Issue

Brief

ISSUE NO. 612
FEBRUARY 2023
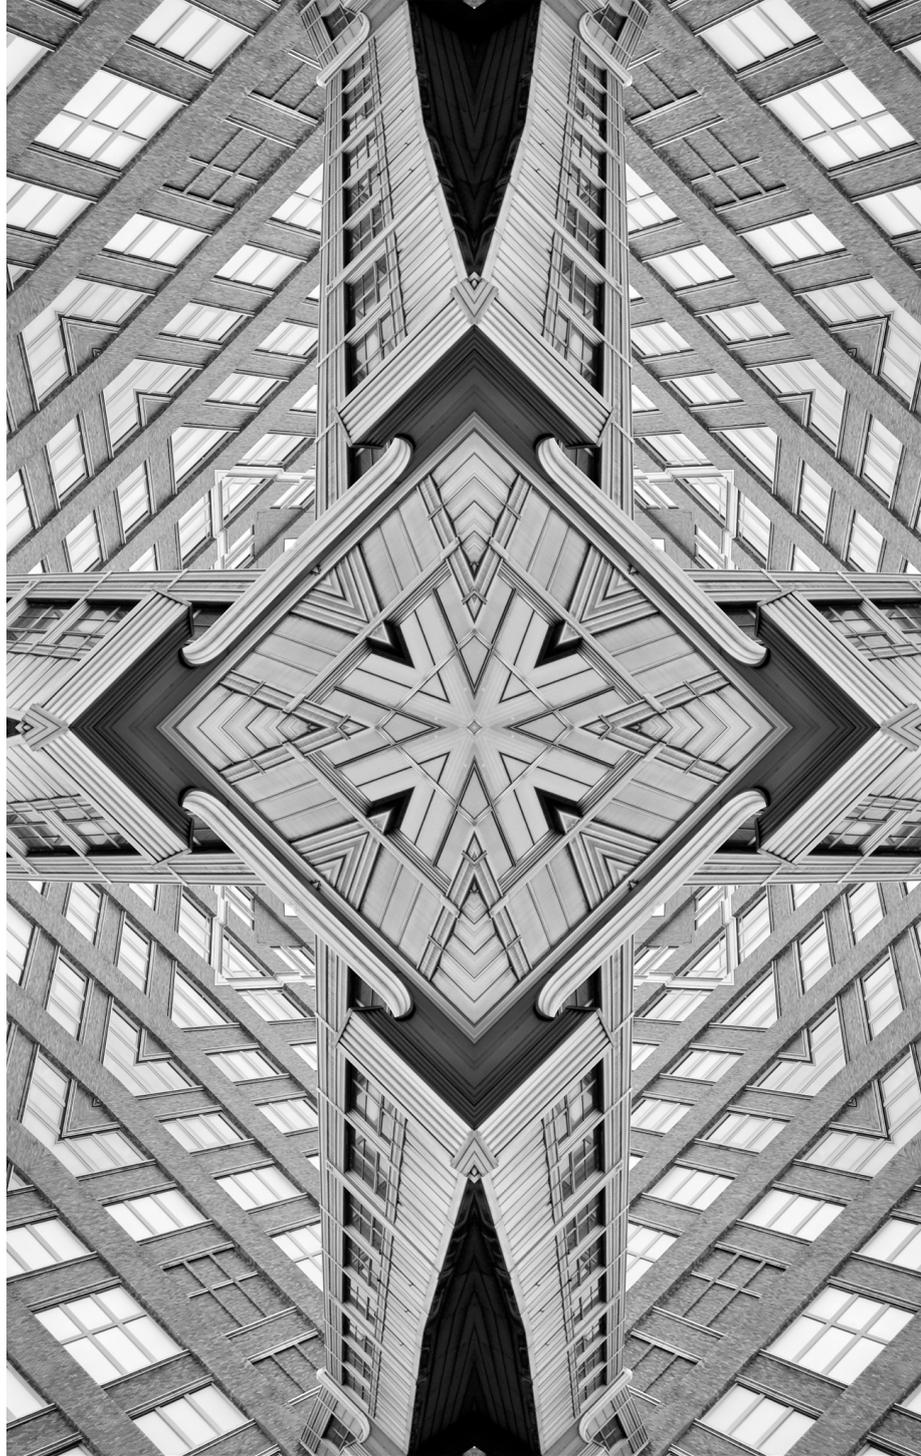
A.I. Systems as Digital Public Goods: Exploring the Potential of Open-Source A.I.

