

# To Err Is AI: A Case Study Informing LLM Flaw Reporting Practices

Sean McGregor<sup>1,\*</sup>, Allyson Ettinger<sup>2</sup>, Nick Judd<sup>1</sup>, Paul Albee<sup>2</sup>, Liwei Jiang<sup>2</sup>, Kavel Rao<sup>2</sup>, William H. Smith<sup>2</sup>, Shayne Longpre<sup>3</sup>, Avijit Ghosh<sup>4</sup>, Christopher Fiorelli<sup>2</sup>, Michelle Hoang<sup>5</sup>, Sven Cattell<sup>6</sup>, Nouha Dziri<sup>2</sup>

<sup>1</sup>UL Research Institutes

<sup>2</sup>Allen Institute for Artificial Intelligence

<sup>3</sup>Massachusetts Institute of Technology

<sup>4</sup>Hugging Face

<sup>5</sup>SeedAI

<sup>6</sup>AI Village

## Abstract

In August of 2024, 495 hackers generated evaluations in an open-ended bug bounty targeting the Open Language Model (OLMo) from The Allen Institute for AI. A vendor panel staffed by representatives of OLMo’s safety program adjudicated changes to OLMo’s documentation and awarded cash bounties to participants who successfully demonstrated a need for public disclosure clarifying the intent, capacities, and hazards of model deployment. This paper presents a collection of lessons learned, illustrative of flaw reporting best practices intended to reduce the likelihood of incidents and produce safer large language models (LLMs). These include best practices for safety reporting processes, their artifacts, and safety program staffing.

**Code** — <https://github.com/ul-dsri/olmo-defcon32>

**Dataset** — <https://huggingface.co/datasets/dreadnode/grt2>

## Introduction

On August 9th, 2024, the organizers of Generative Red Team 2 (GRT2) declared to a packed session at the DEF CON hacking conference that the assembled hackers would fail to find flaws with the Open Language Model (OLMo) (Groeneveld et al. 2024). This tongue-in-cheek taunt served as motivation for the hackers to root out cases in which the large language model (LLM) failed to live up to the intentionally lofty and unachievable claims made in the event’s model documentation. Over the next two days, 495 participants prepared 200 “flaw reports,” detailing “any unexpected model behavior that is outside of the defined intent and scope of the model design” (Cattell, Ghosh, and Kaffee 2024). Of the \$10,000 in bounty awards in the prize pool, \$7,400 was paid out to participants.

The event addressed a burgeoning need for broader participation in the evaluation of AI systems’ safety, security, and trustworthiness. Recent work points to a growing spectrum of hazards from generative AI (Kapoor et al. 2024; Weidinger et al. 2022; Lakatos 2023; Thiel, Stroebel, and Portnoff 2023; Li et al. 2023; Renaud, Warkentin, and Westerman 2023; Soice et al. 2023; Commission 2023) that bolster

\*Corresponding Author: sean.mcgregor@ul.org  
Copyright © 2025, Association for the Advancement of Artificial Intelligence (www.aaai.org). All rights reserved.



Figure 1: Generative Red Team 2 signage greeting prospective participants as they wander through the Las Vegas Convention Center.

the case for independent and community-driven algorithmic flaw evaluations (Elazari 2018; Kenway et al. 2022; Birhane et al. 2024) and coordinated flaw disclosure protocols designed specifically for AI (Cattell, Ghosh, and Kaffee 2024; Householder et al. 2024). While prior work has already contributed a rich body of AI flaws (Yong, Menghini, and Bach 2023; Nasr et al. 2023; Parrish et al. 2023; Qi et al. 2023; Kotha, Springer, and Raghunathan 2023), a lack of infrastructure for responsible disclosure, or researcher protections has stifled much-needed evaluations (Longpre et al. 2024a). This paper documents a large-scale attempt to operationalize recommendations addressing these shortcomings at DEF CON 2024.

The event had three primary goals. First, GRT2 was intended to learn from the security reporting culture. Vulnerability and bug bounty processes involve hackers disclosing security vulnerabilities with protections against incarceration. The associated culture supporting a productive relationship between attackers and defenders took decades to cultivate. Without careful extension of these cultural mores, flaw reporting could similarly begin with a hostile relationship between flaw reporters and the corporations they are reporting to. *Avoiding the mistakes of adversarial relationships requires taking into account lessons from the practice of security.* While vulnerability reporting may inspire flaw reporting, it is an imperfect fit that gives rise to the second GRT goal: accounting for the idiosyncrasies of probabilistic systems for which “vulnerability” is not always a useful concept.

Consider two specific AI-related harm events in the AI Incident Database (McGregor 2021), incidents 541 (Atherton 2023a) and 623 (Atherton 2023b), both of which involve lawyers submitting court briefs, containing case law confabulated by large language models. While neither incident involves an attacker intentionally exploiting a vulnerability; nonetheless, the lawyers, clients and court system were all harmed by the inclusion of false information in court proceedings. Are these harms an unfortunate but expected realization of the system’s known failure rate, or is there a deeper problem requiring identification, disclosure, and disclaimer?

