

Did ChatGPT Help the Tumbler Ridge Shooter?

What verified records show about detection, escalation, account access and the unresolved question of chatbot assistance

Prepared and reviewed by DarkSignals.

Research and evidence cutoff: 26 September 2026 · Public release v1.4

<https://darksignals.org/investigations/did-chatgpt-help-tumbler-ridge-shooter/>

Summary

OpenAI acknowledges that an automated system flagged an account linked to the attacker in June 2025, that people reviewed and banned it, and that police were not notified under the company's then-current threshold. OpenAI later found a second account. The RCMP now says companies receiving requests complied, while public reporting says OpenAI provided chat logs to law enforcement; the public record does not establish their full scope or make them public. Plaintiffs allege chatbot assistance, but DarkSignals has not authenticated the underlying conversations. The evidence establishes a consequential detection-and-escalation failure; it does not establish material assistance, legal causation or preventability.

Bottom line

The independently reviewable record confirms detection, human review, enforcement and non-referral. Whether ChatGPT materially assisted the attacker remains unresolved.

Full investigation

3. Important reader note

DarkSignals did not send this report to OpenAI, the Government of British Columbia, the RCMP or plaintiffs' counsel for comment before publication. No responses are included.

Central evidence limitation. DarkSignals has not obtained the underlying ChatGPT records or OpenAI's internal review file. The RCMP says all platform companies that received requests have complied, and OpenAI told The Wall Street Journal that it provided chat logs to law enforcement. These statements indicate that investigators received requested material, but neither the logs nor their full scope or completeness are public to DarkSignals. The report separates acknowledged facts, attributed reporting, allegations and inference throughout. [S1][S2][S4][S36][S40][S41]

Civil complaints are one party's account. Their allegations are not findings of fact, and OpenAI disputes material allegations. No court has ruled on the merits of the Tumbler Ridge claims reviewed for this draft. [S1][S2][S3]

An external media investigation supplied by the commissioning editor prompted a fresh review of the question. DarkSignals has not relied on its account of the chats because the underlying records and anonymous sourcing cannot be independently tested. It is retained as a background lead only. [S33]

Prepared and reviewed by DarkSignals.

The report concerns a mass shooting in which children died. It names no victims, provides no operational detail and reproduces no violent chatbot output.

4. Executive summary

On 10 February 2026 an 18-year-old killed their mother and 11-year-old half-brother at home in Tumbler Ridge, British Columbia, then killed five students and an education assistant at the town's secondary school before dying by suicide. Official and court-linked sources report 27 people injured. [S1][S2][S7][S18][S35]

OpenAI acknowledges that an automated system flagged an account linked to the attacker in June 2025, that people reviewed it, that the account was banned, and that police were not notified because the activity

did not meet the company's then-current threshold for credible and imminent planning. Sam Altman later apologised for not alerting law enforcement. RCMP court records, as reported by Canadian media, say OpenAI voluntarily disclosed that it had located two accounts believed to be linked to the attacker and relevant to the investigation. A judge later ordered production of account records; OpenAI says it provided chat logs to law enforcement, and the RCMP says companies receiving requests complied. The underlying records remain unavailable to the public. [S4][S5][S6][S36][S40][S41]

The contested questions begin where those admissions end. Plaintiffs and the Province of British Columbia allege that reviewers viewed the material as a credible threat, recommended referral and were overruled. OpenAI disputes claims about who decided and why; its current court filing addresses forum rather than the merits. These are untested allegations. [S1][S2][S3]

The public record supports a finding that OpenAI detected concerning violent-activity misuse and did not refer it before the attack. Newer statements indicate that law enforcement later received requested platform records. They do not make the chat contents public or establish what the chatbot said, whether it encouraged the user, or whether it materially assisted planning. DarkSignals has not reviewed or authenticated the logs. [S1][S2][S4][S36][S40][S41]

OpenAI's later statement that it would refer the June account under its revised protocol shows that its escalation threshold changed. That is relevant to governance and a missed opportunity to seek assessment by authorities; it is not proof that a referral would have prevented the attack or that ChatGPT caused it. [S4][S5]

The strongest public-interest case is institutional: high-risk AI services need documented escalation criteria, accountable decision-makers, durable controls for serious repeat violations, lawful cross-border referral channels and independent audit. A B.C. coroner's inquest has been directed to examine systemic issues, including information-sharing between online platforms, AI and law enforcement. As of this report's research cut-off, no inquest date had been announced in the sources reviewed. The Chief Coroner said he could not compel OpenAI to participate, so the inquest is not a guaranteed route to OpenAI testimony or public access to the logs. [S3][S42][S43]

5. Key judgments

Confidence labels express DarkSignals' assessment of the available evidence, not the seriousness of an allegation.

KJ1. High confidence. OpenAI detected the account, subjected it to human review, banned it and made no police referral in June 2025. Basis: OpenAI's own correspondence and public safety statement, reinforced by Altman's apology. [S4][S5][S6]

KJ2. Moderate confidence. Some staff favoured contacting police. Basis: multiple reports and pleadings describe an internal debate, while OpenAI confirms that referral was considered. Gap: no internal recommendation, review note or testimony is public. [S1][S2][S4][S12]

KJ3. Low confidence. The public record does not establish who made the final decision or whether reputational or commercial considerations affected it. Plaintiffs make those allegations; OpenAI disputes them. [S1][S2][S3][S25]

KJ4. Low confidence. The available primary record does not establish whether ChatGPT merely supplied generic information, reinforced violent ideation, or materially assisted operational planning. Complaints allege assistance, but DarkSignals has not seen the chat records and adopts no finding about their content. [S1][S2]

KJ5. Moderate confidence. Account-level enforcement did not prevent subsequent access. OpenAI acknowledges a second account and says it discovered it only after the attacker's identity became public. The creation date, linkage signals and detection history remain unknown. [S4][S5]

KJ6. Low confidence. A June 2025 referral might have prompted intervention, but the counterfactual cannot be resolved from public material. Police and other institutions already possessed separate warning information, while the content and legal usefulness of an OpenAI referral are unknown. [S3][S8][S20]

KJ7. Moderate confidence. The case exposes a system problem spanning platform safety, law enforcement,

firearms controls, mental-health services and school safeguarding. A B.C. inquest is planned to examine systemic issues, but its date is unset and its Chief Coroner says he cannot compel OpenAI to participate. Assigning legal responsibility to any party requires evidence not yet public. [S3][S7][S42][S43]

KJ8. Low confidence. OpenAI has described post-attack reforms, but DarkSignals found no independent audit establishing their implementation or effectiveness. [S4][S5]

6. Overall confidence assessment

Overall confidence is moderate on the institutional chronology and low on the content of the chats and on causation. OpenAI's own documents establish detection, human review, enforcement, non-referral, subsequent discovery of a second account and a revised referral approach. [S4][S5][S6]

Confidence falls sharply beyond that core. Pleadings are adversarial, much journalism traces to overlapping sources, and repetition is not independent corroboration. Investigators appear to have received requested platform records, but the underlying chat logs, account data and OpenAI decision file remain unavailable to DarkSignals and the public. The coroner's inquest has no announced date in sources reviewed, depends on investigative work concluding, and cannot compel OpenAI according to the Chief Coroner. Strong claims about assistance, encouragement or preventability therefore remain provisional.

