arXiv:2411.10547v2 [cs.CY] 1 Dec 2024

AI Safety Frameworks Should Include Procedure for Model Access Decisions

Edward Kembery
ERA Fellow
edwardkembery@outlook.com

Tom Reed
ERA Fellow

Abstract

The downstream use cases, benefits, and risks of AI models depend significantly on what sort of access is provided to the model, and who it is provided to. Though existing safety frameworks and AI developer usage policies recognise that the risk posed by a given model depends on the level of access provided to a given audience, the procedures they use to make decisions about model access are ad hoc, opaque, and lacking in empirical substantiation. This paper consequently proposes that frontier AI companies build on existing safety frameworks by outlining transparent procedures for making decisions about model access, which we term Responsible Access Policies (RAPs). We recommend that, at a minimum, RAPs should include the following: i) processes for empirically evaluating model capabilities given different styles of access, ii) processes for assessing the risk profiles of different categories of user, and iii) clear and robust pre-commitments regarding when to grant or revoke specific types of access for particular groups under specified conditions.

Preprint. Under review.

1 Executive Summary

Existing frontier AI safety frameworks recognise that the societal impact of an AI model depends on what style of access to the model is provided to which group(s). Some, such as Anthropic's, reference "access controls" explicitly [1]. However, current safety frameworks provide no procedures to ensure that these decisions about who should have what style of access are made responsibly.

Frontier AI companies should expand existing safety frameworks to account for access considerations. We refer to this body of procedure as *Responsible Access Policies* (RAPs). They should set out *general commitments* to follow best practices in governing model access, and *specified procedure* for assessing the risks and benefits of providing different access styles to different groups:

General Commitments

- **Provide *and withhold* styles of access for different groups on the basis of specified procedure**
- **Review & update access decisions and specified procedure as context evolves**

Specified Procedure

- **Evaluate and assess the risks and benefits of different access styles**
- **Categorise and assess the needs and capacities of different user groups**
- **Set significance levels for decisions**

We stress that these decision-making processes should be clear and transparent to stakeholders. To make these decisions legible, we suggest that frontier AI companies use an *Access Assessment Matrix*:

Example Access Assessment Matrix

Values are hypothetical and may not reflect current practices.

User Groups	Chat Interface	Sampling API	Fine-tuning API	Model Weights (Download)
Governments and AISIs				
Trusted Auditors and Partnerships				
Research Community				
Public Release				

Access Styles

"There is strong evidence that this would be safe and highly beneficial"

"The evidence is unclear and further research is needed before proceeding"

"Evidence suggests that this might create undesirable risks that outweigh benefits"

2 Introduction

The rapid development of general-purpose AI has led several frontier developers to implement safety frameworks which help to structure decisions about how to minimise societal harm whilst promoting economic growth and respect for diverse cultural norms [2, 3, 4]. One purpose of these frameworks is to help guide decision-making about *how* and *to whom* developers release their models to [5, 6, 7]. For instance, a recent update to Anthropic's RSP explicitly mentions that "access controls" and a "tiered access system" should be developed as part of commitments to exist safety frameworks [1].

Incorporating access policy into safety frameworks is appropriate. General-purpose AI models can be made available to end users in a wide number of different ways which enable distinct uses cases, each of which may engender different downstream impacts. If access to information or components of a system is provided to users incautiously, it might pose societal risks by increasing the likelihood that powerful AI systems are misused [6, 8, 9] or stolen [10]. Providing clear frameworks for governing decisions would also create certainty for researchers and businesses that rely on publicly available models.

However, existing safety frameworks do not go far enough to support responsible model access governance. Decisions about model access at frontier AI companies are often ad hoc, opaque, and lacking in empirical support.[1] Although Anthropic's recent updates are encouraging, no other frontier AI safety frameworks outline protocols for making informed decisions about how and to whom a model is deployed [3, 4].

Consequently, we contend that frontier AI companies should build on existing safety frameworks by including access considerations. First, we make the case for Responsible Access Policies (RAPs) as part of safety frameworks, building on model access research. Second, we set out a framework for RAPs, suggesting that these should incorporate i) empirical evaluation of access styles, ii) processes for classifying different categories of user and iii) clear pre-commitments on providing or retracting access to specified groups in defined circumstances. We emphasize that frontier AI companies have a real opportunity to set an example here—both for other AI companies and for regulators—helping to develop and implement responsible, empirically-driven model access governance.

