# Mapping IAEA Verification Tools to International AI Governance: A Mechanism-by-Mechanism Analysis

**Christina Krawec**
Earthnote LLC / ERA:AI Fellowship
Washington, DC 20009
cikrawec@gmail.com

## Abstract

*Verification* – the process used to confirm whether actors are complying with agreed-upon rules or obligations – is a central piece of the AI governance puzzle. The International Atomic Energy Agency (IAEA), the United Nations's "nuclear watchdog," safeguards – through its own verification mechanisms – nuclear technology to ensure it is not diverted for military purposes. Therefore, understanding the current international nuclear safeguards regime can inform and potentially guide examination of options for future international AI verification. This report assesses the IAEA at a granular, technical level, conducting a systematic analysis of 24 verification tools used within the Department of Safeguards for State nuclear program evaluations, assessing each for potential applicability to AI-related verification and addressing the technical specifications and challenges of implementation.

## 1. Introduction

At present, the speed of AI advancement outstrips current governance mechanisms. Developing common understanding among all AI stakeholders – from governments to commercial developers to independent researchers – is imperative to the successful integration of assessment and verification measures. The International Atomic Energy Agency (IAEA) provides a mature, technical precedent for nuclear-focused, multilateral verification. This research does not argue for or against an "IAEA for AI" – a topic widely covered in existing literature – but assesses specific verification tools and mechanisms that illustrate historical success within the IAEA nuclear safeguards system, analyzing which could be applicable to AI governance. While many IAEA tools depend on traditional measurements of nuclear material, several – such as open source analysis and voluntary reporting – are not reliant on measures of radioactivity or fissile material. Therefore, a select number of tools could inform short-term, feasible policy changes that various stakeholders could implement in the near future.

### 1.1. Verification for AI: An overview

Successful verification processes ensure that safety commitments (either voluntary or obligatory) are being honored. In an AI governance context, they can enhance trust between developers and other stakeholders, serve as an early warning system for dangerous activity, potentially detect diversion of high-risk hardware, and develop additional groundwork for future international agreements. Without verification, the field of AI safety risks strain from mistrust and unregulated development of systems potentially vulnerable to catastrophic misuse.

AI-focused verification may take many forms, and would involve establishing evidence-based processes to confirm whether an AI developer is meeting standards for creation or dissemination. Existing literature discusses at length the challenges and opportunities for implementing verification measures at both the institutional and technical levels, exploring approaches such as the use of tamper-proof AI chips at data centers, human inspectors for on-site evaluations, and cryptographic signatures for data authentication (Baker et al., 2025). This research intends to add to these discussions, providing additional tools and frameworks through an IAEA-focused lens and exploring whether lessons from the nuclear safeguards State evaluation toolkit can inform the design of trustworthy verification mechanisms for advanced AI.

### 1.2. Why use the IAEA?

The International Atomic Energy Agency – the United Nations's "nuclear watchdog" – supports the use of civilian nuclear technology, safeguards nuclear technology to ensure it is not diverted for military purposes, and develops safety and security standards for nuclear technology. The concept of an "IAEA for AI" has been thoroughly discussed among policymakers, AI company leadership, and technologists, and has been often dismissed as an unworkable analogy as a model for creating a formal regulatory body focused on AI (Horowitz, 2024). These rejections are often made at a high, organizational level, citing general differences in the technology, enforceability, or the role of State actors. This report instead assesses the IAEA at a granular level: it conducts a systematic analysis of IAEA verification tools used within the Department of Safeguards for State nuclear program evaluations, assessing each for potential applicability to AI verification and highlighting the capacity for the integration of these tools in the future at varying levels. The IAEA nuclear safeguards regime specifically – and its array of mechanisms – has shown significant success in the monitoring of nuclear technology development worldwide, and this research provides shared language and a potential toolkit that a range of stakeholders can use to design future verification systems for AI.

## 2. Methodology

The International Atomic Energy Agency's Department of Safeguards is responsible for assessing whether a State has adhered to its international nuclear safeguards-related obligations, ensuring that nuclear technology and material are being used only for civilian purposes and are not being diverted for nuclear weapons programs. This ongoing process of safeguards-related information analysis and the drawing of relevant conclusions is known as *State evaluation* (IAEA, 2022).

