

Community of Practice Recommendations for Assessing AI Systems as Digital Public Goods

Author: Lea Gimpel, Director of AI and Country Engagement (Fellow), DPGA Secretariat

The Community of Practice (CoP) on AI systems as digital public goods, co-hosted by the Digital Public Goods Alliance Secretariat and UNICEF, developed non-binding recommendations for the [DPG Standard Council](#) on updating the DPG Standard for assessing AI systems.

UNICEF and the Digital Public Goods Alliance (DPGA) acknowledge AI's transformative potential, particularly in attaining the Sustainable Development Goals (SDGs), and are committed to ensuring its responsible development and application for the public interest. The recently adopted [UN Global Digital Compact \(GDC\)](#) explicitly mentions digital public goods (DPGs), including open AI models, as a means “to empower societies and individuals to direct digital technologies to their development needs.” The [Ministerial Declaration of the G20 Digital Economy Working Group](#) published in September 2024 also recognises the need for inclusive access to AI technology, particularly for developing countries, and underscores the importance of harnessing open source technologies and open data to deliver AI benefits at scale.

These initiatives build on previous efforts, including the [UN Secretary General's Roadmap for Digital Cooperation](#), which emphasised the importance of digital public goods. The roadmap defines DPGs as open source software, open standards, open data, open AI models, and open content collections that adhere to privacy and other applicable best practices, do no harm by design, and are highly relevant for achieving the SDGs. To put this into practice, the DPG Standard sets the baseline requirements for a solution to be recognised as a DPG.

Against the backdrop of heightened international interest in AI and DPGs, discussing potential changes to the DPG Standard is timely to ensure that AI DPGs can be leveraged responsibly and to the greatest extent. For this reason, the CoP, which brought together experts from diverse sectors and geographies, started work in 2023 to grapple with the evolving discourse on defining open source AI and propose recommendations for evolving the DPG Standard to encompass AI systems better. While the group of more than 50 individuals didn't reach full consensus on every discussion point, this post attempts to capture and summarise key discussion areas and the recommendations submitted to the DPG Standard Council¹ for review and consideration in August 2024.

¹ The DPG Standard Council is the governing body responsible for guiding the development and maintenance of the DPG Standard. It facilitates the process by reviewing, organising, and consolidating feedback from the DPG community, ensuring the standard is developed through a well-planned and transparent approach.

Openness in AI Systems: Binary vs. Tiered

The CoP started by discussing the components of AI systems and examined different release practices for AI models and inspired by, amongst others, [Irene Solaiman's work on a gradient release model for generative AI](#). Core considerations were captured in the [CoP's mid-report](#) published in August 2023. This discussion included debating "meaningful openness" in AI in relation to the DPG definition and its objectives and exploring a tiered approach to the openness of AI systems' components. It also included recognising the current ambiguities surrounding the term "open," especially concerning data extractivism, data colonialism, and the economics of open data.

The idea for a tiered approach to the openness of AI system components was described in the CoP's [blog post from October 2023](#). Core strengths of such an approach include acknowledging current industry practices and emphasising the importance of open AI training data. As the CoP process took place, several leading open source organisations, including the Linux Foundation, Mozilla, and the Open Source Initiative, worked on defining openness in AI systems and an [open source AI definition](#), respectively. The [Linux Foundation's Model Openness Framework](#) (MOF) comprises 16 components. It describes three classes of openness, ranging from an open model and open tooling to open science, depending on the number of open components. [OSI's Open Source Definition 1.0](#) for the sits around the open tooling class, depending on the quality and quantity of available components.

Recommendations for the DPG Standard

There was much discussion on how this should inform a recommendation for the DPG Standard. The DPG definition is rooted in the aim to make DPGs broadly accessible for reuse and adoption, building on the open source definition, championing SDG relevance, and ensuring do no harm by design. The DPGA, as the custodian of the DPG Standard, is also committed to supporting the open movement by establishing stringent rules that support software and AI developers and content creators alike. However, the DPG Standard is not rooted in open science practices but focuses on the applicability of DPGs in real-world contexts that benefit the attainment of the SDGs. Because the Standard is built as a binary - either a product is recognised as a digital public good because it fulfils the DPG Standard indicators or fails this test - the gradient approach was less applicable. For these reasons, the CoP decided to focus on a binary approach of defining AI model releases as either conformant with OSI's open source AI definition or nonconformant. Doing so would ensure the DPGA Secretariat can operationalise this distinction in the DPG Standard. However, some participants also felt that alongside the binary decision, there is value in defining the multiple components of AI models according to the Model Openness Framework and having transparent documentation of what open (including open licensing) means for each individual AI component.