Gaurav Sharma

Abstract

Artificial Intelligence (AI) systems are showing promise in addressing the complex and interrelated challenges facing the world. During the pandemic, for example, AI voice enablement helped in broadcasting advisories in the vernacular and acted as a fact-checking tool. Yet, most AI systems are designed and developed in countries of the Global North. Policymakers in developing economies remain wary of AI systems, especially for use in social sectors, as questions of trust, accountability and privacy remain unanswered. This brief explores the potential of open-source AI as a viable solution to the lacuna of giving AI systems the characteristics of digital public goods—a strategy that would benefit larger populations across the globe.

Introduction

Addressing current global challenges requires innovative solutions, and Artificial Intelligence (AI)-enabled systems could be among them. The design and development of such innovations, however, is occurring mostly in developed economies, as countries of the Global South grapple with questions such as assigning accountability and the protection of privacy.

This brief makes a case for assigning AI systems the characteristics of a Digital Public Good (DPG), while underlining the impediments to such a strategy. In the financial sector, AI tools are already being used for repetitive task automation, fraud detection, and faster processing and response time for various transactions. The ease of use of AI in the financial sector can be attributed to the strict rules and regulations on customer data protection. In India, for example, the Reserve Bank of India (RBI) issues directives that require all banks and payment systems to safeguard customer information. Furthermore, the Public Financial Institutions (Obligation as to Fidelity and Secrecy) Act, 1983[1] prohibits financial institutions from divulging any information relating to the affairs of their clients.

This is not yet the case in social sectors such as healthcare and education. This brief explores the potential of open-source AI to make AI-enabled systems function in the various social sectors and scalable to larger populations. The merit of open-source high-value datasets has been showcased in projects such as the High Energy Physics (HEP) experiments where data are open to public at the Large Hadron Collider at CERN, Geneva.[2] Does open-source provide a similar opportunity to build collaborative advancements in AI for public good? Is the creation of open-source AI initiatives in the social sector feasible, say in healthcare or agriculture?

In the study of economics, public goods are defined as "non-excludable and non-rivalrous".[3] This means that public goods adhere to two broad principles: (i) people cannot be "excluded from consuming" public goods; and (ii) "one person's consumption does not reduce the amount available to other consumers." A Digital Public Good (DPG), meanwhile, is defined as "open-

Introduction

source software, open data, open Artificial Intelligence models, open standards and open content that adhere to privacy and other applicable international and domestic laws, standards and best practices and do no harm and help attain the sustainable development goals (SDGs)."[4] Therefore, integrating the definition of 'public good' with the 'digital' domain, is to suggest free and open availability of digital products and services and free distribution, use, and reuse by all.

The UN high-level panel on digital cooperation sought to simplify the definition of DPGs by providing a number of indicator open standards.[5] These standards state that DPGs are platform-independent; use approved open licenses; produce detailed technical and software documentation such as source codes, use-cases, and functional requirements; have clear ownership and defined mechanisms for extracting data; and adhere to guidelines on the protection of data privacy and security.[6]

Pushing Artificial Intelligence (AI) systems to the domain of DPGs is a difficult task. To begin with, most AI systems are governed by rules on intellectual property (IP), and the free and open use of algorithms and datasets is minimal. As the 'public' notion of 'digital public goods' lies in free use, distribution, and adoption—there is a lacuna in legislation and regulation. This is true in many parts of the globe, and more so in developing countries. The applicability of AI systems as DPGs is related more to their use-case applicability in a particular sector. For example, AI systems in healthcare, to be designated as DPGs, must be able to provide healthcare service accessibility to all under strict data protection guidelines. For AI systems, there is an additional requirement for control over the original code of the digital good, to avoid alteration and misuse. This is also where AI systems suffer, as algorithmic logic and learnings are mostly owned by private enterprises and AI systems evolve as datasets grow. Thus, AI systems are bound by IP rights and are also in a perpetual state of change.

A.I. for Social Good

The principle of 'AI for social good' is interpreted to mean the use of AI technology for applications that redound to the welfare of communities. It encompasses AI applications, design and assessment frameworks and policy initiatives that are focused on benefitting not individuals, but societies as a whole. 'AI for social good' also refers to a set of principles that can inspire the design and assessment of AI systems, and provide a means to advance development of AI policies that prioritise action plans for the adoption of AI in the public interest. Table 1 lists some initiatives that aim to promote 'AI for social good'.

