

A Taxonomic System for Failure Cause Analysis of Open Source AI Incidents

Nikiforos Pittaras^{1,*}, Sean McGregor¹

¹Responsible AI Collaborative

Abstract

While certain industrial sectors (e.g., aviation) have a long history of mandatory incident reporting complete with analytical findings, the practice of artificial intelligence (AI) safety benefits from no such mandate and thus analyses must be performed on publicly known “open source” AI incidents. Although the exact causes of AI incidents are seldom known by outsiders, this work demonstrates how to apply expert knowledge on the population of incidents in the AI Incident Database (AIID) to infer the potential and likely technical causative factors that contribute to reported failures and harms. We present early work on a taxonomic system that covers a cascade of interrelated incident factors, from system goals (nearly always known) to methods / technologies (knowable in many cases) and technical failure causes (subject to expert analysis) of the implicated systems. We pair this ontology structure with a comprehensive classification workflow that leverages expert knowledge and community feedback, resulting in taxonomic annotations grounded by incident data and human expertise.

Keywords

AI Incidents, Failure Analysis, AI Safety

1. Introduction

The pursuit of safe AI is a critical problem of the 21st century which, unless dealt with, harbors dangers likely to define and shape the trajectory and prosperity of the human species [1]. AI Safety research efforts have made progress towards combating AI-induced x-risk in a variety of fronts; these include value loading and refinement by human preferences [2], investigations on inner misalignment manifestations [3], interpretability-oriented analysis on deep network models [4], as well as conceptual work on frameworks and potential harms of superintelligent systems [5, 6]. In contrast, there has been limited work in a different direction: exploiting publicly available data that may provide useful insights in the function, composition, alignment, deployment and application failure causes of real-world AI systems. A prominent example of such data streams are AI incidents, i.e. public articles and reports that describe harms and failures of deployed AI systems in the wild. This study presents early work, proposing the analysis and annotation of AI incidents via the development of a taxonomic system that captures Goals, Methods / Technologies and Failure Causes of a technical nature (abbreviated as GMF), that stem from content-based information in incident reports in conjunction with technical knowledge and expertise in the AI Safety and Alignment community. In the rest of this paper we provide the motivation and contributions of preliminary work conducted to develop

the taxonomic system, including the proposed structure, annotation workflow and development procedures. Finally, we discuss the expected impact of GMF towards the annotation, discovery and analysis of AI systems and their failure causes as they manifest in existing incidents, along with the potential of resulting datasets for future data-driven AI Safety and Alignment research efforts.

2. Related Work


While research efforts have been focused on analyzing and categorizing failures [7, 8, 9, 10, 11] and major components of AI systems [12, 13], there has been limited work on holistically linking multiple aspects of AI systems into a single, interrelated taxonomic model. Additionally, these efforts are often separated from real-world failures and events, resulting in the under-utilization of any research outputs and taxonomic insights via, e.g., shared datasets for operational purposes. Finally, when applying existing taxonomies to real world systems, they can rarely be applied adequately with incomplete or uncertain information, which is the case for the vast majority of AI-related incidents now reported.


Research efforts in the AI Safety community have organized different aspects of the current landscape to advance AI Safety, encompassing organizations, datasets, failures, and harms [14, 15, 7, 10]. Notable works include the CSET taxonomy [16], which provides a broad set of information for AI incident annotation, ranging from high-level descriptions of harm types, severity estimates and distribution among different affected groups, to limited sets of high-level AI functions, causative factors (e.g. robustness failure) and a variety of additional metadata

SafeAI 2023, February 13-14, 2023, Washington D.C., US @ AAAI-23

 pittarasnikif@gmail.com (N. Pittaras);

 smcgregor@seanbmcgregor.com (S. McGregor)

 © 2023 Copyright for this paper by its authors. Use permitted under Creative Commons License

 Attribution 4.0 International (CC BY 4.0).

 CEUR Workshop Proceedings (CEUR-WS.org)

(e.g. system owner, deployment sector, user expertise etc.). While rich, the ontology aims for broad descriptions of incidents and their real-world impact, rather than focusing on the technical aspects and failure causes of the AI system involved. Other efforts focus on the compilation of chronological lists of sourced AI Failures [7], paired with a comparison between AI Safety and cybersecurity viewpoints and concerns around producing safe and reliable systems. The provided incident pool is however limited (< 20), with no analysis performed on the specific incidents cited.