7. Confirmed, alleged, disputed and unknown

Status labels: Confirmed; Acknowledged by OpenAI; Alleged in litigation; Independently reported; Disputed; Unverified; Unknown.

Issue | Current status | What is known | Principal source | OpenAI's position | Confidence | What would resolve it

Attack and casualties | Confirmed | 10 Feb 2026. Eight victims killed (five students aged 12 to 13, an education assistant, the attacker's mother and 11-year-old half-brother). The attacker also died. Twenty-seven injured. | S1, S2, S18 | Not disputed | High | RCMP final report; coroner's findings

June 2025 flag and human review | Acknowledged by OpenAI | Automated system detected the account; it went to human review; the account was banned. | S4 | Acknowledged | High | Redacted audit trail

No police referral in June 2025 | Acknowledged by OpenAI | OpenAI says credible and imminent planning was not identified. The CEO apologised for not alerting law enforcement. | S4, S6 | Acknowledged; explained by threshold | High | Not needed for the fact; needed for the reason

Reviewers recommended referral | Alleged in litigation; Independently reported | About a dozen staff debated; some urged leaders to alert authorities (unnamed sources). Complaints plead a recommendation. | S12, S14, S1, S2 | Considered referral; recommendation not denied in reviewed statements | Moderate | Internal review records; testimony

Leadership or global affairs overruled reviewers | Alleged in litigation; Disputed | September complaints name a chain of command. OpenAI says the named executive was not involved and the CEO was not personally involved. | S24, S25, S3 | Denied | Low | Decision records; depositions

Reputational or commercial motive | Alleged in litigation; Disputed | Pleaded on information and belief; supported by general remarks of a former investigator quoted in a complaint. | S1, S2, S25 | Denied | Low | Internal communications

Account "banned" or only "deactivated" | Disputed | OpenAI calls it a ban. Complaints say deactivation allowed re-registration by ordinary means. Help Centre wording not opened by DarkSignals. | S4, S5, S1 | Says a repeat-violator system existed but did not catch the account | Low | Help Centre version history; account-action logs

A second account existed | Acknowledged by OpenAI | OpenAI says it was found after the attacker's identity became public and then shared with law enforcement. | S4, S5 | Acknowledged | High | Creation date, linkage and detection logs

OpenAI knew of the second account before the attack | Unknown | OpenAI says it discovered the account afterwards. Nothing public contradicts this. | S4 | Denies prior knowledge | Low | Account and detection

logs

What ChatGPT said to the attacker | Unknown to DarkSignals | OpenAI says it provided chat logs to law enforcement and the RCMP says companies receiving requests complied; the logs remain unavailable for independent review. | S1, S2, S4, S40, S41 | Says provided logs to law enforcement | High that the public record is insufficient | Authenticated logs and independent review

ChatGPT materially assisted planning | Alleged in litigation; not established | Plaintiffs allege assistance. DarkSignals has not seen evidence sufficient to determine whether any output was operationally useful or used. | S1, S2 | Disputes responsibility; merits not adjudicated | Low | Authenticated chats, account logs and causal evidence

Prior police contact and firearms history | Independently reported | RCMP said guns were seized and returned after an appeal; unsealed records show 12 police-file entries since 2018. | S20, S8 | Relied on in forum motion | Moderate | RCMP disclosure; inquest

A referral would have prevented the attack | Unknown | Plaintiffs allege it; OpenAI points to other institutions. | S1, S2, S3 | Disputes causation | Low | Full evidentiary record

Promised reforms implemented | Unverified | Written commitments exist. No independent confirmation found. | S4, S5, S2, S31 | Says safeguards strengthened | Low | Audit or regulator review

Any court finding on the merits | Confirmed (none found) | OpenAI's reviewed motion seeks dismissal on forum grounds; no merits ruling was identified by the research cut-off. | S3 | Not applicable | Moderate | Docket monitoring

8. Scope and methodology

The investigation asks four questions: what OpenAI detected; how the company escalated and enforced the case; what the accessible record establishes about the chatbot's role; and what governance lessons follow. It does not assume that ChatGPT caused the attack.

Research for version 1.4 ended on 26 September 2026. Priority was given to court filings, OpenAI's correspondence and policy statements, RCMP releases, British Columbia government statements, official parliamentary evidence and official coroner information. Reputable journalism was used to locate and cross-check primary material, not to convert untested allegations into fact.

The supplied Mother Jones investigation was treated as an editorial lead only. Its descriptions of non-public records were not used as the basis for a factual finding, quoted or reproduced. No violent prompts were run and no chatbot safeguards were tested. [S33]

Each source is classified as primary or secondary and carries access, reliability and legal-review notes. Reports deriving from the same lawsuit, press release or anonymous source chain are not counted as independent confirmation. Section 27 explains the principal limitations.

9. Concise incident background

Tumbler Ridge is a remote mining town of about 2,400 people in north-east British Columbia. [S2] On 10 February 2026 police say an 18-year-old, named by police as Jesse Van Rootselaar, killed their mother and 11-year-old half-brother at home, then attacked Tumbler Ridge Secondary School, killing five students aged 12 and 13 and an education assistant. [S18][S20][S2] The attacker died of a self-inflicted injury. Twenty-seven people were injured, two seriously. [S2]

DarkSignals acknowledges the eight people killed and the survivors, families and community. It names none of the victims and gives no detail of the violence.

Casualty figures vary. Early reports gave ten dead including the attacker. [S19] Later police and court material gives eight victims killed plus the attacker's death, nine deaths in all. [S2][S18] Injuries are given as 25 or 27 depending on whether two seriously injured students are counted separately. [S2][S11] OpenAI's brief describes six killed at the school and two seriously wounded, consistent with the other figures. [S3] DarkSignals uses the RCMP-derived figures in the strongest sources.

