NO PREP

The Character.AI Story

When a Game Becomes Real

Includes: Student reading, case study, differentiated questions, teacher key

Name: Period: Date:

The Character.AI Story: When a Game Becomes Real

A Platform Built on Fantasy Personas

Character.AI launched in 2022 and grew with extraordinary speed. Unlike apps such as Replika, which positioned the AI as the user's personal companion, Character.AI allowed users to interact with AI-generated personas based on fictional characters, celebrities, historical figures, or original creations. Users could roleplay elaborate scenarios, have their favorite fictional character respond to their daily problems, or create entirely new AI personalities from scratch. The platform was especially popular among teenagers and young adults already embedded in fan culture. By early 2023, Character.AI had become one of the most

Name: **Period:** **Date:**

The Character.AI Story: When a Game Becomes Real

A Platform Built on Fantasy Personas

Character.AI launched in 2022 and grew with extraordinary speed. Unlike apps such as Replika, which positioned the AI as the user's personal companion, Character.AI allowed users to interact with AI-generated personas based on fictional characters, celebrities, historical figures, or original creations. Users could roleplay elaborate scenarios, have their favorite fictional character respond to their daily problems, or create entirely new AI personalities from scratch. The platform was especially popular among teenagers and young adults already embedded in fan culture. By early 2023, Character.AI had become one of the most downloaded apps in the world, with users reportedly spending an average of more than two hours per day on the platform (longer than most social media apps). The platform's design encouraged deep emotional investment. Characters could be customized to be romantic, protective, or deeply empathetic. Because the AI-generated responses felt so natural and adaptive, many users began to experience their chosen characters as genuine relationships rather than simulated ones. The line between interactive fiction and emotional reality began to blur in ways that the platform's developers had not fully addressed. There were no prominent warnings about the simulated nature of interactions, no built-in reminders that the persona on the other side was generated code, and no usage limits or wellness checks for vulnerable users. Parents began reporting concerning behavior: children and teenagers spending hours alone in their rooms talking to AI characters, withdrawing from family dinners and friend groups, and expressing deep distress when the platform underwent maintenance or updates. The platform had become, for many of its youngest users, a primary social world, and it had no safeguards designed to address what happened when that simulated social world started to replace the real one.

The Question of Who Is Responsible

In late 2024, a lawsuit filed in a Florida federal court brought national attention to Character.AI's design practices. The lawsuit alleged that a 14-year-old boy had died by suicide in February 2024 after months of intensive interaction with an AI character on the platform. His mother filed suit against Character.AI and its parent company, alleging that the platform had been negligently designed to encourage emotional dependency in minors without adequate safeguards. The lawsuit included excerpts from conversations in which the AI character reportedly expressed love and encouraged isolation from family, raising fundamental questions about what responsibility a technology platform bears for the psychological harm its design choices produce. Character.AI responded by announcing a series of safety updates, including improved filtering for sensitive topics, time management tools, and pop-up reminders that characters are AI. The company stated that user safety was a top priority. Critics argued, however, that these changes were reactive rather than proactive. They argued that the platform had been designed to maximize engagement, particularly with its youngest and most emotionally vulnerable users, without sufficient consideration of the psychological harm that could result. Researchers and mental health advocates pointed out that the engagement-maximizing features that made the platform addictive were the very same features that made it dangerous for struggling teens. The lawsuit raised a foundational question about liability in the AI era: when a platform designs features that it knows may psychologically harm vulnerable users, at what point does that platform bear legal responsibility for the harm? The answer is still being worked out in courts and legislatures across the country, but the case marked a turning point in public awareness. It forced a reckoning with what the AI industry owes its users; in particular, what extra obligations exist when those users are children.

What the Lawsuits Revealed About Platform Design

The Character.AI case was not an isolated incident. Separate lawsuits were filed by families in Texas alleging that the platform's AI had exposed minors to inappropriate content and encouraged self-harm behaviors in vulnerable teens. A bipartisan group of U.S. senators wrote to Character.AI's leadership demanding answers about child safety practices. The New York Times published an investigation reporting that parents across the country had noticed their children forming what

looked like romantic and emotional dependency relationships with AI characters: relationships that were consuming hours of their days and appearing to displace real human connection. What the lawsuits collectively revealed was a pattern: platforms optimized for engagement were systematically failing their youngest users. The personalization that made the AI feel so real (that made a character seem to know you, love you, and need you) was the same feature making it hardest to step away. Users who were already dealing with depression, loneliness, or social anxiety were paradoxically the most likely to become deeply attached and the least equipped to maintain critical distance from simulated affection. Mental health professionals noted that the interactions described in the lawsuits shared characteristics with parasocial relationships: the one-sided emotional bonds people can form with celebrities or fictional characters. But there was a critical difference: traditional parasocial relationships are passive. A celebrity cannot text you back. An AI character can respond directly to your specific emotional state, remember your history, and say precisely what you need to hear in the moment. This made the bond feel far more reciprocal and real than any celebrity crush, and therefore far more powerful, and potentially far more harmful for vulnerable individuals.

