives you a kind of area and a creative guideline." - 15-year-old boy*

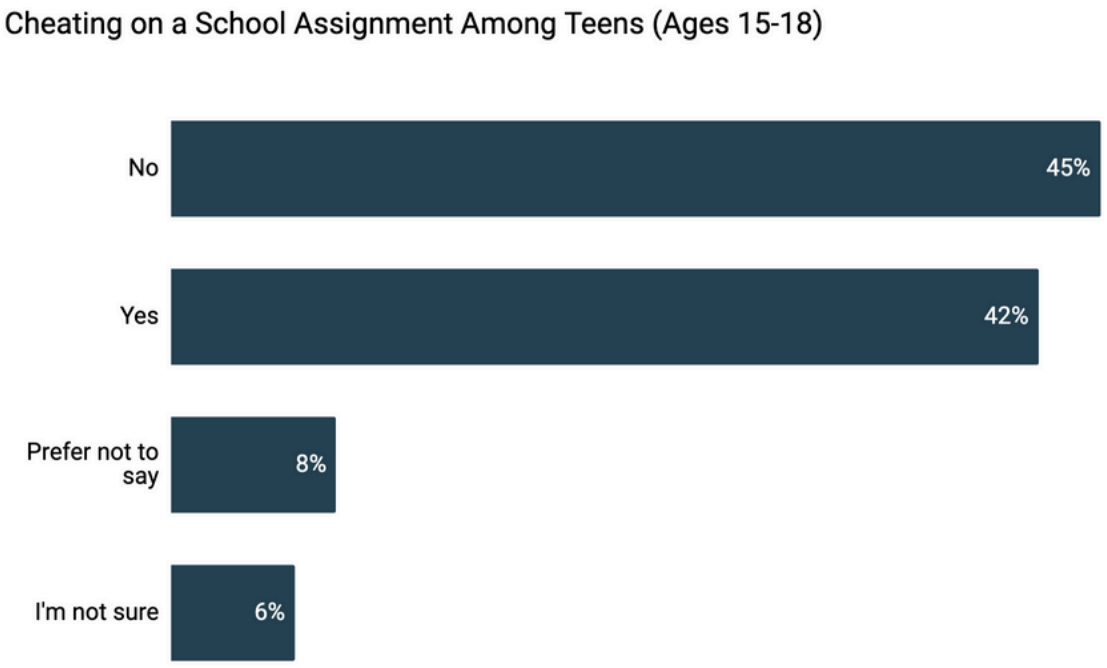

Graph represents answers to the Q: Remember that your answers are completely confidential. Have you ever used GenAI to cheat on a school assignment? Respondents are teen generative AI users ages 15-18. Data are weighted using iterative proportional fitting (raking) to reflect U.S. Census demographic estimates for adolescents by race and gender (weighted n = 1,015). Percentages may not sum to 100% due to rounding.

TEENS' EMOTIONAL ENGAGEMENT WITH GENERATIVE AI

As generative AI becomes more popular, so do its uses. Teens not only use generative AI for schoolwork, but also for emotional support, connection, and companionship. The survey explored the prevalence of teen users' turning to generative AI to discuss their personal feelings with an AI chatbot, and their thoughts about generative AI's "human-like" behavior. Teens' use of generative AI to talk about personal feelings has ignited widespread concern that this behavior could lead to emotional harm.

42% OF TEEN GENERATIVE AI USERS REPORT THAT THEY HAVE TALKED ABOUT THEIR FEELINGS WITH A GENERATIVE AI CHATBOT.

A large share **(42%) of teen generative AI users report that they have talked about their feelings with a generative AI chatbot**. A smaller share (20%) indicated that while they have not talked about their feelings with a chatbot, they know someone who has.

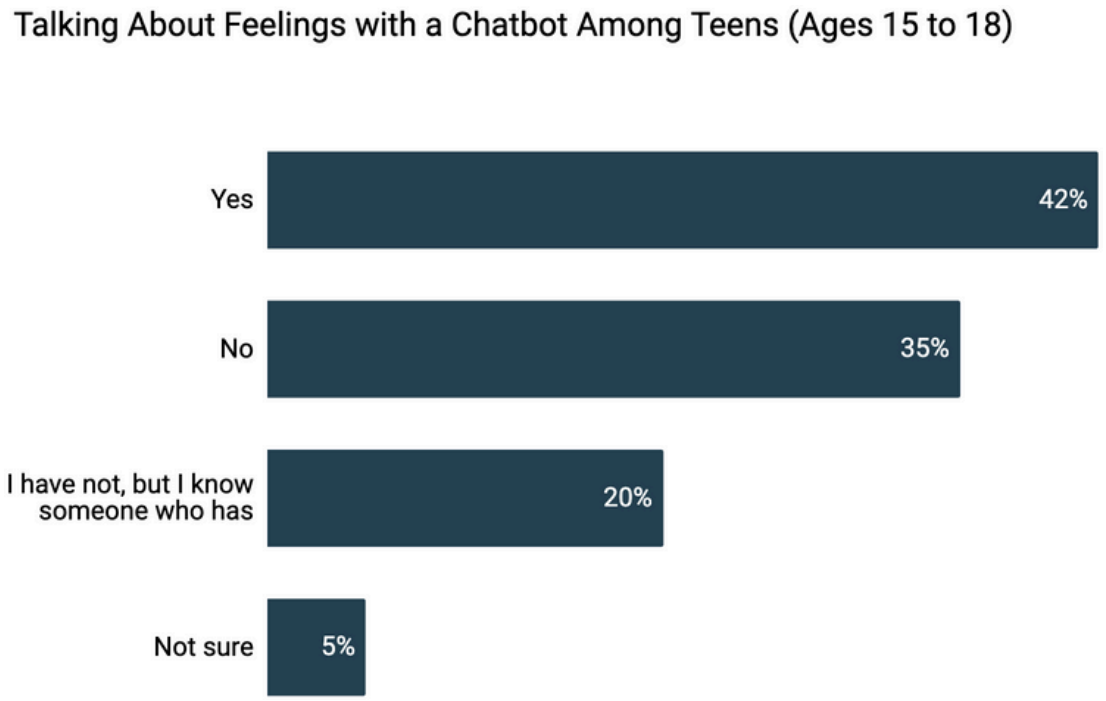

Graph represents answers to the Q: Have you ever talked about your feelings with an AI chatbot? Respondents are teen generative AI users ages 15-18. Data are weighted using iterative proportional fitting (raking) to reflect U.S. Census demographic estimates for adolescents by race and gender (weighted n = 1,015). Percentages may not sum to 100% due to rounding.

PART FOUR - COMPANIONSHIP

When discussing this topic in the focus groups, participants spoke about generative AI's accessibility compared to traditional mental healthcare, or choosing the technology over friends and family. Teens stated that they did not have access to the former and felt like a burden to the latter.

> *"I feel as though our generation is the main target for all of that, because nowadays people don't really care about mental health. So they [other teens] chat with the AI." - 18-year-old girl*

> *"I treat [generative AI] as a therapist, honestly, because I don't have no therapist. Be like, 'What should I do in this situation?' Because they would give, like, an unbiased opinion. Rather than going to your friend who actually knows you and would side with you if there was, like, a situation." - 16-year-old girl*

> *"Being someone who has dealt with like their mental health– I feel like a lot of times you feel like, or you can feel like, you don't want to burden anybody else with your feelings and having like a companion that doesn't necessarily have like genuine emotion, but it's just giving you advice, can be super helpful to not feel like you're like, inconveniencing anyone else." - 15-year-old girl*

Youth focus group participants often spoke about how they did not use generative AI to talk about their feelings, but they knew others who did. Therefore, it felt important to provide teens with the survey option: "I have not talked about my feelings with a chatbot, but I know someone who has." Including this option could help avoid the potential embarrassment that a respondent might feel if they indicated that they had used generative AI in this way. This points to a certain stigma about teens using generative AI to talk about their feelings, with some teens openly judging this particular use of the technology as "bad" or "exploitative."

