

Lessons for Editors of AI Incidents from the AI Incident Database

Kevin Paeth¹, Daniel Atherton^{2,3}, Nikiforos Pittaras¹, Heather Frase^{2,4}, Sean McGregor¹

¹UL Research Institutes

²Responsible AI Collaborative

³Georgetown University

⁴Veraitech

kevin.paeth@ul.org, danny@raicollab.org, pittarasnikif@gmail.com, hnfrase@veraitechus.com, sean.mcgregor@ul.org

Abstract

As artificial intelligence (AI) systems become increasingly deployed across the world, they are also increasingly implicated in *AI incidents* – harm events to individuals and society. As a result, industry, civil society, and governments worldwide are developing best practices and regulations for monitoring and analyzing AI incidents. The AI Incident Database (AIID) is a project that catalogs AI incidents and supports further research by providing a platform to classify incidents for different operational and research-oriented goals. This study reviews the AIID’s dataset of 750+ AI incidents and two independent taxonomies applied to these incidents to identify common challenges to indexing and analyzing AI incidents. We find that certain patterns of AI incidents present *structural ambiguities* that challenge incident databasing and explore how *epistemic uncertainty* in AI incident reporting is unavoidable. We therefore report mitigations to make incident processes more robust to uncertainty related to cause, extent of harm, severity, or technical details of implicated systems. With these findings, we discuss how to develop future AI incident reporting practices.

Introduction

As practitioners of AI incident reporting over the last seven years, we have acted as subject matter experts to researchers, corporations, and governments seeking to develop domain, company, and country-specific incident databases. Such efforts stand to advance safer AI development and deployment along with impact-centered AI governance. Yet throughout these efforts, there are consistent editorial challenges faced by our database and those databases that have yet to launch.

Incident editing challenges require reflection and accommodation, especially in a time where AI incident reporting requirements are being formalized within legislation. For example, the recent European Union AI Act requires developers and deployers of AI systems implicated in “serious” incidents to report to national AI authorities charged with oversight (Council of the European Union 2024). In the United States, the state of Colorado passed legislation requiring AI system developers and deployers to report instances of algorithmic discrimination (Colorado General Assembly 2024), and draft legislation is circulating within the AI policy community for national AI incident reporting.

Copyright © 2025, Association for the Advancement of Artificial Intelligence (www.aaai.org). All rights reserved.

Several corporations are setting up AI-specific incident response programs. These personnel often operate from a history of security incident response. We believe these are positive trends and there is much that can be learned from previously developed incident processes. However, we also believe that the broader community can benefit from a deeper exploration of AI-specific incident cataloging and editing practices.

The contributions of this work are therefore twofold. First, we explore the foundational question of deciding what is and is not an AI incident. While we have contributed to intergovernmental efforts at formalizing the answer, such high level definitions leave many implementation details to the reader when bridging definitions to policy. Second, a collection of incidents is only as useful for preventing incident recurrence as its metadata. Well-motivated analyses of incidents recorded as metadata (i.e., “taxonomies” in our parlance) are necessary for making incidents useful for querying and analysis. Within this work, we explore the challenges and solutions for navigating the evolving technology and practice of AI incident cataloging. These contributions are motivated to both explain our practices and to inform others that may be looking to expand the network of actors collecting and sharing AI incidents.

The AI Incident Database and Related Work

Incident reporting is a practice common in many other fields, from aviation (Federal Aviation Administration 2024) to environmental monitoring (Environmental Protection Agency 2024). Perhaps the field with established incident reporting practices most relevant to AI is cybersecurity, as both fields are rooted in monitoring wide-spread digital technologies where software is associated with both tangible harms (e.g. compromised autonomous vehicle systems resulting in accidents) as well as intangible harms (e.g. compromised data storage resulting in disclosure of private information). Importantly, there exists a widely-adopted practice to handle potential incidents shared by research, industry, and government actors alike, realized in the Common Vulnerabilities and Exposures (CVE) Program (The MITRE Corporation 2020). The CVE Program catalogs reported cybersecurity vulnerabilities for verification and awareness across many stakeholders and industries, providing a “canonical” identification of a vulnerability so that it may be consistently ref-

erenced.

