
MULTINATIONAL AGI CONSORTIUM (MAGIC): A PROPOSAL FOR INTERNATIONAL COORDINATION ON AI

Jason Hausenloy*
UWCSEA East
hause56773@uwcsea.edu.sg

Andrea Miotti*
Conjecture
andrea@conjecture.dev

Claire Dennis*
Princeton University
claire.dennis@princeton.edu

ABSTRACT

This paper proposes a Multinational Artificial General Intelligence Consortium (MAGIC) to mitigate existential risks from advanced artificial intelligence (AI). MAGIC would be the only institution in the world permitted to develop advanced AI, enforced through a global moratorium by its signatory members on all other advanced AI development. MAGIC would be exclusive, safety-focused, highly secure, and collectively supported by member states, with benefits distributed equitably among signatories. MAGIC would allow narrow AI models to flourish while significantly reducing the possibility of misaligned, rogue, breakout, or runaway outcomes of general-purpose systems. We do not address the political feasibility of implementing a moratorium or address the specific legislative strategies and rules needed to enforce a ban on high-capacity AGI training runs. Instead, we propose one positive vision of the future, where MAGIC, as a global governance regime, can lay the groundwork for long-term, safe regulation of advanced AI.

1 Executive Summary

Today, a handful of U.S.-based AI companies are spearheading the development of extremely powerful, general AI systems – what the leading developers refer to as Artificial General Intelligence (AGI). AGI lacks any unified, concrete definition, but has been described in some contexts as “superintelligence” or “highly autonomous systems that outperform humans at most economically valuable work.”[1] While the specific threshold for defining an “AGI” is often contested, there is an emerging scientific and public consensus that the unchecked development of such advanced, autonomous AI systems may pose enormous, societal-scale risks.[2, 3, 4] The challenges posed by advanced AI are sizable enough to necessitate an international response. A growing chorus of policymakers, technologists, and governance experts are thus calling for global AI governance through various proposed international bodies to facilitate global coordination and oversight.[5, 6, 7, 8, 9]

Some recent AI governance proposals rely on *post-hoc* auditing and verification mechanisms – evaluating the system once it has already been developed. [10] Given the potential of advanced AI systems to rapidly improve during training, this method is insufficient. Regulation of most global hazards, such as biological weapons, do not require the technology to be first engineered in order to subsequently control it. To intervene in the pre-development stage, however, requires centralized, tightly controlled, and constant oversight of advanced AI models, with methods in place to abort training if there are signs of dangerous capabilities. To comprehensively halt unchecked AGI advancement therefore requires monopolization of development in a centralized facility, coupled with a simultaneous moratorium on all unofficial development.

To achieve this level of oversight, this paper proposes an **exclusive Multinational AGI Consortium (MAGIC)** – an international governance structure designed to avoid existential risks of AI systems. MAGIC has the following four core characteristics:

*Equal contribution of all authors. Andrea Miotti is the corresponding author.

²These include but are not limited to: UN Secretary-General Antonio Guterres, Microsoft CEO Satya Nadella, OpenAI CEO Sam Altman, UK Prime Minister Rishi Sunak (in conversation with US President Joe Biden) and Chair of the UK Foundation Model Taskforce Ian Hogarth.

1. **Exclusive:** the world’s only advanced AI facility, with a monopoly on the development of advanced AI models, and nonproliferation of AI models everywhere else.
2. **Safety-focused:** focused on the development of AI systems that are safe by design, including development of new architectures and ways to bound existing AI systems.
3. **Secure:** among the most highly secure facilities on Earth, with strict protocols for information security.
4. **Collective:** supported internationally, where the benefits of AI systems are distributed among all member countries.

2 Multinational AGI Consortium (MAGIC)

The Multinational AGI Consortium (MAGIC) relies on four core characteristics: 1) exclusivity, 2) focus on safety, 3) security, and 4) collective action. MAGIC goes further than other proposals for international AI research institutions in its call for an immediate restriction of all external advanced AI development.