> *"Mainly they just create AI that try to get somebody to, like, be with, like in a relationship. It's just really sad, really, when it's just trying to go after like lonely people or young kids. So it's like another example of AI just being used in the wrong way besides, like, cheating and such." - 16-year-old boy*

PART FOUR - COMPANIONSHIP

"I feel like it's exploiting, like, all these lonely people that should be getting, like, real human connection for money. And I feel like it's pretty selfish of the people who are, like, creating these apps." - 17-year-old girl

"I hope this doesn't come off as mean, but I, like, feel kind of bad. Like, I feel like if you have to rely on technology for social interaction, that's so, like, vital for us –I feel like you don't have that in your life. And I kind of, like, feel bad. And I think, like, maybe the resources in, like, the environment that they're in should provide that to them, if they're going to high school or like, have any clubs they're part of." - 17-year-old girl

Teens discussed their feelings with generative AI chatbots regardless of gender and usage frequency. This suggests that using generative AI to discuss emotions is not just for heavy users, or for teens of a specific gender, but appeals to a broad range of teens.

Notably, though, LGBTQ+ teens (53%) are significantly more likely than non-LGBTQ+ teens (35%) to say they have talked about their feelings with generative AI. Members of this community are also more likely to say they know someone who has talked about their feelings with a chatbot (23% of LGBTQ+ vs. 15% of non-LGBTQ+).

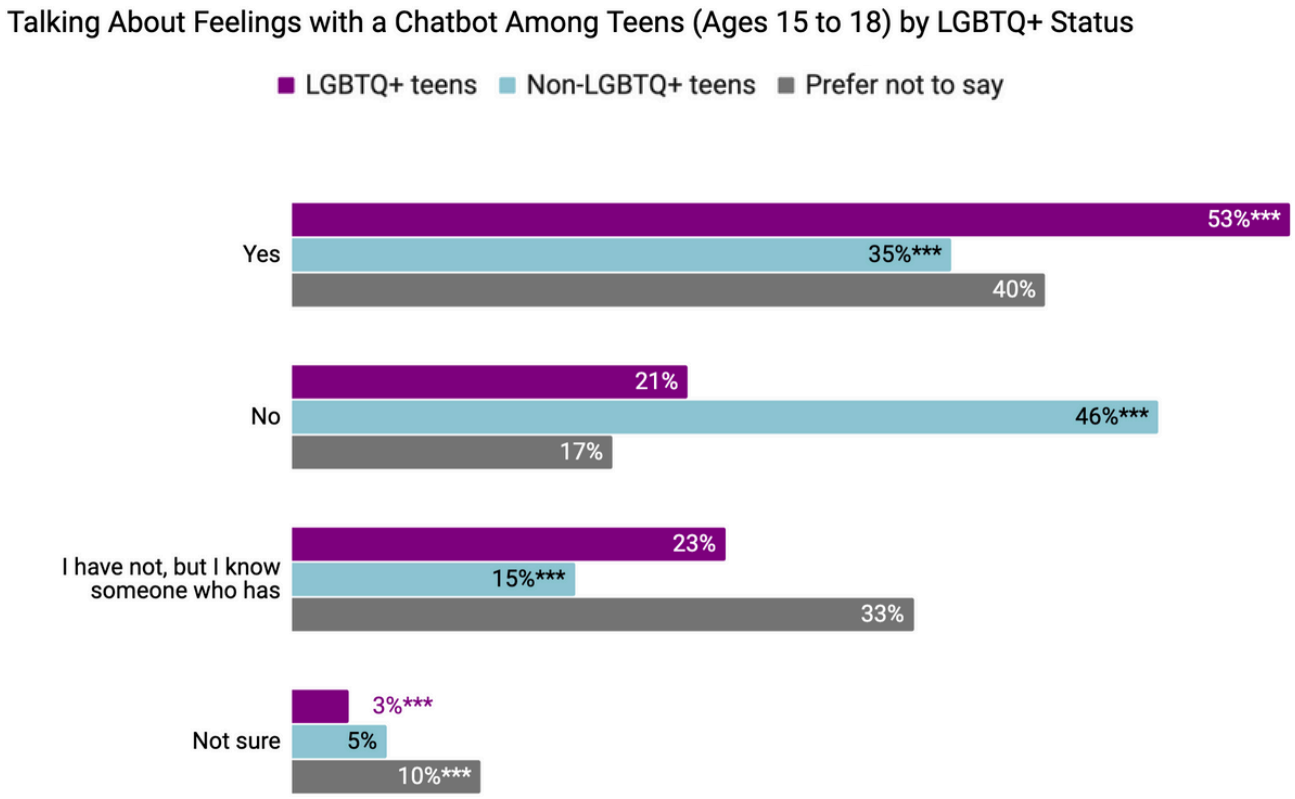

Graph represents answers to the Q: Have you ever talked about your feelings with an AI chatbot?. Bars with asterisks (***) differ significantly within each response option (p < .001). Respondents are teen generative AI users ages 15-18. Data are weighted using iterative proportional fitting (raking) to reflect U.S. Census demographic estimates for adolescents by race and gender (weighted n = 1,015). Percentages may not sum to 100% due to rounding.

Perceived Humanity of Generative AI Chatbots

It is clear that many teens either use or know someone who is using chatbots to talk about their feelings. This brings up questions about generative AI's perceived "humanity." Four-in-ten **(42%) of teen generative AI users agree that talking to generative AI feels like talking to a human.**

FOUR-IN-TEN (42%) OF TEEN GENERATIVE AI USERS AGREE THAT TALKING TO GENERATIVE AI FEELS LIKE TALKING TO A HUMAN.

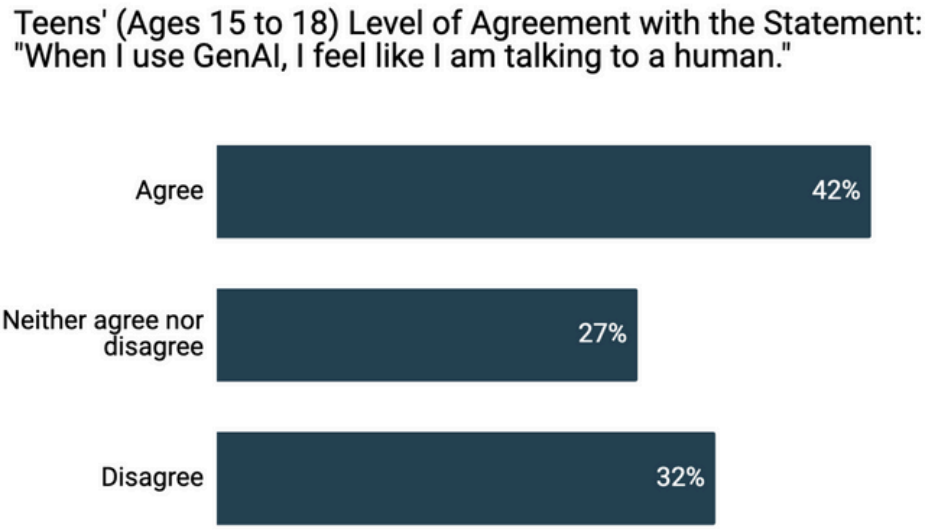

Graph represents answers to the Q: When I use GenAI, I feel like I am talking to a human. Respondents are teen generative AI users ages 15-18. Data are weighted using iterative proportional fitting (raking) to reflect U.S. Census demographic estimates for adolescents by race and gender (weighted n = 1,013). Percentages may not sum to 100% due to rounding.

PART FOUR - COMPANIONSHIP

Teens who view chatbots as human-like have a higher likelihood of discussing their feelings with these technologies. **In fact, teens who agree that talking to generative AI feels like talking to a human are three times as likely to talk to generative AI about their feelings** compared to teens who disagree that generative AI feels human-like (62% vs. 21%).