2.1 Related Literature

Model Deployment Governance. Frontier AI companies that have produced safety frameworks around AI systems include Anthropic [2], Google DeepMind [4] and OpenAI [3]. However, these extent to which these frameworks are comprehensive or feasible is unclear [12], and other than Anthropic's recent update, they do not make explicit commitments to procedures for making decisions about which groups to provide access to under which circumstances. Comapnies like Meta, Mozilla and LAION [13, 14] have associated themselves with 'open-source' access policies, but may continue to do so as AI models develop.

Structured access and model access. Structured transparency—a framework for information sharing designed to facilitate collaboration and information sharing whilst addressing privacy concerns [15]—has been applied to model access decisions around AI systems for several years [16]. Recent work has explored what styles of model access are advantageous for valuable safety research [17] and audits [18, 19, 20, 21, 22]. Technical work has explored the extent to which it is possible to make model access governance robust: for instance, by assessing how easy it is to "steal" the weights of a production language model through interactions [23, 24].

Open vs closed-source debate. Early work related to model access governance transferred a framing of 'openness' from earlier software research (e.g. Millar et al., 2022 [25]). In practice, key intuitions have proved unsuitable for AI systems [26], and there is as yet no official definition of 'open source AI' [27]. Some research has argued that we should move 'beyond the open vs closed debate' [28] towards 'hybrid' access regimes that involve providing access to some aspects of the model but not others [29]. Nonetheless, several papers have usefully focused on models with widely available model weights in order to identify theoretical marginal risks [30] and possible alternatives to open-weight deployment [6]. Others have made a vigorous case for various definitions of open-source [7, 14].

[1] For instance, Meta's CEO stated that if "there's some qualitative change in what the thing is capable of, and we feel like it's not responsible to open source it, then we won't" [11]. Meta has not released any statements as to how these decisions would be made in practice.

2.2 Definitions

Model Aspects: AI models are complex structures composed of various elements, such as code and weights [26]. They are the basis for AI systems, which involve additional infrastructure like GPUs and user interfaces [26, 31]. Organizations developing AI models must decide not only which components of the model to make accessible to different groups but also how to share broader information about the model, such as training data or sample outputs. We refer to this set of information as the *model aspects*.

Access Styles: When a developer makes an aspect of a model available to a user in a certain way (for instance, with certain permissions or restrictions on how the information is to be used), we refer to this as an *access style*. Different access styles might include chat, fine-tuning, weights inspection, or weights modification [17].

Access Regimes: When organizations develop an AI model, they may offer different access styles to various groups. For example, a company might provide access to the weights for an AI Safety Institute, API-based fine-tuning for trusted researchers, and a chat interface for the public. We refer to this collective set of access styles granted to all users as an *access regime*. Different access regimes might create different use cases en masse, leading to different downstream impacts.

3 The Case for Responsible Access Policies

3.1 Misgoverning model access could have serious societal consequences

Incautious access policies could create greater societal risks [8, 6]. Whilst there is already substantial literature investigating the potential of powerful AI systems to cause serious harm, [32, 33], miscalibrated access governance might make these risks significantly worse, for three main reasons. First, it could make it easier to bypass the safeguards of models that would otherwise be robust to jailbreaking [34], either by allowing detuning or making it easier to design custom jailbreaks. Second, unsecured models could spread globally, with no way to retract them, expanding both the geographic risk zone and the time frame of the risk [6]. Thirdly, miscalibrated access could make it harder to have oversight on how models are used, reducing decision-makers' insight into evolving risks, thereby making them less prepared to address these issues.

On the other hand, overly restrictive access policies might incur high opportunity costs [7]. Broadly speaking, these costs fall into three categories. First, providing insufficient access to safety organisations could delay the development of safety research [17, 35], thus resulting in greater future risks from AI systems. Second, miscalibrated access might lead to underutilisation of valuable use cases, hampering the ability of the public to develop novel AI applications without substantially increasing risk [13]. Third, overly restrictive access could lead to a more unequal distribution of decision-making power over advanced AI, perhaps disproportionately affecting users from particular socioeconomic backgrounds [36, 37].