2.1 Exclusive: The World’s Only Advanced AI Facility

MAGIC would be the sole institution authorized to conduct advanced AI research beyond a defined capability level. In this single, secure facility, the world’s leading AI experts would research advanced AI systems and new architectures to design powerful, safe AI systems. MAGIC would be the only organization authorized to train advanced AI models above a certain threshold of computing power. It would attract top global talent through its exclusive resource access, cutting-edge research, reputation, and funding.³

This exclusivity would make any advanced AI development outside of MAGIC illegal, enforced through a global moratorium on training runs using more than a set amount of computing power.^{4,5}[16] Signatory countries to MAGIC would mandate cloud-computing providers, with whom almost all AI companies partner, to prevent any training runs above a specific size within their national jurisdictions.⁶ Signatory countries would be responsible for ensuring mutual verification of the enforcement of this measure.⁷ Domestically, this would not require the creation of any additional monitoring mechanisms, as cloud providers already track the size of training runs to ensure compliance with their terms of service and for routine business operations.⁸

To make the moratorium sustainable in the long-term, MAGIC members would have to further monitor and control hardware to train AI systems. AI accelerators, advanced chips such as Nvidia’s H100s specialized for AI applications, are currently available for purchase to most commercial actors worldwide and are untracked.⁹[20] There are very few actors who have the expertise and financial resources to aggregate sufficient compute to train state-of-the-art

³Other large-scale scientific collaborations similarly maintain control over specific technologies through the physical concentration of resources. CERN hosts the world’s largest particle physics laboratory, including the Large Hadron Collider, and retains exclusive access to these facilities. [11] ITER similarly hosts the world’s largest experimental fusion reactor in France which cannot be replicated by any single nation due to its cost and complexity. [12] [13]

⁴We suggest that one way to implement such thresholds would be to establish two thresholds: a lower “Frontier Threshold” and a higher “Horizon Threshold”. These would be calculated in terms of total computation used for an AI model’s training run (measured in FLOP) and would help classify AI training runs and apply adequate measures to each:

1. Training runs that meet or exceed the “Frontier Threshold” must adhere to mandatory oversight requirements from the national jurisdiction where they are being conducted in, including closer monitoring, stringent reporting, and regular audits.
2. Training runs that meet or exceed the “Horizon Threshold” should be prohibited by default.

These are discussed in a proposal for the priorities of the UK Foundation Models Taskforce, and Dalrymple elaborates on how such thresholds could be defined and calculated. [14, 15]

⁵In some versions, MAGIC could provide additional computational power remotely through a ‘compute bank’. This would be a centralized resource of computing power that can be accessed remotely. The operations would primarily monitor resource usage and applications of this model. Finally, AI safety research, external to MAGIC, would focus on fundamental technical problems. These include the mechanistic interpretability of models, formal proofs of boundedness, and ultimately, value alignment.

⁶We suggest 10^{24} floating point operations (FLOP) as an appropriate threshold given current levels of compute usage in frontier models [17]

⁷This would be similar to some “IAEA-for-AI” models, which serve as a standard and compliance enforcement role through inspecting hardware, which may require direct physical access to data centres, or model evaluations after development to test the competencies of the model. [18, 19]

⁸Data centers are a key entry point for regulators. They could be required to disclose large training runs which could then be audited by an international regulatory regime.[20] [21]

⁹This is true apart from a handful of entities under export controls by the US government.

models.¹⁰ However, accessibility will increase as enormous amounts of capital continue to flow into AI development and available chips become more powerful and less expensive. This moratorium could be supplemented with stricter monitoring of GPU supply chains.¹¹[19] MAGIC member states would intervene in global semiconductor supply chains to restrict the ability to obtain or build the capacity to manufacture AI accelerators, for example by building a ‘chip directory.’[20] Members achieve this by imposing export controls on AI accelerators, their design blueprints, and equipment needed for their production for non-signatory states or non-compliant signatory countries

Why monitor and restrict compute?

Compute is an effective lever for centralized intervention. Compared to alternatives such as algorithms or data, compute is a physical resource that is easily detectable and already monitored. As the large training runs required for frontier model development utilize vast quantities of these chips, compute measurement provides a proxy to measure the capabilities of models for regulators. [23, 24, 25] Most importantly, while the data and algorithms used for the most powerful AI models might vary over time, all AI and all computation will always need computing power, and the amount of computing power will likely continue to be correlated to the power of the AI system. It is universal and agnostic to AI architecture. A cap on compute affects most compute-intensive modern architectures, such as reinforcement learning and deep learning and non-transformers such as diffusion models. This approach is therefore effective in the short-term until significant improvements in algorithmic efficiency necessitate adjustments.