Focus group and reflection session participants' sentiments about generative AI's humanity mirrored the survey data. Some were quite insistent that generative AI was not human-like, while others were open to the idea.

"I know a lot of people use AI, like ChatGPT, for therapy, free therapy. And I know a lot of people have said that they felt heard and seen from, like, AI's response to what they told them. And so I feel like in that sense, it feels like generative AI is human because you're able to like, see, like kind of the sympathy they show towards you. And it may not necessarily be like a real person showing that sympathy, but like being able to feel that behind the screen, I feel like it's enough for that person."
- 17-year-old girl

Teens (Ages 15 to 18) Who Agree that Talking to a Chatbot Feels like Talking to a Human are Much More Likely to Talk to Generative AI About Their Feelings

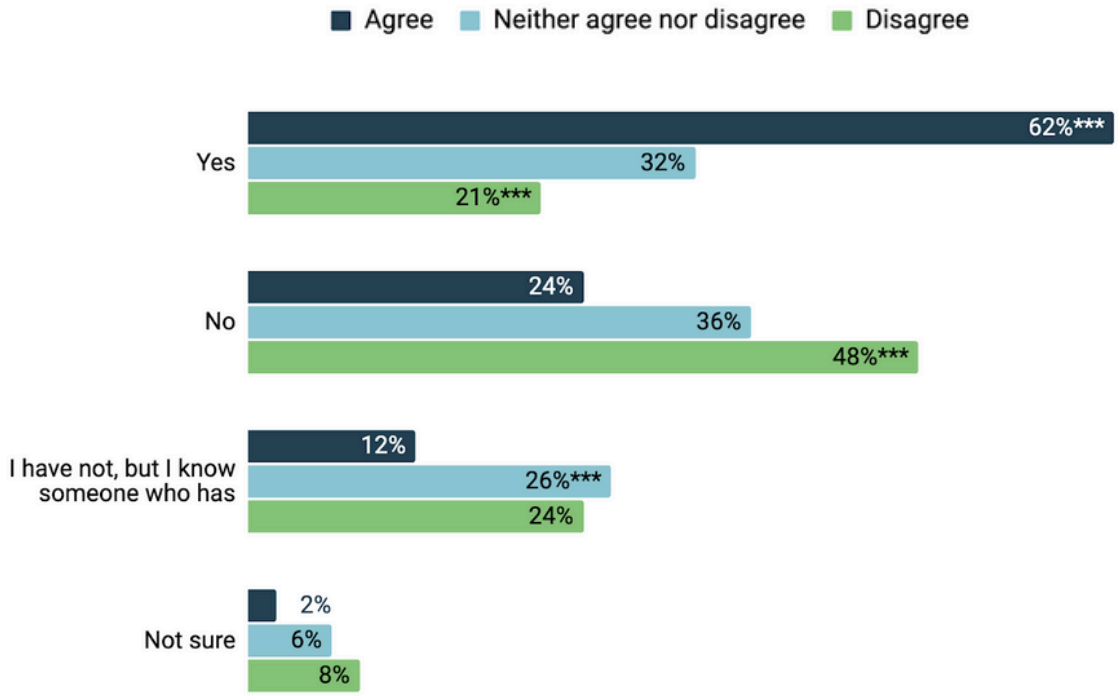

Graph represents answers to the Qs: Have you ever talked about your feelings with an AI chatbot? and: When I use GenAI, I feel like I am talking to a human. Bars with asterisks (***) differ significantly within each response option (p < .001). Respondents are teen generative AI users ages 15-18. Data are weighted using iterative proportional fitting (raking) to reflect U.S. Census demographic estimates for adolescents by race and gender (weighted n = 1,014). Percentages may not sum to 100% due to rounding.

PART FOUR - COMPANIONSHIP

"We did an activity [in class] which was like an imitation game, where we had to ask questions to both a human and the chatbot. And then we compared people's answers, and it's a little surprising that people said that when they're using like, ChatGPT or like AI, that it sounds like they're talking to a human, because I feel like from our little experiment in class, it was like very obviously different answers from the like, the ChatGPT to what like an actual person said." - 18-year-old girl

"AI is not human, so they can't achieve the pure creativity that may come from a human because it's just trying to replicate what we may think, but it's not that advanced yet." - 18-year-old girl

"[Generative AI] can sound very human based off of your data. So, off of how you talk to it. So, sometimes I ask AI or ChatGPT to help me respond to, like, this message someone sent me, not because I'm lazy, it's because sometimes I genuinely don't know, like, what they're saying to me. So I ask [generative AI] to help me with this message, and they will put a message back in like my lingo and how I speak and how I talk. And, you know, it's really weird." - 18-year-old girl

Teen boys (49%) are significantly more likely than teen girls (35%) to agree that talking to generative AI feels like talking to a human. Since teen boys tend to use generative AI more frequently, this increased use might lead them to feel that this technology is more human-like or vice versa.

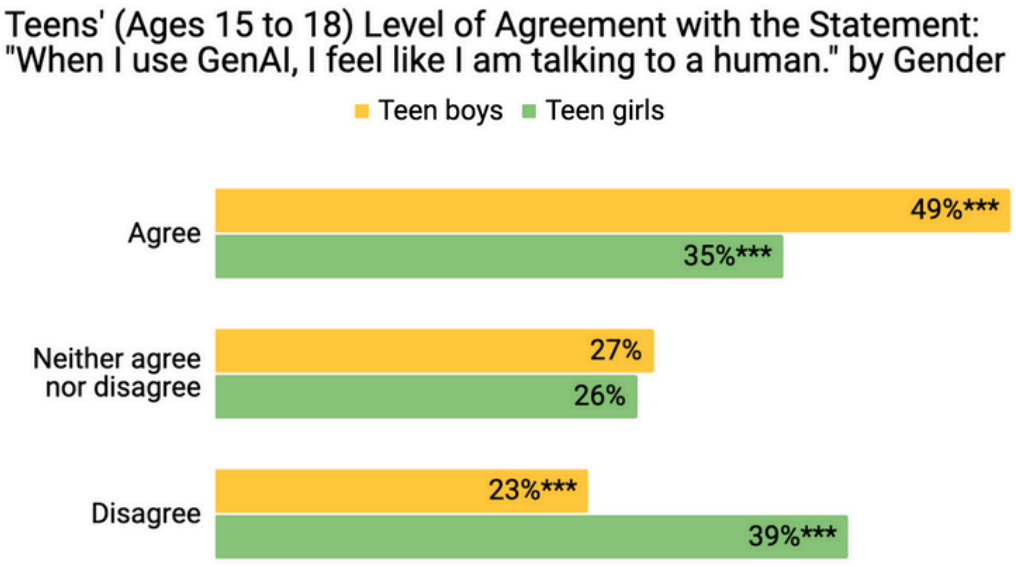

Graph represents answers to the Q: How much do you agree with the following statement: When I use GenAI, I feel like I am talking to a human. Bars with asterisks (***) differ significantly within each response option (p < .001). Respondents are teen generative AI users ages 15-18. Data are weighted using iterative proportional fitting (raking) to reflect U.S. Census demographic estimates for adolescents by race and gender (weighted n = 1,013). Percentages may not sum to 100% due to rounding.

Discomfort and Unease with Generative AI Behavior

While it is relatively common for teens to use generative AI to talk about their feelings, some are uncomfortable with generative AI's behavior. Overall, **44% of teen generative AI users report that generative AI's behavior freaks them out**. Less than a quarter (20%) disagree with this statement.

It should be noted that when considering generative AI's humanity and potentially "freaky" behaviors, a significant number of teen users reported that they neither agree nor disagree with these statements (27% and 35%, respectively).

