

# 12 famous AI disasters

Feature

Jun 9, 2025

11 mins

By Thor Olavsrud

Source: [www.cio.com](https://www.cio.com)

Insights from data and machine learning algorithms can be invaluable, but be warned — mistakes can be irreversible. These recent high-profile AI blunders illustrate the damage done when things don't go according to plan.

In 2017, [The Economist](#) declared that data, rather than oil, had become the world's most valuable resource, and the refrain has been repeated ever since. Organizations across every industry have been investing, and continue to heavily invest, in data and analytics. But like oil, data and analytics have their dark side. According to CIO's State of the CIO Survey 2025, 42% of CIOs say AI and ML are their biggest technology priority for 2025. And while actions driven by [ML algorithms](#) can give organizations a competitive advantage, mistakes can be costly in terms of reputation, revenue, or even lives.

Understanding your data and what it's telling you is important, but it's equally vital to understand your tools, know your data, and keep your organization's values firmly in mind. And with that in mind, here are a handful of high-profile AI blunders from the past decade to illustrate what can go wrong.

## Chicago Sun-Times, Philadelphia Inquirer publish summer reading list of fake books

The Chicago Sun-Times and Philadelphia Inquirer took reputational hits when May 2025 editions featured a special section that included a summer reading list recommending books that don't exist.

The [Chicago Sun-Times explained](#) that the syndicated section, "Heat Index: Your Guide to the Best of Summer," was provided by King Features Syndicate, a unit of Hearst. Marco Buscaglia, the author of the special section, admitted he used AI to assist putting it together, including the recommended reading list, and failed to fact check the output.

The reading list featured many real authors but attributed nonexistent books to them. For example, the first recommendation was Tidewater Dreams by famed Chilean-American writer Isabel Allende, who's written more than 20 novels. But Tidewater Dreams, a "climate fiction novel that explores how one family confronts rising sea levels while uncovering long-buried secrets," isn't one of them. Like most books on the list, it was hallucinated by AI.

The newsrooms of both papers said they had nothing to do with the insert, though neither paper marked it as advertorial content. King Features terminated its relationship with Buscaglia following the incident, noting that his use of AI violated a strict policy.

## McDonald's ends AI experiment after drive-thru ordering blunders

After working with IBM for three years to leverage [AI to take drive-thru orders](#), McDonald's called the whole thing off in June 2024. The reason? A slew of social media videos showing confused and [frustrated customers](#) trying to get the AI to understand their orders.

One TikTok video in particular featured two people repeatedly pleading with the AI to stop as it kept adding more Chicken McNuggets to their order, eventually reaching 260. In a June 13, 2024, internal memo obtained by trade publication Restaurant Business, McDonald's announced it would end the partnership with IBM and shut down the tests.

The restaurant had piloted the AI at more than 100 US drive-thrus, and indicated it still saw a future in a voice-ordering solution.

## Grok AI falsely accuses NBA star of vandalism spree

In an April 2024 post on X, Grok, the AI chatbot from [Elon Musk's xAI](#), falsely accused NBA star Klay Thompson of throwing bricks through windows of multiple houses in Sacramento, California.

Some commentators speculated that Grok may have hallucinated the vandalism story about the Golden State Warriors small forward, who's since moved to the Dallas Mavericks, after ingesting posts about Thompson "throwing bricks," common basketball parlance for badly missed shots. In his final game with GSW, the team suffered a blowout loss as part of the worst postseason performance of Thompson's career.

Grok does display a disclaimer saying, "Grok is an early feature and can make mistakes. Verify its outputs." Still, the incident raises questions about liability when an AI chatbot makes false and defamatory statements.

## NYC AI chatbot encourages business owners to break the law

In March 2024, [The Markup](#) reported that Microsoft-powered chatbot MyCity was giving entrepreneurs incorrect information that would lead to them break the law.

Unveiled in October 2024, MyCity was intended to help provide New Yorkers with information on starting and operating businesses in the city, as well as housing policy and worker rights. The only problem was The Markup found MyCity falsely claimed that business owners could take a cut of their workers' tips, fire workers who complain of sexual harassment, and serve food that had been nibbled by rodents. It also claimed landlords could discriminate based on source of income.

In the wake of the report, indicted New York City Mayor Eric Adams defended the project. The chatbot remains online.

## Air Canada pays damages for chatbot lies

In February 2024, Air Canada was ordered to [pay damages to a passenger](#) after its virtual assistant gave him incorrect information at a particularly difficult time.

