JUNE 2019

Artificial Intelligence in International Development

A Discussion Paper

PREPARED BY The IDIA Working Group on Artificial Intelligence & Development

The International Development Innovation Alliance (IDIA)

Australian Aid
giz Deutsche Gesellschaft für Internationale Zusammenarbeit (GIZ) GmbH
USAID FROM THE AMERICAN PEOPLE
UKaid from the British people
Department for International Development
Canada
Bill & Melinda Gates foundation
World Bank Group
The Rockefeller Foundation
UNDP Empowered lives. Resilient nations.
Grand Challenges Canada / Grands Défis Canada BOLD IDEAS WITH BIG IMPACT®
Sida
Global Innovation Fund
unicef
Results for Development

Table of Contents

Artificial Intelligence in International Development

About this document

This paper has been developed by Thomas Feeny and Olivia Elson at Results for Development in partnership with members of the IDIA Working Group on Artificial Intelligence & Development. It is designed to provide an accessible and concise entry point for actors working in international development who are interested in how Artificial Intelligence (AI) technologies can or will impact their work. This paper draws on a rapid (rather than exhaustive) review of current reports, blogs and commentaries on AI offered by experts around the world and is split into two sections:

- **Part One** explores the history of AI, its current complexity and capabilities, and examples of how it is being used within development to support the Sustainable Development Goals.
- **Part Two** synthesizes contemporary challenges to the deployment of AI in development and outlines some of the key debates that are influencing stakeholders' approaches, alongside a selection of tools and initiatives that are advancing practice in this space.

We thank all those who have contributed to this document, especially member agencies of the IDIA Working Group on Artificial Intelligence & Development who acted as critical reviewers throughout the process.

About the International Development Innovation Alliance (IDIA)

IDIA is a unique collaboration platform that brings together the senior leadership from the innovation teams, labs and departments of some of the world's leading development agencies, with the shared goal of "actively promoting and advancing innovation as a means to help achieve sustainable development". IDIA is committed to the development of new products, services and ways of working, and ensuring that the lessons arising from both success and failure can be disseminated to inform the more efficient adaptation and scaling of innovations within different countries, populations and contexts. For more information, visit www.idiainnovation.org.

About the AI & Development Working Group

The IDIA Working Group on Artificial Intelligence & Development was established in January 2019 to provide a forum through which IDIA members and other actors from the public, private and academic sectors could come together to learn and collaborate around the deployment of AI in development policy and programming.

Technology is a both a source of innovation and a key enabler of the innovation process. For example, the rapid penetration of mobile phones globally has almost instantaneously created a platform for all kinds of mobile-based innovative products and services to reach scale. The new wave of emerging technologies, especially AI, now offers a similar paradigm shift in how we design, implement and scale development innovation, but is not without its risks in terms of its potential to exacerbate, rather than reduce, socio-economic inequalities. To ensure an informed approach, the AI & Development Working Group is creating a range of introductory materials (including this document) examining the responsible use of AI in development, before moving towards potential collaboration in promoting and scaling promising actors, initiatives and innovations which exemplify the responsible use of AI in support of the Sustainable Development Goals.

AT A GLANCE

This short paper provides a high-level introduction to Artificial Intelligence (AI) and its relevance to international development. It is designed for those working in the international development sector — whether in funding agencies, governments, the private sector or other organizations — who are new to the field of AI and would like to understand the basics of emerging technologies and how they may impact their work.

PART ONE:

Recognizing the potential — Focuses on defining AI, providing some background to the field alongside concrete examples of the use of AI in the international development sphere

A BRIEF HISTORY | Page 5

Much common use of and political activity around AI — especially as it relates to the development sector — has only emerged in the last few years. But today's advances rest on almost 70 years of research and development.

DEFINING AI | Page 7

AI is an incredibly broad sphere which is constantly evolving. This section covers some key terms and characteristics of AI which can help demystify this space.

LEVELS OF COMPLEXITY | Page 8

While AI is advancing rapidly, the popular conception of the power of AI is typically far more advanced than is actually the case at present.