When presented with the rather high numbers of respondents who selected "neither agree nor disagree," participants in the survey reflection groups theorized that their peers might view generative AI as somewhere in between a human and an object.

To many teens, generative AI is an ever-evolving, complicated, and convenient robot. Reflection session attendees also stated that generative AI has become part of their lives, so they do not think of its behavior as "freaky" or not "freaky," it just **is**. However, those who use generative AI for a limited number of tasks, for example, only academic work, might not think about generative AI in this context.

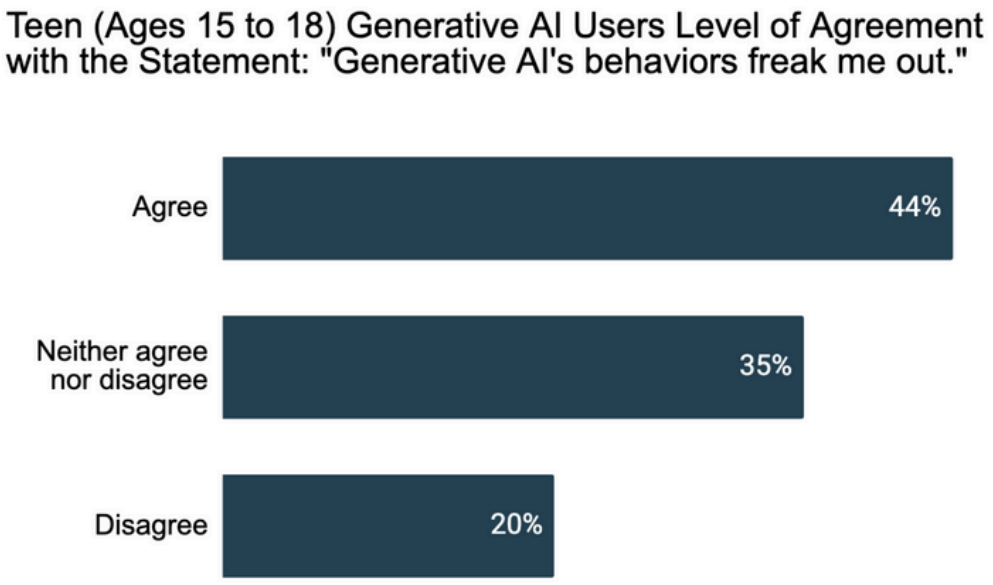

Graph represents answers to the Q: How much do you agree with the following statement: "GenAI's behaviors freak me out." Respondents are teen generative AI users ages 15-18. Data are weighted using iterative proportional fitting (raking) to reflect U.S. Census demographic estimates for adolescents by race and gender (weighted n = 1,013). Percentages may not sum to 100% due to rounding.

PART FIVE - SAFETY & INVOLVEMENT

TEENS' PERCEPTIONS OF SAFETY

As use of these tools becomes more widespread, understanding teen perceptions of generative AI safety is paramount. **Three-in-five (60%) teen generative AI users say they feel safe while using generative AI**. Only 12% of teen generative AI users felt unsafe while using these tools, and 29% reported feeling neither safe nor unsafe.

When discussing the rather high percentage of teen users who reported feeling neither safe nor unsafe, focus group and reflection session participants said that the idea of safety is relative and dependent on the situation. It appears that safety is not a top concern for teens when using generative AI.

"I think me personally, I feel like I'd be under the neither safe nor unsafe [category]. Yes, [generative AI] gives us some like, ideas or like, benefits when we use it. But, like when I think about like, our environment and how AI is being used so much, I would say that it's neither safe nor unsafe." - 18-year-old girl

"When I think of AI, I don't necessarily like, associate it with being safe or unsafe, because I feel like that would depend on what you're using it for." - 18-year-old girl

"Using the word 'safe' is really vague, because if you try to think of 'safe,' like, am I safe if there's a criminal in my house?" - 16-year-old boy

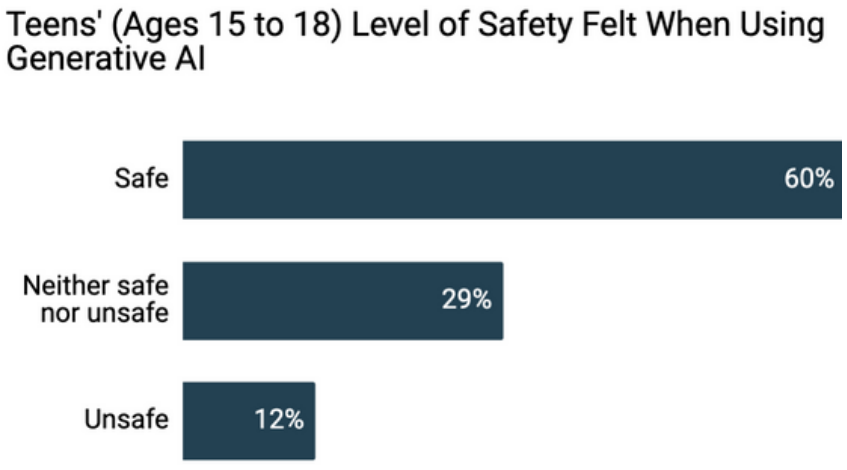

Graph represents answers to the Q: How safe do you feel when using GenAI? Respondents are teen generative AI users ages 15-18. Data are weighted using iterative proportional fitting (raking) to reflect U.S. Census demographic estimates for adolescents by race and gender (weighted n = 1,015). Percentages may not sum to 100% due to rounding.

PART FIVE - SAFETY & INVOLVEMENT

> *"I don't think there's any information that I put into AI where it would, like, be risking my safety or like, my personal data. I guess if it was breached, apart from like, just my accounts that I have, but it's like the same risk that I have when I'm using like Google Docs or anything else like that."*
> *- 18-year-old girl*

Safety Across Gender And Identity

Notable demographic breakdowns emerge among perceptions of safety. Teen boys (15%) are more likely than teen girls (8%) to report feeling unsafe while using generative AI. Teen girls, on the other hand, more often feel neither safe nor unsafe when using generative AI (33%) compared to teen boys (23%).

Perhaps this can be attributed to varying use frequency between genders, or teen boys' increased likelihood to perceive generative AI as human-like.

When comparing LGBTQ+ teens to non-LGBTQ+ teens, smaller shares of LGBTQ+ teens report feeling safe (55%) than non-LGBTQ+ teens (63%). Relatedly, LGBTQ+ teens (16%) are also slightly more likely to report feeling unsafe than their non-LGBTQ+ peers (9%). Despite LGBTQ+ teens being more apt to talk about their feelings with generative AI (see page 27 for finding), this does not mean they perceive these technologies as entirely safe. It should be noted that the overall percentage of those who feel unsafe while using generative AI is quite small for both the gender and LGBTQ+ demographics.