Some research efforts offer causal factors for AI misbehavior [8], paired with harm descriptions and linked to exemplary real-world manifestations, along with general directions for mitigation and handling per category. Although informative, the technical grounding of causal factors is limited, with the majority of the discussion being delivered at a conceptual, high-level framework and no additional views (e.g. categorizations of system objectives or implementation descriptions) being considered. Further recent research focuses on specific domains and models, such as the taxonomy of language model risks [17], which categorizes real-world harms, risks and hazards of large language generators; while the study is comprehensive and provides technical mitigation approaches, the proposed taxonomic grouping focuses on a very high-level view of harmful effects and is restricted to the language domain, without aiming to explore causative factors that potentially generalize across systems and tasks of different modalities.

In light of these efforts, this work describes ongoing work on the proposed Goals, Methods and Failure Causes taxonomic system, which encapsulates three interrelated ontologies: 1) high-level AI system goals, 2) methods and technologies used for system implementation, and 3) technical failure causes that result in misbehavior in the applied system; this structure utilizes inter-taxonomy relationships towards identifying technical failure factors in AI systems. The resource is presented in the context of characterizing real-world AI incidents provided by the AI Incident Database (AIID) [18].

The list of contributions associated with this preliminary body of work includes:

- General, holistic, interrelated taxonomies: We propose three interconnected views of a broad and diverse set of AI incidents, enabling multifaceted data interpretation, analysis and retrieval, yielding various pattern matching avenues towards diagnosing and mitigating harms by parties of different interests, domains and expertise.
- Focus on technical causal factors: While most existing works strive to categorize failures and harms, we focus on AI attributes, approaches, limitations and issues of a technical nature, that

themselves lead to observed harms.

- Grounded to real data: In contrast to conceptual / theoretical analysis or narrow experimental testbeds prone to measurement issues [19], our study centers around annotating real-world AI incidents provided by AIID, which describe real-world observations of AI systems and failures.
- Data-driven, fine-grained and explainable: we propose a workflow that links annotations to specific incident text spans and metadata, enhancing explainability, transparency and validation and enabling further data-driven safety research.

Given this setting, we move on to provide a description of the structure, workflows, development procedures and expected impact of GMF in the sections that follow.

3. The GMF taxonomic system

Here we describe the proposed taxonomic system structure, annotation workflow, development procedure and projected impact of GMF, as reflected by the initial body of research work and early findings.

3.1. Taxonomies

Three interrelated taxonomies are included in GMF, providing a well-rounded view and different discovery avenues for AI systems involved in incidents.

First, the *AI System Goals* taxonomy addresses what the deployed AI system was trying to achieve; it encapsulates high-level goals, objectives and primary use cases pursued in the real world, such as “Translation” or “Face Recognition”. This information enables use case-driven incident discovery and facilitates retrieval of historical AI methods and failure causes of similar systems by interested groups, such as AI developers and safety engineers.

Second, *AI Methods and Technologies* contains information on how the system is built, including learning models, representation construction approaches and other methodological, engineering and implementation-related features, e.g. “Transformer Neural Network”, “Collaborative Filtering”. Incident filtering by elements of this taxonomy provides popularity and utilization trends of different technologies, while highlighting the distributions of harms, historical failures and technical pitfalls associated with specific implementation approaches.

Finally, the *AI Failure Causes* taxonomy consists of technical reasons that lead to the emergence of real-world harms during AI deployment. It involves systemic failures of technical nature which may manifest in system design, engineering, specifications and construction procedure, such as “Concept Drift” or “Distributional Bias”, which are potential causal factors to the observed undesirable AI behavior. Failure cause-based retrieval from

GMF-annotated instances should reveal indicative use cases, methods and technologies where specific failures materialize, enabling data collection for pattern extraction and causal analysis, as well as provide groundwork for research into mitigation efforts.

At the present stage, we reserve enforcing a predefined internal taxonomy structure once GMF development progresses further (see Section 3.3), opting for a flat label organization instead. Taxonomy elements represent general categories (e.g. “Clustering” instead of “K-Means Clustering”) in order to establish high term applicability and are composed of a short descriptive name (e.g. 1-3 words) and a concise description (e.g., up to 30 words) that communicates exact semantic content to annotators.

3.2. Annotation Workflow

Given that AI incident report contents may contain limited amounts of technical information for efficient annotation with GMF, we propose an annotation workflow that additionally leverages taxonomic relationships, historical incident records and technical knowledge in the AI, ML and Safety community.