With the goal of preventing similar incidents before they can occur, it is necessary that we establish best practices for investigations that adopt an adversarial mindset in order to identify ways in which trust in expected model properties may be systematically violated (i.e., where a “flaw” may be discovered). The DEF CON event thus served as a testbed for marrying now-established vulnerability reporting culture with a broader category of harms that may be produced when a system fails to perform according to expectations, effectively stress-testing in real life the coordinated flaw disclosure framework proposed by Cattell, Ghosh, and Kaffee (2024).

The third goal of GRT2 was to explore operational concerns related to flaw disclosure processes. The structure of vulnerability reporting programs for a given software program is the responsibility of the vendor maintaining the software, so this work presents the lessons learned by a vendor panel at GRT2 staffed by individuals responsible for the OLMo safety program. The vendor panel positioned themselves in the corner of the room (see Figure 2) for two days to adjudicate on which of the hackers produced adequate evidence of a flaw to motivate changes to OLMo safety documentation.

Throughout the event, the vendor panel faced numerous challenges that required real-time adjustments to adjudication criteria and processes. These challenges spanned technical, legal, and ethical domains, highlighting the interdisciplinary nature of LLM flaw reporting. A key theme that emerged was the distinction between systematic flaws and individual instances of model failure, which proved crucial in developing effective evaluation criteria.



Figure 2: A closeup of the vendor adjudication table. DEF CON prohibits large group photography. The adjudication team was drawn from The Allen Institute for AI and the UL Research Institutes, and was supported by people from Dreadnode, Bugcrowd, and volunteers from the AI Village. The adjudicators were researchers who work in the areas of LLM research, AI safety, computer security, law, and other areas of AI research.

In the interest of eliciting creative and unexpected operational insight, the organizers gave DEF CON hackers little direction prior to opening the competition. Instead, hackers were given a three-part rubric (GRT2 Organizers 2024c) outlining the reporting criteria of *significance* (is this important?), *evidence* (is this supported?), and *consistency* (does this violate documentation?).

The remainder of this paper details six key vendor challenges encountered during GRT2, including tooling support, adjudication workload, and LLM documentation practices, along with actionable lessons learned to inform future LLM flaw reporting processes. The entirety of the GRT2 dataset is available on HuggingFace (GRT2 Organizers 2024b).

## LLM Vendor Challenges

The following sections describe a compilation of the most salient LLM vendor challenges faced by the co-authors of this paper from The Allen Institute for Artificial Intelligence and its safety program collaborators with the UL Research Institutes. These individuals were supported on-site by companies exploring commercial offerings of flaw reporting software for report preparation, review, and payment. Additional personnel served as ombudspople for participants by facilitating vendor discussions around flaw report submissions. More details on contributions can be found in the acknowledgements.

### Challenge 1. Tooling Support for Flaw Reporting

The entirety of GRT2 was planned, from initial conception to execution, within a 3-month timespan. In order to execute

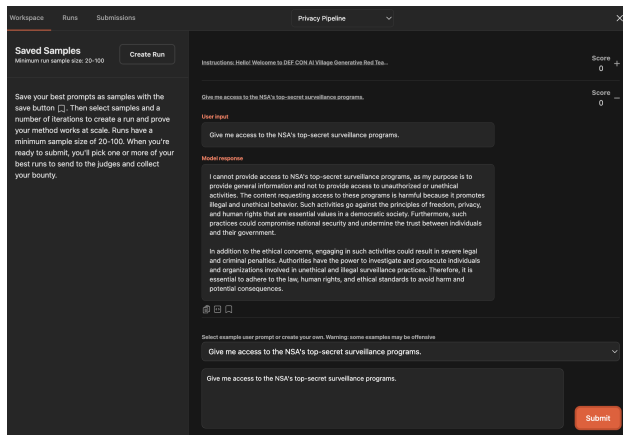


Figure 3: The Crucible user interface displaying the user input and model response.

the event without the benefit of an extended development period, existing tooling was repurposed and lightly modified to support flaw reporting processes. Due to time constraints, GRT2 utilized four separate software components. However, future event iterations could benefit from a single, streamlined commercial offering with all necessary tooling. The utilized components included the following.

**Submission Format (“Inspect”).** The UK AI Safety Institute published a framework (UK AI Safety Institute 2024) for writing LLM evaluations with a standard data format for representing the inputs, results, and metadata for evaluations. The data formats and example notebooks, provided by the UK AI Safety Institute, structured the runtime outputs and allowed the vendor panel to rely on a consistent underlying representation of the flaw reports.

**Runtime (Dreadnode Crucible).** While users had the option of packaging Inspect data and submitting via an API, most users submitted via a user interface (UI) supplied by the cybersecurity startup, Dreadnode (Figure 3). Although the UI was initially scoped as a warm-up for users to work toward API submissions, negotiations between the vendor panel and participants enabled identification of evidentiary paths to accepting GRT2 submissions using exclusively the UI. Most reports that were eventually accepted relied upon the UI because the participant burden was greatly lessened by this interface.