The AI Incident Database. Inspired by the CVE program and other incident reporting practices, the AI Incident Database (AIID) catalogs AI incidents with the goal of reducing the likelihood that similar incidents recur (McGregor 2021). The AIID hosts a human-curated dataset of AI incidents, each of which are substantiated by one or more reports. Currently there are over 750 distinct AI incidents detailed by more than 3000 indexed third-party reports. The nature of the incidents spans various AI systems and contexts, ranging from harms associated with autonomous vehicle accidents to algorithmic discrimination. The AIID’s incidents are maintained by editors who have spent years developing guidelines¹ in response to changing technologies, novel harms, and frequent discussion. The resulting dataset, which is under continual development, is publicly available to download as a database snapshot.²

In addition to serving as a public educational resource on AI safety and ethics (Feffer, Martelaro, and Heidari 2023), the AIID supports deeper AI incident research by providing a platform for further annotating incidents according to different domain-specific analyses and research goals. These are broadly referred to as *taxonomies* in the AIID platform. The two most prominent taxonomies hosted on the AIID are the Center for Security and Emerging Technology (CSET) AI Harm Taxonomy and the Goals, Methods, and Failures Taxonomy, which each address distinct incident analysis goals.

The CSET AI Harm Taxonomy for the AIID (Hoffman and Frase 2023) applies a conceptual framework that *operationalizes notions of harm*. The framework and taxonomy identify differences between “tangible” and “intangible” harms, provide categories of these harms (e.g. harm to physical health or safety, human rights violations, financial loss), and distinguish harms that have occurred versus those that may occur in the future, among other features. At the time of writing, 214 incidents in the AIID contain classifications in the CSET AI Harm Taxonomy for the AIID, but the framework is intended to be illustrative of what is possible in incident annotation more broadly than solve annotation problems within the AIID.³

The Goals, Methods, and Failures (GMF) taxonomy (Pitaras and McGregor 2022) classifies the same underlying incidents in the AIID, but it focuses on failure cause analysis by interrelating the system’s goals, its methods, and their technical causal factors for the observed failure events. The taxonomy structure and methodology encourage *considering what is known or observed* of these aspects along with what is *potential or likely*, guiding how to apply and interpret expert technical knowledge about AI failures in the presence of uncertainty. These “confidence modifiers” are relevant when the incidents in question are based on public

but incomplete third-party reports. At the time of writing, 188 incidents in the AIID contain classifications in the GMF taxonomy.

The efforts behind cataloging incidents and applying these taxonomies present operational challenges beyond the scope of producing research papers. These are open-ended projects for which we expect to never arrive at an end state. As such, we share the operational concerns that companies and countries may also have, herein presented as our research question (**RQ**): what are the current operational challenges to reporting, indexing, and analyzing AI incidents learned from the AI Incident Database as well as the application of its long-term supported taxonomies? To explore the question, we identify incidents in the underlying AIID dataset that “stress” its editing methodology and propose qualitative themes that characterize these cases. We then explore how the downstream taxonomies that further annotate incidents in the AIID either handle or are challenged by these. Throughout, we discuss the implications of these findings as lessons for future AI incident reporting practices.

No such widely-shared and successful practices exist for AI incident reporting despite the growing legislative acts and drafts that would direct agencies to develop AI incident reporting programs. The term “*AI incident*” itself is a relatively new term that presents definitional challenges and ambiguities. Beyond inheriting the issue of defining what is considered to be “AI” across both policy and technical domains (Krafft et al. 2020), defining what can be considered an *AI incident* touches fundamental questions of semantics, values, and law.

Multiple policy-oriented definitions published in 2024 require that an AI incident identifies harm to people, rights, property, or the environment. For example, the definition of “serious incident” in the EU AI Act (Council of the European Union 2024) or definition of “AI incident” by the Organization for Economic Co-operation and Development (OECD 2024) both position “incidents” in relation to their harms. Other perspectives argue a broader sociotechnical approach where AI incidents should be identified as “algorithmic troubles,” events caused by particular algorithmic systems that provoke anxiety or concern about the nature of their functioning (Meunier, Gray, and Ricci 2021). Prior works have focused on similar concerns under the term “AI accidents” (Arnold and Toner 2021), though this term is less popular and nominally excludes purposeful harms.

The construction of these definitions opens up questions integral to reporting practices, such as how to define or scope the nature of harm. An OECD report on approaches to defining AI incidents identifies various potential dimensions of harm – such as the type of harm, severity, reversibility – and highlights difficulties in aligning with existing technical standards and terminology related to severity and risk (OECD 2023). Hoffman and Frase (2023) argue that assessing intangible or derivative harms of AI incidents is inherently difficult, and the difficulty of identifying harm across different incident reporting efforts can lead to inconsistent measurements and recommendations.