Limitations: Enforcing exclusivity, and a moratorium, through compute requires that computing power remains a strong proxy for how advanced an AI model is. In general, scale factors like compute and dataset size are still expected to remain highly relevant to pushing the capabilities of frontier models, but incentives are driving algorithmic efficiency advances which could begin to overturn this assumption.¹²[26, 27] Additionally, some research indicates that, in any case, the current rate of scaling may not be viable in the long-term. If this is the general trend, compute governance may not continue to be the most direct proxy for model size in the long term.[27] In other words, as advanced models become smaller and cheaper to build, establishing a compute-based threshold for monitoring the most powerful AI systems may become unviable. MAGIC’s response to account for this should be to lower the threshold over time, while investing in safety research such that the highest available capabilities threshold is being used, coupled with a monitoring agency to ensure that the best safety techniques are being followed. Flexibility must be built into MAGIC’s monitoring metrics to account for algorithmic progress over time and ensure it is nimble enough to incorporate new capabilities.

2.2 Safety-Focused: Differential Technological Development of Safe Architectures

MAGIC’s core mandate will not be to build the most powerful AI models as soon as possible, but to build them as safely as possible. MAGIC would encourage the differential technological development of architectures that are safer while restricting the training of AI through today’s black-box, foundation model method.

MAGIC would move from a paradigm in which safety is an afterthought, to safety being built into the design from the beginning. This would begin by experimenting with the largest current models (for example, GPT-4 level). Instead of automatically scaling to the next model, MAGIC would pursue an approach that tweaks and improves safer architectures, making them more interpretable and bounded. MAGIC would first focus on developing models that are on par with existing capabilities and provably safe. If MAGIC is successful, it will then move on to slowly expanding its models to be more capable. If it is not proven safe, MAGIC would stop building and wait for new evidence and breakthroughs before restarting its development efforts.

This is important because we do not understand the internal computations of current “black-box” foundation model architectures.¹³[28]They often demonstrate unexpected capabilities, some of which only become evident after

¹⁰Compute researcher Lennart Heim estimates there are “less than 100 actors, probably... 30 actors” who could aggregate enough compute to achieve this.[22]

¹¹In the long run, implementing this would be harder. To maintain a moratorium, a coalition of actors would have to intervene in global semiconductor supply chains to restrict their ability to obtain or build the capacity to manufacture AI accelerators. Members achieve this by imposing export controls on the AI accelerators, their design blueprints, and the equipment needed for their production.

¹²OpenAI in 2020: “Our results suggest that for AI tasks with high levels of investment (researcher time and/or compute) algorithmic efficiency might outpace gains from hardware efficiency (Moore’s Law).” [26]

¹³Prof. Stuart Russell said in a recent Senate Judiciary hearing, of the current paradigm of developing advanced AI models: “They are black boxes... their internals are largely impossible to understand.” [28]

deployment to millions of users.¹⁴ Advanced AI models can thus be considered to have “unbounded” capabilities – we cannot be certain what the limits of any system may be, or how to align these systems to human values and goals. MAGIC circumvents these dilemmas by designing advanced AI to be safe from the ground up, rather than having safety be treated as a post-hoc patch.

Limitations: These safer approaches, however, come with a tradeoff, or a “safety tax”[29]. Technical AI safety research progress could be slow. Connecting AI systems directly to the Internet and allocating greater computing resources will make them more capable and commercially attractive. In a competitive environment, this means that unsafe systems will be favored. Though we currently lack substantial empirical evidence supporting safer architectures, there appears to be a theoretical possibility of developing systems with clear boundaries and coordination to prevent alternative forms of development.¹⁵

Furthermore, if not managed, the pursuit of safer architectures may be dangerous. For example, one strategy MAGIC might adopt involves designing constrained systems that reduce their dependence on black-box models. This would be accomplished by inventing methods to attain the same level of capability using smaller black-boxes. If these techniques were to be released globally, they could potentially be applied to the largest models currently available, resulting in significantly enhanced capabilities. Given this, research undertaken at MAGIC should be thoroughly vetted before being shared with other researchers, including those from signatory countries.

2.3 Secure: Among the Most Highly Secure Facilities on Earth

MAGIC would be one of the most secure facilities on Earth. Security is crucial due to the immense potential for misuse, the risk of powerful models leaking by accident, and interest of malicious actors in acquiring the valuable technological artifacts that MAGIC would produce. Even a system designed with safety measures could pose a threat if exploited by malicious actors, and this security must be rigorously upheld through digital, physical, and personnel measures.