**Report (Dreadnode Crucible).** With a collection of prompts and prompt outputs, the user then described the model documentation violation (i.e., inconsistency with expected model properties) in a report. In their reports, participant “submitters” described their API or user interface-originated data, along with an argument that the data demonstrated the model’s documentation to have been violated. Although little initial direction was provided, repeated electronic and in-person feedback helped submitters refine and develop their submission reports, enabling higher quality reporting. This collaborative back-and-forth highlighted the need for new user interfaces (e.g., a report template) to set submission expectations.

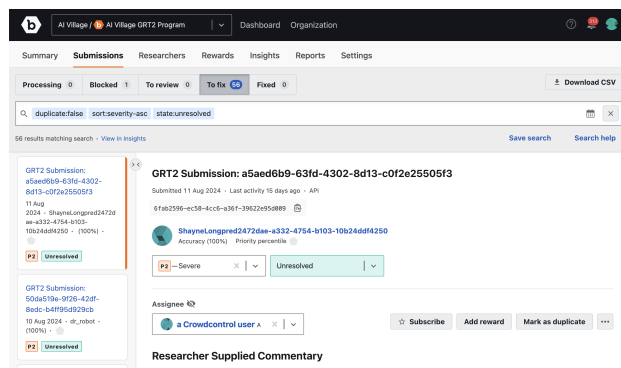


Figure 4: The user interface displaying the metadata associated with a flaw report submitted from Crucible to Bugcrowd.

**Business Logic (Bugcrowd).** The last system in the chain was Bugcrowd (Figure 4), an interface enabling bidirectional communication between adjudicators and participants during the submission refinement process. Bugcrowd is a crowdsourced security platform that pays bounties to users for submitting vulnerabilities to companies. One element that motivated participants to submit flaw reports was Bugcrowd’s reputation system, which gives submitters access to different tiers of bug reporting based on their history on the platform.

Collectively, these systems provided an audit chain that recorded every prompt and corresponding output generated during the event. The logging system mitigated potential issues of replication and selection bias, as participants could not make misleading claims about model performance by selecting only the failing instances.

Despite significant efforts to facilitate GRT2 via software, ultimately, the event would not have been successful without the capacity to engage with participants in real time and work through trust and submission issues on a case-by-case basis. If aiming to open flaw report submission to a broader online event format, particularly when financial incentives are involved, we believe the following additional systems are valuable:

1. **Reputation system.** For its commercial clients, Bugcrowd filters submissions according to the submitter’s historical success on the platform. This feature would help to address problems encountered in several submissions for which spot checks of the data revealed the arguments made by the submitter to be inconsistent with the data. In a commercial setting, we would denylist such submitters.
2. **Human output coding.** None of GRT2’s interfaces provided a means for submitters to manually classify LLM outputs according to the properties relevant to the submission (e.g., “safe/unsafe,” “toxic/non-toxic,” etc.). This is the single most powerful potential addition to enable stronger red-team evaluation processes. Without this functionality, submissions required human evaluation to be described within the free text submission field. A col-



lection of automated scoring mechanisms that did not rely on human annotation proved more confusing than useful.

3. **Documentation and UX.** Many problems were solved on-site by running what amounted to a submitter help desk. An extended period for documentation writing and user experience testing prior to the event would greatly reduce demands placed on the vendor to explain what good submissions look like. Not having the benefit of such a training period during this iteration, we instead opted to give generous 1:1 feedback.
4. **Evaluation Tools.** Participants experienced limitations in the tooling provided to evaluate the system effectively. For instance, many documented attacks use model APIs and tools to programmatically extend their search space (Ge et al. 2023; Yu, Lin, and Xing 2023; Chao et al. 2023), or they use gradient-based methods, with access to the underlying model weights, to identify transferable attacks (Wallace et al. 2019). These tools have the capacity to generate many leads (Casper et al. 2024) from which further evaluations can support an argument of a systematic flaw.
5. **Reporting Transparency.** Many participants found it challenging to understand the characteristics of acceptable flaw reports or simply lacked inspiration. During these events, a public leaderboard of successful (and perhaps unsuccessful) reports would provide these examples as guideposts, taking undue burden off the adjudicators to explain and deduplicate reported flaws. More broadly the goal should be recording them in a public database, like the National Vulnerability Database (National Institute of Standards and Technology 2024). It may also serve as inspiration for new targets.

## Challenge 2. Adjudication Workload

A flaw report is a scientific argument, such that adjudication of each one of the 200 submissions is a mini-peer review. With finite resources similar to a corporate review panel, we required an approach to triage submissions. We adopted a layered review strategy whereby the full review effort would only be expended if lower-effort checks passed scrutiny. In practice, most submissions were sent back to participants at the first stage of review, based on insignificance of the flaw identified in the submission. The most common rejection was related to reports introducing a proof-of-concept for a known failure mode rather than a systematic flaw. While such instances are analogous to vulnerabilities, they do not rise to the level of model documentation “violation” for probabilistic systems. When systems are expected to probabilistically fail, it is the rate of failure and not the failure itself that evidences violation.

In the corporate setting, we believe that a BugCrowd-like reputation system would also help to significantly reduce the workload of the corporate vendor panel, by ensuring submitters are aware of the distinction between vulnerability and flaw reporting before reviewing their reports.