The primary source of information available to annotators for arriving at relevant GMF incident classifications is AI incident contents, i.e. text available in reports that describe the incident. High-level AI system information (e.g., system objective / domain / use-case) required to apply an *AI System Goals* classification should be readily obtainable from incident contents, with no additional investigation and limited speculation. Second, annotators have access to previously annotated incident collections in AIID; labelled incident retrieval provides informative priors via historical incidents, which are useful in the annotation of incidents that include limited technical information, especially for the *AI Methods and Technologies* and *AI Failure Causes* taxonomies, as described in section 3.1. Finally, knowledgeable individuals in AI, Machine Learning (ML) and AI Safety as well as other valuable disciplines (e.g. machine ethics, human-computer interaction, cybersecurity, philosophy, mathematics, etc.) can draw on training, experience and analytic skills to provide diverse, critical insight on methods, technologies and failure causes that produce events and harms reported in the incident, when relevant technical information is not directly available but can be inferred. Notably, such insights can be extracted via crowdsourcing means, when coverage of large amounts of incident data is prioritized, using careful moderation and curation to account for labelling noise that may result from crowdsourced annotators with different backgrounds and levels of expertise.

Given these information streams, the proposed annotation workflow for an incident I is the following:

1. Incident annotation with the *AI System Goals* taxonomy; the appropriate classification G should

be easily identifiable by incident contents alone.

2. Retrieve similar incidents H from AIID that share a goal classification $L = G$, providing a use case-based context of methods and technologies M_H
3. Extract relevant technical community knowledge on the incident and all available context ($I \cup H$)
4. Arrive at likely annotations M from the AI Methods and Technologies taxonomy, $M|I, G, M_H, T$, i.e. by considering given incident contents, current goal classification, historical incident method classifications and relevant technical community knowledge.
5. Update the historical incident pool H by considering a lookup parameter of the method / technology annotation, $L = M$.
6. Arrive at likely annotations F from the AI Failure Causes taxonomy, i.e. with respect to current incident contents and classification (I, G, M), as well as historical failure annotations and expert knowledge from the community (F_H, T).

An illustration of the proposed annotation workflow is presented in Figure 1.

3.3. Development

We now present the proposed GMF development process, which adopts an iterative, bottom-up approach, working from batches of individual incidents from AIID [18] to annotate incidents and populate taxonomy contents. In this setting, the proposed taxonomic system solves an epistemological problem for AIID through an open tagging design that interrelates both high-certainty and speculative classifications. Tags are open in the sense that there is no pre-defined set of goals, methods, or failures so these can develop through time and reflect consensus in the AI Safety research community.

Given an AI incident, the general proposed workflow for an annotator is as follows:

1. Read through incident contents and identify salient passages, e.g. text mentioning technical terms, system use cases, specifications and harms.
2. Pair salient passages with free discussion comments, providing explanation, rationale, additional information and linkage to external resources, if deemed necessary.
3. Create / modify a GMF taxonomy classification C , following the taxonomy definitions, workflow protocols and information sources established in Sections 3.1 and 3.2. Afterwards, link one or more salient passages to the classification, such that they provide reasonable justification and grounding to selecting C .

