When Artificial Intelligence Fails

The Emerging Role of Incident Databases

Rowena Rodrigues,* **Anais Resseguier,**** **Nicole Santiago*****

* Head of Innovation and Research Services, Trilateral Research, London, United Kingdom, e-mail: rowena.rodrigues@trilateralresearch.com

** Research Manager, Trilateral Research, Waterford, Ireland, e-mail: anais.resseguier@trilateralresearch.com

*** Research Manager, Trilateral Research, Waterford, Ireland, e-mail: nicole.santiago@trilateralresearch.com

Abstract: Diverse initiatives promote the responsible development, deployment and use of Artificial Intelligence (AI). AI incident databases have emerged as a valuable and timely learning resource and tool in AI governance. This article assesses the value of such databases and outlines how this value can be enhanced. It reviews four databases: the *AI Incident Database,* the *AI, Algorithmic, and Automation Incidents and Controversies Repository,* the *AI Incident Tracker* and *Where in the World Is AI.* The article provides a descriptive analysis of these databases, examines their objectives, and locates them within the landscape of initiatives that advance responsible AI. It reflects on their primary objective, i.e. learning from mistakes to avoid them in the future, and explores how they might benefit diverse stakeholders. The article supports the broader uptake of these databases and recommends four key actions to enhance their value.

Keywords: Artificial Intelligence, incident databases, repositories, ethical AI, responsible technology, controversies

1. Introduction

Artificial Intelligence (AI) can enhance and support efficient decision-making, optimise resource allocations, reduce human error and risk, improve human safety and service personalisation. However, AI is also a technology with dual use and misuse potential. The potential harms include adverse impact on human rights, manipulative, exploitative, and social control practices and/or further entrenchment of socio-economic inequalities (Rodrigues et al., 2019; Jansen et al., 2020).

AI incident databases have emerged to identify and/or contribute to learning from and mitigate the risks and harms of AI. An AI incident database refers to any platform to gather information and/or collect source of cases or examples or incidents related to AI, including risks and harms, aimed at learning from, anticipating and addressing these.

This article reviews AI incident databases and responds to the following central questions: What is the primary value of AI incident databases, and how can this value be enhanced? The article looks at the primary value of these databases, i.e. learning from mistakes to avoid them in the future and explore their benefits for specific stakeholders. It identifies and analyses four databases: the *AI Incident Database* (AIID),[1] *The AI, Algorithmic, and Automation Incidents and Controversies Repository* (AIAAIC Repository),[2] *The AI Incident Tracker*[3] and the *Where in the World Is AI*.[4] These databases offer an interesting approach to help address potential harms of AI.

This article was developed in 2021–2022. We searched for relevant AI incident databases through desk-based research in April 2021 (with a review for updates in February–October 2022). In June–July 2021, we interviewed founders/creators/people associated with the databases.[5] We analysed the databases using a SWOT framework, i.e. a systematic identification of Strengths, Weaknesses, Opportunities and Threats (Pickton & Wright, 1998; Hofer & Schendel, 1978) in May–June 2021, which was updated between March–November 2022. This article was checked in January–February 2023 and finalised in November 2023 following journal review feedback. The databases changed since November 2022, and the article was updated in November 2023 to reflect changes (interviewees were contacted and feedback was requested by email). The *Where in the World Is AI* dataset is no longer available; therefore, this article reflects information from its past state.

The article is structured as follows: Section 2 presents information on the databases (what information is collected, its organisation, management and statistical information). Section 3 examines their objectives and how they fill gaps in the AI field and support AI developers and managers. Section 4 highlights how different stakeholders (such as policymakers and regulators, researchers, non-governmental organisations and civil society) can benefit from them. The article concludes with four recommendations to enhance their value and use.

---

[1]   See https://incidentdatabase.ai
[2]   See https://www.aiaaic.org/aiaaic-repository
[3]   See https://bit.ly/3GgHo9t
[4]   See https://map.ai-global.org
[5]   Interviews were carried out with representatives of the four databases: Charlie Pownall (AIAAIC), Sean McGregor (AIID), Ashley Casovan and Kara Scully (from the Responsible AI Institute, for Where in the World Is AI) and J. Patrick Hall (AI Incident Tracker). The questions discussed with the interviewees included: Motivation behind the databases? What sustains the database/repository? Coverage of human rights: why limited? How could be broadened? What are the challenges you face/anticipate facing in the future? What support could the databases benefit from? And from whom? What are the plans for further development of the databases? For what would they like the databases to be used in particular? By whom? And how? Specific tailored questions were asked where deemed necessary.