Jake Moffatt consulted Air Canada's virtual assistant about bereavement fares following the death of his grandmother in November 2023. The chatbot told him he could buy a regular price ticket from Vancouver to Toronto and apply for a bereavement discount within 90 days of purchase. Following that advice, Moffatt purchased a one-way CA\$794.98 ticket to Toronto and a CA\$845.38 return flight to Vancouver.

But when Moffatt submitted his refund claim, the airline turned him down, saying bereavement fares can't be claimed after tickets have been purchased.

Moffatt took Air Canada to a tribunal in Canada, claiming the airline was negligent and misrepresented information via its virtual assistant. [According to tribunal member](#) Christopher Rivers, Air Canada argued it can't be held liable for the information provided by its chatbot.

Rivers denied that argument, saying the airline didn't take "reasonable care to ensure its chatbot was accurate," so he ordered the airline to pay Moffatt CA\$812.02, including CA\$650.88 in damages.

## Sports Illustrated may have published AI-generated writers

In November 2023, online magazine Futurism said Sports Illustrated was publishing articles [by AI-generated writers](#).

Futurism cited anonymous sources were involved to create content, and said the storied sports magazine published a lot of articles by authors generated by AI.

Futurism also found the author headshots in question were listed on a site that sells AI-generated portraits. The online magazine then reached out to The Arena Group, publisher of Sports Illustrated, and in a statement, Arena Group said the articles in question were licensed content from a third party, AdVon Commerce.

"We continually monitor our partners and were in the midst of a review when these allegations were raised," Arena Group said in the statement provided to Futurism. "AdVon has assured us that all of the articles in question were written and edited by humans."

The statement added that AdVon writers used pseudonyms in certain articles, noting that Arena Group doesn't condone those actions and subsequently removed the articles in question from the Sports Illustrated website. Responding to the Futurism article, the Sports Illustrated Union [posted a statement](#) that it was horrified by the allegations and demanded answers and transparency from Arena Group management.

"If true, these practices violate everything we believe in about journalism," the SI Union said in its statement. "We deplore being associated with something so disrespectful to our readers."

## iTutor Group's recruiting AI rejects applicants due to age

In August 2023, tutoring company iTutor Group [agreed to pay \\$365,000](#) to settle [a suit](#) brought by the US Equal Employment Opportunity Commission (EEOC). The federal agency said the company, which provides remote tutoring services to students in China, used AI-powered recruiting software that automatically rejected female applicants ages 55 and older, and male applicants ages 60 and older.

The EEOC said more than 200 qualified applicants were automatically rejected by the software.

"Age discrimination is unjust and unlawful," former EEOC chair Charlotte Burrows said in a statement. "Even when technology automates the discrimination, the employer is still responsible."

iTutor Group denied any wrongdoing but did decide to settle the suit. As part of the settlement and consent decree, it agreed to adopt new anti-discrimination policies.

## ChatGPT hallucinates court cases

Advances made in 2023 by LLMs stoked widespread interest in the transformative potential of gen AI across nearly every industry. OpenAI's ChatGPT was at the center of this surge in interest, but the technology still has a long way to go before it can reliably take over most processes, as attorney Steven Schwartz learned when he found himself in hot water with US District Judge Kevin Castel after using it to research precedents in a suit against Colombian airline Avianca.

Schwartz, an attorney with Levidow, Levidow & Oberman in New York, used the OpenAI gen AI chatbot to find precedent to support a case filed by Avianca employee Roberto Mata for injuries he sustained in 2019. But at least six of the cases submitted in the brief didn't exist. In a [document filed in May 2023](#), Castel noted the cases submitted by Schwartz included false names and docket numbers, along with bogus internal citations and quotes. Schwartz's partner, Peter LoDuca, was Mata's lawyer of record and signed the brief, putting himself in jeopardy as well.

In an affidavit, Schwartz told the court it was the first time he used ChatGPT as a legal research source and was unaware of the possibility that its content could be false. He admitted he hadn't confirmed the sources provided by the AI chatbot, and regretted using gen AI to supplement his legal research, adding he'll never do so in the future without verifying authenticity.

In June 2023, Castel imposed a \$5,000 fine on Schwartz and LoDuca, and in a separate ruling in June, he dismissed Mata's lawsuit against Avianca.