Data Component Requirements

The subject of data was an important part of the CoP's discussions. It was informed by [OSI's draft open source AI definition](#), which outlined the openness requirements of the three main components of code, model and data. The requirements describe, for each of these components, the subcomponents that must be made openly available for an AI model to be recognised as compliant with the definition. For the data component, the requirements do not include a strict open data provision for the complete training data set but highlight in the [annexed checklist](#) that open data is the preferred subcomponent for data. Suppose the complete training data can't be distributed openly. In that case, substantial data information for subsets that can't be shared and the full source code and specifications for data processing, model training, and inference are required to facilitate replication. Many participants highlighted the critical importance of having as much data open as possible. They raised concerns that an AI system can't be considered truly open, and a valuable advocacy opportunity might be lost by not requiring the training data to be open. For that reason, in addition to OSI's specifications, the CoP explicitly recommends that as much data is made available as possible, including open subsets of the training data, by providing a data sample, synthetic data modelled on the original dataset or instructions for the gated access to training datasets used. The CoP also suggested the DPGA should encourage developers to submit their open datasets for DPG recognition under the "open data" category.

Openness and Responsibility: the CoP's Take on RAIL Licenses

In addition to their relevance for the SDGs, one aspect distinguishing DPGs from the wider pool of open source solutions is that they must demonstrate how they anticipate, prevent, and do no harm by design. This requirement is addressed in the DPG Standard indicators seven ("Adherence to Privacy and Applicable Laws") and nine ("Do No Harm by Design"). However, licensing (Indicator 2, "Use of approved open licenses") also came up as a topic of discussion in this context.

As laid out in the mid-term report, the CoP also considered recommending responsible AI licenses (RAIL) such as the openRAIL license family or BigScience BLOOM RAIL 1.0 to be recognised as acceptable within the DPG Standard, especially given its core value of do no harm by design. These licenses combine openness with purpose limitations to help safeguard open AI models against harmful use. However, some CoP participants pointed out that purpose limitations contradict the freedom to use an open source artefact freely and without any restrictions. Another concern raised by CoP members was that such licenses could lead to unintended harm by giving model developers a false sense of security, leading them to believe they can replace due process to ensure low-harm AI with a responsible use license. Imposing purpose limitations on a public good while effectively contributing little to AI safety was deemed counterproductive to the DPG Standard's objectives.

Ultimately, the CoP recognised the well-meaning intentions of responsible AI license developers but concluded that licenses as a tool are not fit for the purpose of inhibiting harm, given their lack of enforceability. [Alternative risk mitigation strategies](#) are being

researched and explored, especially for foundation models. Therefore, the CoP decided not to recommend that responsible AI licenses be included within the DPG Standard.

“Do no harm” Recommendations for the DPG Standard

For an update to indicators 7 and 9, which capture the do no harm principle in broader terms, the CoP decided to build on the UNESCO “[Recommendation on the Ethics of Artificial Intelligence](#)” specifically on the principles outlined in section III.2.

The CoP recommended adding to indicator nine a requirement for AI model developers to provide evidence of model testing for bias, fairness, security, resilience, transparency, accountability and that appropriate mitigation measures were implemented if potential harm was identified. Furthermore, this section would indicate that developers have taken measures to explain model outputs and ensure human oversight. A critical discussion point regarding responsible AI development practices touched upon the question of which requirements should be made mandatory versus what should be a recommended practice as part of the DPG Standard, keeping in mind the limited resources, especially of smaller developer teams. Ultimately, the CoP proposed making several risk mitigation measures mandatory, including an AI risk assessment for the use case(s) for which the model was developed, a responsible use guide, and a plan for utilising AI safety by design principles.