2. The databases: What they collect, incident organisation, management and statistical information

Some terminological variations are evident in what each reviewed database collects. One collects 'incidents' (AI Incident Tracker); one collects 'incidents' and 'issues' (AIID), one collects incidents and/or controversies driven by and relating to artificial intelligence, algorithms and automation (AIAAIC Repository), and one collected 'helpful and harmful' AI use cases (Where in the World Is AI). Two databases explicitly include definitions of what they collect.

The *AIID* defines an incident as "an alleged harm or near harm event to people, property, or the environment where an AI system is implicated" (AIID, 2023) and 'issues' as "an alleged harm by an AI system that has yet to occur or be detected" (McGregor et al., 2022). The *AIAAIC Repository* currently classifies entries as 'system' (A technology programme, project, or product and its governance), 'incident' ("A sudden known or unknown event (or 'trigger') that becomes public and which takes the form of, or can lead to a disruption, loss, emergency, or crisis"), 'issue' ("Concerns publicly raised about the nature and/or potential impacts of a System, but without evidence of a public incident or recognised harms") or 'data' ("A public or proprietary dataset/database that has been shown to be inaccurate, unreliable, biased, overly intrusive, etc., and/or that results in issues or incident(s) directly or indirectly associated with the AI, algorithmic, or automation system(s) that draw(s) on it") (AIAAIC, 2023c).

The *AIAAIC Repository* interview clarified that its goal was to provide impartial data to allow people to draw their own conclusion; it did not take a view on what users do with the data, nor seeks to build trust. We note that the databases may reflect the biases of people adding data, but the intent seems to be to make them comprehensive, accurate, fair and balanced as possible and ensure good quality. The *AI Incident Tracker* does not define an incident. Speaking with the founder of this repository clarified that an incident refers to an "identifiable event or series of events", hence, this is not to be confused with a report on a general concern related to AI, such as bias or surveillance. *Where in the World Is AI* collected 'helpful and harmful' AI use cases, but neither was defined, though it elaborated that many of these cases might fall into the grey area of helpful or harmful and that different cultural perspectives, for example, might label cases differently (AI Global, 2020).

All databases organise entries differently. The *AIID* is searchable and features two taxonomies: an AI Harms Taxonomy (in its second edition) developed for the *AIID* by Georgetown University's Center for Security and Emerging Technology (CSET) and a Goals, Methods and Failures (GMF) taxonomy. The *AIAAIC Repository* (2023c) classifies entries as system, incident, issue, or data, by release, when it occurred, country, sector, operator, developer, system name, technology (type/application), purpose, media trigger, risks (e.g. accuracy/reliability, anthropomorphism, bias/discrimination, dual/multi-use, employment, environment, governance, human/civil rights, legal, mis/disinformation, privacy, safety, security, surveillance, transparency) and harms (actual negative impacts). *Where in the World Is AI* categorised by issue, year, domain, location, and whether helpful or harmful. The *AI Incident Tracker* organises entries by date of publication of the article mentioned.

All databases are open to submissions from the public via online forms, though the databases have different approaches to identify submitters. For the *AIAAIC Repository,* an email ID is required for submission. Premium membership (provides free access to hidden data in the Google sheet version of the AIAAIC Repository, including external and internal impacts and the ability to comment on and make minor editorial updates to repository entries) requires the provision of certain information. The *AI Incident Tracker* makes it possible to contribute by sending a pull request on GitHub or filing an issue (contacting the list manager) (AI Incident Tracker, 2023). The *AIID* is the only database that allows a submission to be made anonymously. *Where in the World Is AI* required a Google sign-in.