The IAEA Department of Safeguards conducts State evaluation using three main information streams: State-provided information, information from safeguards activities, and other relevant information. *State-provided information* is the information that a State – either voluntarily or by obligation – reports to the IAEA. *Information from safeguards activities* is data collected by the IAEA itself, generally from inspections or other in-field activities. *Other relevant information* includes additional uncategorized information from open sources, IAEA databases, or information voluntarily provided by States about other States.

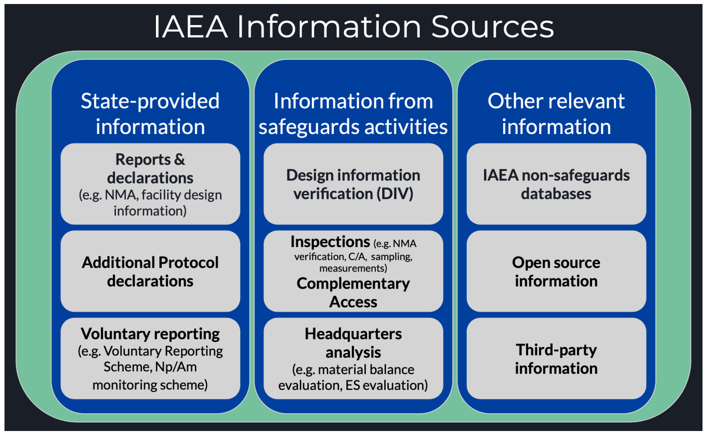


Figure 1: IAEA information streams

These data sources – and associated verification tools – are used in tandem to draw conclusions about State nuclear programs on a roughly annual basis, determining whether a State is in compliance with its IAEA obligations. What makes the IAEA functional is this cross-corroboration of data built into the system: data is compared across the information streams, the built-in redundancy and layering providing a unique method for ensuring information accuracy. For the purposes of this report, these three data streams were converted into a table, providing a comprehensive overview of the specific verification tools used within each category (see Appendix A). These tools were listed with their purpose and technical specifications; following this, potential analogies to AI governance were added.

The IAEA proves that, historically, a variety of overlapping verification methods are useful for determining whether a country is illicitly developing nuclear technology for military purposes. In Appendix A, each of the collected IAEA verification tools was mapped to an equivalent AI analogue for two categories: model development and compute. Each tool was evaluated individually, assessing whether an AI-related analogue could exist. Four categories of feasibility were applied to the tools for evaluation:

- *Technical Feasibility:* How technically complex would implementation of a specific tool be?
- *Institutional Feasibility:* Would the tool's integration require multilateral cooperation, or could it be developed unilaterally?
- *Political Feasibility:* Which stakeholder(s) would be interested, willing, and able to implement the tool?
- *Temporal Feasibility:* Could the tool be implemented in the short term (i.e. within the next year) or would its complexity require multiple years or decades before fruition?

## 3. Discussion of selected verification tools

Appendix A is an in-depth analysis of each IAEA verification tool and its potential applicability to AI governance, addressing the technical specifications and challenges of implementation. Various stakeholders (see Figure 2) may find specific verification methods promising for the future development of international AI governance norms and regimes. Below are examples of three tools that show different levels of promise in an AI governance context.

Table 1: Overview of three selected IAEA verification tools mapped to AI governance

| IAEA Verification Tool | Use in Nuclear Safeguards | AI Analogue & Hypothetical Implementation Scenario | Stakeholder(s) of Interest | General Feasibility Assessment |
|---|---|---|---|---|
| Open Source Information System (OSIS) | Centralized database of open source material relevant to nuclear safeguards State evaluation | *"OSIS for AI":* A database collects and displays news reports, government publications, commercial announcements, information leaks, scientific journal articles, social media posts, and analytical reports relating to model development and deployment. | *Independent researchers, international organizations*, or *governments* could develop a database for communal analysis and completeness of data. | High - Could be implemented outside regulatory or bureaucratic restrictions |
| Voluntary Reporting Scheme (VRS) | States voluntarily share information such as trade data, R&D programs, or expanded site access for inspections | Similar to the established publication of system cards, model developers choose to publicly "declare" additional information such as compute usage summaries, red-teaming participation, and announcements of third-party audit participation. | *Frontier developers* might implement VRS techniques to avoid further limitations from government regulation and "prove" they are taking safety seriously. | Medium - Some motivation already exists for voluntary reporting; further norm establishment necessary |
| Physical Inventory Listing (PIL) | States are required to report a complete, itemized list of all nuclear material at a facility of interest | *Hardware Inventory Reports:* Developers could submit a detailed list of all hardware in use at a data center, including chips, host CPUs, server motherboards/nodes, memory (RAM/HBM), data storage units, networking hardware, and cooling systems. | *International organizations*, with the requisite governing power, would benefit from formal hardware declarations made by frontier developers. | Low - Requires the existence of a formal, centralized verification body |