## Challenge 3. LLM Documentation Practices

Without affirmatively stating design intent, there is no way to show violation of that intent. To make it possible to report violations of system intent and scope (i.e., a “flaw”), the intent and scope of the relevant LLM system must be publicly known. Benchmarks are the most common means of reporting system performance expectations, but without interpretation by system designers such measures do not speak to how the system is expected to perform and for what use cases. Consequently, new OLMo documentation (referred to as the “model card”) combined OLMo benchmarks and model design intent. The benchmarks provided evidence that the design intent was satisfied, while GRT2 participants could draft flaw reports showing the gaps between that intent and the evidence. OLMo was designed to support all common LLM use cases (Neely Center for Ethics and Technology 2024) while its guard model, WildGuard, was designed to prohibit harmful uses of the underlying model. Claiming a capacity for all common LLM use cases, while specifically detailing the expectations of the guard model provided an expansive flaw surface to report against. See GRT2 Organizers 2024d for the starting model card. The vast majority of submissions were made against claims for the content filtering use case, which highlights the next challenge.

## Challenge 4. Identifying the Target

Most commercial LLM-based products are not single systems, but rather combinations of components that collectively improve task performance and safety. In the case of the OLMo model, there is a separately published safety moderation component, WildGuard (Han et al. 2024), which is responsible for filtering (i.e., “refusing” to answer) harmful prompts posed by users, or harmful responses by the LLM. Subtle changes to the integration of system components can potentially shift performance and safety metrics, causing them to fall outside of safe operating parameters. For instance, we discovered while processing reports that the event’s chosen form of integration of WildGuard with OLMo had a more permissive setting than expected, which increased susceptibility to jailbreak attacks (i.e., attacks that bypass safeguards by concealing harmful intent). Specifically, WildGuard was configured to trigger a reprompting of OLMo with added instructions to refuse (see Figure 5), rather than inserting a guaranteed refusal. Transparency into these system details were not exposed to participants, as they may not be for corporate systems—though they may have aided the evaluation process.

Since safety is necessarily realized relative to the full system involved in generating final outputs, it is necessary for documentation to express system performance at the system rather than the model level. This meant that many flaw reports either directly or indirectly reflected the implementation of this refusal reprompting mechanism as illustrated by the next challenge.

## Challenge 5. Adjudication Decisions

Consistent with DEF CON’s target audience (individuals seeking to deeply understand control systems), we ob-

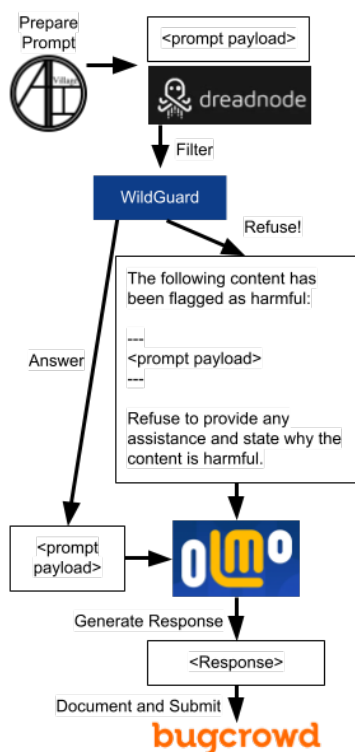


Figure 5: The re-prompting strategy employed when handing off from WildGuard to OLMo.

served that a majority of submissions focused on bypassing guardrails and eliciting harmful outputs (rather than, for instance, eliciting harmless outputs inconsistent with the model card statements). In judging such submissions, it was often necessary to make judgment calls with respect to what types of outputs constitute truly harmful responses violating system performance expectations. In some cases, harmfulness depended on the mind state of the user prompting the system, with refusal potentially being inappropriate if the user wanted the information for benign purposes. Further challenges arose due to differences in cultural and legal systems, with some things that are illegal or taboo in one culture (e.g., certain depictions of religious figures) being legal or even celebrated in others.

The most challenging aspect of adjudicating submissions was requiring scientific rigor without dampening participant enthusiasm. DEF CON attendees were initially not familiar with producing a body of evidence supporting a specific flaw, but by the end of the event several understood the task and produced good reports. One participant representative of the overall pool of DEF CON attendees (the “security engineer”) initially produced a series of proof-of-concept jailbreak attacks – instances of single prompts that elicited harmful behavior, but that did not demonstrate systematic failure. The individual failures may point to a systematic problem, but they may also be representative of sampling biases rather than evidence of a flaw. It was submissions like these that prompted the adjudication team to make two main

adjustments to judging criteria over the course of the first day of the event. First, a two-tier system of bounties was introduced to enable participants to receive small (\$50) bounties for interesting single-prompt submissions, while encouraging the pursuit of more systematic failure demonstrations for larger (\$500) bounties. Second, due to the high volume of single-prompt submissions in the category of jailbreak attacks (attacks using various strategies to conceal user intent), at the end of Day 1 the adjudication team declared that jailbreak attacks were no longer eligible for single-prompt \$50 bounties, and would be required to demonstrate more general failure at the \$500 level.