The RCMP described its criminal investigation as active and ongoing on 7 July 2026 and said it was still

analysing seized devices and collecting data from online accounts. It also said it would participate in the coroner's inquest. [S7] The Chief Coroner directed an inquest on 3 March to examine the circumstances of the deaths and systemic issues, including how mental-health and public-safety systems intersect with firearms oversight and how information is shared between platforms, AI and law enforcement. [S42][S43] As of 26 September, no inquest date was identified in the sources reviewed. Its timing depends on the RCMP investigation and the coroner's own investigative phase. The Chief Coroner told The Canadian Press he could not compel OpenAI, a U.S. company, to participate; he also described the inquest as a non-fault-finding process. It may generate systemic recommendations but cannot be treated as a guaranteed source of OpenAI testimony or public chat records. [S3][S43] Politicians called for a separate public inquiry, but no announcement of one was located by the research cut-off. [S44]

10. Chronological timeline

Date | Event | Status | Sources

From Apr 2018 | Attacker linked to 12 police-file entries over the years; details redacted. Firearms later seized from the home and returned after the owner's appeal. | Independently reported | S8, S20

June 2025 | OpenAI's automated system flags the account; human review; account banned; no police referral. Conversations described gun-violence scenarios over several days (reported). | Acknowledged (flag, ban, no referral); reported (content) | S4, S11, S14

June 2025 | About a dozen staff debate contacting police; some urge it (unnamed sources). Complaints allege leadership overruled reviewers. | Alleged; Disputed | S12, S1, S24, S25

Nov-Dec 2025 | CCDH and CNN test ten chatbots with simulated teen users (not this account). | Published research | S29

10 Feb 2026 | The attack at a family home and the school. | Confirmed | S1, S2, S18

11 Feb | BC declares a state of emergency. An OpenAI representative meets a government contact at a pre-arranged meeting without mentioning the account. | Independently reported | S2, S10, S11

12 Feb | OpenAI asks the province for RCMP contact details. | Independently reported | S10

15 Feb | The FBI tells the RCMP that OpenAI has located two accounts; the sworn ITO later says OpenAI voluntarily disclosed that the accounts were believed linked and relevant. | Reported in unsealed court records | S8, S9, S36

20 Feb | An RCMP corporal learns of the June 2025 internal review from a WSJ story. | Independently reported (court records) | S9

21 Feb | WSJ report dated 21 February (an RCMP record places awareness on 20 February; publication time to be verified). OpenAI confirms the June ban and says the threshold was not met. | Acknowledged | S13, S12, S14

24 Feb | Minister Solomon convenes OpenAI leaders in Ottawa. | Acknowledged | S4, S16

26 Feb | OpenAI letter: second account found after the name was public; revised criteria; commitments. | Acknowledged | S4, S15

9 Mar | Parents of a wounded student sue in the B.C. Supreme Court; later discontinued and refiled in California. | Alleged; Confirmed (filing) | S21, S3

23-24 Apr | CEO letter of apology dated 23 April, published 24 April. Premier says the RCMP investigation is in its final stages. | Acknowledged | S6

28 Apr | OpenAI publishes its community-safety post. | Acknowledged | S5

29 Apr | Seven family and survivor complaints filed in N.D. California (lead: Stacey v. Altman). | Confirmed (filing) | S1, S3, S22

May | A journalist publishes a test of ChatGPT guardrails (excerpt dated 12 May; original date to be confirmed). | Reported | S31

7 Jul | RCMP: investigation active, FBI assisting, online-account data being collected. | Confirmed | S7

13–14 Aug | Heavily redacted RCMP court records become available to media organizations; DarkSignals did not locate public scans of the underlying ITOs or order. | Reported | S8, S9, S36

2 Sep | OpenAI moves to dismiss on forum grounds (hearing 8 Oct). Thirty further complaints filed, adding a global-affairs allegation that OpenAI denies. | Confirmed (filings); Disputed (allegation) | S3, S24, S25

17 Sep | Gun-control groups write to the RCMP seeking transparency about firearms used in the attack. | Independently reported | S28

21 Sep | British Columbia and School District 59 sue in N.D. California (3:26-cv-10743). | Confirmed (filing) | S2, S26

24 Sep 2026 | An external investigation is supplied to DarkSignals as an editorial lead. Its non-public source material is not independently available and its claims are not adopted as findings. | Background lead only | S33

25 Sep | RCMP says companies receiving information requests complied; OpenAI separately tells The Wall Street Journal it provided chat logs to law enforcement. | Public statements reported by media | S40, S41

3–6 Mar | An RCMP investigator discusses record storage with OpenAI counsel; a judge then orders OpenAI to produce records from both accounts within 14 days. | Reported in unsealed court records | S9, S36

3 Mar | Chief Coroner directs an inquest into the deaths. Its remit includes systemic questions about mental health, public safety, firearms oversight, and information-sharing between platforms, AI and law enforcement. No date is announced; timing depends on investigative phases. | Official announcement; reporting on scope and limits | S3, S42, S43

26 Sep | No inquest date identified by the research cut-off. RCMP had described its investigation as active in July; the Chief Coroner says he cannot compel OpenAI to participate. | Status as at report cut-off; absence finding from sources reviewed | S7, S42, S43

Evidence chain

Stage | What is documented | Strongest source | Who makes the claim | Evidence public? | OpenAI position | What is missing | Confidence | Corrective control

User activity | OpenAI says misuse connected with violent activity triggered detection; detailed content is not public. | S4 | OpenAI | Partly | Acknowledged at a general level | Authenticated chats, dates and context | Low on content | Preserve records; independent review

Automated detection | System flagged the account in June 2025. | S4 | OpenAI | Partly (statement) | Acknowledged | Classifier, signal, thresholds | High | Publish detection criteria in aggregate

Human review | Account went to human review. | S4 | OpenAI | Partly | Acknowledged | Reviewers, number, notes | High | Mandatory review for repeated high-risk cases

Internal recommendation | Some staff urged referral (reported and pleaded). | S12, S1 | Unnamed sources; plaintiffs | No | Considered referral | Records of the recommendation | Moderate | Document reasons when recommendations are rejected

Management decision | No referral; OpenAI says threshold not met. | S4 | OpenAI; plaintiffs | Partly | Acknowledged; decision-maker disputed | Who decided; written reasons | High (outcome); Low (who, why) | Named accountable decision-maker

Account action | Account banned or shut down. | S4 | OpenAI | Partly | Acknowledged; ban wording disputed | Exact action; what was retained | Moderate | Identity-level controls for serious violations

Subsequent access | A second account existed and was identified after the attack. | S4 | OpenAI | Partly | Acknowledged | Creation date, linkage and detection history | High on existence; Low on detail | Stronger repeat-violator controls

Attack | Eight victims killed; 27 injured. | S2 | Court filings; police | Yes | Not disputed | Full police findings | High | Not applicable

Post-incident disclosure | OpenAI voluntarily disclosed that it had identified two accounts; a later production order sought records from both. OpenAI says chat logs were provided to law enforcement; RCMP

says companies receiving requests complied. | S8, S9, S36, S40, S41 | ITO reporting; OpenAI statement; RCMP statement | Yes, per OpenAI and RCMP statements | OpenAI says it provided logs to law enforcement | Exact scope and completeness; provincial request is separate | Moderate | Clarify reporting and record-production channels

Government response and litigation | Ministerial engagement, official B.C. statements and federal cases; no merits ruling identified. | S2, S3, S34 | Government and court parties | Yes in part | Contests responsibility and forum | Answers, discovery and rulings | Moderate | Regulatory standard for threat referral

11. What the public record says about the chats

The underlying chats are not public to DarkSignals. A judge signed a 6 March 2026 production order directing OpenAI to provide records from two accounts within 14 days, according to reporting on the unsealed RCMP court materials. The RCMP later said all companies that received requests had complied; OpenAI separately told The Wall Street Journal that it had provided the chat logs to law enforcement. Together, these statements indicate that investigators received the requested records. DarkSignals has not seen the order or logs themselves, so their exact scope, completeness and contents remain unverified here. The coroner's inquest has no announced date in sources reviewed and cannot compel OpenAI to participate according to the Chief Coroner. [S2][S9][S36][S40][S41][S43]

OpenAI's verified account is narrow. It says automated abuse detection identified misuse connected with violent activity; human reviewers examined the account; the company banned it; and the material did not meet its then-current threshold for referral. OpenAI has not publicly described the triggering language or the model's replies. [S4][S5]

The Stacey complaint alleges, on information and belief, that ChatGPT supplied information, instructions or encouragement and that product features reinforced a violent fixation. The Province's complaint advances broader product and failure-to-warn allegations while acknowledging that it does not possess the full conversations. These are allegations, not independently verified facts. [S1][S2]

Three analytically distinct possibilities remain: the system supplied generic information; it reinforced or normalised violent ideation; or it materially assisted planning. The present public record does not allow DarkSignals to choose among them. The fact that OpenAI treated the use as a policy violation does not, by itself, establish intent, assistance or causation.