### 3.1. The Open Source Information System

#### 3.1.1. The IAEA's Open Source Information System (OSIS)

Open sources are used as supplementary information for Safeguards evaluation, corroborating State-declared and IAEA-collected data and validating that all information is consistent. One particular tool of unique applicability to the AI governance field is the IAEA's Open Source Information System (OSIS). The Open Source team within the IAEA Department of Safeguards utilizes OSIS, an internally-developed and centralized database, to automate the collection and processing of open source information (Sköld et al., 2018). A selection of curated keywords related to the nuclear fuel cycle and relevant technologies and materials is used for web scraping: OSIS automatically retrieves news reports, government publications, commercial announcements, social media entries, blog posts, analytical reports, and scientific journal articles of safeguards relevance. In addition, OSIS can collect non-textual media such as photos, videos, or audio data. These items are added to OSIS, and the system is refreshed twice daily with new waves of collected entries. Each entry is tagged automatically with the following information:

- Country or countries of interest;
- Facility or facilities of interest (if applicable);
- Publisher/author (if applicable);
- Relevant state of the nuclear fuel cycle; and
- An AI-generated ranking predicting how safeguards-relevant the content is (Baude et al., 2022).

Following the collection of the entries, analysts must review the entries manually to either confirm or discard them and fix any erroneous tags. All analysts have access to OSIS and keep it up to date. OSIS upkeep is extremely part-time and analysts might spend anywhere from five minutes to an hour each day on curation, depending on their responsibilities. Each analyst is assigned an individual portfolio of countries and is required to maintain an accurate set of entries in OSIS for these assigned States. The collected open source items are then used for annual reviews of State nuclear programs.

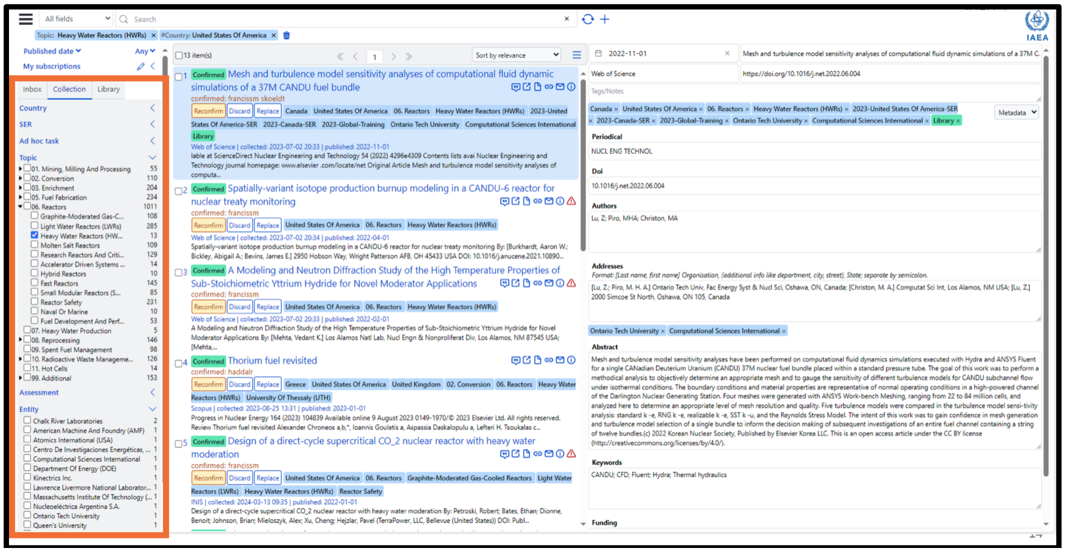

Figure 2: Screenshot of OSIS. The middle column lists the collected articles and their relevant tags, the right column shows metadata for the highlighted article, and the left column displays categories and tags (Baude, 2024).