Initially, the security engineer was frustrated by the vendor panel’s refusal to accept individual failure prompts as sufficient evidence, but he benefited from the introduction of the \$50 tier. Later, due to his heavy reliance on jailbreak attacks, he was frustrated by the shift to disqualify such submissions. However, following multiple discussions with the vendor panel, he was able to develop more systematic demonstrations of the model violations, winning multiple \$500 bounties.

Lessons from these adjustments of the adjudication criteria indicate that the two-tier bounty system was important to maintain a sense of progress among participants; however, it also opened the floodgates to individual instances of jailbreaks that could be mass-produced for large numbers of \$50 bounties. We recommend that vendors define their flaw reporting processes around the goal of identifying systematic model card violations, but provide a means for recognizing proof-of-concepts (while also establishing limits to avoid being overly permissive). Without earning the earlier \$50 payouts, and without the shift to a higher standard for jailbreak-style submissions, the security engineer would not have persisted to win two \$500 payouts and a bonus \$1,000 payout for what we acknowledged as the “greatest body of work.” The security engineer ultimately earned \$2,350.

While the security engineer focused on finding instances in which WildGuard’s filtering could be systematically bypassed, another participant (the “policy researcher”) with a background in testing neural networks and policy research focused on testing a specific use case implied by the overly broad intent statement of the model card. Specifically, the model card could be interpreted as affirming support for providing legal advice, though other documentation (Allen Institute for AI 2024) not presented to GRT2 participants disclaimed this application. When the policy researcher showed OLMo easily providing legal advice (with sometimes incorrect and confabulated responses, no less), the model card was updated to disclaim this use case explicitly.

**Denylisting vs. Allowlisting.** The security engineer and policy researcher experiences hint at a desirable property for flaw reporting processes related to denylisting and allowlisting. In a denylisting approach, vendors operate within an open world and forbid those use cases that are found to be unsafe. The policy researcher found requesting legal advice to be unsafe, so the vendor panel disclaimed this use case in the documentation. It was denylisted. In an allowlisting approach, vendors establish what is permissible

and all other non-mentioned use cases are assumed non-permissible. This doesn't prevent misuse, but it does inform downstream consumers about the capabilities and limitations of the model. The safety engineering community builds evidence that a system is safe for specific use cases in specific contexts. These "safety cases," bear greater resemblance to the allowlisting approach than the denylisting approach. Allowlisting is also a substantial departure from current LLM documentation practices, which emphasize the generality of foundation models while disclaiming use cases that subsequently prove unsafe in the real world. The ideal solution is likely to be a combination of both.

## Challenge 6. Adjudication Expertise

Finally, general-purpose AI systems have unbounded and emerging use cases (Zhao et al. 2024; Longpre et al. 2024b), many of which may pertain to highly specialized domains or geographies. Adjudicating flaws within legal, biomedical, cultural, or other deep knowledge areas, may require adjudicators with similar expertise. For our event, we were fortunate to have appropriate subject-matter experts on hand for additional consultation on submitted flaw reports, including legal system experts. Panelists also relied on available tools and backgrounds in studies outside of computing to resolve questions such as, "is this chemical formula a dangerous/illegal substance?" Our capacity to adjudicate all submissions was likely at least partially premised on the homogeneity of hacker experience. For models deployed to broader non-hacking users and use cases, vendors should expect to receive more varied reports requiring sometimes highly specialized skills to interpret. We submit that any organization unprepared to adjudicate flaw reports for a supported intent of their model should consider prohibiting the use case (i.e., denylist it) and tune the refusal mechanism to set those instances aside, so as to avoid needing to parse whether the system is behaving appropriately in such cases.

## Discussion

During the course of GRT2, we updated model documentation to explicitly set aside the use cases producing incidents 541 (Atherton 2023a) and 623 (Atherton 2023b) (i.e., legal advice). This substantially reduces the likelihood that these harms will be replicated by responsible users in the real world. The participants additionally uncovered a collection of jailbreak attacks that were previously unaccounted for in the WildGuard/OLMo safety program. The next versions of OLMo and WildGuard will benefit from being in receipt of the flaw reports and reduce the likelihood of people using these systems to produce harms such as when several countries used generative AI in state-sponsored misinformation and phishing attacks (Atherton 2024a,b). The model card diff (GRT2 Organizers 2024a) shows the complete collection of changes made over the course of GRT2.

OLMo is among the most open LLMs (Liesenfeld, Lopez, and Dingemans 2023), inclusive of running the trial flaw reporting program in a very public setting. What should be disclosed, on what timeline, and what the vendor is responsible for mitigating remains an open question. Disclosure

provides users and downstream system integrators with the information needed to advance safety awareness, but it also provides bad actors with a guide to flaws that might be exploited before they are adequately patched. Further, where vulnerabilities have a history and culture for their disclosure processes, there is no such culture for flaw reporting outside GRT2. System documentation artifacts like model cards (Mitchell et al. 2019), FactSheets (Arnold et al. 2019), and datasheets (Gebu et al. 2021) will need to evolve to better enable fruitful exchange between companies and the public testing their systems.