With sufficient motivation and resources, secret information can stay secret.¹⁶ Like nuclear weapons systems, MAGIC would have specialized hardware to remain isolated from the Internet. Further physical access devices could include biometric access controls, video surveillance, security patrols, and physical isolation of the facility. Scientists would operate on a compartmentalized basis, with strict background checks and lifestyle monitoring.¹⁷ Digital protections would ensure that sensitive systems are air-gapped from external networks and use specialized encrypted hardware.¹⁸ Scientists working at MAGIC would be required to sign strict non-disclosure agreements, protecting classified information under threat of legal consequences, including imprisonment. Additionally, non-compete clauses would ensure these scientists do not engage in similar AI projects after leaving MAGIC. It would additionally need to be located in a geopolitically and geographically neutral and secure location.¹⁹

Limitations: Centralization and suppression of AI safety research could prove to be inefficient and limit innovation. A diversity of approaches and research groups would likely yield faster progress in advanced AI safety. Additionally, while hardware resources could be consolidated, it would be unnecessary to physically relocate researchers to a single site given the remoteness of the work. MAGIC can, however, take a more adaptable approach, involving a small (<10) network of coordinated labs, each with independent scientific direction, though this would present further tradeoffs for security and control. [31, 32]

¹⁴Anthropic’s CEO said that, “You have to deploy [the model] to a million people before you discover some of the things that it can do.”

¹⁵For instance, one approach for creating safer AI systems is through ‘Cognitive Emulations’, a method that tries to build predictably limited systems by mimicking human-like processes. Another naively ‘safe’ system is the Open Agency model, which separates world models from data and planning and acting in real-time. More generally, from first principles, there do exist human-level systems that do not unilaterally result in human extinction: humans. All of these proposals share the property that these systems cannot compete with black box systems with increased amounts of compute.

¹⁶The US military kept the nature of the Manhattan project secret from 1942-46, despite 130,000 people working on it at its peak.[30]

¹⁷This is common in other projects. At the Wuhan Institute of Virology’s BSL-4 lab, researchers are divided into groups that operate independently; CERN scientists working with classified information require SECRET-level clearance. The IAEA monitors scientists at nuclear facilities to identify insider threats early, including scrutinizing finances, psychological fitness, and personal connections.

¹⁸ITER has an isolated control system secured through cryptography.

¹⁹One example of a location for such a facility would be in a country like Singapore because it has technological expertise, technocratic governance, geopolitical stability, and a strong relationship with both the US and China.

2.4 Collective: Supported Internationally, Benefits Distributed To All Signatories

MAGIC would be a multinational, cooperative effort enacted through an international treaty, with the goal of reducing risks from advanced AI systems. The project could be championed by individual powerful states like the U.S. and China, along with an initial coalition of supporting countries. It would, however, also need to make early credible commitments that all other nations can participate. Inclusivity is important to ensure benefits of AI advancement such as scientific breakthroughs are shared equitably among nations, and ensure that a global moratorium is effective in practice.

MAGIC's inclusive membership would be a core incentive for states. States which would otherwise be sidelined in a technology race can meaningfully influence the direction of AI development. MAGIC would further give those states an opportunity to partake in the benefits of safe, powerful AI as developed. Currently, individual developing states and emerging market economies face a significant challenge in reining in the behaviors of powerful global technology corporations on their own. Meanwhile, they are likely to suffer some of the greatest damage of dangerous AI systems. Through inclusive membership, MAGIC is already an improvement compared to the current closed-door, highly informationally asymmetric situation in which a small group of companies continue with fast-paced, unfettered development.

Inclusivity also ensures that MAGIC is viewed as a global effort for humanity's benefit by institutionalizing benefit-sharing into its governance structure. This would involve the creation of a mechanism to distribute the benefits of AI research and development equitably among participating member states. We expect the first outputs of MAGIC to be scientific breakthroughs, such as an extension of the capabilities achieved by AlphaFold. [33] This is because as a secure global research institution, MAGIC can begin by focusing on narrow AI systems focused on solving specific scientific tasks, such as medical breakthroughs like a cure to cancer or longevity. The outputs of these would have to be verified to pose no safety risks before being released. Because this information can be provided at low or zero marginal cost to an additional actor, MAGIC's advancements can be distributed to all members equitably, acting as a further incentive.