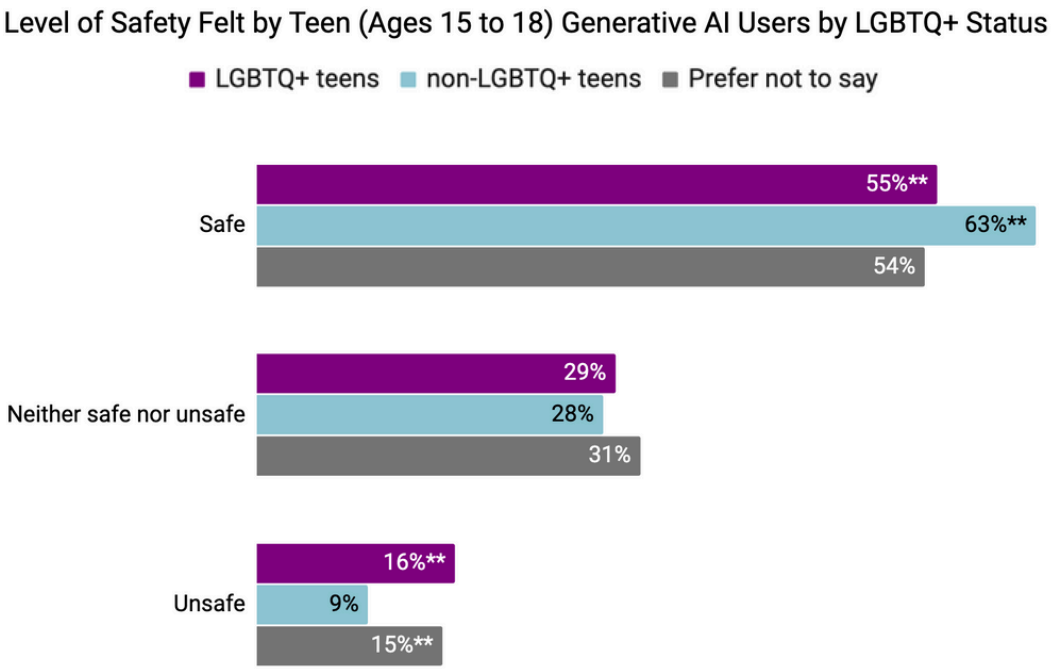

Graph represents answers to the Q: How safe do you feel when using GenAI? Respondents are teen generative AI users ages 15-18. Bars with asterisks (**) differ significantly within each response option (p < .01). Data are weighted using iterative proportional fitting (raking) to reflect U.S. Census demographic estimates for adolescents by race and gender (weighted n = 1,015). Percentages may not sum to 100% due to rounding.

PART FIVE - SAFETY & INVOLVEMENT

Exposure to Chatbot Advertisements Among Teen Users

As teens decide how to use, navigate and feel about generative AI, most of them are receiving targeted advertisements, particularly on social media, encouraging them to interact with chatbots on a deeper level, and a more frequent basis.

A vast majority of teen generative AI users (81%) have received advertisements encouraging them to interact with a chatbot. About a third (34%) report seeing these ads more than once a week. Less than a quarter of users (19%) have never seen these ads.

A VAST MAJORITY OF TEEN GENERATIVE AI USERS (81%) HAVE RECEIVED ADVERTISEMENTS ENCOURAGING THEM TO INTERACT WITH A CHATBOT.

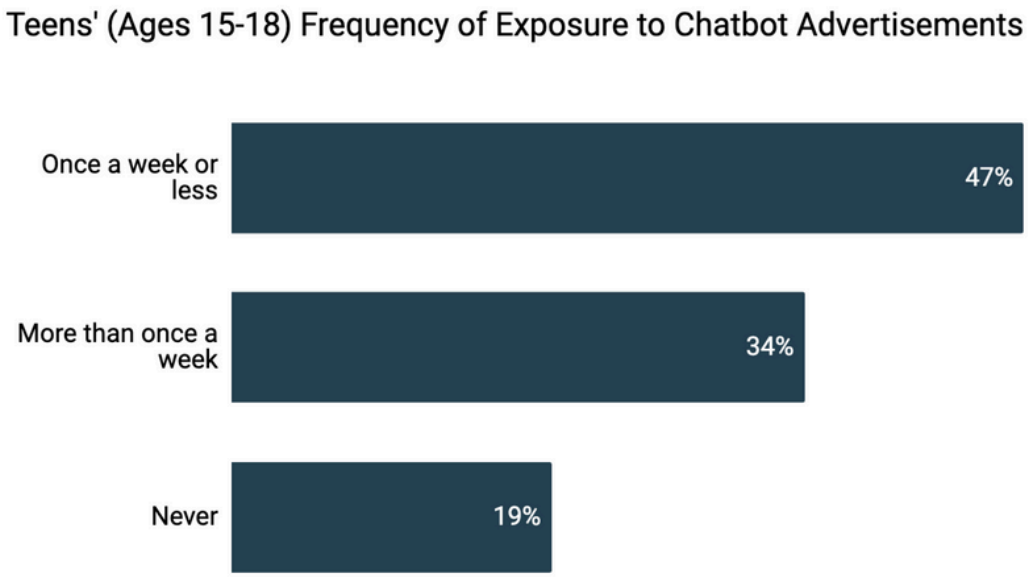

Graph represents answers to the Q: Do you get advertisements (i.e., on social media platforms) encouraging you to interact with a chatbot? Respondents are teen generative AI users ages 15-18. Data are weighted using iterative proportional fitting (raking) to reflect U.S. Census demographic estimates for adolescents by race and gender (weighted n = 1,015). Percentages may not sum to 100% due to rounding.

PART FIVE - SAFETY & INVOLVEMENT

"I'm definitely getting these ads mainly from TikTok, but usually they're of like, anime characters or just or just like fictional characters where their personalities have already been written. So turning it into an AI, it's not that hard. But, I usually see them as like a, like a plus one, or like someone could just pull up the AI and see how they would react to certain things." - 17-year-old boy

"It makes me, like, feel uncomfortable when I get these ads. Like, I know that your phone, like, uses your information that you like, search up on Google. And I'm, like, not even looking for any of that kind of stuff. Like, I'm just trying to do my work and like, move on with my life. And it kind of made me feel a little better when I heard that a bunch of other people were getting these ads." - 17-year-old girl

BOYS ARE SIGNIFICANTLY MORE LIKELY TO GET CHATBOT ADVERTISEMENTS THAN GIRLS (84% VS. 78%).

Exposure to Chatbot Advertisements Across Gender and Identity

Certain teen generative AI users are more likely to see these ads than others. Boys are significantly more likely to get these ads than girls (84% vs. 78%), and girls are more likely to have not seen such ads (23% vs. 16%).

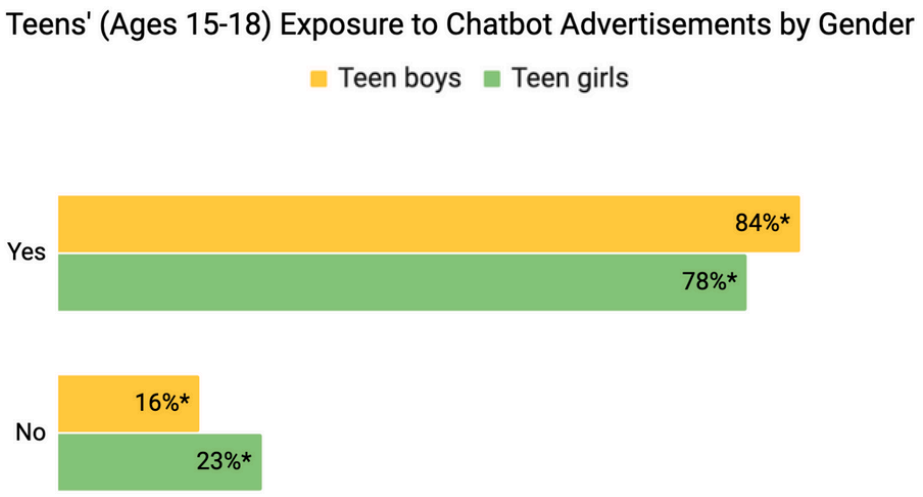

Graph represents answers to the Q: Do you get advertisements (i.e., on social media platforms) encouraging you to interact with a chatbot? Respondents are teen generative AI users ages 15-18. Bars with an asterisk (*) differ significantly within each response option ($p < .05$). Data are weighted using iterative proportional fitting (raking) to reflect U.S. Census demographic estimates for adolescents by race and gender (weighted n = 1,015). Percentages may not sum to 100% due to rounding.

PART FIVE - SAFETY & INVOLVEMENT

Similarly, slightly larger shares of LGBTQ+ teens (89%) than non-LGBTQ+ teens (75%) see these ads.

Larger shares of non-LGBTQ+ teens are significantly more likely to not get these ads at all when compared to LGBTQ+ teens (25% vs. 11%).