The databases indicate that submitted incidents are reviewed or verified before inclusion. The *AIID* developed an editor's guide on reviewing submitted incidents for inclusion in the database (AIID 2023). The *AI Incident Tracker* (2023) states that "a maintainer" may ask the contributor to edit his/her "Pull Request" before it is included, due to spelling error or non-compliance with the guidelines or Code of Conduct. *Where in the World Is AI* did not indicate any process for verifying incidents, but our interview revealed that this was done manually and instructions were to ensure that these came from a credible source (e.g. well-known publication) – this had to be approved by the Managing Editor. The *AIAAIC Repository* is edited and managed by the Managing Editor and contributors (https://www.aiaaic.org/aiaaic-repository/governance); incidents and controversies are also submitted by researchers, NGOs and others via social media or through the AIAAIC Repository incident report form. The *AIAAIC Repository* sets out a six-step process for consideration and processing of entries (detection, assessment, classification, summarisation, approval by the Managing Editor and publication) (2023d). Tools and techniques used to identify and collect incidents are Google Alerts, RSS and subscriptions to high quality newsletters and websites and the AIAAIC Repository incident report form. The incident/issue needs to have been covered by high-quality mainstream, trusted media sources (e.g. Reuters, Financial Times) not just the technical media and criteria for inclusion are fundamental. Where things are not clear/in assessment for inclusion, it is put on the *pending* page and after assessment added to main page.

None of the databases outline an explicit detailed process to contest an incident entered in the database. The *AIAAIC Repository* states it "is committed to handling complaints in a fair and transparent manner". An email is provided. Complaints are assessed by the Managing Editor and decisions are published on the *AIAAIC* website (2023d). It appears that the *AI Incident Tracker* stopped adding references to the list of incidents in June 2021. *Where in the World Is AI* indicated it was updated weekly.

The number of incidents reported vary. The *AIAAIC Repository,* in November 2023, lists almost 1,200 entries from 8,100+ reports. The *AIID* has 2,872 reports pertaining to 583 incidents and 215 issues. *Where in the World Is AI* had 415 entries when checked in February–October 2022 and the *AI Incident Tracker* had 251 incidents listed. The types of incidents included in each database also vary. Categorisation of types varies from detailed (specific) to casual and/or uncategorised *(Where in the World Is AI).* The *AI Incident Tracker* included issues related to discrimination, privacy, security, social polarisation, organisational culture in technology companies, etc. The *AIID* taxonomies break

down incidents by sector of deployment, harm distribution and AI tool, among others. The *AIAAIC Repository* covered a wide variety of issues, including, e.g. accuracy, bias, confidentiality, dual use, ethics, fraud, governance, hype, intellectual property, privacy and transparency (2023f).

With regards to the number of types of incidents covered, the databases varied when analysed in 2022. The *AIAAIC Repository* had featured accuracy/reliability – 307 times, privacy – 361, bias/discrimination – 206, safety – 178 times when analysed. *Where in the World Is AI* had 323 Harmful cases and 97 helpful cases.[6] *The AI Incident Tracker* did not specify types. The *AIID* is no longer searchable by 'type' of incident.

All the databases studied are initiatives from the past six years. Where specified,[7] their dates of creation range from 2018 *(Where in the World Is AI*[8]*)* to 2019 *(AIAAIC Repository, AIID).* One database was created by a non-profit organisation *(Where in the World Is AI),* and one *(AIID)* by an individual as part of a project supported by the Partnership on AI (PAI), a non-profit partnership of academic, civil society, industry and media organisations (2023a). Two were created by individuals or groups of individuals (*AIAAIC Repository* and the *AI Incident Tracker*). The *AIAAIC Repository* was founded and is managed by a managing editor, Charlie Pownall and contributors (AIAAIC, 2023g). The *AI Incident Tracker* is curated by J. Patrick Hall. *Where in the World Is AI* was aggregated by AI Global. The *AIID,* which was created by an individual (Dr Sean McGregor) with funding from the PAI, is currently managed by the Responsible AI Collaborative (led by Scott Allen Cambo).

The formats of the databases vary. The *AIAAIC Repository* (2023e) has a web interface supported by a Google Sheets repository. It also has a search engine, and the full repository can also be downloaded with Premium membership. The *AIID* is a searchable website; the full database can be downloaded. The *AI Incident Tracker* is a bulleted list of links on GitHub, part of other resources dedicated to machine learning interpretability. *Where in the World Is AI* also used Google Sheets but is no longer accessible.

Regarding funding, the *AIAAIC Repository* is a privately funded, free database working with a non-profit model. Initial funding for the *AIID* came through the PAI until 2021; as of 2022, its funding came through the Responsible AI Collaborative, which was financed by grants. Two databases did not provide any information about funding sources on their websites (*AI Incident Tracker* and *Where in the World Is AI*). *Where in the World Is AI* did not seem to have dedicated funding but was resourced and managed ad hoc internally. It was done in partnership with an individual who ran a blog on AI and updated it on a weekly basis. They had a design partner and in interview had stated they were invested in keeping it going for the future, as they did not expect a shortage in use cases.