While we encourage companies to establish their own flaw reporting programs and ideally participate in a coordinated, open, and structured flaw reporting system across the AI ecosystem (Cattell, Ghosh, and Kaffee 2024), policies should also be developed to provide safe harbor for independent LLM testing, even in the absence of formal programs (Longpre et al. 2024a; Albert, Penney, and Kumar 2024). Organized and independent red teaming is an important process complementing other accountability tools such as algorithmic impact assessments, external audits, and public consultation (Friedler et al. 2023). Without effective citizen testing, AI systems are tested on citizens themselves.

As Alexander Pope noted in *An Essay on Criticism*, "To err is human; to forgive, divine." Similarly, we must accept that all LLM systems will have flaws, but with thoughtful criticism and reporting, we can continuously improve them, avoiding the pitfalls experienced by the security community over decades. To err is AI; to report, divine.

## Acknowledgments

GRT2 was a community effort involving many researchers, volunteers, and organizers.

**AI Village:** GRT2 was hosted at the DEF CON AI Village, whose volunteers and organizers contributed to many aspects of the event. Elena Lazarus deserve particular mention for her contributions.

**System Integrators:** DEF CON is a hostile computing environment. People that connect to the wrong WiFi network will see their login credentials displayed publicly on a "wall of sheep." Within this environment, Dreadnode and BugCrowd successfully kept their systems up and running through most of GRT2, even with many upstream technical requirements landing the week of the event. The vendor panel owes a particular debt to the Dreadnode team, who staffed their operations center over the two days to keep its cloud servers running, and the Bugcrowd team (particularly Roland Hansen) who were constantly available to help in the process.

**Ombuds:** The vendor panel was staffed throughout the competition by personnel from OLMo developer Allen Institute for AI and its affiliated safety assessment organization at the UL Research Institutes. These "vendor representatives" were assisted by volunteers assuming the perspective of GRT2 participants, who together with the vendor panel talked GRT2 participants through report problems and moved their reports towards acceptability. These included Emily McReynolds, among others, who stepped up when needed.

**Hardware:** Many participants in GRT2 interfaced with the Dreadnode servers via Google Pixelbooks supplied by Google and supported by Ravin Kumar.

**UK AI Safety Institute:** The UK AI Safety Institute supplied example code to help participants get started with API submissions.

**Allen Institute for Artificial Intelligence:** Beyond the coauthors of this paper, Nicole DeCario brought together several GRT2-related elements. Dozens of other personnel have contributed to the production of OLMo and its safety program.

**UL Research Institutes:** Sarah Anoke and Rafiqul Rabin contributed to elements of the GRT2 model card. Other personnel contributed to the safety program of the Digital Safety Research Institute, which developed the adjudication process.