LARGER SHARES OF LGBTQ+ TEENS (89%) THAN NON-LGBTQ+ TEENS (75%) SEE CHATBOT ADVERTISEMENTS.

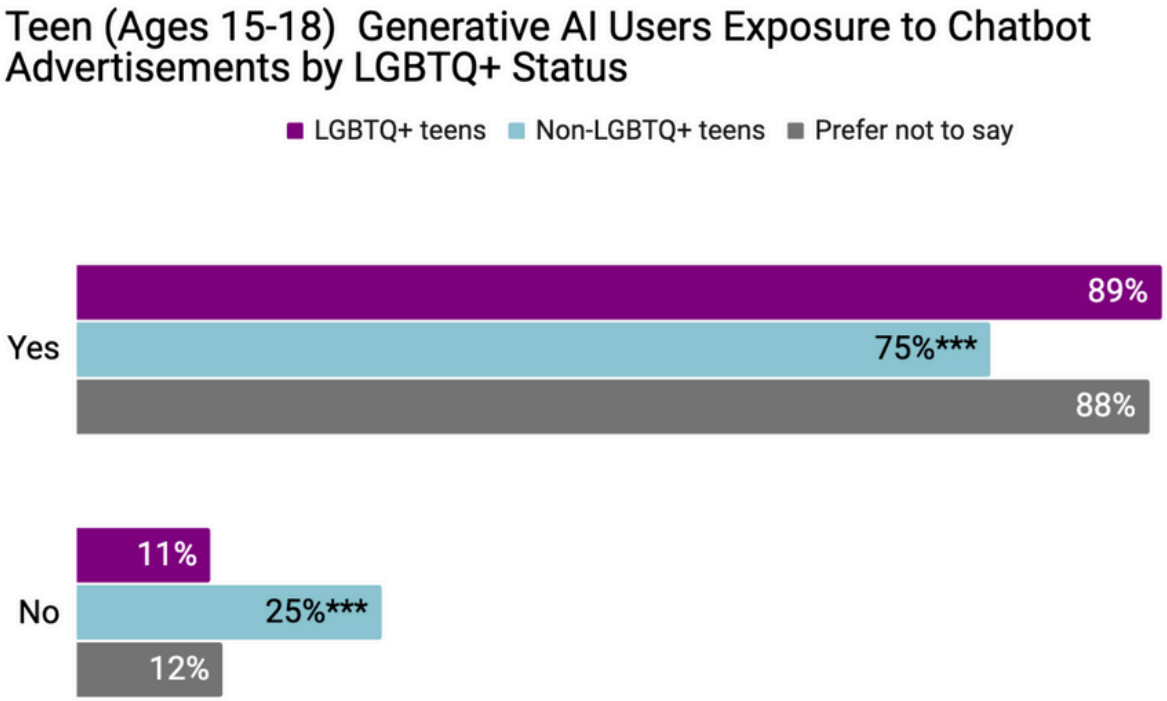

Graph represents answers to the Q: Do you get advertisements (i.e., on social media platforms) encouraging you to interact with a chatbot? Respondents are teen generative AI users ages 15-18. Bars with asterisks (***) differ significantly within each response option ($p < .001$). Data are weighted using iterative proportional fitting (raking) to reflect U.S. Census demographic estimates for adolescents by race and gender (weighted n = 1,015). Percentages may not sum to 100% due to rounding.

PART FIVE - SAFETY & INVOLVEMENT

Teen Perspectives on Participation in Generative AI Design

This report makes evident teen users' many perspectives on generative AI. The ultimate goal of this research is to give a voice to young people who often are the first to try new technologies, but are also the first to experience potential harms. With this in mind, it was imperative to ask teens about their desired involvement in the development and design of generative AI tools. **Over half of teen generative AI users (54%) believe that young people should be involved in the design of generative AI tools.** Only 17% said that young people should not be involved.

Notably, the term "involvement" had a different meaning for different teens. In the focus groups, some teens spoke about involvement from a trust and safety perspective (i.e., shaping generative AI tools to be an inclusive, safe space). Other teens were more interested in playing a role in the engineering/technology perspective (i.e., providing guidance on updates, programming Large Language Models, etc.).

> *"[These tools] are being developed by a very certain demographic, which is white older men. And I think that being the main demographic developing these things, as it grows, it can become very one-sided."*
> *- 18-year-old girl*

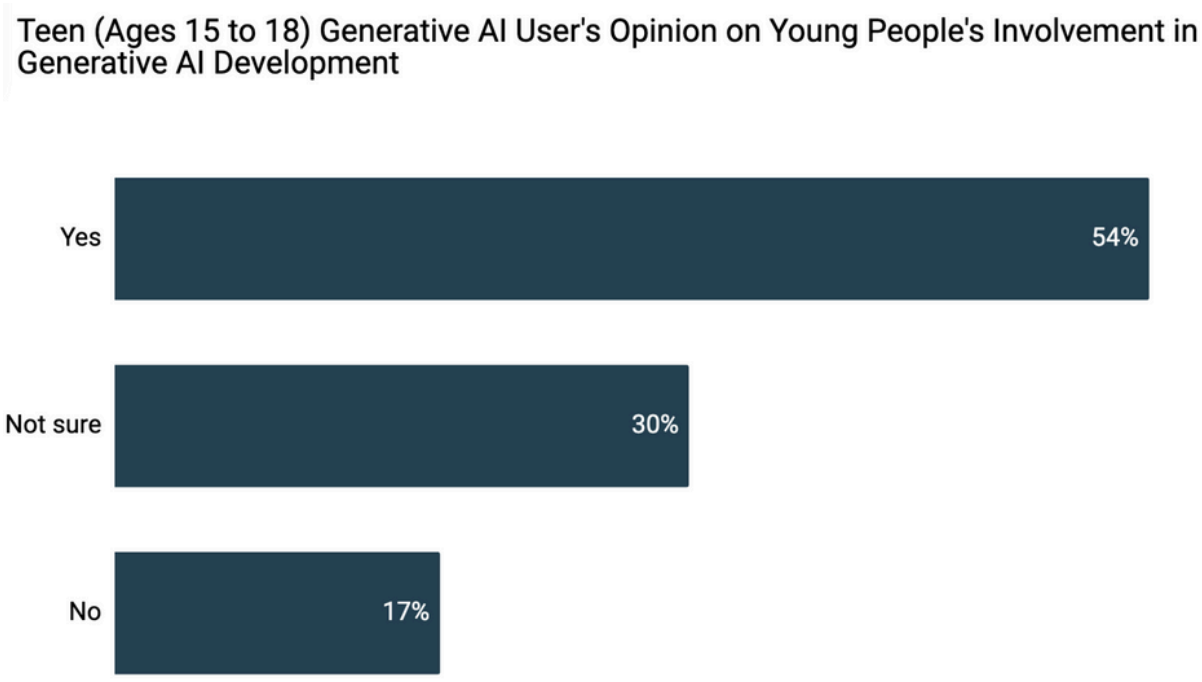

Graph represents answers to the Q: Do you think young people should be involved in the design of GenAI tools? Respondents are teen generative AI users ages 15-18. Data are weighted using iterative proportional fitting (raking) to reflect U.S. Census demographic estimates for adolescents by race and gender (weighted n = 1,015). Percentages may not sum to 100% due to rounding.