## AI algorithms identify everything but COVID-19

Since the [COVID-19 pandemic](#) began in 2020, numerous organizations have sought to apply ML algorithms to help hospitals diagnose or triage patients faster. But according to the UK's Turing Institute, a national center for data science and AI, the predictive tools made little to no difference.

MIT Technology Review chronicled a number of failures, most of which stemmed from errors in the way the tools were trained or tested. The use of mislabeled data, or data from unknown sources, was a common culprit.

Derek Driggs, an ML researcher at the University of Cambridge, together with his colleagues, [published a paper in Nature Machine Intelligence](#) that explored the use of deep learning models to diagnose the virus. But the paper determined the technique wasn't fit for clinical use. For example, Driggs' group found their own model was flawed because it was trained on a data set that included scans of patients that were lying down while scanned, and patients who were standing up. The patients who were lying down were much more likely to be seriously ill, so the algorithm learned to identify COVID risk based on the position of the person in the scan.

A similar example included an algorithm trained with a data set with scans of chests of healthy children. The algorithm learned to identify children, not high-risk patients.

## Zillow wrote down millions, slashed workforce due to algorithmic home-buying disaster

In November 2021, online real estate marketplace Zillow told shareholders it would wind down its Zillow Offers operations and cut 25% of the company's workforce — about 2,000 employees — over the next several quarters. The home-flipping unit's woes were the result of the error rate in the ML algorithm it used to predict home prices.

Zillow Offers was a program through which the company made cash offers on properties based on a "Zestimate" of home values derived from an ML algorithm. The idea was to renovate the properties and flip them quickly. But a Zillow spokesperson [told CNN](#) the algorithm had a median error rate of 1.9%, and could be as high as 6.9% for off-market homes.

CNN then reported that Zillow bought 27,000 homes through Zillow Offers since its launch in April 2018, but sold only 17,000 through the end of September 2021. Black swan events like the COVID-19 pandemic and a home renovation labor shortage contributed to the algorithm's accuracy troubles.

Zillow said the algorithm led it to unintentionally purchase homes at higher prices rather than its current estimates of future selling prices, resulting in a \$304 million inventory write-down in Q3 2021. In a conference call with investors following the announcement, Zillow co-founder and CEO Rich Barton said it might have been possible to tweak the algorithm, but ultimately it was too risky.

## Dataset trained Microsoft chatbot to spew racist tweets

In March 2016, Microsoft learned that using Twitter interactions as [training data for ML](#) algorithms can have dismaying results.

Microsoft released Tay, an AI chatbot, on the social media platform, and the company described it as an experiment in conversational understanding. The idea was the chatbot would assume the persona of a teenage girl and interact with individuals via Twitter using a combination of ML and natural language processing. Microsoft seeded it with anonymized public data and some material pre-written by comedians, then set it loose to learn and evolve from its interactions on the social network.

Within 16 hours, the chatbot posted more than 95,000 tweets, and those tweets rapidly turned overtly racist, misogynist, and anti-Semitic. Microsoft quickly suspended the service for adjustments and ultimately pulled the plug.

"We are deeply sorry for the unintended offensive and hurtful tweets from Tay, which do not represent who we are or what we stand for, nor how we designed Tay," Peter Lee, corporate VP, Microsoft Research & Incubations (then corporate VP of Microsoft Healthcare), wrote in a post on Microsoft's official blog following the incident.

Lee noted that Tay's predecessor, Xiaoice, released by Microsoft in China in 2014, had successfully conducted conversations with more than 40 million people in the two years prior to Tay's release. What Microsoft didn't take into account was that a group of Twitter users would immediately begin tweeting inappropriate comments to Tay. The bot quickly learned from that material and incorporated it into its own tweets.

## Amazon AI-enabled recruitment tool only recommended men

Like many large companies, Amazon is hungry for tools that can help its HR function screen applications for the best candidates. In 2014, Amazon started working on AI-powered recruiting software to do just that. There was only one problem: The system vastly preferred male candidates. In 2018, [Reuters broke the news](#) that Amazon had [scrapped the project](#).

Amazon's system gave candidates star ratings from 1 to 5, but the ML models at the heart of the system were trained on 10 years' worth of résumés submitted to Amazon — most of them from men. As a result of that training data, the system started penalizing phrases in résumés that included the word "women's," and even downgraded candidates from all-women colleges.

At the time, Amazon said the tool was never used by Amazon recruiters to evaluate candidates. The company tried to edit the tool to make it neutral but ultimately decided it couldn't guarantee it wouldn't learn some other discriminatory way of sorting candidates and ended the project.