Table 1:
A.I. With a Social Impact: Examples

Organisation	Area	Social Purpose	Impact Sectors
Climate Change AI[7]	Climate Change	Use the power of AI and machine learning to help reduce greenhouse gas emissions (GHG)	Sectors such as energy and urban infrastructure development, and scalable to other sectors
Organization for Economic Co-operation and Development (OECD).[8]	AI Policy guidelines	Recommendation of the Council on Artificial Intelligence	All sectors
Bill and Melinda Gates Foundation (BMGF) and the German Development Cooperation (GIZ)[9]	Vernacular Languages – To democratise voice technology.	Save local languages and availability of AI services in local languages	Multiple sectors – healthcare, education, agriculture, financial inclusion
Lacuna Fund[10] – multiple international partners: The Rockefeller Foundation, Google.org, International Development Research Institute (IDRC), FAIR Forward: AI for All initiative of German Development Cooperation	Open Datasets: Funding creation of open datasets for social impact	World's first collaborative effort to provide data scientists, researchers, and social entrepreneurs in low- and middle-income contexts globally with the resources they need to produce labeled datasets that address urgent problems in their communities.	Current: Language, Agriculture and Health. Expandable to other sectors.
UNESCO: Ethics of AI[11]	Policy Recommendations on the Ethics of AI	The very first global standard-setting instrument on ethical use of AI	All sectors

A.I. for Social Good

Healthcare is one important area where AI-enabled systems hold promise. For an AI system to be used in healthcare, the first imperative is to recognise that people are at both the supply side and the demand side of the healthcare sector. This means that people are generating the datasets e.g., through X-rays, ECG reports, Retina scans, among others—or around whose environment the data is generated; and people would also be at the receiving end of the output data generated by AI systems—for example, predictive analysis undertaken by an AI system, based on X-ray data, whether a patient has tuberculosis (TB) or not. Simply put, the AI system for social good would inculcate 'data creation' as a human-centric process, as most data would be *created by* people or would be implicitly *about* people; or is created by people *for* other people; or is a measurement of the environment that people live in (e.g. tracing a virus outbreak).[12]

At present, most large AI models are scraping datasets from the internet for voice, text, images, videos, and other data. This has little use for social-impact sectors such as healthcare. This is because for targeted solutions such as detection of TB, X-ray datasets are required and this in itself is a meticulous exercise demanding the organised collection of datasets from people who donate their health data based on trust and with appropriate data-protection mechanisms in place. A country that has experience in this regard is Israel, with its MIMIC-III[13] – a critical care database or over 40,000 patients admitted to intensive care units at the Beth Israel Deaconess Medical Center (BIDMC). The access to this critical information is via cloud platforms, but requires becoming a credentialed user of PhysioNet and adherence to strict data-use agreement. However, such detailed, secure and massive efforts of data standardisation, coupled with strict personal data protection schemes, are often absent in countries of the global South.

The application of AI systems for social good in sectors that collect datasets from human subjects—such as healthcare, education, nutrition, and poverty alleviation—is labour-intensive. The task of comprehensive data collection, tailored to solve a particular problem, has yet to gain adequate attention in countries of the global South.

The Missing Global South Representation

There are a number of large AI models existing at present, with the most prominent ones being developed in the 'Silicon Valley' of the United States. These are: GPT-3 and chatGPT (natural language models), by OpenAI in San Francisco;[14] BERT (natural language models) and ViT/G14[15] (computer vision models), by Google based in California;[16] Galactica by Meta (Facebook) also in California. These AI models are focused on the supply side of data generation, feeding enormous amounts of data by crawling the internet for training and fine-tuning. As data generation is a generic activity by online users, and data such as text, images, audio, and video are widely generated and shared openly by users on the internet, it is easy for large AI systems to crunch these digital datasets and train the big AI models.

The problem, however, is that these current, large AI model datasets represent mostly the populations of the global North. These exclude the three billion people who do not have access to internet, 96 percent of whom are residing in developing countries.[17] Therefore, there is a yawning wealth gap in inclusion in large AI models. This is one reason why very large AI models are difficult to classify as DPG: there is simply no equity in data representation.