Two of the databases are maintained by people in the United States (*AI Incident Tracker* – Washington, D.C.); the AIID has paid personnel in California, Rhode Island, Massachusetts, Uruguay and Greece, and volunteers all over the world. One has team members in the United States and Canada *(Where in the World Is AI)* and one in the U.K.

---

[6]   As checked on 21 February 2022.
[7]   We could not find it specified for the AI Incident Tracker.
[8]   Has cases dating back to 2018.

(*AIAAIC Repository*). The databases are publicly available (except for *Where in the World Is AI* which is now offline), although some had access restrictions in terms of functionality.

3. Main objective: Learn from harms and raise awareness

3.1. Stated aims: Learning from harms

The AI incident databases studied share a common aim: *raising awareness of current AI harms to prevent future harm*. These databases also have a similar primary audience: developers and managers of AI systems.

The *AIAAIC Repository* was created in June 2019 "to better understand the reputational and other risks" of AI, algorithms and automation and "to get a better handle on how to design, develop, deploy, and regulate them" (AIAAIC, 2023a). It currently (November 2023) states that "by collecting, dissecting, and surfacing incidents and issues from a non-technical, 'outside-in' perspective in an objective and balanced manner, the Repository enables users to identify, examine, and understand the nature, risks, and impacts of AI, algorithms, and automation" (AIAIC, 2023a). The *AIAAIC Repository* is reportedly used by researchers, academics, advocates, journalists, lawyers, policymakers, industry experts at universities, business schools, NGOs, non-profits, think tanks, media organisations, industry associations and businesses globally (AIAAIC, 2023b).

The *AI Incident Tracker* aims to provide a "blueprint for a more human-centred, lower-risk machine learning" (The AI Incident Tracker, 2023). It does not specify its target user, but its founder, Patrick Hall, states the main motivation in building this database is for developers and data scientists to be able to learn from past mistakes. The *AIID* presents itself as a "repository of problems experienced in the real world as a result of AI" that "can help AI researchers and developers mitigate or avoid repeated bad outcomes in the future" (PAI, 2023b). Sean McGregor, the creator of the *AIID,* developed the database "to enable AI incident avoidance and mitigation", noting the present lack of "collective memory of failing" (McGregor, 2020). *Where in the World Is AI* intended to map out "interesting examples where AI has been harmful and where it's been helpful" (2020).

For AI developers, deployers and providers who are the most obvious and the primary targeted audience, AI incident databases are a valuable resource of real-world examples. This knowledge can help reflect on, anticipate problems in their own systems and explore how harms are/could be addressed and mitigated. Such databases can help developers, deployers and providers find concrete examples and insights of problematic and risky AI use, what is/might be harmful, how systems have caused harms and had consequences, types of harms and help address concerns early on. It might also help anticipate what mitigation measures are required to be embedded into the system, at development, deployment or use stages. This would improve the transparency, accountability, safety and reliability of AI systems. Such databases are good teaching and mutual learning resources, and their use should be encouraged in AI/AI ethics and STEM (Science, Technology, Engineering and Math) curricula.

In summary, the main objective of these AI incident databases is to learn from risks and harms to avoid these in the future. As such, they offer a valuable resource support for the design, development and promotion of safe, transparent and accountable AI systems.

3.2. Growing the field's maturity by supporting prudence

The 'move fast and break things' approach has often accompanied the growth of the digital economy (Taplin, 2017) including AI and big data technologies. Cassie Kozyrkov, Head of Decision Intelligence at Google, stated it thus, "go for it, see what happens. Get it wrong, because you're going to have to do it over and over again. You're going to fail, fall down over and over again. Pick yourself up. Try again until finally, it works" (Google Cloud Tech, 2019). In too many cases, it is after the harm is done, that action is taken to prevent it. As McGregor put it: "Technology companies are famous for their penchant to move quickly without evaluating all potential bad outcomes" (McGregor, 2020). The specific technical characteristics of AI as a technology that is based on the trial-and-error method might have further contributed to the lack of prudence in deploying potentially harmful AI systems. AI incident databases are useful tools to help grow the field's maturity in supporting prudent decision-making.

Policy makers share the responsibility for the fast deployment of AI without always sufficiently accounting for potentially harmful outcomes early-on (though of late there is a greater recognition of the need to regulate AI). Globally and in the European Union, there has been a strong push to fund, increase and hasten the deployment of AI. The European Commission's White Paper on Artificial Intelligence (2020) illustrated this push from the policy-making side, although it also acknowledged the need for trustworthy technology.