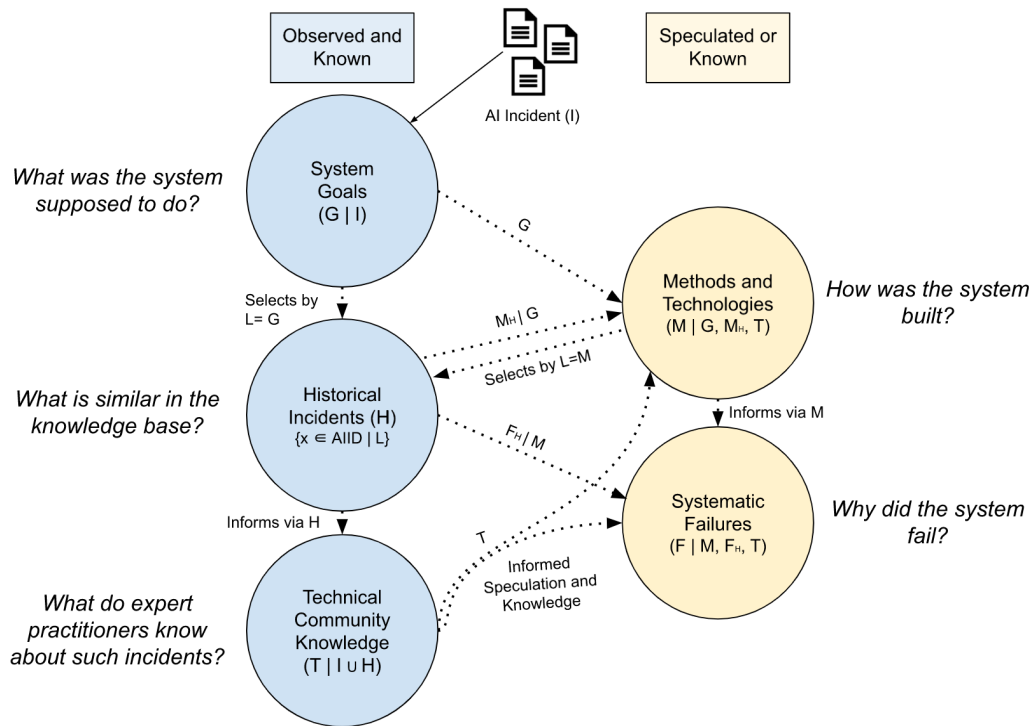


Figure 1: The proposed annotation workflow of AIID instances with GMF, illustrating relationships between taxonomic elements and related sources of information (AIID data, technical AI / ML / Safety community). System Goals are generally known for all incidents in the AIID. The system goal G then determines which candidate similar incidents H are pulled from the AIID and presented to the technical community. Then, the combination of previous methods and technologies M_H from the AIID conditioned on the Goals and Technical Community Knowledge T helps determine candidate Methods and Technologies (M) whenever they are not stated in available incident reports. Finally, the combination of historical incidents for similar system goals and methods help determine the candidate Systemic Failures (F) with the input of the Technical Community.

4. Pair each classification with a confidence modifier, i.e. “known” or “potential”, conveying near-certain or above-average degrees of certainty that C is relevant to the incident. The totality of accumulated information (selected snippet, content terms / technical information / ambiguity, historical incidents, technical background and community knowledge) should determine the most relevant modifier according to the annotator.
5. Pair C with free discussion comments, which supply adequate reasoning for why the classification and confidence modifier are fitting / relevant to the linked passages, given the totality of accumulated information available to the annotator. This discussion should be able to reveal the decision-making process, rationale and evidence used, serving as documentation to other interested third parties (annotators, evaluators, editors etc.). Notably, such comments are especially important in annotations where non-trivial

amounts of experience, intellectual work and information gathering were marshalled to produce C , which are classification cases that will be characterized by higher uncertainty, on average.

An illustration of the overall GMF structure, workflow and development protocol for classification of a real-world incident ¹ is available in Figure 2, providing details on information flow, resource utilization and labelling operations for a concrete annotation example in an end-to-end fashion.

The proposed annotation configuration provides a number of notable desired features to the taxonomy development process. First, improved transparency and validation of classifications is achieved by grounding annotations with relevant passages and free discussion comments. Passages list supporting input evidence, while comments may elaborate on rationale, sources, reasoning and intuition across different annotators, levels of