3.2 Frontier AI companies might not necessarily make good decisions by default

The current absence of frameworks for model access governance makes it difficult for developers to make good decisions about how and to whom they release their models. There are at least three reasons this is the case. First, decision-makers are currently very uncertain about key issues, such as the relative increase in risk of harm caused by releasing a particular access modality to a particular audience. Empirical evidence as to potential risks and benefits of different access regimes is generally lacking [28].

Second, current safety frameworks fall short in governing model access. While they outline how companies manage risks in developing and deploying frontier AI systems [12], they lack procedures for decisions on model access. For example, many of them do not address how the 'risk category' of a model may vary by access style [3, 4]. Nor do they clarify how user categories are classified or how these relate to threat models informing risk thresholds for deployment. OpenAI's Preparedness Framework, for instance, mentions restricting dangerous models to trusted parties but lacks criteria for defining a trusted party [3]. (Furthermore, if the risk category of models was to depend significantly on what style of access was provided to particular parties, conventional safety frameworks may prove insufficient for reducing risks—however, this is beyond the scope of this paper [12]).

Finally, organisational incentives might not be naturally aligned with responsible model access governance. For instance, pursuing one type of access regime might increase the number of customers willing to pay for a model and so be desirable for the company, even if this regime creates an uplift in the capabilities of malicious actors.

3.3 Governments and societies need more clarity about how this will develop in future

Even if miscalibrated access regimes posed no marginal risk of harms, or if most frontier AI companies were already well-placed to make good decisions, there are still at least three reasons why developers should adopt RAPs. First, governments and industry that rely on AI models are likely to require clear guidance on the sort of model access they can expect to have in the future. Second, these policies should be transparent so as to encourage governments and researchers to audit them, improve them, and contest disagreements in a open manner. Third, whilst some frontier AI companies might be well situated to govern model access effectively, sharing best practises between companies, and around the world, might provide incentives towards developing good model access governance practises across industry.

4 Incorporating Responsible Access Policies

4.1 Deployers should pre-commit to responsible model access governance processes

Model governance pre-commitments, as implemented in safety frameworks, offer at least three key advantages for model access governance. They increase the likelihood of compliance, inform governments about the effectiveness of self-regulation, and provide clarity for those who depend on particular model access styles for their everyday use cases. RAPs should involve two types of pre-commitments: general commitments about the procedures that will be followed related to model access governance, and specified procedures related to access style 'roll-forward' and 'roll-back' decisions:

General Commitments. Frontier AI companies should provide general commitments as to how they intend to govern model access. This could include clauses to i) pre-commit not to release any access styles without a clear safety case to do so, ii) pre-commit to evaluating different access styles for frontier models in the context of different user groups, and iii) pre-commit to reviewing and potentially updating these decisions (whether for or against a particular access regime) on an ongoing basis as use-cases and technologies change.

Specified Procedure. Frontier AI companies should offer explicit commitments regarding decisions about whether to 'roll-forward' or 'roll-back' different access styles. These commitments may include: i) detailed protocols for conducting evaluations, including defined significance levels; ii) procedures for assessing the different use cases of groups concerned and how they might be affected by the decision; and iii) governance frameworks that demonstrate how this information will be integrated into decision-making processes, as well as identifying those groups accountable for the resulting decisions.

We suggest that frontier AI companies approach these decisions with caution, especially in regard to situations where providing access might mean that companies are not able to retract it in future (for instance, if a group was able to download model weights). In some cases, it may be appropriate to investigate mitigating strategies for reducing the risks associated with model access regimes (for instance, providing more robust watermarking for models with open-source access regimes), or reducing the burden on users of retracted access styles.