Aristotle considered prudence *(phronesis)* a key value and attitude of ethical behaviour. This value appears to be lacking in the AI field – this is concerning not just because there is much greater aversion to risk but also because there is much more work on, and awareness of the risks and harms that might result from the improper application or use of AI technologies. The imprudence that can be observed in the AI field reveals some degree of immaturity of the field and a certain hype and fascination with AI that lowers the threshold of caution. There seems not to have been too many lessons learnt from other risk-heavy domains such as pharmaceuticals and air transportation. AI incident databases, therefore, constitute a powerful tool to support prudent decision-making.

3.3. Fostering accountability

The public and visible nature of such databases can help develop accountability, i.e. require AI developers, deployers and users to be accountable for harms their systems might generate. By flagging a potential harm, the databases are a resource that could put pressure on a specific actor to act to prevent and mitigate harm. The "move fast and break things" approach of Silicon Valley is no longer acceptable (Taneja, 2019).

The more these databases are known, publicised and used, the greater this pressure will be. As AI technology matures, preventing potential harms to individuals and the society is critical. Furthermore, as the regulation of AI in the EU increasingly moves toward a risk-based approach as framed in the AI Act proposal, these databases constitute a useful resource to identify and assess these risks (European Commission, 2021).

The databases also present a means to encourage public accountability by raising public awareness of AI harms. The public often lacks awareness of the risks and harms of AI systems. Considering the complexity and opacity of AI systems, lay people may fail to recognise that an issue exists or might have originated there. This makes it difficult for the public to ask for accountability. As Charlie Pownall, Founder of the *AIAAIC Repository* put it, opacity "erodes confidence and trust". These databases can help further promote transparency,[9] although it might prove costly for the organisation developing, deploying, or using the AI that has led to an incident. By improving public awareness, the databases foster accountability. They make it possible to "share power" by sharing knowledge beyond the experts. As the databases refer to media sources that are for a non-technical audience, this further contributes to this.

4. How other stakeholders can benefit from the databases

In addition to AI developers and managers, a broad range of other stakeholders can benefit from accessing and using AI incident databases, including policymakers and regulators, researchers, non-governmental organisations and civil society.

For policymakers and regulators, AI incident databases provide examples of risks and harms within AI ecosystems, which can inform policy and decision-making related to the governance of AI. Policymakers and regulators working on AI policy need to better understand AI and be sensitised to the risks and harms of this technology (a view supported in the *AIAAIC* interview). In that sense, the databases could be a great tool to inform their work. Information could also be used for training and/or awareness-raising. Databases could become a legal requirement for the industry, incentivised, for example, as due diligence requirements. If required, funding support should be forthcoming from policymakers and regulators, helping address sustainability challenges and enabling the databases to be maintained.

For example, an AI incident database could support implementation of the proposed European AI Act, which may require certain AI providers to inform national competent authorities about serious incidents or malfunctioning that constitute a breach of fundamental rights (European Commission, 2021). AI incident databases would complement these formal efforts in a more non-formal and non-legalistic manner, especially in capturing a wider variety of instances that might not fall within the scope of the regulation and even occur beyond the European Union.

---

[9] Jobin et al. 2019 analysed 84 AI ethics guidelines and found that the principle of transparency is the one that occurs the most often: it is mentioned in 73 out of the 84 guidelines studied.

Policy makers and regulators also play a key role for the promotion of accountability of AI systems. As Naurin (2006) states: "Being held accountable involves 'paying the price' for one's actions. Accountability therefore involves something more than just having one's actions publicly exposed. In case of misconduct some kind of sanction should be imposed on the actor." We encourage the use of these databases by decision makers to ensure accountability for harms caused by AI.

Researchers on ethics, SSH (social sciences and humanities), and human rights can utilise the information collected in the databases to better understand the AI ecosystem and identify incidents (or incident profiles) for further studies. While soft law initiatives (e.g. ethics guidelines) have emerged since 2018 to respond to potential harms caused by AI (Ethical ML, 2023) there is increasing discontent with this form of response (Rességuier & Rodrigues, 2020; Wagner, 2018) and the growing recognition that human rights should play a role in addressing these harms (Siemaszko et al., 2020). These databases could further support the work of the human rights community by framing the incidents, where relevant, as human rights violations. Human rights provide well-established lenses to consider harms to individuals or communities. Framing of harms as human rights violations make them more directly relevant to the human rights community to address these harms. It could also be a useful lens for developers themselves to better understand the human rights risks of their systems to individuals, society and the environment.