¹<https://incidentdatabase.ai/cite/72>

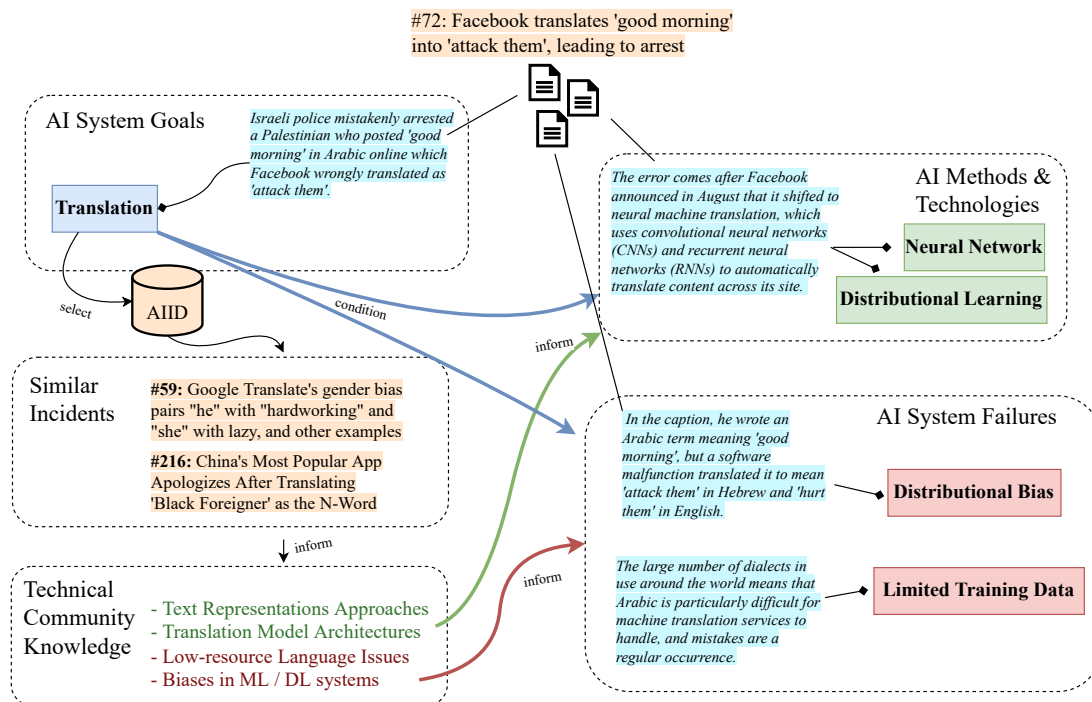


Figure 2: An end-to-end example of the proposed GMF classification workflow for the real-world AIID incident #72. Incidents titles and report passages are illustrated with an orange and cyan highlight respectively, while colored rectangles represent GMF classifications and diamond edges showcase their grounding. The “Translation” AI System Goal is directly attainable from relevant passages of the incident in question (top), which enables similar incident retrieval from AIID (mid-left); these are utilized to obtain technical knowledge about translation methods, models, implementations (top-right) as well as known limitations and issues, from experts in the AI / ML / Safety community (bottom-left). This feedback in turn informs and introduces priors to AI Methods and Technologies and AI Failure Causes (bottom-right), in conjunction with useful incident contents. Comments and some interaction details are omitted from the visualization for the sake of readability.

expertise and points in time.

Second, the linkage stated above comes with built-in potential for data-driven automation, enabling the development of Machine Learning workflows to enhance, automate and accelerate future manual annotation efforts via, e.g., classification recommendations, salient passage extraction, keyword extraction, etc.

Finally, the proposed workflow can rapidly generate annotated dataset versions of variable levels of classification grounding and quality control; given the large cost of manual technical annotation, quality can be improved by iteratively refining existing versions of ground truth. For example, an initial release might include noisy crowd-sourced annotations (e.g. with limited classification grounding and a small annotator pool). Subsequent versions that undergo multiple steps of correction, verification and validation should produce datasets eligible for research-grade utilization (e.g. via imposing minimum levels / thresholds for classification grounding statistics,

annotator pool size and agreement, etc.).

The proposed setting regards uncertainty as an explicit element in the taxonomy, correlating it with consensus; in that sense, incidents with high annotator disagreement stem from their corresponding AI systems and their failure causes being opaque, ambiguous and generally hard to diagnose. We plan to utilize this feature in conjunction with improving transparency by providing incoming annotators access to labels assigned by people with different expertise and points of view, in order to encourage a holistic consideration of incidents. Additionally, we aim to improve classification reliability and quality by initially iterating the exposure of the annotation interface to selected groups of limited but increasing size, validating new users via moderation (e.g. spam detection, verification of relevance / precision of classification groundings) and using community-based trustworthiness mechanics (e.g. in a wiki-style knowledge management), etc.

We have currently conducted a series of classification

exercises on a representative set of incidents, to explore the information and context available to back various taxonomy designs, with the findings from this investigation being consolidated in the interrelated structure and workflows herein introduced. A summary of the status of GMF ontology development is provided in Table 1: at this stage, we have iterated the taxonomic system over the first $\approx 12\%$ percent of the database with expert annotators (i.e. PhD-level ML / AI Safety researchers and engineers). Annotation confidence breakdown is approximately evenly split, highlighting the crucial uncertainty aspect of the task due to the limited amount of technical information available in incident reports. This fact, along with cases of identification of multicausal failures (e.g. the system may exhibit “Distributional Bias” that cannot be correctly diagnosed or fixed due to “Lack of Transparency” of the underlying model, etc.), complex model architectures (e.g. multimodal / ensemble approaches) and multitask systems, leads to the generation of more than one annotation per taxonomy for each incident.