PART FIVE - SAFETY & INVOLVEMENT

"If you type in a certain prompt, the responses you would get were very, like stereotypical, or even offensive towards certain groups –like mocking the way that they speak or stuff like that. So, I think when we talk about the younger generation having a voice, I think since we tend to be more like aware about, you know, the effect of our words and how we use them, I think it's important to have those controversial topics be brought to, like, awareness and have more like regulations in place."
- 17-year-old girl

"I also think that younger people tend to be more empathetic or at least aware of, for instance, like, the environmental issues or ethical issues. So, I think it could be like a slippery slope if you have like these massive tech corporations with, like, only money in mind, like kind of leading the scene. So I think it's great to have younger voices, you know, take a role in that creation." - 17-year-old girl

"I was thinking more about, like, the debugging side of it because a lot of youth use it, and they hear a lot of the problems from their peers. So, maybe more than the design, more like problem solving with it and troubleshooting." - 15-year-old girl

"I do think we should have like, an influence on the future of AI since we will be the ones still using it down the line. And I think, like this should apply to like, future generations as well, since it does affect our future. And I do think as we like, get more information on, like its limits and its effects on like the environment and also like ethical issues — I feel like each generation will gain more information and be able to input that into AI." - 17-year-old girl

"I think we should be more involved. I've created some AI models before. And I'd say that it's nice to have different inputs because different people have their different thoughts on different parts of information. So, working with different people like, helps with the reliability and accuracy of AI." - 17-year-old boy

Regardless of how teens think about ways to get involved with generative AI, they want a seat at the table by playing a role in shaping the generative AI tools that will impact their futures.

CONCLUSION

GROWING UP WITH GENERATIVE AI

It has been over three years since generative AI first began publicly making waves. Since then, several tools have been released promising to support people with tailored generated content for just about any problem, question, or skill they need assistance with. The promise of convenient, easily accessible solutions just a click away may appeal to teens who are undergoing unfamiliar formative changes biologically, cognitively, and socially within a short time frame. Those without access to academic, social, psychological, or financial resources may be drawn to the purported promises of generative AI to provide an equitable playing field. While cognitively maturing, some teens may jump into the technology head first without considering or understanding the risks.

Teens can misstep when emotions and pressures are high, especially around peers. When all their peers are using the technology, teens may feel the need to use it to keep up. While teens were adopting and testing out this technology, adults began asking questions: Is generative AI safe? What does it mean to cheat on an assignment using generative AI? Who is most responsible for teaching teens about generative AI? Teens were asking these questions too. This research shows a glimpse into their world with generative AI, as they navigate how and when to use it.

WHILE COGNITIVELY MATURING, SOME TEENS MAY JUMP INTO THE TECHNOLOGY HEAD FIRST WITHOUT CONSIDERING OR UNDERSTANDING THE RISKS.

This study also provides insight into an incredibly unique time and place –when teens are using generative AI relatively frequently, amidst uncertainty and confusion about how to provide guidance and guardrails.

Teen generative AI users are concerned about critical thinking, job loss and mis/disinformation. They are also aware of generative AI's benefits around convenience and speed. Teens report a lack of household rules about the technology, but a large share of teens find their school rules to be at least somewhat clear. Many teens find generative AI to be human-like, while others find its behaviors "freaky." As these contradictions swirl, most teens believe people their age should be involved in generative AI development and design.

The purpose of this research and subsequent report is to provide insights to stakeholders such as parents, policymakers, and industry leaders, so that teens' perspectives on generative AI are taken into consideration during these tumultuous times. The sections below will provide suggestions to each of these three stakeholders.

HOW PARENTS CAN SUPPORT SAFE GENERATIVE AI USE

Parents often play a key role in their teens' online safety education. However, there can be a disconnect between parents and teens in this area. This is especially true with newer technologies such as generative AI. This report provides insights into caregivers' current role in their teens' generative AI journeys. Such insights provide key takeaways for parents hoping to engage their teens in generative AI safety education.

Start a conversation.

This research demonstrates that most teens do not view their parents as their primary source of generative AI education. Teens are more likely to turn to themselves, their peers, or their teachers to learn about these tools. If caregivers want to be more involved in their teens' generative AI use, they can start by showing an interest in this technology. To start, parents can ask their children why they use generative AI, what's appealing about it, and if they have any concerns. Showing interest is a great way for parents to demonstrate to their teens that they are aware and want to learn more about generative AI.

Set household rules and boundaries.

Currently, less than a third of parents have established household rules about generative AI. Like any other technology, generative AI use should come with guidelines for teens. These guidelines do not have to be so different from previously established rules about technology. Parents can: designate certain areas of the house and times of day when their teens are permitted to use generative AI; emphasize the importance of honesty and integrity when it comes to AI-generated text, photos, and videos; and determine what platforms are appropriate (and which are inappropriate) for their teens.

IMPLICATIONS FOR PARENTS

Understand the array of possibilities.

This report shows that teens use generative AI for a variety of purposes. Though schoolwork is the most common use of this technology, substantial numbers of teens also utilize generative AI for search, entertainment, social activities, health, and more. It is important that parents understand the breadth and depth of generative AI, so they know how to best support their teen generative AI users.

Work together, learn together, bond together.

As with many new technologies, the teens in this report indicate they feel they have an upper hand when it comes to understanding and using generative AI. Rather than let this fact be disheartening, caregivers should ask their teens to help teach them about generative AI.

IT IS IMPORTANT THAT PARENTS UNDERSTAND THE BREADTH & DEPTH OF GENERATIVE AI, SO THEY KNOW HOW TO BEST SUPPORT THEIR TEEN GENERATIVE AI USERS.

Parents can turn this into a bonding experience by asking their teens to show them how to use generative AI to support shared interests. For example, parents could ask teens to show them how to generate a recipe which they could then cook together. If caregivers make it a habit to ask questions, teens will feel empowered, and caregivers will learn more about this new technology.

HOW INDUSTRY CAN BUILD SAFER, MORE RESPONSIBLE GENERATIVE AI

AI companies have an opportunity and a responsibility to design products and services that empower young users while minimizing risks. As the data shows, teens are already active users of generative AI, often with little guidance from the adults in their lives. Industry leaders can and must play a critical role in shaping a safer, more ethical, and more inclusive digital environment for this generation of early adopters.

Design with teens, not just for them.

Young people want a voice in how these tools are built, and more than half believe that teens should be involved in generative AI design. Companies should establish youth advisory boards that reflect diverse backgrounds and lived experiences. These groups can help developers understand how teens actually use generative AI –what motivates them, what confuses them, and what concerns them. Incorporating direct teen feedback throughout the design and testing process will ensure that new features support positive learning, creativity, and wellbeing.

Make safety the default.

Generative AI products should adopt "safety by design" principles, particularly for teen users. Default settings should minimize risks of overreliance, misinformation, and exposure to inappropriate or manipulative content. Companies can introduce "teen modes" or adjustable risk tiers that guide users toward responsible, age-appropriate uses of generative AI. For chatbots advertising emotional connection, it is especially important to include built-in guardrails, such as clear disclaimers, escalation pathways to trained human support, and boundaries around conversations about mental health or relationships.

Prioritize transparency and explainability.

Teens should be able to understand how a generative AI system works –what data it uses, what limitations it has, and how its responses are produced. Developers can integrate transparency features such as "AI influence summaries," visible source citations, and explanations of why a certain response was generated. Clarity builds trust and helps teens distinguish between helpful automation and human judgment, fostering, rather than replacing, critical thinking.

Protect privacy and reduce manipulation.

Given teens' heightened exposure to targeted chatbot advertising, industry standards should prohibit emotionally manipulative or parasocial chatbot marketing aimed at minors.

IMPLICATIONS FOR INDUSTRY

Developers should collect the minimum data necessary to operate these tools and avoid using teen interactions for behavioral advertising or unrestricted model training. Privacy dashboards written in plain language can empower teens to see what data is collected, how long it is stored, and how to delete it.

Build interoperable, proportional protections.

Effective age assurance and content safeguards do not have to mean blanket restrictions (see **FOSI's 2025 Age Assurance White Paper**). Companies should explore interoperable, privacy-preserving verification systems—such as wallet-based or third-party credentials—to enable graduated access to generative AI features based on maturity and context.

These proportional measures protect younger users while preserving educational and creative benefits for older teens.

Collaborate for accountability.