After automation of scientific research in narrow domains, if MAGIC's systems are deemed to be sufficiently safe, research could move onto general-purpose, powerful AI systems. Once these controllable systems are achieved, it is likely they could be used to generate enormous economic value, as they will be able to automate a large fraction of economic tasks currently carried out by people. If these are deployed by member states, profits and economic gains from these applications will necessitate equitable economic distribution.

Limitations: In practice, inclusion may pose a direct tradeoff with security. Though a globally inclusive recruitment model is ideal in theory, pragmatic concerns around preventing misuse of advanced AI may necessitate more selective participation criteria in practice. The goal should be assembling a world-class multinational team of AI experts, but with safeguards to keep sensitive research under control. However, this tradeoff can be mitigated through governance structures that enable broad participation in the moratorium, but narrow participation in the most advanced research.

Achieving this balance would require a rigorous selection process for candidates to join the research collective based on AI expertise, while also enabling broad international representation. This selection process would need to consider: a) necessary level of expertise b) conflicts of interest c) security checks d) international representation as possible. This would be operable because there are still strong incentives for nations to join even if they are not given access to MAGIC's internal research, including economic benefits and greater global safety. The broader the participation, the more effective the moratorium would become.

3 Conclusion

MAGIC has precedence as an institution. When scientists discovered chlorofluorocarbons (CFCs) in man-made chemicals were creating a hole in the ozone layer, which posed a significant threat to human health and the environment, the international community agreed to phase out their production. The global community shifted from CFCs to the safer HFC technology through the Montreal Protocol which is considered the UN's first universally ratified treaty.²⁰ On matters of information security, the U.S. military maintained strict secrecy around the Manhattan Project from 1942 to 1946, despite involving approximately 130,000 people at its peak[30]. Additionally, the Trade-Related Aspects of Intellectual Property Rights (TRIPS) Agreement has allowed developing countries to access and produce lifesaving medicines researched and developed in more economically advanced countries. [34] These historical examples show

²⁰Additionally, in the case of unleaded gasoline, policy interventions led to the deliberate restriction and eventual phasing out of the incumbent technology – leaded gasoline. This facilitated the rise and development of a safer alternative, unleaded gasoline, which was initially less competitive but eventually surpassed the performance of leaded gasoline due to targeted advancements and societal preference for a less harmful option. The same can be seen in other industries, such as synthetic colorings in food production.

that in the face of comparable global health or security risks, countries have come together to restrict the proliferation of dangerous technologies and carefully steward their safe and equitable development.

MAGIC would not hinder the majority of current AI development. Narrow systems focused on specific domains like medical imaging or fraud detection would remain unaffected, and could continue to benefit economies and technological progress. While this proposal does not seek to minimize or dismiss the misuse risks of models below the frontier threshold, these models take supervised actions, not autonomous ones, and are largely controllable through human oversight and evaluation.

This paper does not delve into the political feasibility of implementing such an international regime or address the specific legislative strategies and rules that would need to be put in place to enforce a ban on high-capacity AGI training runs. We envision this process could plausibly occur through unilateral enforcement of national moratoriums in the United States and a small group of other nations with oversight of the majority of advanced AI development operations. Collectively, these national moratoria could serve as the foundation for an international treaty, and the basis for creating MAGIC.

This paper does not seek to endorse MAGIC as the only solution, but rather one possible scenario which can effectively halt unchecked advanced AI advancement and significantly reduce large-scale risks of advanced AI. Governance models need not be mutually exclusive. It seems inevitable that advanced AI will require international governance at numerous levels, and other models including a scientific observatory²¹ or an auditing mechanism²² could be complementary to MAGIC. [35]

As AI capabilities accelerate, the risks of uncontrolled development also multiply. This paper outlines MAGIC as one potential governance framework to institute inclusive oversight and responsible development of advanced AI models. By consolidating the most advanced research within a single international body, MAGIC can enact a global moratorium on unauthorized models while enabling safety-focused progress within tight parameters. The aim is not absolute prohibition, but careful stewardship. With coordinated action, the potential of advanced AI models can be harnessed and equitably distributed, while risks are monitored and managed.