The CoP also recommended that the DPGA Secretariat provide templates for these recommendations to ease implementation and allow for comparability between submissions applying for DPG recognition, given the many different responsible AI frameworks and toolboxes. The CoP was also mindful of the environmental impact of AI systems and suggested adding a carbon footprint estimation for model training as a mandatory requirement, for instance, by using established footprint estimators such as [Code Carbon](#) or the [ML CO2 Impact Estimator](#). It was motivated by the desire to provide higher transparency and encourage efficient model development. However, it was also noted that the carbon footprint heavily depends on the energy sources used in the country where training occurs (for example, countries whose energy mix is primarily based on renewables would fare better than those using coal plants) and the hardware used. Older hardware is usually less energy efficient, which might disadvantage smaller developers and those in under-resourced contexts. Thus, carbon footprint estimations are of limited use, and energy consumption as a benchmark should also be considered - both for training and inference, given that the Standard encourages the adoption and reusability of DPGs.

Next Steps: the DPG Standard Update Governance Process

The CoP’s work concludes with these recommendations to the DPG Standard Council. The DPG Standard Council will consider these contributions, in line with the [DPG Governance framework](#), and conduct its own deliberations, considering technical feasibility, alignment with organisational goals, and feedback from the DPGA leadership and community to reach a proposal for updating the DPG Standard. This may take the form of mandatory requirements and recommended practices. Far-reaching changes will be transparent on the [DPG Standard’s GitHub page](#) and open for public comment for four weeks.

Given the dynamics of AI development, we also expect the need to revisit the revised DPG Standard regularly. For this reason, the DPGA Secretariat will continue to work with experts at the intersection of the openness agenda and responsible AI to inform our work. Specifically, the DPGA Secretariat recognises that various perspectives on open AI training data have recently emerged. This issue requires further understanding and work with stakeholders to establish norms, standards, and governance models that support responsible data sharing in the public interest.

Acknowledgements

We would like to thank the members of the AI DPG Community of Practice listed below for their time and expertise in shaping our understanding of AI DPGs, bringing essential discussions to the table, and helping to shape these recommendations. A special thanks goes to Dr. Irina Mirkina for representing UNICEF as co-host of the CoP for the past year and a half.

- Abdoulaye Diack, Google
- Aditeswar Seth, Indian Institute of Technology
- Alexander Hradecky, UNDP
- Alek Tarkowski, Open Future
- Cable Green, Creative Commons
- Darlington Akogo, minoHealth AI Labs
- Daniel Brumund, GIZ
- David Manset, ITU
- Davide Storti, UNESCO
- Deb Bryant, OSI
- Emmanuel Letouzé, Data Pop Alliance
- Fabro Steibel, ITS Rio
- Filippo Pierozzi, Office of the UN Tech Envoy
- Freya Hewett, Humboldt Institute for Internet and Society
- Friederike von Franqué, Wikimedia Deutschland e.V.
- Govind Shivkumar, Omidyar Network
- Hadi Asghari, Humboldt Institute for Internet and Society
- Ignatius Ezeani, Lancaster University
- Irina Mirkina, UNICEF
- Jaan Li, One Fact Foundation and University of Tartu
- Jameson Voisin, DPGA Secretariat
- Jerry John Kponyo, Kwame Nkrumah University of Science and Technology
- Judith Faßbender, Humboldt Institute for Internet and Society
- Kuuku Abraham Sam, GIZ
- Luis Villa, Tidelift
- Merl Chandana, LIRNEasia
- Miko Canares, Independent Consultant
- Moritz Fromageot, Office of the UN Tech Envoy
- Prateek Sibal, UNESCO
- Renata Ávila, Open Knowledge Foundation
- Ricardo Mirón, DPGA Secretariat
- Ruth Schmidt, GIZ
- Sami Nenno, Humboldt Institute for Internet and Society
- Stefano Maffulli, OSI
- Susan Ariel Aaronson, George Washington University
- Tanvi Lall, people+AI
- Tanuj Bhojwani, people+AI
- Tarunima Prabhakar, Tattle Civic Tech
- Theresa Züger, Humboldt Institute for Internet and Society
- Urvashi Aneja, Digital Futures Lab
- Vipul Siddharth, UNICEF
- Yasmine Hamdar, UNDP
- Zuzanna Warso, Open Future