Moreover, there is little sharing of knowledge regarding the deployment and use of AI systems in the socio-economic sectors. This can be attributed to the fact that most AI systems use cases are narrow in their application and designed to solve a specific problem, and therefore are difficult to replicate in other contexts. For example, the project 'Sunroof',[18] which helps in the estimation of rooftop solar potential in different parts of the United States, may be working well in those places but may not be easily deployable, for example, in unorganised settlements in cities like Delhi. The question remains, therefore, whether these successful global North AI systems are scalable and replicable.

Furthermore, the current knowledge and technical research about most AI systems, studies contexts in developed countries. These scholarly research, demonstrations, and conferences do not account for the impact of AI-enabled systems on populations of developing economies.

AI Systems as Digital Public Goods: Obstacles

As discussed briefly earlier, AI systems are being used extensively in banking and financial services and insurance in countries like the United Kingdom and Switzerland, where regulatory and legislative frameworks are in place. The European Union, with its General Data Protection Regulation in place, is also advancing the use of AI in the financial sector. Thus, AI does already credit itself, to a certain extent, as a 'digital public good' in the banking and financial sectors. The applicability in social sectors such as healthcare remains largely unexplored.

Much of primary research on AI systems and their development and deployment is taking place in countries of the global North, via either public sector institutions or the industrialised laboratories of Big Tech companies such as Google, Meta, Apple, IBM, Huawei, and Microsoft. These institutions are generally governed by proprietary practices and work with closed datasets and proprietary AI models; there is very limited open-sharing, if at all, of these AI models or datasets. Most AI systems, therefore, do not qualify as a 'digital public good'. The following points outline the most crucial challenges that inhibit the infusion of DPG characteristics into AI systems.

1. **Unavailability of Open Data:** In the absence of frameworks that govern local data and privacy in the global South, it is a difficult task to work with personal data and use it for AI systems. As the pillars of a 'digital public good' are open datasets and open AI models, the scarcity of datasets from countries in the global South, in general, presents a roadblock to define AI systems as a DPG. Furthermore, current AI systems are trained on datasets from the global North and may fail to account for social and cultural sensitivities when applied to developing countries. As discussed briefly earlier, quality and standardised datasets are absent in the healthcare sector, for example, of developing countries. Second, in sectors such as banking, finance and insurance where quality data is applicable and used by AI systems, the data is not available to be used freely and openly by start-ups and small and medium-scale enterprises (SMEs), as most datasets are proprietary in nature. Furthermore, large AI models such as GPT-3[19] and chatGPT are trained on internet-crawled datasets. The accessibility of high-quality filtered, refined, and fine-tuned datasets rests with the organisation developing these large AI models. Thus, if one wants to work on a health issue use-case in a developing country and is looking for finely-tuned

and curated datasets from large AI models, these datasets are not freely and openly available. As most AI systems being developed fall under this category and do not subscribe to open data and open AI models, they are far from being a public good.

2. **Aligning 'Public Good' Definitions:** The 'black box' nature of AI systems has been a fundamental roadblock to their adoption by the public sector in developing countries. The complexity of existing AI systems makes it a *privileged good* accessible only to a few and, in this respect, understood by a few. For public sector administrators in developing countries, making use of non-excludable AI systems means investing in commercial AI solutions to narrow inequity, providing access to services such as subsidies, and improved administrative planning, for example on climate change events such as floods and famines. The 'black box' nature of AI systems makes the comprehension and understanding of AI products and services difficult. The adoption of AI systems in public spheres, therefore, is less attractive as government is accountable for citizens' trust and there is no proven level of trust yet in AI systems. Moreover, the current work in the field of explainable AI (XAI)[20] remains primarily focused on contexts in the West. Current XAI techniques and tools focus on explainability at the level of a software developer and not a citizen user.