References

- [1] *OpenAI Charter*. en-US. Apr. 2018. URL: <https://openai.com/charter> (visited on 09/21/2023).
- [2] Alan Chan et al. “Harms from Increasingly Agentic Algorithmic Systems”. In: *2023 ACM Conference on Fairness, Accountability, and Transparency*. arXiv:2302.10329 [cs]. June 2023, pp. 651–666. DOI: 10.1145/3593013.3594033. URL: <http://arxiv.org/abs/2302.10329> (visited on 09/21/2023).
- [3] Richard Ngo, Lawrence Chan, and Sören Mindermann. *The alignment problem from a deep learning perspective*. arXiv:2209.00626 [cs]. Sept. 2023. DOI: 10.48550/arXiv.2209.00626. URL: <http://arxiv.org/abs/2209.00626> (visited on 09/21/2023).
- [4] Joseph Carlsmith. *Is Power-Seeking AI an Existential Risk?* arXiv:2206.13353 [cs]. June 2022. DOI: 10.48550/arXiv.2206.13353. URL: <http://arxiv.org/abs/2206.13353> (visited on 09/21/2023).
- [5] *Secretary-General Urges Security Council to Ensure Transparency, Accountability, Oversight, in First Debate on Artificial Intelligence* | UN Press. July 2023. URL: <https://press.un.org/en/2023/sgsm21880.doc.htm> (visited on 09/21/2023).
- [6] Stephen J. Dubner. *Satya Nadella’s Intelligence Is Not Artificial*. en. June 2023. URL: <https://freakonomics.com/podcast/satya-nadellas-intelligence-is-not-artificial/> (visited on 09/21/2023).
- [7] Sam Altman, Greg Brockman, and Ilya Sutskever. *Governance of superintelligence*. en-US. May 2023. URL: <https://openai.com/blog/governance-of-superintelligence> (visited on 09/21/2023).
- [8] George Parker. “Rishi Sunak to lobby Joe Biden for UK ‘leadership’ role in AI development”. In: *Financial Times* (June 2023).
- [9] Ian Hogarth. “We must slow down the race to God-like AI”. In: *Financial Times* (Apr. 2023).
- [10] Toby Shevlane et al. *Model evaluation for extreme risks*. arXiv:2305.15324 [cs]. May 2023. DOI: 10.48550/arXiv.2305.15324. URL: <http://arxiv.org/abs/2305.15324> (visited on 09/21/2023).
- [11] LHcb collaboration. en. In: *CERN* (Sept. 2023). URL: <https://home.cern/news/news/knowledge-sharing/lhcb-releases-first-set-data-public>.

²¹Other proposals have called for an Intergovernmental Panel on Climate Change model for AI.

²²Similar to the The International Atomic Energy Agency.