¹<https://incidentdatabase.ai/editors-guide/>

²<https://incidentdatabase.ai/research/snapshots/>

³More details about the taxonomy’s annotation methodology and application to AIID can be found in a publicly-available guidelines: <https://github.com/georgetown-cset/CSET-AIID-harm-taxonomy>

Actual Harm vs. Potential Harm

The AIID maintains an operating definition of an AI incident in its editor guidelines as “an alleged harm or near harm event to people, property, or the environment where an AI system is implicated.” Notably, this articulation includes “near harm” events as AI incidents, which characterize incidents such as Incident 8: “Uber Autonomous Cars Running Red Lights” (Anonymous 2014). This construction captures behaviors of AI systems that would have caused harm if not for other environmental factors – a particular subset of alleged “potential” harm. Complementary to this term and intended to capture other types of potential harm, the AIID also ingests what it terms as “AI issues”, which are “an alleged harm or near harm by an AI system that has yet to occur or be detected.” In practice, AIID’s AI issues are often “incidents waiting to happen” and are collected evidence that understanding of where AI presents the potential for harm. This articulation of *actual* versus *potential* harm is also operationalized in the OECD’s AI incident definition and guidelines, which propose the term “AI incident hazard” to capture analogous types of potential and future harm caused by AI systems. The term “hazard” is more common in the product safety community and is likely to be the dominant term moving forward. In the future, we will be harmonizing the AIID “issues” with the OECD “hazards.”

Reporting and Monitoring Projects

In addition to the AIID, other projects that seek to index AI incidents include the OECD AI Incidents Monitor (AIM)⁴ and the AI, Algorithmic, and Automation Incidents and Controversies (AIAAIC) repository,⁵ both of which index similar public reports of incidents albeit by different methodologies or tools. More recently, incident databasing activities have begun to target specific types of harms and incidents, such as the Political Deepfakes Incidents Database (Walker, Schiff, and Schiff 2024), or specific domains, such as the Database of AI Litigation (DAIL)⁶. Other related efforts focus on aggregating and recording AI system vulnerabilities, such as the AI Vulnerability Database,⁷ also inspired by the CVE Program.

Taxonomies for Deeper Analysis

Detailed analysis and organization of AI incident data and AI risks more broadly is often done using taxonomies or other structured annotation. These taxonomies vary in their purposes, often targeting specific technical AI research and development, public policy, humanities-oriented, or educational audiences.

AI incident- and harm-specific taxonomies can seek to generally annotate incidents and harms and opt for inclusion, such as the aforementioned CSET AI Harm and GMF taxonomies. Other taxonomies focus on risks or harms implicating particular use cases of systems, such as speech generation system harms (Hutiri, Papakyriakopoulos, and Xi-

ang 2024), or analyze specific contexts in how the systems are used, such as “misuses” of generative AI (Marchal et al. 2024). Other recent works have compared AI incident- and harm-focused taxonomies in more detail and propose additional taxonomies that seek to annotate AI incidents for broader public consumption (Abercrombie et al. 2024).

The abundance of taxonomies seeking to characterize risks, harms, and incidents caused by AI and other algorithmic systems is so great that meta-studies of the recent literature seek to compose more abstract or more systematically inclusive taxonomies. Based on a scoping review of 172 relevant papers and resources, Shelby et al. proposed a high-level taxonomy of sociotechnical harms and reflect on challenges in unifying taxonomy development and applications across various computing disciplines (Shelby et al. 2023).

Nevertheless, despite the awareness raised by incident cataloging projects, the challenge of capturing the dynamic nature of the AI incidents remains. Further taxonomic analysis and techniques applied to incident data may be required to better understand and prevent future harms (Turri and Dzombak 2023). In the following section, we present the challenges of ingesting and organizing incident data, especially as it impacts the ability to aggregate and analyze data.

Editorial Challenges to Indexing AI Incidents

We tentatively group challenging editing themes as “key AI incident editing challenges.” These themes are not mutually exclusive; AI incidents that challenge intuitive definitions often exhibit one or more of these patterns in practice.

Finally, we review how these challenges impact downstream analysis of AI incidents, focusing on the CSET AI Harm Taxonomy and the GMF Taxonomy.

Key AI Incident Editing Challenges

(1) Temporal Ambiguities. While some harms caused by AI systems are the results of discrete events with defined and intuitive event timelines, other AI incidents can span ambiguous amounts of time, present as series of events that are difficult to neatly separate and aggregate in a database setting, or obscure when exactly harm can be considered to have occurred. This challenges incident reporting models that require knowing the exact or even approximate times of harm occurrence, which are desirable when seeking to analyze incident trends.