DarkSignals discusses the existence and reported handling of the records without reproducing operational instructions, casualty estimates or graphic material. The fact that details have appeared in reporting does not make the underlying conversations independently authenticated. [S33][S40]

12. What the public evidence does not establish

- The authenticated text, sequence and context of the chatbot exchanges.
- Whether the model refused, redirected, complied with or encouraged particular requests.
- Whether any output was novel, operationally useful or used in the attack.
- When the second account was created, how it was linked to the first, or whether it was flagged before the attack.
- The reviewers' written assessment, the form of any recommendation, and the final decision record.
- Whether commercial, political, privacy or reputational considerations affected the decision; those motives are alleged and disputed. [S1][S2][S25]
- What information a referral would have contained, what legal powers police could have used, and whether intervention would have followed.
- Whether OpenAI's promised reforms have been independently implemented and tested.

13. How OpenAI detected the activity

OpenAI says an automated system detected the account, which was then sent to human review to decide

whether usage policies were violated and whether the account warranted referral to law enforcement. [S4] Its public statement in February said the account was identified through abuse detection combining automated tools and human investigations aimed at misuse in furtherance of violent activities. [S14]

OpenAI's April post describes classifiers, reasoning models, hash-matching and blocklists, and says human reviewers assess flagged content in context, including patterns over time. [S5] It does not say which tool flagged this account.

Policy violation or credible threat? OpenAI found a policy violation and banned the account. It says it did not find credible and imminent planning at the time. [S4] The two are different tests. A user can break a rule on violent content without presenting an imminent threat. The dispute is whether the June conversations crossed the second line.

How many reviewers? The Wall Street Journal reported that about a dozen staff debated the matter. [S12] [S14] Plaintiffs' counsel and the Stacey complaint use the figure twelve. [S23][S1] The figure comes from unnamed sources, and OpenAI has not confirmed it in any reviewed statement.

14. Human review and the internal warning

OpenAI confirms that people reviewed the account and considered whether to refer it. [S4] It has not confirmed that anyone recommended a referral.

The Wall Street Journal reported that some staff read the conversations as an indication of real-world violence and urged leaders to alert authorities, and that leaders decided against it. [S14][S12] The complaints go further and say the team found a credible, specific threat. [S1][S2] The September complaints attribute this to the Intelligence and Investigations team. [S24]

These accounts are not independent. The British Columbia complaint says the facts became public because whistleblowers spoke to the Wall Street Journal. [S2] The Stacey complaint also cites remarks by a former investigations-team member, as reported by other outlets, that OpenAI was reluctant to make referrals because they would expose how the product may worsen the threat environment. [S1] DarkSignals has not seen that reporting and the remarks are general, not about this case.

The Stacey complaint says the safety team urged notification. Counsel said publicly that as many as twelve people did so. [S1][S23] Counsel's statements are advocacy, not evidence.

15. Why police were not alerted

OpenAI's stated reason is that the activity did not meet its threshold, then described as credible and imminent planning. [S4] The company has also said it weighs privacy and the distress caused by unnecessary police involvement, as quoted in the Stacey complaint. [S1] Its letter says it must protect the privacy and security of the many people who use its tools responsibly. [S4]

Was referral discussed and who supported it? Yes, according to OpenAI and the reports. [S4][S14] Who supported it beyond unnamed staff is not public.

Who decided? The September complaints allege that the global affairs team led by Chris Lehane overrode the safety team and that Lehane and Sam Altman approve referral decisions. [S24] OpenAI's Chief Strategy Officer, Jason Kwon, has said it is false that Lehane was involved in the original decision or that investigators report to him. [S25] OpenAI's court filing says Altman was not personally involved in the June 2025 review or decision. [S3] These are allegations and denials. No decision record is public.

Why did the activity allegedly fail the threshold? OpenAI later said its revised criteria account for the fact that a user may not state target, means and timing, and that under them the account would be referred today. [S4][S5] The admission implies that the June threshold was narrower than the current one. It does not say the June decision was unreasonable at the time.

Commercial or reputational motives. The complaints allege that OpenAI avoided a hard referral rule to protect its safety narrative and a possible public offering. [S1][S2] They plead this on information and belief. OpenAI denies that political or public-relations factors were at play. [S25] DarkSignals found no public document showing such a motive.

What OpenAI has not addressed: the content of the conversations, the identity of the decision-makers,

whether the decision was documented, and what it told police directly after the attack. [S10]

16. Account action and subsequent access

Terms. OpenAI's letter says the account was shut down and banned. [S4] Its April post says enforcement can include disabling an account, banning other accounts of the same user, and steps to detect and stop new accounts. [S5] The Stacey complaint says OpenAI does not ban users but deactivates them, which shuts down the email address used and leaves the person free to register again. [S1] DarkSignals found no authoritative OpenAI definition of suspension, deactivation and ban and did not open the Help Centre article the complaints cite.

When? The initial account review was in June 2025. OpenAI says it later identified a second account after the attack. RCMP court records, as reported by Canadian media, state that OpenAI voluntarily disclosed that it had located two accounts believed to be linked to the attacker and relevant to the investigation. A judge then ordered OpenAI to produce records from both accounts. OpenAI says it provided chat logs to law enforcement, and the RCMP says all companies that received requests complied. The logs remain unavailable to DarkSignals; the public record reviewed here does not establish their completeness or the precise material provided. [S4][S8][S9][S36][S40][S41]

The second account. OpenAI says it maintained systems intended to identify repeat policy violators but found the second account only after the attacker's name became public. It then shared information with law enforcement. The public record does not establish when the account was created, what linkage signals existed, or whether any earlier detection occurred. [S4][S5]

Did OpenAI know of the second account before the attack? OpenAI says not. [S4] Nothing public contradicts that. What safeguards existed? OpenAI committed to strengthening repeat-offender detection, which implies the existing safeguard was insufficient for this case. [S4]

17. The wider network of warning signs

OpenAI was not the only body with information. Unsealed RCMP records link the attacker to 12 police-file entries dating from 2018, and the RCMP has said it once seized guns from the home and returned them after the owner successfully appealed. [S8][S20] A search warrant records insecure firearms at the residence after the attack. [S8] The main firearm believed used at the school was never previously seized, and its origin has not been disclosed. [S28]

OpenAI's court filing, citing media reports, says the attacker had been involuntarily hospitalised more than once, had left school, and had held a firearms licence that expired in 2024. [S3] DarkSignals has not seen those exhibits and treats them as advocacy pending verification.