#### 3.1.2. An "OSIS for AI"

A database for AI-relevant open source information could be a feasible next step for the AI governance community. A centralized repository of AI-focused open source information, made public, could consist of model release announcements, research publications, corporate disclosures, research collaborations, conference proceedings, social media activity, patent filings, and shipping information on compute hardware. Tags in this database could include country of interest, company, specific AI model of interest, and relevant hardware.

This is an attractive short-term step for a variety of reasons: Firstly, creating a public database does not require state or developer consent to be implemented. With proper funding, it could be created by individual researchers much more quickly and independently than relying on the bureaucratic State government or international organization environment. This tool utilizes publicly-available information and sidesteps any issues of access or sovereignty. It is a low-barrier entry to verification work, and, if made public, would make governance and verification research more streamlined and accessible for the wider AI community. This would be a credible, independent layer of societal verification operating outside the scope of government.

### 3.2. The Voluntary Reporting Scheme

#### 3.2.1. The IAEA's Voluntary Reporting Scheme (VRS)

IAEA Member States may enter into a range of agreements with the Agency, such as a Comprehensive Safeguards Agreement (CSA), Voluntary Offer Agreement (VOA), Additional Protocol (AP), and Small Quantities Protocol (SQP). The nature and implementation of these agreements is outside the scope of this report, but all establish formal mechanisms by which a State is obligated to report specific quantities of material and types of technology to the IAEA. These vary by State. However, the IAEA's Voluntary Reporting Scheme (VRS) provides a holistic, informal system by which States can share additional information that falls outside the scope of their obligatory declarations. Depending on the type of agreement in force, VRS information could include the following:

- Imports and exports of natural uranium ore concentrates (which have not typically been converted to forms suitable for fuel fabrication, which would be covered under, for example, an Additional Protocol agreement);
- Supplementary trade data, to include shipping routes; and
- Transactions of material falling outside minimum obligatory thresholds.

In addition, the VRS allows for greater frequency of reporting. For example, under an AP, a State must make formal reports once per year. However, with the expanded VRS, a State may choose to send reports much more frequently (i.e. on a quarterly or monthly basis). This gives the IAEA greater visibility and the ability to assess information and take follow-up actions in a more timely manner. All of this data, if reported to the IAEA, aids in comprehensive assessments of State nuclear material and technology development.

### 3.2.2. A "VRS for AI"

In the same way that States are encouraged to report to the IAEA trade data, nuclear material quantities, and technology falling outside the scope of their legal obligations, frontier AI developers might be encouraged by the international community informally to publicly and voluntarily report additional information. At present, there would be significant differences between a "VRS for AI" and the IAEA's scheme: A VRS for AI may not fall under the auspices of a formal governing body; instead, it could take the form of norm development and general advocacy for commercial firms to publicly report elements of their models, showing goodwill and verifying that they are developing their systems with a clear focus on preventing misuse (i.e. the development of chemical, biological, radiological and nuclear weapons).

Developers have already begun to implement these norms. AI model developers regularly publish system cards outlining their model's updated capabilities, features, risks, and evaluation practices. These system cards have, since 2019, become widespread among the frontier AI companies, and are an example of transparency best practices continuing to be established within the field. In addition, the recent European Union's General Purpose AI Code of Practice outlines voluntary measures that model developers can use to demonstrate compliance to the EU AI Act. A "VRS for AI", thus, could take the form of an informal agreement between companies to, for mutual trust-building, to expand the quantities of information publicly reported. Types of additional information to be reported could include:

- Incident reports and known cases of misuse;
- Announcements of participation in third-party audits; or
- Specific results of red-teaming activities.

A "VRS for AI" thus was assessed to have a mid-tier level of feasibility. While its existence does not rely on the creation of a formal regulatory body, and could be created informally among frontier AI companies on a short-term basis, this scheme does inherently rely on the willingness of AI companies to continue to publicly declare information about their models. This is somewhat feasible in the short term, but depends on the consent and buy-in of these firms. AI researchers can continue to advocate for expanded transparency from a research and lobbying perspective. In addition, State governments (particularly the US) can use warnings of impending legislation, which may be a motivator for AI companies to expand system cards: these companies would prefer to prove they are verifying their systems internally to avoid the possibility of additional restrictive legislation being passed.