A well-known example of an AI incident with temporal ambiguities is known as the “Dutch Childcare Benefits Scandal,” wherein Dutch families were wrongfully accused of tax fraud in part due to a discriminatory algorithm (Anonymous 2018; Peeters and Widlak 2023). In this incident, it is unclear precisely when the problematic algorithmic components in the decision-making process were introduced, which would constitute a potential “start” of the incident. It is also difficult to ascertain when the harms of the incident can be considered to have “ended” given that the harm to families was deep and social; as a result of the algorithmic fraud accusations, many children were taken into foster care and the financial burdens were severe. How does an AI reporting system or database capture the complexity of such a timeline of

⁴<https://oecd.ai/en/incidents>

⁵<https://www.aiaaic.org/>

⁶<https://blogs.gwu.edu/law-eti/ai-litigation-database/>

⁷<https://avidml.org/>

harms? How can ambiguous records be used in data-driven approaches to regulating AI? There is a tension between unifying all of the harms experienced under a single umbrella event and considering the harm to each family or victim as its own distinct AI incident.

It is also clear that AI incidents can be *ongoing*. Those that are a part of disinformation campaigns that leverage AI technologies across platforms, for example. In these situations, the individual pieces of disinformation often defy simple categorization as they evolve over time and reach across multiple jurisdictions while often being described as part of wider, sustained influence operations.

For instance, incidents like “*Kremlin-Linked Entities Allegedly Using Generative AI to Spread Russian Disinformation in Latin America*” (Atherton 2023b), “*Russia Using AI in Disinformation Campaigns to Erode Western Support for Ukraine*” (Atherton 2023c), and “*TikTok AI System Used to Amplify Election Disinformation by Foreign Networks*” (Starkweather 2024) are not isolated events; they certainly contain many discrete harm events, but are often discussed collectively as long-running problems that are continuously adapting and expanding.

Reporting on incidents of this nature involves navigating the inherent tension of capturing the specificity of each incident while also acknowledging the overarching nature of these disinformation campaigns. In some cases, sufficient reporting exists on a singular harm event to formulate one incident ID, such as “*Deepfake Targets Olena Zelenska in Russian Disinformation Campaign*” (Atherton 2024b), but in the aforementioned cases, reporting the wider contours of ongoing and long-running incidents becomes necessary. That complexity is further compounded by the ways in which editorial decisions at the journalistic level might merge incidents into larger narratives in such a manner that eschews the detailed specifics that are necessary for discrete incident reports.

(2) Multiplicity. Most AI incidents are not unique and un-reproducible events but are often caused by systems that are involved in our everyday lives. This means that patterns of harm events can develop wherein a fixed AI system can harm different parties in nearly the same causative manners or repeat the same failures. For example, there are multiple distinct entries for autonomous vehicle accidents implicating the same vehicle models and systems. While the causative factors can differ between incidents (e.g. unexpected highway braking (Atherton 2022) vs. stop sign misidentification (Anonymous 2021)), it is often the case that each incident report bears no distinguishing technical details, resulting in reports of harm and investigations that group together such accidents without a means to distinguish individual harm events (Atherton 2024a).

Other patterns of AI incidents can be more trivially repeatable. Consider the potential for widely-available machine translation systems to separately fail many individuals in the same or similar manners, such as perpetuating gender bias (Anonymous 2017) or creating offensive translations (Suriyawongkul 2020). The same can be said for popular large language model (LLM) applications which have mil-

lions of monthly active users, for which even a relatively small probability of model “hallucination” can materialize to impact thousands of individuals.

In this way, **AI incident reporting systems require a strategy to handle multiplicity of incidents**, especially if the system has a mandate to record every instance of harm (as opposed to being selective). As a potential way to index such incidents, the AIID proposed identifying and cataloging “AI incident variants”, defined as an “*AI incident that shares the same causative factors, produces similar harms, and involves the same intelligent systems as a known AI incident*” (McGregor, Paeth, and Lam 2022). The need for incident reporting strategies to relate or cluster similar incidents and harms in a lightweight and performant manner is increasingly important, especially as popular and generalizable AI systems are increasingly made available and deployed as commercial products (such as LLMs).

Aside from generalized AI incident reporting systems, the multiplicity of certain types of AI incidents also **motivates the opportunity for more specialized incident reporting and database practices that restrict and apply domain expertise to particular types of incidents**. For example, the Political Deepfakes Incidents Database (PDID) focuses explicitly on synthetic media that relate to political and public discourse given the proliferation of such deepfakes and the generative AI systems that produce them (Walker, Schiff, and Schiff 2024).

(3) Aggregate and Societal Harms. All of the AI incident definitions discussed would generally consider harms to groups of people caused by AI systems as incident-worthy. But what are AI incident reporting systems to do when the harm in question is an emergent effect distributed across thousands of people and potentially unobservable at the individual level? Put differently, what happens when an AI system is seen to have harmed society? Such incidents are perhaps among the most pressing to rigorously identify but also the most difficult, as it is often the case that “aggregate” and society-impacting incidents do not indicate tangible harms to individuals and require a holistic perspective. In such cases, the AIID opts to take this broader perspective and index such incidents, but such incidents are perhaps the most difficult to analyze and compare with others.