Reporting based on the Wall Street Journal also mentioned a mall-shooting game on Roblox and Reddit posts about guns. [S12] These are unverified and may have been apparent only in hindsight.

Which signs were known before the attack? The police files and the firearms seizure were known to police. OpenAI's flag was known to OpenAI. No public record shows any body holding both. Which appear only with hindsight is uncertain because the timing of each contact is redacted or unstated.

Authority to intervene sat mainly with police and firearms officials, and with health providers under provincial law. As a private provider, OpenAI's practical options were ending service and reporting. Focusing responsibility exclusively on OpenAI is therefore not supported. Focusing on everyone but OpenAI is not supported either, because OpenAI held conversation content that no other body appears to have seen. The failures interacted: an OpenAI referral would have reached police who already held files, which is the plaintiffs' argument. [S2] It would also have raised the question of why earlier contacts had not restricted access to guns, which is OpenAI's. [S3]

Note that the Province, which is now a plaintiff, is also a party whose institutions are under scrutiny. [S2][S3]

18. Causation versus missed opportunity

Would a referral have identified the individual? Plausibly. OpenAI held an account and later shared the second account with law enforcement. [S4] The complaints say it held name and location data. [S2]

Unverified.

Could police have acted? Unknown. No public analysis shows what Canadian police could lawfully have done on a private company's warning. The RCMP already had files and had dealt with firearms at the home. [S8][S20]

What could OpenAI lawfully have disclosed? US law permits a provider to disclose communications to a governmental entity where it believes in good faith that an emergency involving danger of death or serious physical injury requires disclosure without delay. [S32] Whether that provision applies to ChatGPT conversations, or to a foreign police force, has not been assessed by DarkSignals and is a question for counsel. OpenAI in practice worked through the FBI after the attack. [S8][S9]

Would the information have met the threshold for intervention? Unknown. Could it have affected access to firearms? Plausibly, because police had previously seized guns, though they returned them and the main school firearm was never seized. [S20][S28] The main firearm's origin is unknown.

Where does prevention sit on the scale? DarkSignals assesses with low confidence that prevention is plausible but unproven, and not probable on the current record. It is not unknowable in principle; it is unknowable today.

Legal or factual causation would require the chat content, OpenAI's decision records, what OpenAI would have told police, police procedures for such a tip, and proof that the resulting steps would probably have stopped the attack. None is public.

19. OpenAI's public position and subsequent changes

After the attack, OpenAI said it proactively contacted the RCMP. The unsealed-record reporting gives more detail: the FBI notified the RCMP on 15 February that OpenAI had located two accounts; the RCMP's sworn ITO, as quoted by The Canadian Press, says OpenAI voluntarily disclosed that it had found two accounts believed linked to the attacker and relevant to the investigation. On 3 March an RCMP investigator spoke with OpenAI lawyer Jeff Shih about record storage, and on 6 March a judge ordered production of account records. OpenAI later told The Wall Street Journal it provided chat logs to law enforcement; the RCMP said companies receiving requests complied. This indicates the requested records were provided, but DarkSignals has not seen the logs or the underlying order and cannot assess scope or completeness. [S4][S9][S36][S40][S41]

OpenAI's 26 February letter says it made its referral criteria more flexible, brought in mental-health, behavioural and law-enforcement experts, and would refer the June account today. It promised direct contacts with Canadian law enforcement, stronger detection of repeat violators and periodic review of detection thresholds. [S4] Its chief executive apologised in April. [S6] The April post restates the imminent-and-credible standard and says criteria do not require the user to state target, means and timing. [S5]

OpenAI has said it has strengthened its safeguards. [S22] DarkSignals found no independent confirmation that these changes were implemented or that the referral threshold changed in practice. British Columbia alleges that promised reforms were not followed through despite efforts by its Attorney General. [S2][S26] A journalist's test in May reported that guardrails could be bypassed by reframing; this is one unreplicated test and not an audit. [S31]

British Columbia's Attorney General has publicly called for accountability, stronger safeguards and greater transparency, while the Province's filed complaint seeks relief through the courts. These are government positions and litigation claims, not findings on the merits. [S2][S34]

20. The legal cases and what they actually allege

Proceeding | Court and number | Filed | Status (per sources)

Gebala family claim | B.C. Supreme Court, S-261734 | 9 Mar 2026 | Discontinued; refiled in California [S21][S3]

Seven family and survivor complaints (lead: Stacey v. Altman) | N.D. California, 3:26-cv-03701 to -03707; Judge Corley | 29 Apr 2026 | Three complaints directly reviewed; consolidated forum motion filed 2 Sep; hearing listed for 8 Oct 2026 [S1][S3][S37][S38]

Thirty further complaints | N.D. California | 2 Sep 2026 | Reported filed; procedural posture requires docket verification [S24]

Province and School District 59 | N.D. California, 3:26-cv-10743 (PACER header: 4:26-cv-10743-KAW) | 21 Sep 2026 | Complaint directly reviewed; no merits answer or ruling identified [S2]

Claims. Three directly reviewed April family complaints—Stacey, Mwansa and Chisanga, and Lampert—each plead the same eleven principal causes of action: negligent failure to warn law enforcement, negligent entrustment, aiding and abetting a mass shooting, negligent failure to warn, negligent undertaking, negligent and strict-liability design defect, strict-liability failure to warn, unfair competition, wrongful death and survival. [S1][S37][S38] The British Columbia and School District complaint pleads the first eight. [S2]

Source-chain caution. The three April complaints were filed on the same date by the same counsel and use substantially the same structure and OpenAI allegations; the plaintiff-specific accounts of loss differ. Their repetition therefore shows a coordinated litigation theory, not three independent confirmations of what happened inside OpenAI or what the chatbot said. [S1][S37][S38]

Remedies. Family plaintiffs seek damages, punitive damages and injunctive relief including identity-level bans, mandatory referral and independent monitoring. [S1][S37][S38] The Province seeks extraordinary public costs for policing, health care, victim services, education and a replacement school, together with structural safeguards and independent monitoring. [S2] It pleads no sum.