Companies should not act in isolation. Cross-industry cooperation can establish shared safety benchmarks, standardized disclosures, and independent audits focused on youth wellbeing. By participating in transparent, multi-stakeholder frameworks—including researchers, educators, and youth voices—industry can ensure that innovation and responsibility progress together. Teens are shaping the future of generative AI; industry must ensure that the future they inherit is one that supports learning, agency, and trust.

IMPLICATIONS FOR POLICYMAKERS

HOW POLICYMAKERS CAN PROMOTE SAFE & ETHICAL GENERATIVE AI USE

As generative AI becomes increasingly popular, lawmakers have an opportunity to avoid the mistakes made with social media. When developing legislation, policymakers should encourage companies to build safety into the design of their products, craft policies that are rooted in research, and invest in media and digital literacy education.

Collaborate with stakeholders to build safety into systems.

Protecting young people online involves building safety into systems. Generative AI should not be considered protected speech. This way, it will be easier to install baseline protections into the products without running into constitutional challenges. Ensuring generative AI remains a product allows industry to build safety guardrails for young people. Examples include platforms accurately responding to alarming prompts and enabling companies to be held responsible for the output of their products.

Building safe generative AI platforms begins with knowing the age of a user in order to provide age appropriate experiences.

PROTECTING YOUNG PEOPLE ONLINE INVOLVES BUILDING SAFETY INTO SYSTEMS.

Effective age assurance methods exist and can be deployed to give kids and teens safer experiences without fully denying access to these transformational technologies.

Now is the time to think through how best to provide baseline protections for the most vulnerable users. Requiring safety settings to be on by default sets a stronger foundation for all users as compared to parental controls, which are opt-in and rely on busy and technologically literate parents/caregivers to set up. Additional baseline protections could include clear transparency and disclosure requirements that could be assessed by third parties, as well as data privacy considerations such as collection, retention, and use limitations.

Craft evidence-based policy with young people in mind.

Understanding the appeal and attraction to these technologies is helpful when thinking through policy solutions.

IMPLICATIONS FOR POLICYMAKERS

This report highlights that teens appreciate the convenience and speed of generative AI. Notably, the majority of teens say they feel safe while using generative AI. That said, teens are not blind to the negative impacts of technology, in fact, they cite a loss of critical thinking skills and the potential impact on the future generations as top concerns about the technology. A significant share even say that generative AI's behavior freaks them out.

Teens are already regularly interacting with this technology, with many reporting that they use a generative AI tool more than once a week, and do so currently with limited regulatory protections. Ensuring young people have a safe experience online begins by understanding how and why they use this technology, and then developing policies to ensure they can use it safely.

Invest in media and digital literacy rather than blanket bans.

While blanket bans of apps and platforms are often viewed as a cure-all, this research shows that a more nuanced approach is a better path forward. For example, a significant share of teen boys say that they view the technology as human-like. This has the potential to lead to unhealthy relationships with generative AI that should be addressed. Rather than banning these tools out of fear, meet young people where they are with investments in media and digital literacy education.

Consistent guidance for integrating generative AI into the classroom will be helpful for students, teachers, and parents. Policymakers should collaborate with educators and school districts to understand the needs of teachers and students, and work together to deliver quality guidance for all groups.

RESEARCH DESIGN AND DATA COLLECTION

Qualitative

Focus group and reflection session participants were recruited by **In Tandem**. Participants were members of In Tandem's Youth Voice Fellowship which is a panel of teens across the United States. Focus groups were conducted from May 12 - 15th, 2025. Four focus groups were conducted with a total of 18 participants.

Three reflection sessions were conducted with a total of 10 participants. The timeline is as follows:

June 23, 2025: Focus group analysis reflection session

July 10, 2025: Survey design reflection session

September 23, 2025: Survey analysis reflection session

Focus group demographics are as follows: 12 participants identified as female and six as male. Thirteen participants attended public school and the other five attended charter school. Household income ranged from $40,000 to $200,000/year.

Ten participants identified as Black or African American, three as White, two as Asian or Asian American, one as Hispanic or Latino, and two as Mixed Race (Asian or Asian American and Black or African American; Middle Eastern or North African and White).

Reflection session demographics were as follows: Seven participants identified as female and three as male. Eight participants attended public school and the other two attended charter school. Household income ranged from $40,000 to $200,000/year. Three participants identified as Black or African American, two as White, two as Asian or Asian American, one as Hispanic or Latino, and two preferred not to say.

Both focus group and reflection session participants were 15 to 18-years-old, lived in the United States, and had used generative AI in the past.

Qualitative analysis was conducted in NVivo. Researchers used inductive analysis by identifying key codes, which were later aggregated into main themes. These themes helped guide survey design.

METHODOLOGY

Quantitative

The survey was fielded by **TeenVoice** through the platform **EvolveMe**. TeenVoice and EvolveMe are products supported by **American Student Assistance**.

The survey was conducted from July 28 - August 14th, 2025. Quotas were set for age, race, gender, and LGBTQ+ status. This survey had 1,000 respondents. Because this survey was fielded in the summer months, when many U.S. teenagers were on summer break from school, respondents were asked to think about their generative AI use in the past six months.

Survey data were weighted using iterative proportional fitting (raking) to reflect U.S. Census demographic estimates for adolescents by race/ethnicity and gender. The sample size (n) varies slightly due to this weighting. This raking process was reviewed via a quality assurance check.

Thank you to the **Digital Wellness Lab** for their creation of the survey question: “How much do you agree with the following statement: GenAI's behaviors freak me out.”

DEFINITIONS

Sample Characteristics and Demographics

AGE:

Respondents were between 15 and 18-years-old. Each age category reflects a nationally representative sample. Age distribution was monitored through quotas but not weighted.

RACE:

Respondents could select from White, Black, Hispanic/Latino, AAPI (Asian American Pacific Islander), and Other. Iterative proportional fitting (raking) was conducted so each race category reflects a nationally representative sample.

GENDER:

Respondents could select from Female, Male, or Non-binary. When weights were applied to ensure a nationally representative sample of race, gender was skewed to overrepresent female respondents, and underrepresent male ones. To correct for this, iterative proportional fitting (raking) was conducted to reflect a nationally representative sample. Due to small sample size (n = 20), non-binary respondents are included in overall statistics but detailed subgroup analyses focus on male and female respondents.

LGBTQ+ STATUS:

LGBTQ+ teens are overrepresented in this report, making up 34% of all respondents. While this is important to note, an oversampling of this underrepresented group is ultimately a strength of this project.

SCHOOL TYPE:

Respondents could select from the following school types: public school, charter school, private school, homeschooled, and other. However, quotas were not set to reflect a nationally representative sample. As a result, there is an overrepresentation of non-traditional schools (schools other than public schools) which could be reflected in the responses.

STATISTICAL SIGNIFICANCE:

When making comparisons between subgroups, all data were tested for statistical significance. Unless otherwise noted, these findings are described in the text in a comparative manner (e.g., more likely, less likely) only if the differences are statistically significant at the level of $p < .05$.

Bars with one asterisk (*) are statistically significant at the $p < .05$ level.

Bars with two asterisks (**) are statistically significant at the $p < .01$ level.

Bars with three asterisks (***) are statistically significant at the $p < .001$ level.

Chi-square tests were used to assess differences between demographic groups. All statistical significance is at the 95% confidence level ($p < .05$). For crosstabulations showing pairwise comparisons, Bonferroni correction was applied to adjust for multiple comparisons and control for Type I error. Percentages may not sum to 100% due to rounding.

NOTES

About FOSI: The Family Online Safety Institute is an international, non-profit organization that works to make the online world safer for kids and their families. FOSI convenes leaders in industry, government, and the non-profit sectors to collaborate and innovate new solutions and policies in the field of online safety. Through research, resources, events, and special projects, FOSI promotes a culture of responsibility online and encourages a sense of digital citizenship for all. FOSI's membership includes many of the leading internet and telecommunications companies around the world.