Algorithmic recommendation systems are particularly of interest, as they are often deployed globally and are potential sources of societal harm such as reinforcing existing inequalities (Saurwein and Spencer-Smith 2021) or amplifying misinformation and extremist content (Whittaker et al. 2021). Consider Incident 348: “*YouTube Recommendation Reportedly Pushed Election Fraud Content to Skeptics Disproportionately*,” where YouTube’s recommendation algorithm allegedly disproportionately displayed election fraud content to users most skeptical of the election’s legitimacy compared to less skeptical users (Lam 2020). While it is unclear if individuals receiving these recommendations can be considered directly harmed, there is a considerable danger to society presented by the associated system if we maintain that polarizing political discourse online is a harm to society in itself.

Despite the ambiguity present in such an incident, other society-impacting incidents can positively indicate individual harms. Consider again the Dutch Childcare Benefits Scandal (Anonymous 2018), wherein specific individuals and families were directly impacted by algorithmic decision-making, and the scale of which was so significant that it resulted in resignation of the then executive government (Peeters and Widlak 2023).

As incident reporting practices formalize, **they must consider how they will substantiate and investigate harm to groups of people and society absent clear harm to individuals.**

(4) Epistemic Uncertainty in Incident Reporting. In the prior section, particular dimensions of ambiguity can be considered to be artifacts or flaws introduced by the way AI incidents are defined, modeled, and compared within incident reporting systems. However, the role of uncertainty in incident reporting is more fundamental than artifacts introduced by modeling an incident for the sake of indexing it. Rather, uncertainty about facts of AI incidents can be *epistemic* – the knowledge of whether certain facts about an AI incident are true or false is unavailable or unattainable. Such facts might concern the nature of the implicated AI system (such as exact machine learning model weights or configuration of the broader AI system), downstream impacts (such as the number of people definitively harmed or exposed), or any number of other details.

Much of this information is simply *unavailable* from the perspective of the incident reporter or collector. As the AIID relies primarily on “open source” incident information to identify incidents (i.e. that which is publicly available, such as third-party media reports, legal allegations, and other public disclosures), editors frequently encounter reports of incidents that lack insight into particular facts (Pittaras and McGregor 2022).

Further, while it is generally true that developers and deployers of implicated AI systems naturally have more privileged access and knowledge about technical details and operating contexts, their knowledge of relevant facts is not necessarily perfect or even sufficient. Key technical information (such as AI system architecture or model weights) may not be fixed in particular AI systems that allow downstream users to customize aspects of the system or integrate it with other systems (e.g., for large language models that operate on a context window of data provided by the users); operational information (such as device logs of system inputs/outputs) may not sufficiently monitor the system or be maliciously destroyed or corrupted. What are AI incident reporting systems to do when information is – for all intents and purpose – irretrievably lost? In this way, AI incident information can be *unattainable*. This illustrates that even under a mandatory reporting regime, an AI incident reporting system will grapple with missing information to some extent.

Another example of re-occurring (and perhaps unavoidable) uncertainty in AI incident analysis is establishing *causality* in the link between the AI system and the harm event. That is, the harm event in question would not have

occurred if not for the behavior of the AI system implicated. This construction of causality is inherently counterfactual and here the answer can also be considered *unattainable* and subject to epistemic uncertainty.

Regardless, such uncertainty should not deter AI incident reporting practices – just as similar uncertainty has not stopped aviation or cybersecurity incident reporting. **Rather, the takeaway is that AI incident reporting systems must embrace methodologies and cultures that are conducive to investigation in the presence of uncertainty.** While the core AIID editing methodology does not explicitly address such uncertainty, downstream taxonomies and consumers of this information can apply additional methods that do so (e.g., “confidence modifiers” in the GMF Taxonomy, discussed next).

Impacts on and Further Lessons From AI Incident Taxonomies

Not only do the aforementioned challenges impact the decision of whether and how to include candidate AI incidents in the database, they also impact the ability of third-party taxonomies in the AIID platform to analyze the data. Depending on the particular methods and goals of a taxonomy, these issues can be mitigated or they can compound.

CSET AI Harm Taxonomy. A goal of the CSET AI Harm Taxonomy for AIID⁸ is to create a structure for extracting AI harm information that helps AI stakeholders to understand risk and harm from AI systems, track trends in AI harm incidents, and identify emerging types and vectors of AI harm (Hoffman and Frase 2023). The taxonomy differentiates between tangible and intangible harm. The taxonomy requires tangible harm to have observable injury, loss, or damage. In contrast, intangible harm usually cannot be directly observed and often depends on societal norms or cultural context. While the taxonomy attempts to capture details about the AI system, sector, environment, entities, locations, dates, and type of harm that were involved in the AI incident, it was primarily developed to identify if tangible or intangible AI harm was involved; distinguish between harm events, near-misses, and issues; and decrease inter-annotator variation.