### 3.3. Physical Inventory Listings

#### 3.3.1. The IAEA's Physical Inventory Listings (PILs)

A Physical Inventory Listing (PIL) is a report that provides a complete, itemized list of all nuclear material stored at a facility or material balance area (MBA) during a specific time period. A Facility must count and measure all material, record this data, and submit it to the IAEA. Information required in a PIL includes:

- Material type (i.e. plutonium, low-enriched uranium, or spent fuel);
- Quantity of material;
- Specific location;
- Form (i.e. pellets, powders, gases, or solutions); and
- Any changes noted since the last report.

Information from a PIL is then integrated in the wider and more comprehensive Material Balance Evaluation (MBE), which is the analytical process the IAEA uses to verify that all declared nuclear material is accounted for. This set of calculations and analyses detect potential loss or diversion and compare State declarations with data collected physically from inspections or site visits (IAEA, 2022). The PILs are integrated into MBE calculations, with the final output being a Material Balance Report (MBR), which is a formal document outlining current materials and any changes at a State or facility of interest.

#### 3.3.2. "PILs for AI"

A PIL for AI might take the form of a "Hardware Inventory Report" in which private companies – either the frontier developers themselves or intermediary entities such as data center operators – are required by a centralized authority to submit a detailed list of all hardware in use at a data center. This could include chips, host CPUs, server motherboards/nodes, memory (RAM/HBM), data storage units, networking hardware, and cooling systems. Crucially, an HIR would not be a public-facing document. Mirroring how State declarations to the IAEA are kept confidential, HIRs could be submitted privately to a centralized verification authority: an AI developer might be more willing to participate in a verification scheme that ensures the confidentiality of their systems. However, this tool was assessed to have a low level of feasibility at present. Its success heavily relies on a centralized reporting regime and the existence of a monitoring authority with the power to mandate that a firm reports this data. The stakeholders who would be involved in HIRs would be the declaring entity – the AI model developers – and the verifying entity, which would take the form of a centralized international organization responsible for verification. In short, this is a conceptually sound but low-feasibility AI analogue for the short-term, as compliance would require the existence of, for example, an "IAEA for AI" – a centralized AI verification body with the power to compel AI developers to declare this type of information.

## 4. Policy recommendations

In this section, the author provides several options moving forward based on the efficacy and feasibility of applying IAEA-inspired verification tools to the AI governance field.

### 4.1. Tool-specific recommendations

1) ***Create an Open Source Information System (OSIS) for AI.*** This appears to be one of the strongest candidates for a short-term and high-feasibility verification tool that could add to the public sphere of AI governance research work. An independent AI research organization, with the proper resources and funding, could develop an OSIS focusing on collecting, tagging, and otherwise centralizing all open source information related to model updates or compute-related milestones. While coverage would be unavoidably incomplete, its creation would improve situational awareness, help identify trends in model release best practices, and report on safety and security incidents.

2) ***Building off the Voluntary Reporting Scheme, advocate for frontier developers to publicly declare information about model updates.*** While this was assessed to have mid-level feasibility, due to the required buy-in of model developers, there are at least two paths forward to promote commercial transparency:
    a) The higher-feasibility recommendation is a less formal route, building civil society pressure at a grassroots level. The AI researcher community and governments could encourage model developers to expand the information included in system cards and advocate for the publication of additional information on training runs, updates, and significant model changes.
    b) The lower-feasibility option would be for AI model developers to establish a consortium to communally agree and publicly commit to formalizing and normalizing the publication of model data. While this would require motivation and trust from developers, it would give them the power to communally agree upon disclosure standards and potentially avoid further regulation development at the State level. Because commercial entities will prioritize their own best interests, multi-stakeholder participation advocating for voluntary reporting would be ideal.

### 4.2. General recommendations

1) ***Promote overlapping information streams for verification.*** Following the IAEA's three-information-stream model for data collection and analysis, the AI governance community should focus on verification using three categories of information: open sources, voluntary disclosures, and third-party activities such as measurements, inspections, or compute monitoring. In the future of AI verification, this approach provides greater confidence in data assessments, as information can be corroborated across the different streams.