CAPABILITIES | Page 10

What can AI do? Some of its present capabilities include computer vision, content generation, natural language processing, robotics, expert decision-making and speech/audio processing.

GOVERNMENT RESPONSE | Page 11

Many countries have announced significant financial commitments to advance their agendas in AI. Strategies are varied in scope, with some focusing more on research and development, others on policy, legal and regulatory frameworks, some on digital infrastructure and others on education and skills. It is notable that so far only a few Global South countries have developed targeted AI agendas.

APPLICATIONS BY SDG | Page 11

AI can help to accelerate and/or increase impact and we are already witnessing examples of this across many fields including healthcare, education, agriculture and more.

PART TWO:

Navigating the challenges — Describes the key risks that AI presents in development and highlights many promising initiatives that are emerging to support the responsible application of these technologies

POOR AVAILABILITY, ACCESSIBILITY AND QUALITY OF DATA | Page 15

AI is only as good as the data that it is based on — and in the development world, challenges around the accessibility, quality, depth, diversity and volume of data are particularly common, especially in fragile / conflict-affected contexts.

LIMITED CAPACITY TO ENGAGE WITH OR USE AI | Page 16

Globally, there is a shortage of knowledge and skills in the field of AI, as well as a concentration of AI expertise in the hands of a select few, which is limiting its spread across more geographically and ethnically diverse groups.

DATA OWNERSHIP, PRIVACY AND SECURITY Page 17

Development actors and their partners working in sectors that involve large amounts of individually-identifiable data (such as health and education) are at particular risk of undermining privacy and ownership, even (or perhaps especially) when their use of this data is for apparently benign / 'public good' purposes.

PERPETUATING BIAS AND DISCRIMINATION Page 18

The integrity of an AI application is wholly dependent on training data that it learns from, such that it will replicate and perpetuate any biases that may exist within that data, and controlling for this is extremely difficult. Given the current scale of this 'diversity disaster' within AI, the task for anyone using AI in development becomes more one of identifying and mitigating bias rather than eliminating it.

PART ONE

Artificial Intelligence in International Development

Recognizing the Potential

"Emerging ML/AI applications promise to reshape healthcare, agriculture, and democracy in the developing world... At the same time, the very nature of these tools — their ability to codify and reproduce patterns they detect — introduces significant concerns alongside promise."[1]

Artificial Intelligence (AI) is evolving rapidly and is making substantial impacts in homes, businesses and political processes worldwide. Despite emerging over half a century ago, the field of AI is still relatively new, especially as it relates to its application within international development. As the driving force behind the Fourth Industrial Revolution, it is bringing deep and far-reaching changes to the way in which people live, work and play. But while advances in AI are generating positive outcomes at scale across sectors and geographies, it is vital to consider how to ensure the responsible deployment of this technology. This is especially important in the context of the AI divide, in which the Global South is arguably more vulnerable to the risks and challenges inherent in the design, development and implementation of AI technologies, and far less likely to reap the benefits..

Part one of this paper has been designed to provide a high-level introduction to AI and its relevance to international development; highlight some practical examples of its application within development programming; and signpost the reader to further materials and more in-depth resources on AI in the development sector.

1. What is Artificial Intelligence?

1.1 A brief history

Much common use of and political activity around AI has only emerged in the last few years, with Canada the first country in the world to announce a national strategy for AI in March 2017. But today's advances rest on almost 70 years of research and development. The timeline on the following page highlights some key moments in the evolution of AI throughout the 20th and 21st centuries, which help to contextualize what is happening in the field of AI today.

Of note is the fact that coordinated discussion and efforts around AI in international development only began in around the last five years, with a call for a data revolution at the United Nations' High-Level Panel on the Post-2015 Development Agenda, the creation of the World Economic Forum's Centre for the Fourth Industrial Revolution in 2016, and the inaugural AI for Good Global Summit in 2017.

There are of course many more milestones and examples[2] of technologies developed since the early 20th century, and for those interested in exploring the theoretical foundations of the field in more depth, the history of AI can be traced back much further to its early origins in mathematics and mechanics in the 1300s.[3]

Milestones in the development of AI

1950 Computer scientist Alan Turing publishes a paper, Computing Machinery and Intelligence in which he discusses how to build intelligent machines and test their intelligence.