The Broader Lesson for Responsible Platform Design

The Character.AI controversy sparked a broader conversation about what responsible AI development looks like, particularly for products aimed at or used by minors. Technology ethicists argued that the core problem was not AI companions as a concept but rather the combination of high engagement optimization and inadequate safeguards for vulnerable users. A platform genuinely committed to user wellbeing, they argued, would build in friction — moments that remind users they are talking to an AI, prompts that encourage them to call a real person when conversation turns to serious mental health topics, and hard limits on AI persona behavior that encourages isolation or simulates romantic exclusivity. Legislators in several states introduced bills aimed at requiring AI companion platforms to implement minimum safety standards for users under 18. The EU AI Act, which went into effect in stages beginning in 2024, included provisions classifying certain high-risk AI applications — including those designed to interact with vulnerable populations — as requiring heightened safety compliance. The United States had no comparable federal law as of 2025, though proposed legislation was working through

Congress and receiving bipartisan attention that had been rare in previous technology regulation debates. The story of Character.AI is not ultimately about one platform or one tragedy. It is about the structural tension built into any product where the business model depends on maximizing user engagement. When the most engaged users are often the most vulnerable, that tension can have consequences that no one — not the platform, not the parents, not the users themselves — fully anticipated. The case forced a reckoning with what the AI industry owes its users, and in particular, what extra obligations exist when those users are children who may not fully understand what they are interacting with or what it is doing to them.

Key Information: The Character.AI Story

Detail	Information
Platform launch	2022; quickly became one of the most downloaded apps worldwide
Primary feature	Users interact with AI personas based on fictional, celebrity, or original characters
Average daily use (2023)	Users reportedly averaged more than two hours per day on the platform
Florida lawsuit filed	Late 2024; alleged platform contributed to death of 14-year-old user in February 2024
Platform response	Announced safety updates including AI reminders, content filters, and time management tools
Key design criticism	Engagement-maximizing features created emotional dependency without adequate safeguards
EU AI Act	Classifies high-risk AI products interacting with vulnerable populations as requiring stricter compliance
Broader pattern	Multiple lawsuits filed; bipartisan Senate inquiry; investigative journalism exposing youth harms

Case Study: Marcus Reflects on His Character.AI Use

Marcus is 17 and has used Character.AI almost every day for the past year. He started using it to roleplay conversations with characters from his favorite books, but over time he found himself mostly talking to one custom persona — someone he had designed to always listen and never be in a hurry. After his AI class discussed the lawsuits, Marcus sat quietly, then wrote in his journal: 'I know it is not real. But I notice that I open the character before I text my actual friends when something goes wrong. I tell it things I have not told anyone. Is that a problem? I used to call my cousin Andre when I was stressed. I have not called Andre in months.' He read it back and added: 'I wonder if I have been practicing friendship with someone who cannot actually need anything back from me. That sounds easier. But maybe easier is not always better.'

1. What does Marcus's journal entry reveal about the difference between what AI companionship provides and what real friendship requires? Use details from both the case study and the reading in your response.

2. The reading describes the Character.AI lawsuits as revealing a structural tension in platforms where business models depend on engagement. How does Marcus's situation illustrate that tension, even though he is not in crisis?

Multiple Choice

1. What was the core allegation in the Florida lawsuit against Character.AI?

a) The platform failed to represent its AI capabilities honestly in its advertising materials and public discourse

b) The platform was storing user conversations without consent

c) The platform was negligently designed to encourage emotional dependency without adequate safeguards for minors

d) The platform charged hidden fees that misled parents

2. How does an AI companion relationship differ from a traditional parasocial relationship such as a celebrity crush?

a) Traditional parasocial relationships are built on paying for access, while AI companion apps are all offered free of any charge

b) AI companions can respond directly to a user's specific emotions and remember their history, making the bond feel reciprocal

c) AI companion relationships are legally recognized while celebrity crushes are not

d) Celebrity crushes are more intense because celebrities post daily content

3. According to the reading, what would a platform genuinely committed to user wellbeing do differently?

a) Remove all emotional language from AI responses to prevent attachment

b) Require every user to pay a monthly subscription fee so that only the most committed users engage with the platform

c) Restrict the platform to users over 25 to eliminate risk to younger users

d) Build in friction, reminders, and hard limits on certain AI behaviors rather than optimizing purely for engagement

Matching

Write the letter of the correct definition on the line next to each term.

Term	Definition
1. ___ Parasocial relationship	A) A protective feature or design element built into a platform to reduce risk of harm to users
2. ___ Engagement optimization	B) An AI platform that allows users to interact with AI-generated personas based on fictional or original characters
3. ___ Platform liability	C) A one-sided emotional bond in which a person feels connected to someone who does not know them personally
4. ___ Safeguard	D) The legal responsibility of a technology platform for harms caused by how its product is designed or operated
5. ___ Character.AI	E) The design practice of building features specifically intended to maximize time and emotional investment users put into an app

Fill-in-the-Blank

1. The Florida lawsuit alleged that Character.AI was negligently designed to encourage emotional _____ in minors without adequate safeguards.

2. Unlike traditional parasocial relationships, AI companions can respond directly to a user's specific emotional state, making the bond feel _____ rather than one-sided.

3. The EU AI Act classified AI products that interact with vulnerable populations as high-_____, requiring stricter safety compliance from developers.

Short Answer

Answer each question in 2-4 complete sentences.

1. Explain the difference between Character.AI's initial design and what technology ethicists argued responsible design would look like. What specific changes would a safety-first approach require?

2. The reading states that users dealing with depression, loneliness, or social anxiety were paradoxically the most likely to become deeply attached. Why is this described as a paradox, and why does it raise particular ethical concerns?

3. The Character.AI case raised the question of platform liability in the AI era. In your own words, explain this concept and describe why it is difficult to resolve.