What British Columbia adds. The government plaintiffs advance a public-cost theory and seek structural reform. Their complaint also says the attacker was a minor under provincial law, where the age of majority is nineteen, while using ChatGPT. These remain pleaded positions, not findings. [S2]

Shared allegations. The complaints allege a June 2025 flag, a reviewer recommendation, a leadership decision not to refer, later access through a second account and chatbot assistance. The directly reviewed pleadings ultimately trace those platform allegations to the same asserted information rather than separate authenticated records. [S1][S2][S37][S38]

OpenAI's position. OpenAI has not filed a merits answer in the seven April actions reviewed here. Its consolidated motion seeks dismissal in favour of British Columbia, arguing that Canadian witnesses and evidence are central and that causation and comparative fault will be core issues. The motion does not establish or concede the disputed platform allegations. [S3]

Docket numbering note. The filed British Columbia complaint carries 4:26-cv-10743-KAW in the PACER header and 3:26-cv-10743 in its internal caption. The source register records both to prevent the same filing being mistaken for two proceedings. [S2]

Issues that could decide liability include duty, product status, causation, comparative fault, recoverability of public costs and the legal significance of platform safety promises. No reviewed court has decided those issues on the merits. [S1][S2][S3][S37][S38]

Still untested: that reviewers recommended referral, that leadership overruled them, that ChatGPT assisted planning, that commercial or reputational interests affected the decision, and that a referral would have prevented the attack. [S1][S2][S37][S38]

21. The public-safety and privacy dilemma

When should an AI company report a user? DarkSignals' analysis is that fantasy, fiction, research, anger and curiosity can resemble intent, so words alone cannot decide. Reviewers need context: plan specificity, access to means, any named place or person, a timeframe and repeated behaviour. OpenAI's revised criteria accept that target, means and timing need not all appear. That approach may reduce missed threats while widening false-positive risk. [S4][S5]

On 23 March 2026, University of Ottawa law professor Michael Geist cited Tumbler Ridge once in opening testimony before the House of Commons Standing Committee on Industry and Technology, as an example in support of a proposed AI Transparency Act. The hearing was a broad AI-policy study, not an inquiry into the shooting; the transcript records no case-specific questioning or finding. It adds evidence of the transparency debate, not new facts about the incident. [S39]

False positives matter. OpenAI serves hundreds of millions of people. [S4] A low threshold would send police

to many people who posed no danger, with harm to them and to trust, and could deter people from talking about violent feelings. A high threshold risks missing the rare case. The trade-off is real, and no evidence in the public record settles where the line should sit.

Protections that follow: human review of every referral; minimum necessary disclosure; deletion and retention rules; notice to the user where safe; an appeal route after serious bans; and audit of aggregate outcomes. Whether users banned for serious violations should be identifiable across accounts depends on privacy law, which DarkSignals has not assessed. An independent regulator could audit decisions, and a common cross-border protocol would remove doubt about which police force to contact. [S2]

22. Practical safeguards and recommendations

Each recommendation ties to a documented failure point: the referral decision, the ban that did not hold, the cross-border gap, and the absence of oversight.

AI companies

- Use multi-factor threat assessment, not a single imminence test. [S4][S5]
- Require human review for repeated high-risk conversations.
- Publish written escalation thresholds and name the accountable decision-maker.
- Record reasons whenever a safety recommendation is rejected.
- Apply identity-level controls after serious violent-activity bans and hold new linked accounts for review. [S1][S2]
- Preserve relevant records and publish aggregate referral figures.

Governments and regulators

- Set a legal standard for credible-threat referral by AI services, with privacy safeguards.
- Create a lawful cross-border route for emergency disclosure. [S32]
- Require independent audit of safety decisions.
- Clarify whether AI chat services are covered by online-safety law. [S3]

Law enforcement

- Agree a standard intake and triage protocol for platform referrals.
- Link referral intake to firearms and mental-health checks under the law.
- Publish how returned or seized firearms decisions are reviewed. [S28]

Schools and local safeguarding bodies

- Include AI-chat warning signs in threat-assessment training.
- Set a clear channel from platforms to school and health safeguarding leads.

Independent auditors and researchers

- Test referral and ban-evasion controls against a published methodology, not ad hoc prompts. [S29][S31]
- Publish false-positive and false-negative rates in aggregate.

23. Alternative explanations and counterarguments

Violent discussion does not establish intent. Correct. Neither OpenAI nor the complaints say the words alone proved intent. The question is whether the pattern justified review by police. OpenAI's later view is that it would. [S4]

Platforms process vast volumes of ambiguous chat. True, and it explains a high threshold. It does not explain the absence of a written, accountable process for the rare case. [S4]

Reporting ambiguous users can breach privacy and chill speech. This is the strongest counterweight. It supports safeguards, not silence. [S5]

Police may not have been able to act. Possible. Nothing public shows what police would have done, and police already held files. [S8][S20]

The attacker might have proceeded without ChatGPT. Possible. Firearms were reportedly accessible, and no public record shows what ChatGPT added. [S3][S28] This cuts against material assistance claims but not against the referral question.

The information may have been available elsewhere. Likely for generic material. It does not address a referral, which concerns the user, not the information.

Other institutions held more direct warning signs and authority. Supported in part by police records and OpenAI's filing. [S8][S3] But OpenAI held content others did not.

Civil complaints are adversarial, not neutral. Accepted throughout this report. Allegations are attributed and untested. [S1][S2]

Hindsight makes weak signals look obvious. A real risk, especially for unnamed-source accounts of a debate. It is partly answered by OpenAI's own decision to widen its criteria. [S4]

Failure to report does not establish legal causation. Correct. The litigation would still need to test duty, breach, factual causation, legal causation and comparative responsibility against a complete evidentiary record. [S1][S2][S3]

24. Indicators to watch

- The 8 October 2026 forum hearing identified in OpenAI's reviewed motion, and any later forum or consolidation orders. [S3]
- Whether the records provided to police are complete, what date range they cover, what material was produced under the March order, and whether any portion will become public or be disclosed in civil discovery. The inquest date, scope of evidence, and whether OpenAI will participate voluntarily also remain unknown. [S9][S36][S40][S41][S42][S43]
- The coroner's inquest: a date, evidence disclosures, OpenAI's voluntary participation, and any systemic recommendations. [S3][S7][S42][S43]
- OpenAI answers to complaints, and any release of redacted review records.
- Evidence of implemented changes: published thresholds, a Canadian point of contact, audits. [S4]
- Canadian legislation on online harms cited in OpenAI's brief (Bill C-34, not opened by DarkSignals). [S3]
- Any independently authenticated release of the account records or internal review file.

25. Intelligence gaps

- The content, completeness and context of the ChatGPT logs reportedly provided to law enforcement, and the OpenAI internal review file.
- The account opening and deactivation dates and what OpenAI retained.
- The reviewers' written assessment and any recommendation.
- The identity of the decision-maker and written reasons.
- What OpenAI told police directly, and when.
- The RCMP's handling of earlier contacts and firearms returns.
- Canadian and US disclosure law as applied to ChatGPT.
- The 30 complaints filed on 2 September and the Gebala B.C. claim, not read in full.
- Independent evidence of implemented reforms.

26. Final assessment

The headline asks: Did ChatGPT help the Tumbler Ridge shooter? On the present independently reviewable record, the answer is unresolved.