4.2 Decisions about access should be informed by empirical evaluations of the use cases and risks across diverse access styles and end-user categories

The impact profile of a model is related both to its capabilities and the particular access regime it is subject to — that is, both *what* sort of access is provided and *whom* they provided to. In the *Access Assessment Matrix*, we refer to these as 'styles of access' and 'user groups' respectively. RAPs should incorporate evaluations that account for both before considering whether to provide particular styles of access to particular groups:

Evaluating Access Styles. Researchers already assess how different access styles might uplift different use cases for different groups [17]. Researchers should continue to explore how particular capabilities change between access styles such as a chat-based interface, secure fine-tuning access, and open weights release. They might also explore how different access styles might change in future (for instance, as synthetic data for fine-tuning models becomes better at eliciting greater capabilities from smaller models). Key priorities might be investigating this change across domains like Chem/Bio, Cyber, and Autonomous Replication, but researchers should also investigate the extent to which different styles of access might make it easier to bypass model safeguards and enable harmful use cases related to CSAM, NCII and disinformation. These evaluations might be supplemented by research into the more general long-term downstream effects of particular access styles by third-party organisations.

Accounting for User Profiles. Frontier AI companies already assess how different capabilities might differently enable different groups [38]. Groups that frontier AI companies might want to provide access to include general research organisations, specified trusted AI evaluation organisations like AISIs, governments, and the public. These groups might have different goals, resources, time, and levels of skill at eliciting capabilities from models given particular access styles. Frontier AI companies should have reasonable views as to which groups have what level of skill, time and resources, and should, if necessary, assess how these are changing on an ongoing basis (for instance,

as prompt-engineering skills improve or sophisticated scaffolds proliferate online). These user profiles should be paired with evaluations of the ranging capabilities afforded by different access styles (as described above), in order to inform decisions about whether providing a given access modality to a given category of end-user is high-risk, or alternatively, societally desirable. If evidence is unclear, it may be useful to incorporate a phase in which a model is tested for a sustained period in a secure and closely monitored API to understand how the model is used by a group.

We expect that modelling different malicious actors using new technologies over different time frames with different resources will be a significant challenge. Frontier AI companies should collaborate with different actors in government to get better assessments of these as models become more powerful, and refrain from making definitive statements in lieu of sparse evidence or inconclusive results. We suggest that frontier AI companies should be especially wary of dangerous capabilities when considering irreversible access styles such as downloadable weights release [8].

4.3 Decision-making procedures should be clear and transparent

If RAPs are implemented but not operationalized effectively by frontier AI companies, they might diminish appetite for legislation without ensuring that model access is governed responsibly. If they are not transparent to the public and key stakeholders, they might be less useful for coordinating governance, increasing the possibility of unnecessary legislation, and creating uncertainty for groups who depend on particular access styles. Policies should therefore be clear and transparent:

Clarity. Making RAPs clear and operationalizable will require precise definitions which avoid unfairness, prevent confusion, and eliminate the potential for the creation of legal loopholes that could be exploited. It will be important to define access styles carefully, paying special attention to the extent to which it might be possible to elicit information about some parts of the models by accessing others [10]. It will also be important to define key terms such as 'research organisations', 'trusted third-parties' and 'public' carefully so that these decisions can be adhered to and evaluated [5]. Definitions should be robust enough to provide external watchdogs clear signals when access procedures are not being properly followed, and to hold decision-makers accountable.

Transparency. Making RAPs transparent will require documentation on both general principles and specified procedures at various levels of detail. Appropriate documentation for processes might include i) a longer-form access policy document circulated internally detailing general commitments and specified procedure, ii) a medium-length summary of both aspects published online and iii) a short memo on critical issues tailored for government decision-makers. Appropriate reporting related to a specific decision to release a model for a particular access regime (say, to release fine-tuning access to the public) might include a long-form report detailing the evaluations procedure undertaken, the risks and benefits of the proposed access regime, possible implications of the regime, and the set of circumstances which would have to be met for the decision to change. It might also include a shorter document with key headlines for public release, in the style of a model card. The *Access Assessment Matrix* (see Executive Summary) might be a useful way to build on existing data representation techniques used by safety frameworks to represent some of this information concretely.

5 Conclusion

Whilst frontier AI companies have taken several important and necessary steps to ensure that they govern model deployment responsibly, model access governance procedures are still nascent, and unlikely to be sufficient to avoid societal risks and opportunity costs as models become more capable. Incorporating Responsible Access Policies into existing safety frameworks effectively will be crucial for ensuring that frontier AI companies make responsible decisions about model access.