- [12] en. <http://www.iter.org>. ITER is the world’s largest fusion experiment. Thirty-five nations are collaborating to build and operate the ITER Tokamak, the most complex machine ever designed, to prove that fusion is a viable source of large-scale, safe, and environmentally friendly energy for the planet.
- [13] *Defending ITER’s cyberspace*. en. URL: <http://www.iter.org/newsline/-/2843>.
- [14] Andrea Miotti. *Priorities for the UK Foundation Models Taskforce*. en. July 2023. URL: <https://www.conjecture.dev/uk-taskforce-priorities/>.
- [15] davidad. “Compute Thresholds: proposed rules to mitigate risk of a “lab leak” accident during AI training runs”. en. In: (). URL: <https://www.lesswrong.com/posts/Zfk6faYvcf5Ht7xDx/compute-thresholds-proposed-rules-to-mitigate-risk-of-a-lab>.
- [16] Markus Anderljung et al. *Frontier AI Regulation: Managing Emerging Risks to Public Safety*. arXiv:2307.03718 [cs]. Sept. 2023. DOI: 10.48550/arXiv.2307.03718. URL: <http://arxiv.org/abs/2307.03718> (visited on 09/21/2023).
- [17] *Computation used to train notable artificial intelligence systems*. <https://ourworldindata.org/grapher/artificial-intelligence-training-computation>. [Accessed 21-09-2023].
- [18] Jason Hausenloy and Claire Dennis. “Towards a UN Role in Governing Foundation Artificial Intelligence Models”. en. In: *United Nations University* (July 2023). URL: <https://unu.edu/cpr/working-paper/towards-un-role-governing-foundation-artificial-intelligence-models>.
- [19] Mauricio Baker. *Nuclear Arms Control Verification and Lessons for AI Treaties*. arXiv:2304.04123 [physics]. Apr. 2023. DOI: 10.48550/arXiv.2304.04123. URL: <http://arxiv.org/abs/2304.04123> (visited on 09/21/2023).
- [20] Yonadav Shavit. *What does it take to catch a Chinchilla? Verifying Rules on Large-Scale Neural Network Training via Compute Monitoring*. arXiv:2303.11341 [cs]. May 2023. DOI: 10.48550/arXiv.2303.11341. URL: <http://arxiv.org/abs/2303.11341> (visited on 09/21/2023).
- [21] Konstantin Pilz and Lennart Heim. *Compute at Scale - A broad investigation into the data center industry*. en-US. July 2023. URL: <https://www.konstantinpilz.com/data-centers/report>.
- [22] *Lennart Heim on the compute governance era and what has to come after*. en-US. URL: <https://80000hours.org/podcast/episodes/lennart-heim-compute-governance/> (visited on 09/21/2023).
- [23] OpenAI. *AI Compute Trends: A Decade of Accelerated Progress*. <https://openai.com/research/ai-and-compute>. 2018.
- [24] Jaime Sevilla et al. “Compute Trends Across Three Eras of Machine Learning”. In: (Mar. 2022). arXiv:2202.05924 [cs]. DOI: 10.48550/arXiv.2202.05924. URL: <http://arxiv.org/abs/2202.05924>.
- [25] Will Henshall. “4 Charts That Show Why AI Progress Is Unlikely to Slow Down”. en. In: *Time* (Aug. 2023). URL: <https://time.com/6300942/ai-progress-charts/>.
- [26] OpenAI. *AI and efficiency*. May 2020. URL: <https://openai.com/research/ai-and-efficiency>.
- [27] Andrew Lohn and Micah Musser. *AI and Compute*. en-US. Jan. 2022. URL: <https://cset.georgetown.edu/publication/ai-and-compute/> (visited on 09/21/2023).
- [28] *Oversight of A.I.: Principles for Regulation | United States Senate Committee on the Judiciary*. en. July 2023. URL: <https://www.judiciary.senate.gov/committee-activity/hearings/oversight-of-ai-principles-for-regulation> (visited on 09/21/2023).
- [29] Steven Byrnes. “Safety-capabilities tradeoff dials are inevitable in AGI”. en. In: (Sept. 2023). URL: <https://www.alignmentforum.org/posts/tmyTb4bQqi7C47sde/safety-capabilities-tradeoff-dials-are-inevitable-in-agi> (visited on 09/21/2023).
- [30] *Manhattan Project Background Information and Preservation Work*. en. URL: <https://www.energy.gov/lm/manhattan-project-background-information-and-preservation-work> (visited on 09/21/2023).
- [31] Yoshua Bengio. “One of the “godfathers of AI” airs his concerns”. In: *The Economist* (July 2023). ISSN: 0013-0613. URL: <https://www.economist.com/by-invitation/2023/07/21/one-of-the-godfathers-of-ai-airs-his-concerns> (visited on 09/21/2023).
- [32] Yoshua Bengio. *FAQ on Catastrophic AI Risks*. en-CA. June 2023. URL: <https://yoshuabengio.org/2023/06/24/faq-on-catastrophic-ai-risks/>.
- [33] John Jumper et al. “Highly accurate protein structure prediction with AlphaFold”. en. In: *Nature* 596.7873 (Aug. 2021), pp. 583–589. ISSN: 1476-4687. DOI: 10.1038/s41586-021-03819-2. URL: <https://www.nature.com/articles/s41586-021-03819-2>.
- [34] Ellen’t Hoen. “TRIPS, Pharmaceutical Patents, and Access to Essential Medicines: A Long Way from Seattle to Doha”. In: *Chicago Journal of International Law* 3 (2002), p. 27. URL: <https://heinonline.org/HOL/Page?handle=hein.journals/cjil3&id=33&div=&collection=>

- [35] Lewis Ho et al. “International Institutions for Advanced AI”. In: (2023). DOI: 10.48550/arXiv.2307.04699. arXiv: 2307.04699 [cs.AI]. URL: <http://arxiv.org/abs/2307.04699>.