Regarding scalability and throughput, current average per-user time requirements reach 20 – 30 minutes for the development workflow (i.e. both incident annotation and generation of new taxonomy labels); incidents for which relevant taxonomy elements are already defined can be annotated in $\approx 3 - 5$ minutes, depending on the amount and size of reports citing the incident. Annotation time is thus expected to decline heavily once a core majority of relevant taxonomy elements is established and can be utilized by users. Additionally, planned tools, UI aids and utilities are set to introduce further speed-ups by streamlining and automating the annotation workflow.

Given the above, we are now preparing to comprehensively apply the taxonomy across the database on an ongoing basis. At the same time, we are working on developing, designing and integrating toolsets, interfaces and procedures to support accelerated, crowd-sourced operations in the near future; namely, the proposed developmental workflow will be supported by a dedicated user interface designed to facilitate annotation at a faster pace and support automation functionalities (e.g. auto-completion, recommendation, highlighting, etc.) for improved user experience, reduced boilerplate and higher throughput. Apart from annotation, this infrastructure will support information extraction, incident retrieval and navigation for the inspection of possible failure modes, causes and risks in existing AI incidents, as well as exploratory analysis on deployment descriptions of new systems, by experts and laypeople alike.

3.4. Expected Impact

The proposed taxonomic system will complement existing ontologies by providing applicability to a broad set of real-world AI incidents (and textual content of AI mis-

Statistic	Value
Number of annotated AIID incidents	41
AIID coverage %	12.40
<i>Known</i> annotations %	48.68
<i>Potential</i> annotations %	51.32
Goals per incident (avg.)	1.20
Methods / Technologies per incident (avg.)	2.43
Technical Failure Causes per incident (avg.)	3.78

Table 1

Status of initial GMF development, with respect to AI incident instances in AIID.

behavior in general), with a strong focus on grounded technical descriptions of failure causes. We expect that the set of annotation workflows, tools and resources will enable rapid annotation, curation and verification of incidents with GMF labels, leverage the support of AI communities and experts alike, and generate a wide variety of useful data-driven applications for further automation and development of related research.

At this early stage of GMF construction, we center our offering on real-world AI harm events and posit that the proposed work can help researchers identify, analyse and contextualize open problems in AI Safety and related research domains (e.g. Human-Centered AI, Cybersecurity, Machine Ethics, etc), aid policymakers in effectively regulating the most damaging systems, and empower corporations to identify when systems under development are subject to previously experienced failure modes.

4. Conclusion and Future Work

In this work, we presented preliminary work on the GMF taxonomic system, a proposed set of taxonomies that capture AI System Goals, AI Methods and Technologies and AI Failure Causes. We outlined recommendations for taxonomy composition, annotation workflow, development procedure and future development plans, in the context of incident reports in AIID, listing the rationale, benefits and expected impact of the resulting resources to the research, policy and industry sectors.

At present, we have applied classifications from the perspective of machine learning research engineers, under an iterative taxonomy development procedure on an initial batch of incident data. Given that incidents are typically multi-faceted in their causes and safe system design calls for a variety of organizational processes in addition to design accommodations, in the future we plan to augment the classification set with additional perspectives (e.g., scientists and researchers with expertise in philosophy, ethics, AI governance and various human factors). Additionally, we plan on publishing the taxon-

omy within the platform interface of AIID² and invite external participation in the incident classification and taxonomy development process in a crowdsourced setting. Further planned steps include enriching, structuring and interlinking GMF as annotation coverage of AIID increases, via, e.g., consolidating taxonomy elements into a structured hierarchy emerging from collections of independently assessed annotations, or connecting GMF taxonomies with other AI Safety-related resources and ontologies (e.g. CSET [16], AVID³, etc.). We believe that the broad inclusion of AI Safety-related research communities will provide invaluable feedback, constructive revisions and significant expansions to existing annotations and ontology contents, lending critical technical insights about the holistic description of AI system failures, towards building up mitigation measures of future AI-induced harms in the real world.

Acknowledgements

In addition to participating in the 2018 discussions that launched the AIID, Richard Mallah gave valuable feedback on the taxonomy. The AIID is an effort by many people and organizations operating under the banner of the Responsible AI Collaborative. It is through their collective efforts that the ontological perspectives presented above have meaning and real world importance.