Doing so will have challenges. The different use cases afforded by different access styles may be unclear and change substantially over time as people find new ways to use the system, or technologies develop. New modalities like agents may affect the extent to which different access styles enable different use cases, requiring access regime changes to minimise risks. Specific decisions, such as retracting styles of access, may be difficult to implement and harmful to particular user groups if not executed with care. However, these challenges only serve to emphasise the importance of forethought, research, responsible pre-commitments and clear decision-making processes that could be supported by RAPs.

6 Acknowledgements

This research was supported by the ERA Fellowship. The authors would like to thank the ERA Fellowship for its financial and intellectual support, and Ben Bucknall, Guarav Sett and Michael Chen in particular for their support and suggestions.

7 Errata

An earlier version of this paper mistakenly included style components indicating that this paper had been accepted to the Socially Responsible Language Modelling Research (SoLaR) Workshop at NeurIPS 2024. This was not the case. We apologise.

8 Bibliography

References

[1] Anthropic. *Responsible Scaling Policy Updates — anthropic.com*. https://www.anthropic.com/rsp-updates. [Accessed 09-11-2024].

[2] METR. *Responsible Scaling Policies (https://metr.org/blog/2023-09-26-rsp/)*. Tech. rep. METR, 26 September 2023.

[3] OpenAI. *OpenAI Preparedness Framework*. https://cdn.openai.com/openai-preparedness-framework-beta.pdf. [Accessed 17-09-2024]. 2023.

[4] Anca Dragan, Helen King, and Allan Dafoe. *Introducing the Frontier Safety Framework — deepmind.google*. https://deepmind.google/discover/blog/introducing-the-frontier-safety-framework/. [Accessed 17-09-2024]. 17 May 2024.

[5] Edward Kembery, Ben Bucknall, and Morgan Simpson. *Towards Model Access Governance*. [Accessed 13-09-2024]. Forthcoming.

[6] Elizabeth Seger et al. *Open-Sourcing Highly Capable Foundation Models: An evaluation of risks, benefits, and alternative methods for pursuing open-source objectives*. 2023. arXiv: 2311.09227 [cs.CY]. URL: https://arxiv.org/abs/2311.09227.

[7] Francisco Eiras et al. "Near to Mid-term Risks and Opportunities of Open Source Generative AI". In: *arXiv preprint arXiv:2404.17047* (2024).

[8] David Evan Harris. *How to Regulate Unsecured "Open-Source" AI: No Exemptions*. Tech. rep. Tech Policy Press, December 4 2023.

[9] Sayash Kapoor et al. *On the Societal Impact of Open Foundation Models*. 2024. arXiv: 2403.07918 [cs.CY]. URL: https://arxiv.org/abs/2403.07918.

[10] Nicholas Carlini et al. *Stealing Part of a Production Language Model*. 2024. arXiv: 2403.06634 [cs.CR]. URL: https://arxiv.org/abs/2403.06634.

[11] Dwarkesh Patel. *Mark Zuckerberg - Llama 3, Open Sourcing $10b Models, & Caesar Augustus — dwarkeshpatel.com*. https://www.dwarkeshpatel.com/p/mark-zuckerberg. [Accessed 26-09-2024]. 2024.

[12] Jide Alaga, Jonas Schuett, and Markus Anderljung. *A Grading Rubric for AI Safety Frameworks*. 2024. arXiv: 2409.08751 [cs.CY]. URL: https://arxiv.org/abs/2409.08751.

[13] Mark Zuckerberg. *Open Source AI Is the Path Forward | Meta*. https://about.fb.com/news/2024/07/open-source-ai-is-the-path-forward/. [Accessed 27-07-2024]. 2024.

[14] Misc. *NTIA AI Open Model Weights Request for Comments*. https://www.regulations.gov/document/NTIA-2023-0009-0001/comment. [Accessed 17-09-2024]. 26 February 2024.

[15] Andrew Trask et al. *Beyond Privacy Trade-offs with Structured Transparency*. 2024. arXiv: 2012.08347 [cs.CR]. URL: https://arxiv.org/abs/2012.08347.

[16] Toby Shevlane. "Structured access: an emerging paradigm for safe AI deployment". In: *arXiv preprint arXiv:2201.05159* (2022).

[17] Benjamin S. Bucknall and Robert F. Trager. *Structured access for third-party research on frontier AI models: Investigating researchers' model access requirements*. Tech. rep. Oxford Martin School of Governance, Oct. 2023.