The possibility cannot be dismissed: civil complaints allege assistance and the account was serious enough to trigger detection, human review and enforcement. But those facts do not reveal what the model said or prove that any output influenced the attack. [S1][S2][S4]

Five propositions must remain separate. OpenAI detected concerning use: acknowledged. It banned an account and did not alert police: acknowledged. A second account existed: acknowledged after the attack. ChatGPT reinforced or materially assisted violence: alleged but not independently established. OpenAI or ChatGPT legally caused the attack: untested and not established. [S1][S2][S3][S4]

Moral, operational and legal responsibility are different. Altman apologised for the non-referral, and OpenAI says its current protocol would produce a different escalation outcome. Neither statement concedes legal causation or liability. No merits judgment has been made. [S3][S4][S6]

DarkSignals assesses with moderate confidence that OpenAI had information it now says would be referred and did not refer it under the 2025 threshold. It assesses with low confidence that referral would have changed the outcome and makes no finding that ChatGPT materially assisted the attack. The planned coroner's inquest may examine systemic issues but has no announced date in sources reviewed; its Chief Coroner says he cannot compel OpenAI to participate. No public release of the logs is scheduled in the materials reviewed. Civil discovery and future court orders could change that position, but access and timing remain unknown. [S3][S7][S42][S43]

27. Methodology and limitations

- Open-source research only; no interviews, document requests or forensic examination were conducted.
- DarkSignals has not authenticated or independently reviewed the underlying ChatGPT conversations, account logs or OpenAI decision records.
- Several secondary reports rely on overlapping anonymous sources or repeat allegations from the same pleadings; they are not independent corroboration.
- Three April family complaints were read in full. They were filed by the same counsel and contain substantially similar platform allegations, so they are treated as one litigation source chain rather than independent confirmation. [S1][S37][S38]
- The supplied Mother Jones article is recorded as a background lead but is not relied upon for the report's factual conclusions. [S33]
- Some dates and docket developments remain subject to confirmation; the source register identifies each access limitation.
- Legal questions are described rather than answered and require qualified Canadian and US counsel.

Prepared and reviewed by DarkSignals.

DarkSignals did not seek comment from OpenAI, the Government of British Columbia, the RCMP or plaintiffs' counsel before publication. This report therefore contains no responses from them.

29. Full source register

Source register · S1-S44

S17 and S27 are background only and are not relied upon in the assessment. Some publisher links may require a subscription or restrict access.

S1 — Complaint, Stacey v. Altman et al. — Edelson PC (counsel for plaintiff Mark Stacey) (US District Court, N.D. California (via CourtListener RECAP), 2026-04-29)

<https://storage.courtlistener.com/recap/gov.uscourts.cand.468851/gov.uscourts.cand.468851.1.0.pdf>

S2 — Complaint — Province of British Columbia and School District No. 59 v. OpenAI et al. (N.D. California, 21 September 2026)

<https://storage.courtlistener.com/recap/gov.uscourts.cand.479418/gov.uscourts.cand.479418.1.0.pdf>

S3 — Consolidated motion to dismiss seven family and survivor actions for forum non conveniens — OpenAI defendants (2 September 2026)

https://storage.courtlistener.com/recap/gov.uscourts.cand.468851/gov.uscourts.cand.468851.29.0_1.pdf

S4 — Letter from Ann M. O’Leary to Minister Evan Solomon — OpenAI (Vice President, Global Policy) (OpenAI (cdn.openai.com), 2026-02-26)

<https://cdn.openai.com/pdf/8e938d69-0b67-4994-b9ff-683733ed587e/openai-letter-minister-solomon.pdf>

S5 — Our commitment to community safety — OpenAI (OpenAI, 2026-04-28)

<https://openai.com/index/our-commitment-to-community-safety/>

Link access restricted; independent retrieval not confirmed.

S6 — Letter from Sam Altman to the community of Tumbler Ridge — Sam Altman (OpenAI) (Tumbler RidgeLines (local outlet publishing the letter; PDF linked on page), Dated 2026-04-23; published 2026-04-24)

<https://tumbleridgelines.com/2026/04/24/openai-apologizes-to-tumbler-ridge/>

S7 — Statement on Tumbler Ridge criminal investigation and recent announcement by the BC Attorney General — Royal Canadian Mounted Police (BC) (RCMP, 2026-07-07)

<https://rcmp.ca/en/bc/tumbler-ridge/news/2026/07/4354810>

S8 — Unsealed documents reveal details of RCMP’s response, investigation into Tumbler Ridge shooting — CBC News (Michelle Ghoussoub) (CBC News, 2026-08-13)

<https://www.cbc.ca/news/canada/british-columbia/tumbler-ridge-unsealed-documents-9.7305135>

S9 — Court records detail chaos, fearful calls to RCMP during Tumbler Ridge mass shooting — The Globe and Mail (The Globe and Mail, 2026-08-13)

<https://www.theglobeandmail.com/canada/article-bc-rcmp-handling-tumbler-ridge-mass-shooting-court-documents-released/>

S10 — OpenAI did not mention Tumbler Ridge shooter’s posts in meeting with B.C. officials day after mass shooting: province — Mike Hager, Jesse Winter, Joe Castaldo (The Globe and Mail) (The Globe and Mail, 2026-02-21)

<https://www.theglobeandmail.com/canada/article-tumbler-ridge-openai-rcmp-disclosure/>

S11 — OpenAI’s handling of Tumbler Ridge shooter info opens regulation questions — Global News (Global News, 2026-02-24)

<https://globalnews.ca/news/11687903/openai-tumbler-ridge-shooting-duty-to-inform/>

S12 — OpenAI debated calling police about suspected Canadian shooter’s chats — TechCrunch (TechCrunch, 2026-02-21)

<https://techcrunch.com/2026/02/21/openai-debated-calling-police-about-suspected-canadian-shooters-chats/>

S13 — OpenAI Employees Raised Alarms About Canada Shooting Suspect Months Ago — Georgia Wells (The Wall Street Journal) (The Wall Street Journal, 2026-02-21)

<https://www.wsj.com/us-news/law/openai-employees-raised-alarms-about-canada-shooting-suspect-months-ago-b585df62>

Paywalled / access restricted.

S14 — Tumbler Ridge suspect’s ChatGPT account banned before shooting — BBC News

<https://www.bbc.com/news/articles/cn4gq352w89o>

S15 — OpenAI says Canada mass shooter evaded ban with second ChatGPT account — Associated Press (AP (via Yahoo News), 2026-02-26)

<https://www.yahoo.com/news/articles/openai-says-canada-mass-shooter-015159042.html>

S16 — OpenAI connection to Tumbler Ridge tragedy puts balance between privacy and public safety in the spotlight — BetaKit (BetaKit, 2026-02-27)

<https://betakit.com/openai-connection-to-tumbler-ridge-tragedy-puts-balance-between-privacy-and-public-safety-in-the-spotlight/>

S17 — Canada orders OpenAI safety review after grilling Sam Altman over security lapses — Wire-derived report (original outlet not shown) (Yahoo News, 2026 (c. March; exact date not shown in excerpt))

<https://www.yahoo.com/news/articles/canada-orders-openai-safety-review-161611193.html>

Background only; not relied upon.