In addition to confirming the key modeling and uncertainty issues in AI incidents, key lessons learned from the development and application of the CSET AI Harm Taxonomy include acknowledging the challenge of mitigating **inter-annotator variation**. Inter-annotator variation can be reduced but not removed through iterative development of definitions, examples, and annotation processes. For instance, due to the iterative development process, the CSET AI Harm Taxonomy includes questions about the domain or conditions in which the potential AI harm occurred and sought to focus annotators on essential details. These questions helped reduce inter-annotator variation but did not eliminate it, in part because each annotator brings their distinct cultural and personal experiences.

In addition, the taxonomy’s focus on differentiating tangible and intangible harms highlights a particular source of

⁸<https://incidentdatabase.ai/taxonomy/csetv1/>

uncertainty. Especially in the case of **recording intangible harms**, it is inherently difficult to know if intangible harm definitively occurred or if there was simply a nonzero probability that it could have occurred.

Goals, Methods, and Failures (GMF) Taxonomy. As a failure cause analysis taxonomy, the GMF taxonomy⁹ seeks to interrelate the goals of the system deployment, the system’s methods, and technical causal factors for the observed failures in AI incidents. Notably, the GMF taxonomy provides multiple features that intend to mitigate the effects of uncertainty in AI incident reporting. Firstly, the taxonomy records both rich snippets of text from reports as well as annotator rationale for resulting annotations, which serve to address potential uncertainty introduced by public but incomplete third-party reports.

Secondly, the taxonomy also specifically identifies which annotations represent “potential” or “known” facts of the incident. In this way, the GMF taxonomy **explicitly addresses epistemic uncertainty**. Downstream consumers of GMF annotations can compare and aggregate incidents with respect to these confidence modifiers, and these modifiers enable a process where expert and editor communities can iteratively review and upgrade (or downgrade) annotations (Pitaras and McGregor 2022).¹⁰

While the taxonomy mitigates uncertainty by explicitly recording it, it still faces certain ambiguities. In GMF, temporal ambiguity issues transfer from incident delineation to the failure cause analysis domain. Incidents that span an extended period of time may involve multiple AI systems that change over time (e.g. different model iterations, updates to problem-solving approaches, or new mitigation measures), which may in turn result in new methods to annotate or sources of technical failure to consider (it is less likely that the relevant goal of the AI system will change). As a result, incidents that exhibit temporal ambiguity increase complexity of failure cause analysis.

In addition, the development and application of the GMF taxonomy to the AIID dataset surfaces the concern of which incidents are “in scope” for its analysis. The GMF taxonomy focuses on technical failure causes (i.e. methodological or implementation-related steps) that lead to the manifestation of harm. This is in tension with fact that the AIID also indexes incidents where the application of AI systems leads to harm without technical failings. This not only includes cases of malicious use of AI systems (e.g. “deepfakes” (Atherton 2023a)) but also extends to incidents that have aggregate or society-level harms that are not attributed to technical failures (as discussed previously). A simple takeaway here is that the GMF taxonomy and others need not be applied to such incidents, but this highlights the variability of incident editing methodologies where the underlying dataset is more broad.¹¹

⁹<https://incidentdatabase.ai/taxonomy/gmf/>

¹⁰Today, the GMF taxonomy powers the AIID’s “Risk Checklists” feature, which enables users to retrieve incidents by potential goals, methods, and failures of underlying AI systems in order to illustrate risks.

¹¹Ongoing discussion of how to discern non-technical failure

Conclusion

This study presents operational insights into the challenges of cataloging AI incidents faced by the AI Incident Database. They are rooted in seven years of deliberative work to faithfully describe AI incidents for researchers, developers, policy makers, and the general public. We highlight potential mitigations and solutions to these challenges with the goal of empowering the growing AI incident community. Just as the development of AI is dynamic, so must be AI safety practices. AI incident reporting will continue to confront new ambiguities with new technologies, but safer systems will become possible by a careful examination and presentation of the past.

Acknowledgments

This work benefits from years of discussion and concept development among a collection of incident and taxonomy editors. We are particularly indebted to Mia Hoffman, Patrick Hall, Khoa Lam, Zachary Arnold, Helen Toner, Janet Schwartz, and Kate Perkins, but we also would have much less data with which to work if not for the 96 (and counting) named submitters to the AI Incident database submitting the reports of 3,300+ on AI incidents across 1,200 different source domains. Finally, we are lucky to work with a talented team of engineers building the computer systems behind the AIID, including César Varela, Pablo Costa, Clara Youdale, and Luna McNulty.