While we encourage the human rights community to use these resources, we acknowledge it is not the primary audience of the databases. Our interviews flagged concerns that using the human rights lens might undermine the databases' use and may scare away the intended audience from supporting and proactively using them. We do not fully agree. As outlined in the UNESCO Recommendation on the Ethics of Artificial Intelligence (2021) respect for human rights is essential throughout the life cycle of AI systems and should be considered in the collation of data for these databases. AI is fundamentally multi-stakeholder – it relies on diverse stakeholders for its development, deployment and use and has a wide impact on individuals and society, including human rights. Therefore, the use of the human rights lens or the engagement of the human rights community with these databases would be highly valuable. The databases are a resource that could target the diversity of stakeholders in the development, deployment and use of AI.

Civil society organisations (CSOs) and non-governmental organisations (NGOs), including those that work in human rights, may also find AI incident databases valuable to identify incidents for further investigation and follow-up. As exemplified in the case of the *AIAAIC,* it has been endorsed by researchers, academics, advocates, journalists, policymakers, industry experts, NGOs and non-profits, think tanks, media organisations, industry associations and businesses (AIAAIC, 2023b). A privacy-focused CSO, for example, may find specific examples of AI-related privacy violations to support a campaign for policy reform or increased public awareness. CSOs working with specific demographic groups, including vulnerable groups, may use such databases to discover incidents with similar characteristics to help build a network of similar victims and leverage the strength of a larger collective to advocate for harm mitigation. A legal services organisation or law school legal clinic could also use an AI incident database to identify specific victims for individual or collective (i.e. class action) representation.

5. Conclusion

As AI is pervasive and cuts across sectors and disciplines, it requires oversight from people with a diversity of experience and expertise. An AI incident database can be an especially useful resource that is worth sharing and using across a diversity of communities. Such databases have a valuable role to play in helping AI design, development and use to learn from what's gone wrong. They complement other governance measures aimed at mitigating potential harms. To this effect, we support their broader uptake and recommend:

1. Ensuring funding and sustainability by linking the databases to other AI and sectoral initiatives and making them available for wider use (e.g. teaching, research and policy decision-making).
2. Enhancing/adopting measures for better vetting, incident reporting accuracy and improving discoverability of incidents. This would boost the reputational value of the databases and improve trust and accountability.
3. Analysing the types of incidents occurring and re-occurring over a period to help understand key AI risks and harms, whether these are being addressed, and the remaining challenges.
4. Ensuring accountability, by addressing understandability, accuracy, verifiability, comparability, timeliness, and completeness and resourcing issues (Ullah et al., 2021).

AI incident databases can be useful to raise awareness and provide a resource base for diverse stakeholders. It is important to make these databases more accessible to audiences non-familiar with harms of AI, especially those that are not given the opportunity to provide their inputs during the research and development phases. This will promote a civil society that can ask for redress in case of harm caused by AI and require accountability.

Acknowledgements and disclaimer

References

AI Global (2020). Responsible AI Design Assistant. Online: https://bit.ly/3uyI6MK

AIAAIC (2023a, November 4). About the Repository. Online: https://bit.ly/47tl5th

AIAAIC (2023b, August 3). Repository users and endorsements. Online: https://bit.ly/49XER1N

AIAAIC (2023c, November 22). Classifications and Definitions. Online: https://bit.ly/3sTmk5T

AIAAIC (2023d, November 22). Governance. Online: https://www.aiaaic.org/aiaaic-repository/governance

AIAAIC (2023e, November 22). Repository. Online: https://www.aiaaic.org/aiaaic-repository

AIAAIC (2023f, November 22). Research and Advocacy Citations and Mentions. Online: https://www.aiaaic.org/about-aiaaic/research-citations-mentions

AIAAIC (2023g, November 22). Team. Online: https://www.aiaaic.org/about-aiaaic/team

AIID (2023). Editors Guide. Online: https://incidentdatabase.ai/editors-guide

Ethical ML (2023). Awesome Artificial Intelligence Guidelines. Online: https://github.com/EthicalML/awesome-artificial-intelligence-guidelines

European Commission (2020). White Paper on Artificial Intelligence – A European Approach to Excellence and Trust. Online: https://bit.ly/40V34BA

European Commission (2021). Proposal for a Regulation of the European Parliament and of the Council Laying Down Harmonised Rules on Artificial Intelligence (Artificial Intelligence Act) and Amending Certain Union Legislative Acts. Online: https://bit.ly/47V0oX7

Google Cloud Tech (2019). Making Friends with Machine Learning: Regression. Online: https://www.youtube.com/watch?v=WNvOtwP_yf4

Hofer, C. W. & Schendel, D. (1978). Strategy Formulation: Analytical Concepts. West Publishing Company.