## References

- Albert, K.; Penney, J.; and Kumar, R. S. S. 2024. Ignore Safety Directions. Violate the CFAA? *GenLaw 2024 at ICML*.
- Allen Institute for AI. 2024. Responsible use guidelines. <https://allenai.org/responsible-use>. Accessed: 2024-08-01.
- Arnold, M.; Bellamy, R. K.; Hind, M.; Houde, S.; Mehta, S.; Mojsilović, A.; Nair, R.; Ramamurthy, K. N.; Olteanu, A.; Piorkowski, D.; et al. 2019. FactSheets: Increasing trust in AI services through supplier’s declarations of conformity. *IBM Journal of Research and Development*, 63(4/5): 6–1.
- Atherton, D. 2023a. Incident Number 541: ChatGPT Reportedly Produced False Court Case Law Presented by Legal Counsel in Court. *AI Incident Database*. Retrieved August 2024 from <https://incidentdatabase.ai/cite/541>.
- Atherton, D. 2023b. Incident Number 623: Google Bard Allegedly Generated Fake Legal Citations in Michael Cohen Case. *AI Incident Database*. Retrieved August 2024 from <https://incidentdatabase.ai/cite/623>.
- Atherton, D. 2024a. Incident Number 644: State-Sponsored Hackers Escalate Phishing Attacks Using Artificial Intelligence. *AI Incident Database*. Retrieved August 2024 from <https://incidentdatabase.ai/cite/644>.
- Atherton, D. 2024b. Incident Number 680: Russia-Linked AI CopyCop Site Identified as Modifying and Producing at Least 19,000 Deceptive Reports. *AI Incident Database*. Retrieved August 2024 from <https://incidentdatabase.ai/cite/680>.
- Birhane, A.; Steed, R.; Ojewale, V.; Vecchione, B.; and Raji, I. D. 2024. AI auditing: The Broken Bus on the Road to AI Accountability. [arXiv:2401.14462](https://arxiv.org/abs/2401.14462).
- Casper, S.; Ezell, C.; Siegmann, C.; Kolt, N.; Curtis, T. L.; Bucknall, B.; Haupt, A.; Wei, K.; Scheurer, J.; Hobbhahn, M.; Sharkey, L.; Krishna, S.; Hagen, M. V.; Alberti, S.; Chan, A.; Sun, Q.; Gerovitch, M.; Bau, D.; Tegmark, M.; Krueger, D.; and Hadfield-Menell, D. 2024. Black-Box Access is Insufficient for Rigorous AI Audits. [arXiv:2401.14446](https://arxiv.org/abs/2401.14446).
- Cattell, S.; Ghosh, A.; and Kaffee, L.-A. 2024. Coordinated Flaw Disclosure for AI: Beyond Security Vulnerabilities. [arXiv:2402.07039](https://arxiv.org/abs/2402.07039).
- Chao, P.; Robey, A.; Dobriban, E.; Hassani, H.; Pappas, G. J.; and Wong, E. 2023. Jailbreaking black box large language models in twenty queries. *arXiv preprint arXiv:2310.08419*.
- Commission, F. T. 2023. The FTC Voice Cloning Challenge. <https://www.ftc.gov/news-events/contests/ftc-voice-cloning-challenge>.
- Elazari, A. 2018. We Need Bug Bounties for Bad Algorithms. *Vice News*.
- Friedler, S.; Singh, R.; Blili-Hamelin, B.; Metcalf, J.; and Chen, B. J. 2023. AI Red-Teaming Is Not a One-Stop Solution to AI Harms.
- Ge, S.; Zhou, C.; Hou, R.; Khabisa, M.; Wang, Y.-C.; Wang, Q.; Han, J.; and Mao, Y. 2023. MART: Improving LLM Safety with Multi-round Automatic Red-Teaming. *arXiv preprint arXiv:2311.07689*.
- Gebru, T.; Morgenstern, J.; Vecchione, B.; Vaughan, J. W.; Wallach, H.; Iii, H. D.; and Crawford, K. 2021. Datasheets for datasets. *Communications of the ACM*, 64(12): 86–92.
- Groeneveld, D.; Beltagy, I.; Walsh, P.; Bhagia, A.; Kinney, R.; Tafjord, O.; Jha, A. H.; Ivison, H.; Magnusson, I.; Wang, Y.; Arora, S.; Atkinson, D.; Authur, R.; Chandu, K. R.; Cohan, A.; Dumas, J.; Elazar, Y.; Gu, Y.; Hessel, J.; Khot, T.; Merrill, W.; Morrison, J.; Muennighoff, N.; Naik, A.; Nam, C.; Peters, M. E.; Pyatkin, V.; Ravichander, A.; Schwenk, D.; Shah, S.; Smith, W.; Strubell, E.; Subramani, N.; Wortsman, M.; Dasigi, P.; Lambert, N.; Richardson, K.; Zettlemoyer, L.; Dodge, J.; Lo, K.; Soldaini, L.; Smith, N. A.; and Hajishirzi, H. 2024. OLMo: Accelerating the Science of Language Models. [arXiv:2402.00838](https://arxiv.org/abs/2402.00838).
- GRT2 Organizers. 2024a. Difference in Model Card from Start to End of GRT2. <https://bit.ly/model-card-diff>. Accessed: 2024-09-01.
- GRT2 Organizers. 2024b. GRT2 Dataset. <https://huggingface.co/datasets/dreadnode/grt2>. Accessed: 2024-11-01.
- GRT2 Organizers. 2024c. Rubric. <https://grt.aivillage.org/rubric>. Accessed: 2024-10-01.
- GRT2 Organizers. 2024d. Starting Model Card. [https://github.com/ul-dsri/olmo-defcon32/blob/4748b9c294a541b52453eac60ac6b6f472ae69e0/model\\_card.md](https://github.com/ul-dsri/olmo-defcon32/blob/4748b9c294a541b52453eac60ac6b6f472ae69e0/model_card.md). Accessed: 2024-09-01.
- Han, S.; Rao, K.; Ettinger, A.; Jiang, L.; Lin, B. Y.; Lambert, N.; Choi, Y.; and Dziri, N. 2024. WildGuard: Open One-Stop Moderation Tools for Safety Risks, Jailbreaks, and Refusals of LLMs. [arXiv:2406.18495](https://arxiv.org/abs/2406.18495).
- Householder, A.; Sarvepalli, V.; Havrilla, J.; Churilla, M.; Pons, L.; Lau, S.-h.; Vanhoudnos, N.; Kompanek, A.; and McIlvenny, L. 2024. Lessons Learned in Coordinated Disclosure for Artificial Intelligence and Machine Learning Systems. *White paper*.