References

- Abercrombie, G.; Benbouzid, D.; Giudici, P.; Golpayegani, D.; Hernandez, J.; Noro, P.; Pandit, H.; Paraschou, E.; Pownall, C.; Prajapati, J.; Sayre, M. A.; Sengupta, U.; Suriyawongful, A.; Thelot, R.; Vei, S.; and Waltersdorfer, L. 2024. A Collaborative, Human-Centred Taxonomy of AI, Algorithmic, and Automation Harms. ArXiv:2407.01294 [cs].
- Anonymous. 2014. Incident Number 8: Uber Autonomous Cars Running Red Lights. *AI Incident Database*. Retrieved August 2024 from <https://incidentdatabase.ai/cite/8>.
- Anonymous. 2017. Incident Number 59: Gender Biases in Google Translate. *AI Incident Database*. Retrieved August 2024 from <https://incidentdatabase.ai/cite/59>.
- Anonymous. 2018. Incident Number 101: Dutch Families Wrongfully Accused of Tax Fraud Due to Discriminatory Algorithm. *AI Incident Database*. Retrieved August 2024 from <https://incidentdatabase.ai/cite/101>.
- Anonymous. 2021. Incident Number 145: Tesla’s Autopilot Misidentified the Moon as Yellow Stop Light. *AI Incident Database*. Retrieved August 2024 from <https://incidentdatabase.ai/cite/145>.
- Arnold, Z.; and Toner, H. 2021. AI Accidents: An Emerging Threat. Technical report, Center for Security and Emerging Technology.
- Atherton, D. 2022. Incident Number 434: Sudden Braking by Tesla Allegedly on Self-Driving Mode Caused Multi-Car causes in the GMF Taxonomy can be found here: <https://github.com/responsible-ai-collaborative/aiid/issues/2704>