Jansen, P., Brey, P., Fox, A., Maas, J., Hillas, H., Wagner, N., Smith, P, Ouoch, I., Lamers, L., van Gein, H., Resseguier, A., Rodrigues, R., Wright, D. & Douglas, D. (2020). SIENNA D4.4: Ethical Analysis of AI and Robotics Technologies. Zenodo. Online: https://doi.org/10.5281/zenodo.4068082

Jobin, A., Ienca, M. & Vayena, E. (2019, September 1). The Global Landscape of AI Ethics Guidelines. Nature Machine Intelligence, 1(9), 389–399. Online: https://doi.org/10.1038/s42256-019-0088-2

McGregor, S. (2020). Preventing Repeated Real World AI Failures by Cataloging Incidents: The AI Incident Database. Online: https://bit.ly/47AWliL

McGregor, S., Paeth, K. & Lam, K. (2022). Indexing AI Risks with Incidents, Issues, and Variants. Online: https://doi.org/10.48550/arXiv.2211.10384

Naurin, D. (2006). Transparency, Publicity, Accountability – The Missing Links. Swiss Political Science Review, 12(3), 90–98. Online: http://hdl.handle.net/1814/661

Partnership on AI (2023a). About. Online: https://partnershiponai.org/about/

Partnership on AI (2023b). AI Incident Database. Online: https://partnershiponai.org/aiincidentdatabase/

Pickton, D. W. & Wright, S. (1998). What's SWOT in Strategic Analysis? Strategic Change, 7(2), 101–109. Online: https://doi.org/10.1002/(SICI)1099-1697(199803/04)7:2%3C101::AID-JSC332%3E3.0.CO;2-6

Rességuier, A. & Rodrigues, R. (2020). AI Ethics Should Not Remain Toothless! A Call to Bring Back the Teeth of Ethics. Big Data & Society, 7(2). Online: https://doi.org/10.1177/2053951720942541

Rodrigues, R., Siemaszko, K. & Warso, Z. (2019). SIENNA D4.2: Analysis of the Legal and Human Rights Requirements for AI and Robotics In and Outside the EU (Version V2.0). Zenodo. Online: https://doi.org/10.5281/zenodo.4066812

Siemaszko, K. Rodrigues, R. & Slokenberga, S. (2020). SIENNA D5.6: Recommendations for the Enhancement of the Existing Legal Frameworks for Genomics, Human Enhancement, and AI and Robotics (V2.0). Zenodo. Online: https://doi.org/10.5281/zenodo.4121082

Taplin, J. (2017). Move Fast and Break Things: How Facebook, Google, and Amazon Have Cornered Culture and What It Means for All of Us. Pan Macmillan.

Taneja, H. (2019). The Era of "Move Fast and Break Things" Is Over. Harvard Business Review. Online: https://bit.ly/47NUqag

The AI Incident Tracker (2023). Awesome Machine Learning Interpretability. Online: https://bit.ly/40ZnVUA

Ullah, M. H., Hazelton, J. & Nelson, P. F. (2021). Can Databases Facilitate Accountability? The Case of Australian Mercury Accounting Via the National Pollutant Inventory. Accounting. Auditing & Accountability Journal, 34(1), 164–193. Online: https://doi.org/10.1108/AAAJ-11-2017-3232

UNESCO (2021). Recommendation on the Ethics of Artificial Intelligence. Paris. Online: https://bit.ly/3sIvGS9

Wagner, B. (2018). Ethics as an Escape from Regulation: From Ethics-Washing to Ethics-Shopping. In E. Bayamlioğlu, I. Baraliuc, L. Janssens & M. Hildebrandt (Eds.), Being Profiled: Cogitas Ergo Sum: 10 Years of 'Profiling the European Citizen' (pp. 84–89). Amsterdam University Press.

Where in the World Is AI (2021). Global Map. Online: https://map.ai-global.org