3. **Issues of Opacity and Privacy:** Most AI global value chains are 'opaque' and 'unaccountable', as AI systems research and development is still being undertaken by the Big Tech companies and big public sector enterprise laboratories. There are no legislations yet regarding AI systems R&D that make it mandatory for Big-Tech companies to adhere to the privacy protection of personal data. For example, Google Health is licensing the mammography AI research model to iCAD to validate and incorporate the AI technology for use in clinical practices to improve breast cancer detection.[21] There are two privacy challenges in this case: first, iCAD and Google are undertaking proprietary research and thus data use, processing, and fine-tuning to improve AI algorithmic outputs is opaque; and second, as iCAD's AI algorithms are utilising the medical datasets in the United States, the learnings of the AI system are not usable for other

AI Systems as Digital Public Goods: Obstacles

geographies, especially in the developing countries. Given how current AI global value chains are opaque and datasets are subject to privacy challenges, existing AI systems are not a public good. Until AI systems become open and transparent and use of internet-based public datasets is appropriately regulated for Big Tech companies across the globe, AI systems cannot be public goods. Furthermore, as shown in the above example, AI projects funded in Western countries may not necessarily be responsive and applicable to the actual needs of developing economies. For example, the US Department of Agriculture generates annual 'Cropland Data Layer' using Landsat and other satellite data,[22] which helps in crop yield estimation. This, however, would not be applicable to countries such as India, for example, where farmlands are smaller, farm-level species are more diverse, agricultural practices are more variable, and intercropping and crop rotations are predominant.

> "Current AI systems are trained on datasets from the global North and may fail to account for social and cultural sensitivities in developing countries."

Recommendations

This brief makes the following recommendations to push AI systems into the realm of Digital Public Goods.

1. **Open-Source AI:**[23] An open-source AI is any AI technology that is publicly available for commercial and non-commercial use under various open-source licenses.[24] Open-source AI includes datasets, pre-built algorithms, and ready-to-use interfaces that can help developer communities worldwide to take advantage of open AI development.[25] Open-source AI is different from freeware AI applications—in open-source, the underlying code is exposed to the user and open for modifications and implementations in scenarios other than the ones originally intended.

Current AI initiatives are use-case specific. Open-source AI could identify common needs and gaps in disparate disciplines, abstract them and create commodity generic tools, which can act as references to similar challenges in other parts of the globe. For example, it is important to define sufficiently the minimum threshold generic AI-enabled, usable and inclusive datasets that are of high quality and standardised for use in specific purposes across multiple domains. For example, is a dataset of 1,000 hours of good quality speech enough for AI language-based applications? If yes, then this should become a global standard.

The 'Global Forest Watch'[26] based in the United Kingdom attempting to track and protect forests worldwide—is an example. Its findings and databases should have a commitment to be shared in the public domain for their practical reusability and understanding of the content and context. This will enable transparency and reproducibility in making AI systems open-source. Open-source AI presents massive applicability to developing countries, which are generally deficient in key AI resources such as computing power, and technical and domain expertise. Open-source AI could act as one of the key drivers for accelerated adoption of AI systems in addressing social sector challenges in developing countries.

Innovation can thrive by opening-up large AI models such as GPT-3 (Open AI) and ViT (Google) as most AI start-ups do not have the financial resources to build such large AI models. Opening up large AI models can be easily adapted and interpreted, providing equity of code-access. An open-source software[27] is fundamental to the definition of a DPG and the success of regulatory mechanisms inspired by initiatives such as Wikipedia suggests that it is possible to integrate and monitor globally shared technical resources. As open-source, by definition, does not discriminate against persons or groups, it is a good starting point for AI systems. To be sure, open-source AI is not completely risk-free: AI datasets and AI models in the public sector may be vulnerable to cyberattacks and open access could lead to privacy violations. Open-source AI is also subject to trust, reliability and accountability issues regarding AI system output.[28] Its overall social gains, however, exceed the potential demerits.

2. **Global Partnerships**: A multi-stakeholder inter-governmental global partnership between wealthy and developing countries—such as the global partnership on AI (GPAI)—[29] could set in motion the initial steps in deliberating the strategy of making AI systems, public goods. The current mandate of GPAI is to promote responsible development and use of AI while respecting human rights and democratic values.

As data governance is already the key theme of GPAI, it can lead efforts and create dialogue and research space in moving AI systems to the space of digital public goods. Similarly, the Digital Public Goods Alliance (DPGA) has already set standards to ascertain if a digital solution conforms to the definition of a DPG and open data and open AI models are part of that set of specifications.

The DPGA is thus rightfully placed to lead discussions on AI system-specific standards, especially non-personal data AI systems. It can echo the voices of developing countries supporting Open-Source AI. Similar partnerships can empower AI governance directives to embed AI systems within the public goods ethos of 'non-exclusion and non-rival' and gather support for the development of AI system platforms with a DPG mandate.