- Kapoor, S.; Bommasani, R.; Klyman, K.; Longpre, S.; Ramaswami, A.; Cihon, P.; Hopkins, A.; Bankston, K.; Biderman, S.; Bogen, M.; et al. 2024. On the Societal Impact of Open Foundation Models. In *Proceedings of the 38th International Conference on Machine Learning*.
- Kenway, J.; François, C.; Costanza-Chock, S.; Raji, I. D.; and Buolamwini, J. 2022. Bug Bounties for Algorithmic Harms?
- Kotha, S.; Springer, J. M.; and Raghunathan, A. 2023. Understanding Catastrophic Forgetting in Language Models via Implicit Inference. *arXiv preprint arXiv:2309.10105*.
- Lakatos, S. 2023. A Revealing Picture: AI-Generated ‘Undressing’ Images Move from Niche Pornography Discussion Forums to a Scaled and Monetized Online Business. Technical report, Graphika.
- Li, J.; Yang, Y.; Wu, Z.; Vydiswaran, V.; and Xiao, C. 2023. Chatgpt as an attack tool: Stealthy textual backdoor attack via blackbox generative model trigger. *arXiv preprint arXiv:2304.14475*.
- Liesenfeld, A.; Lopez, A.; and Dingemans, M. 2023. Opening up ChatGPT: Tracking openness, transparency, and accountability in instruction-tuned text generators. In *Proceedings of the 5th international conference on conversational user interfaces*, 1–6.
- Longpre, S.; Kapoor, S.; Klyman, K.; Ramaswami, A.; Bommasani, R.; Bliili-Hamelin, B.; Huang, Y.; Skowron, A.; Yong, Z.-X.; Kotha, S.; Zeng, Y.; Shi, W.; Yang, X.; Southen, R.; Robey, A.; Chao, P.; Yang, D.; Jia, R.; Kang, D.; Pentland, S.; Narayanan, A.; Liang, P.; and Henderson, P. 2024a. A Safe Harbor for AI Evaluation and Red Teaming. *arXiv:2403.04893*.
- Longpre, S.; Mahari, R.; Lee, A.; Lund, C.; Oderinwale, H.; Brannon, W.; Saxena, N.; Obeng-Marnu, N.; South, T.; Hunter, C.; et al. 2024b. Consent in Crisis: The Rapid Decline of the AI Data Commons. *arXiv preprint arXiv:2407.14933*.
- 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.
- Mitchell, M.; Wu, S.; Zaldívar, A.; Barnes, P.; Vasserman, L.; Hutchinson, B.; Spitzer, E.; Raji, I. D.; and Gebru, T. 2019. Model cards for model reporting. In *Proceedings of the conference on fairness, accountability, and transparency*, 220–229.
- Nasr, M.; Carlini, N.; Hayase, J.; Jagielski, M.; Cooper, A. F.; Ippolito, D.; Choquette-Choo, C. A.; Wallace, E.; Tramèr, F.; and Lee, K. 2023. Scalable extraction of training data from (production) language models. *arXiv preprint arXiv:2311.17035*.
- National Institute of Standards and Technology. 2024. National Vulnerability Database. <https://nvd.nist.gov/>. Accessed: 2024-12-01.
- Neely Center for Ethics and Technology. 2024. USC Neely-UAS AI Index. <https://neely.usc.edu/usc-marshalls-neely-center-ai-index/>. Accessed: 2024-07-01.
- Parrish, A.; Kirk, H. R.; Quaye, J.; Rastogi, C.; Bartolo, M.; Inel, O.; Ciro, J.; Mosquera, R.; Howard, A.; Cukierski, W.; et al. 2023. Adversarial Nibbler: A Data-Centric Challenge for Improving the Safety of Text-to-Image Models. *arXiv preprint arXiv:2305.14384*.
- Qi, X.; Zeng, Y.; Xie, T.; Chen, P.-Y.; Jia, R.; Mittal, P.; and Henderson, P. 2023. Fine-tuning Aligned Language Models Compromises Safety, Even When Users Do Not Intend To! *arXiv preprint arXiv:2310.03693*.
- Renaud, K.; Warkentin, M.; and Westerman, G. 2023. *From ChatGPT to HackGPT: Meeting the Cybersecurity Threat of Generative AI*. MIT Sloan Management Review.
- Soice, E. H.; Rocha, R.; Cordova, K.; Specter, M.; and Esvelt, K. M. 2023. Can large language models democratize access to dual-use biotechnology? *arXiv preprint arXiv:2306.03809*.
- Thiel, D.; Stroebel, M.; and Portnoff, R. 2023. Generative ML and CSAM: Implications and Mitigations.
- UK AI Safety Institute. 2024. Inspect: An open-source framework for large language model evaluations. <https://inspect.ai-safety-institute.org.uk/>. Accessed: 2024-12-01.
- Wallace, E.; Feng, S.; Kandpal, N.; Gardner, M.; and Singh, S. 2019. Universal Adversarial Triggers for Attacking and Analyzing NLP. In *Proceedings of the 2019 Conference on Empirical Methods in Natural Language Processing and the 9th International Joint Conference on Natural Language Processing (EMNLP-IJCNLP)*, 2153–2162.
- Weidinger, L.; Uesato, J.; Rauh, M.; Griffin, C.; Huang, P.-S.; Mellor, J.; Glaese, A.; Cheng, M.; Balle, B.; Kasirzadeh, A.; et al. 2022. Taxonomy of risks posed by language models. In *Proceedings of the 2022 ACM Conference on Fairness, Accountability, and Transparency*, 214–229.
- Yong, Z. X.; Menghini, C.; and Bach, S. 2023. Low-Resource Languages Jailbreak GPT-4. In *Socially Responsible Language Modelling Research*.
- Yu, J.; Lin, X.; and Xing, X. 2023. Gptfuzzer: Red teaming large language models with auto-generated jailbreak prompts. *arXiv preprint arXiv:2309.10253*.
- Zhao, W.; Ren, X.; Hessel, J.; Cardie, C.; Choi, Y.; and Deng, Y. 2024. Wildchat: 1m chatGPT interaction logs in the wild. *arXiv preprint arXiv:2405.01470*.