References

- [1] E. Yudkowsky, Agi ruin: A list of lethalties, <https://intelligence.org/2022/06/10/agi-ruin/>, 2022. Accessed October 28th, 2022.
- [2] L. Ouyang, J. Wu, X. Jiang, D. Almeida, C. L. Wainwright, P. Mishkin, C. Zhang, S. Agarwal, K. Slama, A. Ray, et al., Training language models to follow instructions with human feedback, arXiv preprint arXiv:2203.02155 (2022).
- [3] K. Cobbe, O. Klimov, C. Hesse, T. Kim, J. Schulman, Quantifying generalization in reinforcement learning, in: International Conference on Machine Learning, PMLR, 2019, pp. 1282–1289.
- [4] N. Elhage, N. Nanda, C. Olsson, T. Henighan, N. Joseph, B. Mann, A. Askell, Y. Bai, A. Chen, T. Conerly, N. DasSarma, D. Drain, D. Ganguli, Z. Hatfield-Dodds, D. Hernandez, A. Jones, J. Kernion, L. Lovitt, K. Ndousse, D. Amodei, T. Brown, J. Clark, J. Kaplan, S. McCandlish, C. Olah, A mathematical framework for transformer circuits, Transformer Circuits Thread (2021). <https://transformer-circuits.pub/2021/framework/index.html>.
- [5] K. E. Drexler, Reframing superintelligence, The Future of Humanity Institute, The University of Oxford, Oxford, UK (2019).
- [6] R. V. Yampolskiy, Agi control theory, in: International Conference on Artificial General Intelligence, Springer, 2021, pp. 316–326.
- [7] R. V. Yampolskiy, M. Spellchecker, Artificial intelligence safety and cybersecurity: A timeline of AI failures, arXiv preprint arXiv:1610.07997 (2016).
- [8] D. N. Banerjee, S. S. Chanda, AI failures: A review of underlying issues, arXiv preprint arXiv:2008.04073 (2020).
- [9] R. V. Yampolskiy, Predicting future AI failures from historic examples, foresight (2018).
- [10] R. Mallah, The landscape of AI safety and beneficence research. input for brainstorming at beneficial AI 2017, Beneficial AI 2017 (2017).
- [11] D. Manheim, S. Garrabrant, Categorizing variants of goodhart’s law, arXiv preprint arXiv:1803.04585 (2018).
- [12] V. Honavar, Artificial intelligence: An overview, Artificial Intelligence Research Laboratory (2006) 1–14.
- [13] S. Samoil, M. L. Cobo, E. Gomez, G. De Prato, F. Martinez-Plumed, B. Delipetrev, AI watch. defining artificial intelligence. Towards an operational definition and taxonomy of artificial intelligence (2020).
- [14] J. Porway, A taxonomy for AI / data for good, <https://data.org/news/a-taxonomy-for-ai-data-for-good/>, 2022. Accessed October 28th, 2022.
- [15] M. Zheng, H. Robbins, Z. Chai, P. Thapa, T. Moore, Cybersecurity research datasets: taxonomy and empirical analysis, in: 11th USENIX Workshop on Cyber Security Experimentation and Test (CSET 18), 2018.
- [16] S. McGregor, The first taxonomy of AI incidents, <https://incidentdatabase.ai/blog/the-first-taxonomy-of-ai-incidents>, 2021. Accessed October 28th, 2022.
- [17] L. Weidinger, J. Uesato, M. Rauh, C. Griffin, P.-S. Huang, J. Mellor, A. Glaese, M. Cheng, B. Balle, A. Kasirzadeh, et al., Taxonomy of risks posed by language models, in: 2022 ACM Conference on Fairness, Accountability, and Transparency, 2022, pp. 214–229.
- [18] S. McGregor, Preventing repeated real world AI failures by cataloging incidents: The AI incident database, in: Proceedings of the AAAI Conference on Artificial Intelligence, 17, 2021, pp. 15458–15463.
- [19] J. Wentworth, You are not measuring what you think you are measuring, <https://www.lesswrong.com/posts/9kNxxhKWvixtKW5anS/you-are-not-measuring-what-you-think-you-are-measuring>, 2022. Accessed October 28th, 2022.

²<https://incidentdatabase.ai>

³<https://www.avidml.org>