Recommendations

Recommendations

3. **FAIR Principles Approach:** FAIR principles (Findability, Accessibility, Interoperability, and Reusability)[30] can play a role in designing FAIR means of practice for AI systems and not only support innovation and research but help in better data representation and in making provisions for access to AI models and data in the context of no-harm. Human-centric AI systems make AI systems socially beneficial.

AI systems development must become fundamentally human-centric by default (human-centric by design, human-centric during development, human-centric in ease of adoption and use). This might demand a minimal viable governance mechanism to be actioned by the United Nations such as a 'Universal Declaration on AI' that sets parameters as moral 'code of conduct' that supports fundamental rights, democratic values, public opinions and operational capacity via multi-stakeholder engagement, inclusion, and compatibility in the public domain.

FAIR Principles for AI would also contribute to the sustainability of AI models applicability inter and intra domains. FAIR AI models can facilitate comparison of benchmarks results across models and support explainability of AI. This will, in turn, support education in AI for interpretability of AI models, uncertainty quantification, and ease of access of data and AI models for key use cases.

Conclusion

Advancements made by AI systems, especially with the explosion of large AI models, are bound to have a more pervasive effect on significant populations across the globe. Infusing a notion of DPG for AI systems demands greater assurances and standards of care and governance structures. These begin from the conception stage of designing an AI system for social sectors. The notion of open AI datasets and open AI models, open sharing and cross-transfer of knowledge rooted in responsible integration of AI systems into the public sector—can pave the way for AI systems alignment as a digital public good. Open-source AI is one alternative approach that can be explored, coupled with governmental accountability, as open-source AI could provide a greater level of transparency and make it easier for AI systems to be accepted as a DPG by end users.

Furthermore, ethical and responsible AI practices must be put in place that are inclusive and would incorporate the aspirations of developing countries. A thoughtful process that embeds the societal beliefs and undertakes risk assessment of AI systems is imperative in all discussions. The public goods foundation of 'non exclusion and nontrivial' and support for an open-source AI definition can serve as the pillars to push discussions on AI systems as digital public goods. The use-case-based approach, displaying examples from developing countries, could further strengthen the use of open-source AI. There is a need for the AI community to invest in open-source AI from a public-use perspective. ORF

Gaurav Sharma *is Artificial Intelligence Fellow at the Academy of International Affairs.*

Endnotes

1   The Public Financial Institutions (Obligation as a Fidelity and Secrecy) Act, 1983; December 1983, https://legislative.gov.in/sites/default/files/A1983-48.pdf

2   The European Organization for Nuclear Research (CERN) provides open datasets from particle physics. See: http://opendata.cern.ch/

3   Stanford Encyclopedia of Philosophy, 'Public Goods', https://plato.stanford.edu/entries/public-goods/#NonRivaNonExcl

4   United Nations, Report of the Secretary-General, 'Roadmap for Digital Cooperation', June 2020, https://www.un.org/en/content/digital-cooperation-roadmap/assets/pdf/Roadmap_for_Digital_Cooperation_EN.pdf

5   Digital Public Goods Alliance, "Digital Public Goods Standard", https://digitalpublicgoods.net/standard/

6   Digital Public Goods Alliance, "Digital Public Goods Standard"

7   Climate Change AI is a global non-profit that catalyses impactful work at the intersection of climate change and machine learning, https://www.climatechange.ai/

8   The Organisation for Economic Co-operation and Development (OECD) – Legal Instruments, "*Recommendation of the Council on Artificial Intelligence,*" 2019, https://legalinstruments.oecd.org/api/print?ids=648&lang=en

9   Mozilla Foundation blog, "Mozilla Common Voice Received $3.4 Million Investment to Democratize and Diversity Voice Tech in East Africa," May 24, 2021, https://foundation.mozilla.org/en/blog/mozilla-common-voice-receives-34-million-investment-to-democratize-and-diversify-voice-tech-in-east-africa/ (accessed December 28, 2022)

10   Lacuna Fund is a global collaborative effort to fund labelled data for social impact in various domains – Language, Agriculture, Healthcare. See: https://lacunafund.org/about/

11   *Ethics of Artificial Intelligence*, UNESCO, https://www.unesco.org/en/artificial-intelligence/recommendation-ethics

12   Rishi Bommasani, et al., "On the Opportunities and Risks of Foundation Models," Stanford University Human-Centered Artificial Intelligence and Centre for Research on Foundation Models, 2022, https://arxiv.org/abs/2108.07258