S18 — Canadian police say 8 killed in British Columbia's Tumbler Ridge shooting — Al Jazeera (Al Jazeera, 2026-02-11)

<https://www.aljazeera.com/news/2026/2/11/canadian-police-say-10-dead-in-british-columbia-school-shooting>

S19 — Canada in mourning as rare mass shooting claims ten lives in British Columbia town — Reuters (via Malay Mail) (Malay Mail, 2026-02-11)

<https://www.malaymail.com/news/world/2026/02/11/canada-in-mourning-rare-mass-shooting-claims-ten-lives-in-british-columbia-town/208858>

Reader access may vary.

S20 — Carney tells Tumbler Ridge residents that Canadians will always be with you following mass shooting — Reuters (via Yahoo News) (Yahoo News, February 2026 (day not shown in excerpt))

<https://www.yahoo.com/news/articles/carney-other-leaders-mourn-victims-110651329.html>

S21 — OpenAI sued by parents of girl critically wounded in Canada school shooting — Associated Press (Fortune (AP wire), 2026-03-10)

<https://fortune.com/2026/03/10/openai-mass-shooting-canada-lawsuit/>

S22 — Families sue OpenAI over Tumbler Ridge mass shooter's use of ChatGPT — NPR (Geoff Brumfiel) (NPR, 2026-04-29)

<https://www.npr.org/2026/04/29/nx-s1-5798896/tumbler-ridge-mass-shooting-chat-gpt-lawsuit>

S23 — Families of Tumbler Ridge, B.C., mass shooting victims suing OpenAI in California — CBC News (Lauren Vanderdeen) (CBC News, 2026-04-29 (per S2 fn 3; search index shows later update))

<https://www.cbc.ca/news/canada/british-columbia/tumbler-ridge-b-c-mass-shooting-families-suing-openai-9.7181214>

S24 — New lawsuits claim OpenAI execs put image ahead of safety in Canadian mass shooting — NPR (NPR, 2026-09-02)

<https://www.npr.org/2026/09/02/nx-s1-5953021/openai-tumbler-ridge-mass-shooting>

S25 — Teachers, students file new wave of lawsuits against OpenAI over Tumbler Ridge shooting — CBC News (Matt Preprost, Lauren Vanderdeen) (CBC News, 2026-09-02)

<https://www.cbc.ca/news/canada/british-columbia/tumbler-ridge-shooting-open-ai-lawsuits-9.7328382>

S26 — Canada's BC sues OpenAI over ChatGPT role in Tumbler Ridge school shooting — Al Jazeera (Al Jazeera, 2026-09-22)

<https://www.aljazeera.com/news/2026/9/22/canadas-bc-sues-openai-over-chatgpt-role-in-tumbler-ridge-school-shooting>

S27 — B.C. sues OpenAI: a duty to warn police that no court has yet tested — Parminder Kumar Sharma (P.K. Sharma (independent practitioner briefing), 2026-09-22)

<https://www.pk-sharma.com/briefing/bc-sues-openai-tumbler-ridge-duty-to-warn>

Background only; not relied upon.

S28 — Gun control groups call for transparency of firearms used in Tumbler Ridge shooting — Terrace Standard (Black Press) (Terrace Standard, 2026-09-17)

<https://terracestandard.com/2026/09/17/gun-control-groups-call-for-transparency-of-firearms-used-in-tumbler-ridge-shooting/>

S29 — Killer Apps: how popular AI chatbots are enabling the next generation of school shooters and extremists — Center for Countering Digital Hate (with CNN) (CCDH, 2026-03-11)

<https://counterhate.com/research/killer-apps/>

S30 — AI chatbots helped teen users plan violence in hundreds of tests — CNN (CNN, 2026-03-11)

<https://www.cnn.com/2026/03/11/americas/ai-chatbots-help-teen-test-users-plan-violence-tests-intl-invs>

S31 — ChatGPT Gave Me Chilling Advice, as I Simulated Planning a Mass Shooting (excerpt) — Mark Follman (Mother Jones), excerpted by Longreads (Longreads (excerpting Mother Jones), 2026-05-12)

<https://longreads.com/2026/05/12/ai-chatbot-violence-guardrails-investigation/>

S32 — 18 U.S.C. section 2702, Voluntary disclosure of customer communications or records — United States Congress (US Government Publishing Office (US Code), Current text (2024 edition))

<https://www.govinfo.gov/content/pkg/USCODE-2024-title18/html/USCODE-2024-title18-partI-chap121-sec2702.htm>

S33 — External investigation supplied as an editorial lead — Mother Jones (24 September 2026)

<https://www.motherjones.com/media/2026/09/chatgpt-tumbler-ridge-mass-shooter-openai/>

S34 — Attorney General's statement on B.C. filing legal action against OpenAI — Government of British Columbia

https://archive.news.gov.bc.ca/releases/news_releases_2024-2028/2026AG0067-001105.htm

S35 — Tumbler Ridge investigative update — British Columbia RCMP (13 February 2026)

<https://rcmp.ca/en/bc/tumbler-ridge/news/2026/02/4350292>

S36 — Report on unsealed Tumbler Ridge court records — CityNews / The Canadian Press (13 August 2026)

<https://vancouver.citynews.ca/2026/08/13/court-documents-reveal-confusion-around-number-of-shooters-in-tumbler-ridge->

[killings/](#)

S37 — Complaint, Mwansa and Chisanga v. Altman et al. — N.D. California (29 April 2026)

<https://storage.courtlistener.com/recap/gov.uscourts.cand.468853/gov.uscourts.cand.468853.1.0.pdf>

S38 — Complaint, Lampert v. Altman et al. — N.D. California (29 April 2026)

<https://storage.courtlistener.com/recap/gov.uscourts.cand.468856/gov.uscourts.cand.468856.1.0.pdf>

S39 — Evidence, INDU Committee, Meeting 29

<https://www.ourcommons.ca/DocumentViewer/en/45-1/INDU/meeting-29/evidence>

S40 — Report on Tumbler Ridge shooter’s ChatGPT conversations shocks B.C. attorney general

<https://globalnews.ca/news/12072859/what-did-chatgpt-tell-tumbler-ridge-shooter/>

S41 — British Columbia Sues OpenAI, Alleging ChatGPT Aided Mass School Shooting

<https://www.wsj.com/tech/ai/british-columbia-sues-openai-alleging-chatgpt-aided-mass-school-shooting-aac66568>

Paywalled / access restricted.

S42 — Chief coroner’s statement on Tumbler Ridge

https://archive.news.gov.bc.ca/releases/news_releases_2024-2028/2026PSSG0017-000210.pdf

S43 — Focus on mental health, public safety at inquest into Tumbler Ridge, B.C., deaths

<https://ottawa.citynews.ca/2026/03/03/bc-coroner-announces-tumbler-mass-shooting-inquest/>

S44 — B.C. Conservative MPs Zimmer, Caputo ask for public inquiry into Tumbler Ridge deaths

<https://vancouver.citynews.ca/2026/03/01/b-c-conservative-mps-zimmer-caputo-ask-for-public-inquiry-into-tumbler-ridge-deaths/>

DarkSignals did not seek comment for this report; it contains no responses. Publication authorized following human editorial review.