2) ***Establish a global verification research initiative.*** Several initiatives within the UN and wider international organization system already focus on AI safety and ethics. However, none is focused specifically on AI verification. A new verification-centric initiative could be launched within the architecture of one of these existing initiatives. For example, this might fit under the Organisation for Economic Co-operation and Development (OECD)'s Global Partnership for AI (GPAI). Rather than develop an entirely new verification treaty body, a multi-stakeholder initiative focused on verification research might be a feasible step towards institutional legitimacy and global coordination. In addition, if a new organization were to be established, it could function as an alliance between private and public organizations, promoting collaboration in particular between the US and China. Within any of these frameworks, research could explore compute measurement, model behavior auditing, and reporting protocols. This entity could also work to develop formal standards for verifying AI development, which also do not exist: verification only truly works within a system that has established formal benchmarks.

## 5. Opportunities for further research

This report has outlined and provided a guide to specific tools that inform the future of AI verification. Additional research could involve conducting case studies of more of these different tools, assessing them in further detail for their applicability to AI governance. Future research could involve a deep dive into one specific tool. For example, one project could involve the actual implementation and development of an "OSIS for AI", using IAEA open source tools and processes to collect, analyze, and disseminate up-to-date public information about AI developments, sharing findings with the wider public in a centralized database or other forum.

Another branch of research could involve exploring the IAEA's Technical Cooperation (TC) Programme, which provides IAEA Member States access to nuclear technology, material, and training in exchange for remaining in good standing with the Agency and complying to their nuclear safeguards obligations. There may be an analogue for AI here as well.

In addition, further research could involve the development of an OS case study specifically studying the efficacy of advanced remote sensing tools for the assessment of data center operations. This could include traditional electro-optical image assessments, synthetic aperture radar (SAR) imagery analysis, and thermal imagery heat signatures estimating energy usage.

## 6. Conclusion

Verification is a necessary foundation for future enforcement and trust-building for AI governance. While no international regulatory body for AI currently exists, understanding the existing international nuclear safeguards regime at a granular, technical level provides an overview of the options moving forward for the future of international AI verification. The AI governance community should develop a solid verification toolkit and can be inspired by the IAEA's methods and its accompanying history of successful implementation.


## Acknowledgements

This research was supported by the ERA Fellowship and mentored by Kevin Kohler and Belinda Cleeland. The author would like to thank the ERA Fellowship for its financial and intellectual support.


## References


Allen, Greg. *AI Chips: What They Are and Why They Matter*. Center for Security and Emerging Technology (CSET), Georgetown University, January 2023. https://cset.georgetown.edu/publication/ai-chips-what-they-are-and-why-they-matter/

Anderljung, Markus and Lennart Heim. *Computing Power and the Governance of AI*. Centre for the Governance of AI, 2023. https://www.governance.ai/analysis/computing-power-and-the-governance-of-ai

Baker, Mauricio. *Nuclear Arms Control and Lessons for AI Verification Treaties*. Stanford University, 2023. arXiv. https://arxiv.org/pdf/2304.04123

Baker, Mauricio; Kulp, Gabriel; Marks, Oliver; Brundage, Miles; and Lennart Heim. *Verifying International Agreements on AI: Six Layers of Verification for Rules on Large-Scale AI Development and Deployment*. Stanford University, July 2025. arXiv:2507.15916. https://arxiv.org/pdf/2507.15916

Baude, Stephane. *The IAEA's Innovative Approach to Address the Challenges in the Collection and Analysis of Safeguards Relevant Information*. Presentation slides, INMM 2024 Annual Meeting, Portland, Oregon, July 4, 2024. https://event.fourwaves.com/65inmm/abstracts/5ff9bed9-411e-4aad-b994-6f7a306807e3/slides

Baude, S., P. Schneeweiss, and T. Stojadinovic. *The IAEA's Innovative Approach to Address the Challenges in the Collection and Analysis of Safeguards Relevant Information*. Vienna: International Atomic Energy Agency, 2022. https://cdn.fourwaves.com/static/media/formdata/b12613f9-28b7-4568-a069-5c7cfeff3c8e/8d5ee6c8-9afc-4bf0-812c-2f57cf07b9d9.pdf

Cha, Seokki. *Towards an international regulatory framework for AI safety: lessons from the IAEA's nuclear safety regulations.* Humanities and Social Sciences Communications, April 12, 2024. https://www.nature.com/articles/s41599-024-03017-1