1955 Newell, Shaw, and Simon create the Logic Theorist, a program designed to emulate human problem solving skills and considered by many to be the first AI program.

1956 John McCarthy and Marvin Minsky host a landmark conference for the Dartmouth Summer Research Project on Artificial Intelligence. It catalyzes the next twenty years of AI research.

1960s The development of AI accelerates as computers become more powerful and more accessible, and machine learning algorithms improve. A natural language processing chatbot named ELIZA is created at the MIT AI Laboratory in 1964 and in 1966 the Stanford Research Institute creates the first mobile robot, Shakey, that could process reasoning about its actions and surroundings.

"In three to eight years we will have a machine with the general intelligence of an average human being."
— MARVIN MINSKY, 1970

1970s A period of skepticism and impatience around the advances made so far sets in. The end goals of a machine that could exhibit intelligence, abstract thinking and self-recognition are hampered by a lack of computational power, and funding and research slow down.

1980s A boost in funding for AI sees a number of advances. Just a few examples include: Hopfield and Rumelhart popularize "deep learning" techniques which are inspired by the human brain and enable computers to learn using experience; Douglas Lenat's Cyc is developed to codify the knowledge that composes human common sense; and Feigenbaum's expert systems, which mimic human decision-making processes to provide advice to non-experts, are introduced in various industries.

1990s Many landmark goals are achieved. In 1995 advances in natural language processing are demonstrated through the chatbot ALICE which engages in conversation; in 1997, IBM's Deep Blue defeats the reigning world chess champion and grand master; and in 1998 Dr. Breazeal develops Kismet, a robot that recognizes and displays emotions.

"Within thirty years, we will have the technological means to create superhuman intelligence. Shortly after, the human era will be ended." — VERNOR VINGE, 1993

2000s The development of more sophisticated robots continues (such a driverless cars, industrial robots and drones) and AI capabilities are mainstreamed through applications such as Apple's Siri, Google's Assistant and Amazon's Alexa. The application of AI is gathering pace through the availability of big data which is used to shape industries such as telecoms, marketing and financial services.[4]

Specific attention is brought to AI in the sphere of international development, with a call for a "data revolution" at the United Nations' High-Level Panel on the Post-2015 Development Agenda, the creation of the World Economic Forum's Centre for the Fourth Industrial Revolution in 2016, and the inaugural AI for Good Global Summit in 2017.

1.2 Defining Artificial Intelligence

The term 'artificial intelligence' first emerged in the 1956 Dartmouth College research proposal which described "the conjecture that every aspect of learning or any other feature of intelligence can in principle be so precisely described that a machine can be made to simulate it". Since then, definitions of AI have continued to surface, emphasizing various different aspects of the field (see the International Development Research Centre's 2018 paper Artificial intelligence and human development: toward a research agenda for examples).

It is worth noting that many commentators consider the broader field of AI to be a fluid and constantly-evolving concept. Andrew Moore (Former-Dean of the School of Computer Science, Carnegie Mellon University) describes AI as "the science and engineering of making computers behave in ways that, until recently, we thought required human intelligence". That is to say, our understanding of what is considered AI is time-bound and changing as machine capabilities become increasingly sophisticated and functions that were once considered AI (to take examples from several decades ago, basic calculation or the application of if-then rules) become part of basic computer functionality.

For the purposes of this paper, we use the definition proposed in the 2018 USAID report Making AI Work for International Development, which is a useful starting point for non-AI-experts in the development community, and helpfully distinguishes between AI and machine learning:

***Artificial intelligence** uses computers for automated decision-making that is meant to mimic human-like intelligence. Automated decisions might be directly implemented (e.g. in robotics) or suggested to a human decision-maker (e.g. product recommendations in online shopping); the most important thing for our purpose is that some decision process is being automated. AI often incorporates machine learning (when using data-driven predictions to make better decisions) but doesn't have to. For shorthand, you can think of AI as "smart automation."*

***Machine learning** is a set of methods for getting computers to recognize patterns in data and use these patterns to make future predictions. For shorthand, you could think of machine learning as "data-driven predictions."*

WHAT IS AI?