13   Alistair Johnson, Lucas Bulgarelli, et al., "MIMIC-IV," *PhysioNet – The Research Resource for Complex Physiologic Signals* (2022), https://physionet.org/content/mimiciv/2.1/

14   GPT-3 is a set of large AI models that can understand and generate natural language. See: https://beta.openai.com/docs/models/overview

Endnotes

15   Xiaohua Zhai, et al., Google Research, Brain Team, Zürich, *Scaling Vision Transformers*, 2022, https://arxiv.org/pdf/2106.04560v2.pdf

16   "BERT 101: State Of The Art NLP Model Explained," *Huggingface blog*, comment posted March 2, 2022, https://huggingface.co/blog/bert-101 (accessed November 25, 2022)

17   International Telecommunication Union, "Facts and Figures 2021: 2.9 Billion People Still Offline," https://www.itu.int/hub/2021/11/facts-and-figures-2021-2-9-billion-people-still-offline/#:~:text=An%20estimated%2037%20per%20cent,still%20never%20used%20the%20Internet.

18   Project Sunroof, https://sunroof.withgoogle.com/

19   GPT3 is a third-generation large-scale language model that can understand and generate natural language human-like text output. Not only can it produce text, but it can also generate code, stories, poems, and others. GPT-3 is trained on nearly 45 Tera Bytes (TB) of text data. GPT-3's training data is still primarily English (93 percent by word count). See: https://arxiv.org/pdf/2005.14165.pdf, p. 14

20   Chinasa T. Okolo, Nicola Dell, and Aditya Vashistha, "Making AI Explainable in the Global South: A Systematic Review," Paper presented at the Conference on Computing and Sustainable Societies, June 29 – July 01, 2022, https://dl.acm.org/doi/proceedings/10.1145/3530190

21   Gred Corrado, "Partnering with iCAD to Improve Breast Cancer Screening," *Google Blog*, November 28, 2022, https://blog.google/technology/ai/icad-partnership-breast-cancer-screening/

22   Daniel Kpienbaareh, et al., "Crop Type and Land Cover Mapping in Northern Malawi using the Integration of Sentinel-1, Sentinel-2, and PlanetScope Satellite Data," *Special Issue, Environmental Mapping Using Remote Sensing*, 13 (4), 700, (2021), https://www.mdpi.com/2072-4292/13/4/700

23   Open-source initiative, "Open- Source software is software that can be freely accessed, used, changed, and shared (in modified or unmodified form) by anyone," https://opensource.org/osd

24   Chiradeep Basu Mallick, "Top 10 Open Source AI Software in 2021," *Spice Works*, https://www.spiceworks.com/tech/innovation/articles/top-open-source-artificial-intelligence-software/#:~:text=Open%20source%20AI%20is%20defined,started%20with%20AI%20app%20development

25   Mallick, "Top 10 Open Source AI Software in 2021"

26   Global Forest Watch offers the latest data, technology and tools that empower people everywhere to better protect forests. See: https://www.globalforestwatch.org/

Endnotes

27   Open-Source Initiative, "Frequently Answered Questions," https://opensource.org/faq#osd

28   Alexandra Theben, et. al., "Challenges and Limits of an Open Source Approach to Artificial Intelligence," *Artificial Intelligence in a Digital Age*, May 2021, https://www.europarl.europa.eu/RegData/etudes/STUD/2021/662908/IPOL_STU(2021)662908_EN.pdf

29   The Global Partnership on Artificial Intelligence (GPAI) is a multi-stakeholder initiative which aims to bridge the gap between theory and practice on AI by supporting cutting-edge research and applied activities on AI-related priorities. See: https://www.gpai.ai/about/

30   Mark D. Wilkinson, et al., "The FAIR Guiding Principles for Scientific Data Management and Stewardship," *National Library of Medicine: National Centre for Biotechnology Information USA,* Sci Data, 3:160018, (2016), https://www.ncbi.nlm.nih.gov/pmc/articles/PMC4792175/

Images used in this paper are from Getty Images/Busà Photography.

ORF

OBSERVER RESEARCH FOUNDATION

Ideas . Forums . Leadership . Impact

20, Rouse Avenue Institutional Area,
New Delhi - 110 002, INDIA
Ph. : +91-11-35332000. **Fax :** +91-11-35332005
E-mail: contactus@orfonline.org
Website: www.orfonline.org