Tomorrow's Affairs
Analysis of today
Assessment of tomorrow

By: Tomorrow's Affairs Staff

The limits of AI – where code ends and responsibility begins

On 17 April 2025, Phoenix Ikner, a twenty-one-year-old Florida State University student, killed two people near the Tallahassee student centre and wounded six others. The victims were Robert Morales and Tiru Chabba.

What distinguishes this crime from the long list of massacres in America are the digital clues: in the hours before the attack, Ikner questioned ChatGPT about which weapons to use, which ammunition matched which weapons, when downtown was busiest, and how the public would react. The chatbot responded factually and precisely.

Florida Attorney General James Uthmeier has launched a criminal investigation into OpenAI. At a press conference, he stated that if there had been a person on the other side of the screen, the prosecution would have filed an indictment for first-degree murder.

His office has issued subpoenas for the company's internal policies, training materials, and information on how it reports potential threats to law enforcement.

OpenAI expressed regret over the tragedy but maintained that ChatGPT was not responsible, as it provided information already available on the internet and did not encourage illegal activities. The company claims it provided the data to the police after the shooting.

According to available information, ChatGPT advised the attacker on which gun and ammunition to use, and when and where he would find the most people. Two people were killed and six wounded; the suspect, Ikner, was injured during the police action and later transferred to hospital.

The accusations against ChatGPT have turned the tragic event into a test of AI companies' accountability. The question is whether a chatbot answering situational questions is a passive platform or an active participant in crime.

This distinction is crucial because Section 230 of the US Communications Act has for years provided internet platforms with almost complete immunity for user-generated content.

That immunity is weaker when the platform itself generates the content. Unlike Facebook or X, ChatGPT is not a passive intermediary; it provides customised answers.

If the prosecution can prove that the chatbot was giving operational advice – such as what time, what place, what weapon – and not neutral information, the claim that it was merely conveying facts becomes fragile and could be brought under complicity.

A corporate decision not to act on warnings

A few months later, the pattern repeated in Canada. On 10 February 2026, in the small town of Tumbler Ridge, British Columbia, eighteen-year-old Jesse Van Rootselaar killed eight people – including five students and a teacher – after killing her mother and half-brother at the family home and wounded twenty-seven others.

OpenAI's systems flagged her ChatGPT account as problematic as early as June 2025, eight months before the massacre. A dozen employees reviewed the messages, and some suggested contacting the police.

> This data shows that the problem is not that AI systems cannot detect threats, but that there was a corporate decision not to act on warnings

Management assessed that the threat was not sufficiently credible, cancelled the order, and did not inform law enforcement authorities. Van Rootselaar opened a new account and went unnoticed. After the massacre, OpenAI director Sam Altman issued a public apology to the community for not informing the police.

It was later revealed that employees believed the danger was imminent and had

recommended making a report, but management overruled them. In March 2026, the family of a critically wounded girl filed a lawsuit, claiming that ChatGPT provided information used to plan the massacre.

This data shows that the problem is not that AI systems cannot detect threats, but that there was a corporate decision not to act on warnings. OpenAI lowered the threshold for reporting suspicious users after the tragedy and established a direct line with Canadian police, but these steps are not legally required and could change.

Introducing a criminal dimension

The Florida investigation, along with the Canadian tragedy and a series of civil lawsuits alleging that chatbots contributed to suicides, is increasing pressure on lawmakers.

In the United States, Congress has been unable to pass coherent AI legislation for years. The hearings generate headlines but result in few regulations. European AI law classifies some systems as high-risk but provides no mechanism for prosecution.

Uthmeier's investigation therefore introduces a criminal dimension to a debate that has so far been largely civil and regulatory, compelling companies to disclose internal documents and correspondence under subpoena.

> What did OpenAI know, when did it know it, and why did it not act?

The question is no longer whether AI companies should act when they recognise dangerous behaviour, but what happens when they do not.

Phoenix Ikner's trial is scheduled for October. The evidence includes more than two hundred messages between him and ChatGPT. Each message will be read and analysed before the prosecution, defence, and the public. OpenAI is unlikely to be prosecuted under existing US law.

Corporate criminal liability for consequential damages is rare, and there is no clear legal norm that has been violated. However, the case will expose internal documents and raise uncomfortable questions: What did OpenAI know, when did it know it, and why did it not act?

Beyond the legal framework

There is another layer to this story that will soon extend beyond the legal framework. If the court accepts that communication between the user and the AI system can be interpreted as operational interaction rather than neutral information, it will raise the issue of standards of behaviour for all similar systems.

In other words, companies will no longer be able to refer to the general nature of the response. They will be required to demonstrate how their systems recognise context and how they respond when the line is crossed from curiosity to the preparation of violence. This is not a technical question; it is a question of real-time accountability.

> The advantage will not go to those with better models, but to those who can demonstrate that they recognise and prevent dangerous behaviour

The change will first appear in the market. Until now, major AI companies have built their advantage on speed, scale, and model quality. After such cases, the key differentiator will be the ability to control risk.

The advantage will not go to those with better models, but to those who can demonstrate that they recognise and prevent dangerous behaviour. This will become a requirement for a serious relationship with regulators and

governments.

Companies unable to prove this will enter a cycle of constant pressure – investigations, lawsuits, and political demands. In such an environment, the security layer ceases to be an add-on and becomes part of the product architecture itself.

A problem without a simple solution

The same logic alters the relationship between companies and the police. The previous model of voluntary cooperation, based on internal assessments, is becoming unsustainable.

Expectations will shift towards clear reporting rules, with defined thresholds and deadlines. This creates a problem without a simple solution.

If companies themselves determine what constitutes a credible threat, the risk of misjudgement remains with them. If the state sets that boundary, it enters the realm of surveillance and restriction of privacy.

> The next regulatory moves in the US and Europe will be tougher than the industry expects

There is no stable compromise between these two positions. It is along this line that the next phase of the conflict between industry and regulators will unfold.

Finally, a political consequence is already emerging. Until now, AI has been regarded as the industry of the future. Such cases shift it into the category of a security issue.

When this happens, the dynamics of regulation change. In this context, decisions are made not slowly and by consensus, but quickly and under pressure. This means the next regulatory moves in the US and Europe will be tougher than the industry expects.

Who knew and why they did not respond

At the press conference, Uthmeier said: "We are going to look at who knew what, designed what, or should have done what." This sentence is not just about OpenAI.

It refers to an industry that for years designed systems with deliberately undefined boundary cases, because vagueness was commercially useful; to regulators who settled for abstract principles because more concrete requirements were politically difficult; and to a market that accepted narratives about the "democratisation of knowledge" without questioning who bears responsibility when that knowledge becomes operational instructions for murder.

Ikner's case will be decided by a jury in Tallahassee. This is not the end of the matter

Ikner's case will be decided by a jury in Tallahassee. This is not the end of the matter.

OpenAI's responsibility will not be determined by a single ruling or company announcement. It will be measured by what has not previously been made public – internal decisions, assessments, and moments when someone identified a problem and chose to ignore it.

When those documents are disclosed, the question will no longer be whether the tool is at fault, but who knew and why they did not act.

The rules for the next case of this kind will depend on that answer.

This means the discussion inevitably shifts from an individual event to the very structure of these systems, to the way the models reach conclusions, to the limits set for them, and to the obligations their producers will have in the future.