[18] Shayne Longpre et al. “A safe harbor for ai evaluation and red teaming”. In: *arXiv preprint arXiv:2403.04893* (2024).

[19] Jakob Mökander et al. “Auditing large language models: a three-layered approach”. In: *AI and Ethics* (2023), pp. 1–31.

[20] Jakob Mökander et al. “Conformity assessments and post-market monitoring: a guide to the role of auditing in the proposed European AI regulation”. In: *Minds and Machines* 32.2 (2022), pp. 241–268.

[21] Jakob Mökander and Luciano Floridi. “Ethics-based auditing to develop trustworthy AI”. In: *Minds and Machines* 31.2 (2021), pp. 323–327.

[22] Jakob Mökander and Luciano Floridi. “Operationalising AI governance through ethics-based auditing: an industry case study”. In: *AI and Ethics* 3.2 (2023), pp. 451–468.

[23] Florian Tramèr et al. *Stealing Machine Learning Models via Prediction APIs*. 2016. arXiv: 1609.02943 [cs.CR]. URL: https://arxiv.org/abs/1609.02943.

[24] Nicholas Carlini et al. “Stealing part of a production language model”. In: *arXiv preprint arXiv:2403.06634* (2024).

[25] Stuart Millar. *Vulnerability Detection in Open Source Software: An Introduction*. 2022. arXiv: 2203.16428 [cs.CR]. URL: https://arxiv.org/abs/2203.16428.

[26] Adrien Basdevant et al. *Towards a Framework for Openness in Foundation Models: Proceedings from the Columbia Convening on Openness in Artificial Intelligence*. 2024. arXiv: 2405.15802 [cs.SE]. URL: https://arxiv.org/abs/2405.15802.

[27] Open Source Initiative. *The Open Source AI Definition – draft v. 0.0.3*. [Accessed 01-08-2024]. Oct. 2023.

[28] Jon Bateman et al. *Beyond Open vs. Closed: Emerging Consensus and Key Questions for Foundation AI Model Governance*. https://carnegieendowment.org/research/2024/07/beyond-open-vs-closed-emerging-consensus-and-key-questions-for-foundation-ai-model-governance?lang=en. [Accessed 01-08-2024]. July 2024.

[29] Irene Solaiman. *The Gradient of Generative AI Release: Methods and Considerations*. 2023. arXiv: 2302.04844 [cs.CY]. URL: https://arxiv.org/abs/2302.04844.

[30] Rishi Bommasani et al. *Considerations for Governing Open Foundation Models*. Tech. rep. Stanford HAI, 2024.

[31] Council of European Union. *Regulation 2024/1689 of the European Parliament and of the Council*. http://eur-lex.europa.eu/legal-content/EN/TXT/?qid=1416170084502&uri=CELEX:32014R0269. June 2024.

[32] Markus Anderljung et al. *Frontier AI Regulation: Managing Emerging Risks to Public Safety*. 2023. arXiv: 2307.03718 [cs.CY]. URL: https://arxiv.org/abs/2307.03718.

[33] Laura Weidinger et al. “Ethical and social risks of harm from language models”. In: *arXiv preprint arXiv:2112.04359* (2021).

[34] Xiangyu Qi et al. *Fine-tuning Aligned Language Models Compromises Safety, Even When Users Do Not Intend To!* 2023. arXiv: 2310.03693 [cs.CL]. URL: https://arxiv.org/abs/2310.03693.

[35] Stephen Casper et al. *Black-Box Access is Insufficient for Rigorous AI Audits*. 2024. DOI: 10.1145/3630106.3659037. arXiv: 2401.14446 [cs.CY]. URL: https://arxiv.org/abs/2401.14446.

[36] Kunat Pipatanakul et al. *Typhoon: Thai Large Language Models*. 2023. arXiv: 2312.13951 [cs.CL]. URL: https://arxiv.org/abs/2312.13951.

[37] Damian Blasi, Antonios Anastasopoulos, and Graham Neubig. “Systematic Inequalities in Language Technology Performance across the World’s Languages”. In: *arXiv preprint arXiv:2110.06733* (2021).

[38] OpenAI. *Building an early warning system for LLM-aided biological threat creation*. Tech. rep. OpenAI, 2024.