Dennis, Claire, and Jason Hausenloy. *Towards a UN Role in Governing Foundation Artificial Intelligence Models*. United Nations University, September 2023. https://unu.edu/sites/default/files/2023-09/Working%20paper%20-%20Towards%20a%20UN%20Role%20in%20Governing%20Foundation%20Artificial%20Intelligence%20Models_0.pdf

Harack, Ben; Trager, Robert F.; and Reuel, Anka. *Verification for International AI Governance*. Oxford Initiative for AI Governance and Inspection (AIGI), July 2025. https://aigi.ox.ac.uk/wp-content/uploads/2025/07/Verification_for_International_AI_Governance.pdf

Hobbs, Christopher. *The Nuclear Governance Model Won't Work for AI*. Chatham House, June 27, 2023. https://www.chathamhouse.org/2023/06/nuclear-governance-model-wont-work-ai

Horowitz, Michael C., and Lauren A. Kahn. *Nuclear Non-Proliferation Is the Wrong Framework for AI Governance*. AI Frontiers, 2024. https://ai-frontiers.org/articles/nuclear-non-proliferation-is-the-wrong-framework-for-ai-governance

IAEA (International Atomic Energy Agency). *IAEA Safeguards Glossary: 2022 Edition*. IAEA, 2022. https://www-pub.iaea.org/MTCD/Publications/PDF/PUB2003_web.pdf

IAEA (International Atomic Energy Agency). *Safeguards Techniques and Equipment: 2011 Edition*. IAEA, 2011. https://www-pub.iaea.org/MTCD/Publications/PDF/nvs1_web.pdf

Krawec, Christina. "Open Source Collection and Analysis for IAEA Safeguards." Presentation at *Nonproliferation and International Nuclear Safeguards in a Changing World*, Brookhaven National Laboratory, June 26, 2025.

Law, Harry, and Lewis Ho. *The IAEA's Dual Mandate and AI*. Nature Reviews Physics, 2023. https://www.nature.com/articles/s42254-023-00670-4

LinkedIn – Control AI. "Demis Hassabis Warns of the Risk of Bad Actors." July 2024. https://www.linkedin.com/posts/controlai_demis-hassabis-warns-of-the-risk-of-bad-actors-activity-7339317760813944832-KK7O/

Pilz, Konstantin and Lennart Heim. *Compute at Scale: A Broad Investigation into the Data Center Industry*. Cornell University, November 2023. arXiv:2311.02651. https://arxiv.org/abs/2311.02651

Reuel, Anka and Trond Arne Undheim. *Generative AI Needs Adaptive Governance*. 2024. arXiv:2406.04554. https://arxiv.org/abs/2406.04554

Reuel, Anka; Soder, Lisa; Bucknall, Ben; and Undheim, Trond Arne. *Position Paper: Technical Research and Talent is Needed for Effective AI Governance*. 2024. arXiv:2406.06987. https://arxiv.org/abs/2406.06987

Sköld, T.; Courbon, F.; and Spence, K. *OSIS 2.0: Optimizing Analyst-Driven Automation of Open Source Information Collection and Processing for Safeguards State Evaluation*. Vienna: International Atomic Energy Agency, 2018. https://media.superevent.com/documents/20181030/0c98adb085f34db27ceaff2468a60b72/id059_skoeld_paper.pdf

Stewart, Ian J. "Why the IAEA Model May Not Be Best for Regulating Artificial Intelligence." The Bulletin of the Atomic Scientists, June 2023. https://thebulletin.org/2023/06/why-the-iaea-model-may-not-be-best-for-regulating-artificial-intelligence/

Trager, Robert; Reuel, Anka; Bucknall, Ben; Casper, Stephen; Fist, Tim; Soder, Lisa; Aarne, Onni; Hammond, Lewis; Ibrahim, Lujain; Chan, Alan; Wills, Peter; Anderljung, Markus; Garfinkel, Ben; Heim, Lennart; and others. *Open Problems in Technical AI Governance*. 2024. https://cdn.governance.ai/Open_Problems_in_Technical_AI_Governance.pdf

Wasil, Akash. "Do We Want an 'IAEA for AI'?" Lawfare, 2024. https://www.lawfaremedia.org/article/do-we-want-an--iaea-for-ai

Winecoff, Amy A. and Miranda Bogen. *Improving Governance Outcomes through AI Documentation: Bridging Theory and Practice*. 2024. arXiv:2409.08960. https://arxiv.org/abs/2409.08960