- Pileup in Tunnel. *AI Incident Database*. Retrieved August 2024 from <https://incidentdatabase.ai/cite/434>.
- Atherton, D. 2023a. Incident Number 544: Deepfakes and AI-generated disinformation in the 2023 presidential elections of Turkey. *AI Incident Database*. Retrieved August 2024 from <https://incidentdatabase.ai/cite/544>.
- Atherton, D. 2023b. Incident Number 585: Kremlin-Linked Entities Allegedly Using Generative AI to Spread Russian Disinformation in Latin America. *AI Incident Database*. Retrieved August 2024 from <https://incidentdatabase.ai/cite/585>.
- Atherton, D. 2023c. Incident Number 602: Russia Using Artificial Intelligence in Disinformation Campaigns to Erode Western Support for Ukraine. *AI Incident Database*. Retrieved August 2024 from <https://incidentdatabase.ai/cite/602>.
- Atherton, D. 2024a. Incident Number 711: NHTSA Opens New Probe into Tesla's Autopilot Following More than a Dozen Fatal Accidents. *AI Incident Database*. Retrieved August 2024 from <https://incidentdatabase.ai/cite/711>.
- Atherton, D. 2024b. Incident Number 755: Deepfake Targets Olena Zelenska in Russian Disinformation Campaign. *AI Incident Database*. Retrieved August 2024 from <https://incidentdatabase.ai/cite/755>.
- Colorado General Assembly. 2024. Consumer Protections for Artificial Intelligence (SB24-205) Concerning consumer protections in interactions with artificial intelligence systems. <https://leg.colorado.gov/bills/sb24-205>. Accessed: 2025-01-20.
- Council of the European Union. 2024. Regulation (EU) 2024/1689 of the European Parliament and of the Council of 13 June 2024 laying down harmonised rules on artificial intelligence (Artificial Intelligence Act). <https://eur-lex.europa.eu/eli/reg/2024/1689/oj>. Accessed: 2025-01-20.
- Environmental Protection Agency. 2024. About the Incident Data System (IDS). <https://www.epa.gov/pesticide-incidents/about-incident-data-system-ids>. Accessed: 2025-01-20.
- Federal Aviation Administration. 2024. Accident and Incident Data. https://www.faa.gov/data_research/accident_incident. Accessed: 2025-01-20.
- Feffer, M.; Martelaro, N.; and Heidari, H. 2023. The AI Incident Database as an Educational Tool to Raise Awareness of AI Harms: A Classroom Exploration of Efficacy, Limitations, & Future Improvements. In *Equity and Access in Algorithms, Mechanisms, and Optimization*, 1–11. Boston MA USA: ACM. ISBN 9798400703812.
- Hoffman, M.; and Frase, H. 2023. Adding Structure to AI Harm. Technical report, Center for Security and Emerging Technology.
- Hutiri, W.; Papakyriakopoulos, O.; and Xiang, A. 2024. Not My Voice! A Taxonomy of Ethical and Safety Harms of Speech Generators. In *The 2024 ACM Conference on Fairness, Accountability, and Transparency*, 359–376. Rio de Janeiro Brazil: ACM. ISBN 9798400704505.
- Krafft, P. M.; Young, M.; Katell, M.; Huang, K.; and Bugingo, G. 2020. Defining AI in Policy versus Practice. In *Proceedings of the AAAI/ACM Conference on AI, Ethics, and Society*, AIES '20, 72–78. New York, NY, USA: Association for Computing Machinery. ISBN 9781450371100.
- Lam, K. 2020. Incident Number 348: YouTube Recommendation Reportedly Pushed Election Fraud Content to Skeptics Disproportionately. *AI Incident Database*. Retrieved August 2024 from <https://incidentdatabase.ai/cite/348>.
- Marchal, N.; Xu, R.; Elasmr, R.; Gabriel, I.; Goldberg, B.; and Isaac, W. 2024. Generative AI Misuse: A Taxonomy of Tactics and Insights from Real-World Data. ArXiv:2406.13843 [cs].
- McGregor, S. 2021. Preventing Repeated Real World AI Failures by Cataloging Incidents: The AI Incident Database. *Proceedings of the AAAI Conference on Artificial Intelligence*, 35(17): 15458–15463. Number: 17.
- McGregor, S.; Paeth, K.; and Lam, K. 2022. Indexing AI Risks with Incidents, Issues, and Variants. <https://doi.org/10.48550/arXiv.2211.10384>. arXiv:2211.10384.
- Meunier, A.; Gray, J.; and Ricci, D. 2021. A New AI Lexicon: Algorithm Trouble: Troublesome Encounters with Algorithms that go Beyond Computational Processes.
- OECD. 2023. Stocktaking for the development of an AI incident definition. Technical report, Organisation for Economic Co-operation and Development.
- OECD. 2024. Defining AI incidents and related terms. Technical report, Organisation for Economic Co-operation and Development.
- Peeters, R.; and Widlak, A. C. 2023. Administrative exclusion in the infrastructure-level bureaucracy: The case of the Dutch daycare benefit scandal. *Public Administration Review*, 83(4): 863–877.
- Pittaras, N.; and McGregor, S. 2022. A taxonomic system for failure cause analysis of open source AI incidents. In *Proceedings of the Workshop on Artificial Intelligence Safety 2023 (SafeAI 2023)*. CEUR-WS.
- Saurwein, F.; and Spencer-Smith, C. 2021. Automated Trouble: The Role of Algorithmic Selection in Harms on Social Media Platforms. *Media and Communication*, 9(4): 222–233.
- Shelby, R.; Rismani, S.; Henne, K.; Moon, A.; Rostamzadeh, N.; Nicholas, P.; Yilla-Akbari, N.; Gallegos, J.; Smart, A.; Garcia, E.; and Virk, G. 2023. Sociotechnical Harms of Algorithmic Systems: Scoping a Taxonomy for Harm Reduction. In *Proceedings of the 2023 AAAI/ACM Conference on AI, Ethics, and Society*, AIES '23, 723–741. New York, NY, USA: Association for Computing Machinery. ISBN 9798400702310.
- Starkweather, C. 2024. Incident Number 761: TikTok AI System Used to Amplify Election Disinformation by Foreign Networks. *AI Incident Database*. Retrieved August 2024 from <https://incidentdatabase.ai/cite/761>.
- Suriyawongkul, A. 2020. Incident Number 415: Facebook Provided Offensive Translation for King of Thailand's Birthday Ceremony. *AI Incident Database*. Retrieved August 2024 from <https://incidentdatabase.ai/cite/415>.

The MITRE Corporation. 2020. CVE - Common Vulnerabilities and Exposures. <https://cve.mitre.org/>. Accessed: 2025-01-20.

Turri, V.; and Dzombak, R. 2023. Why We Need to Know More: Exploring the State of AI Incident Documentation Practices. In *Proceedings of the 2023 AAAI/ACM Conference on AI, Ethics, and Society*, AIES '23, 576–583. New York, NY, USA: Association for Computing Machinery. ISBN 9798400702310.

Walker, C. P.; Schiff, D. S.; and Schiff, K. J. 2024. Merging AI Incidents Research with Political Misinformation Research: Introducing the Political Deepfakes Incidents Database. *Proceedings of the AAAI Conference on Artificial Intelligence*, 38(21): 23053–23058. Number: 21.

Whittaker, J.; Looney, S.; Reed, A.; and Votta, F. 2021. Recommender systems and the amplification of extremist content. *Internet Policy Review*, 10(2).