This **online tool** helps to determine whether or not a technology incorporates AI. It is based on a flowchart originally published in the **MIT Technology Review**.

But beyond the textbook definitions, descriptions of AI are numerous and diverse, reflecting the vast array of functions that AI is able to perform (e.g. data processor, decision-maker, navigator); the benefits AI may bring to an individual, a business or society; and/or the various goals which AI systems are built to accomplish, which range from simple systems that can replicate discrete human behaviors, to highly advanced systems that truly simulate human cognition (i.e. that think exactly like humans do).

The diagram below[5] helps to situate the fields of Deep Learning and Machine Learning, each a subset of the broader field of AI.

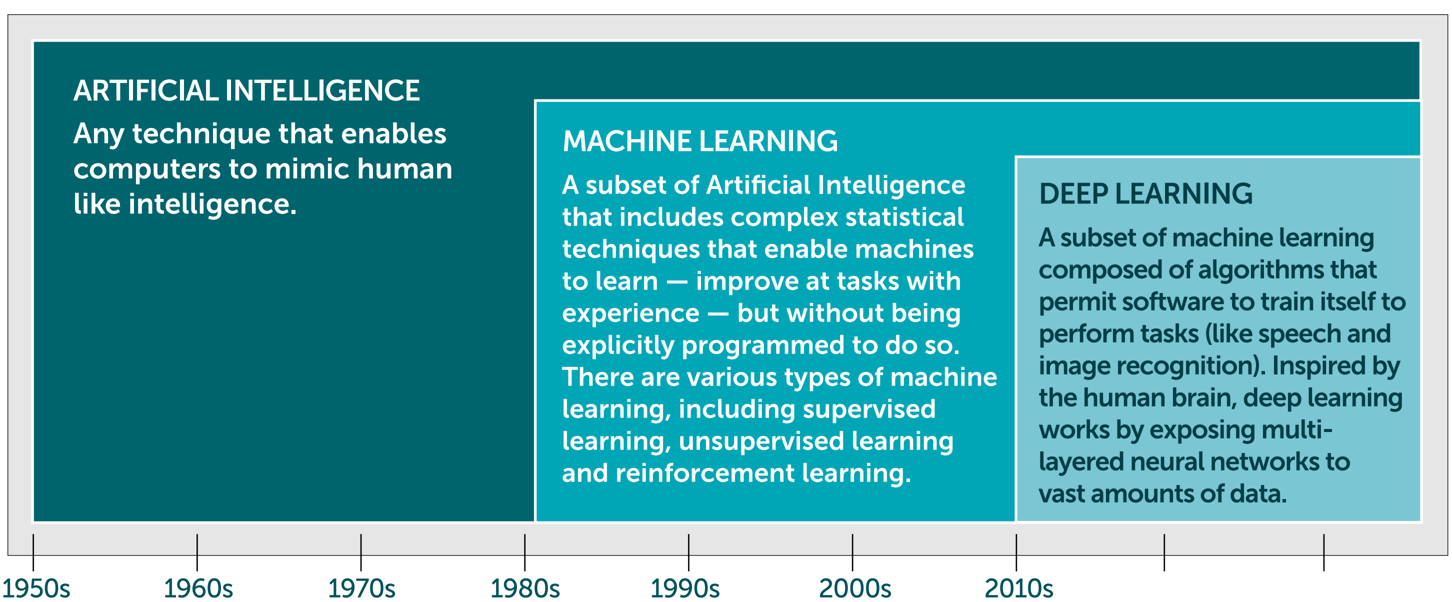

Notably, any given AI solution is likely to be made up of a combination of different AI techniques/ models / machine learning tasks, each involving vast quantities of data and innumerable decisions in how they are built and how they are subsequently applied in practice. This provides great potential for AI to help to address an array of complex problems, but also magnifies many of the risks and challenges of the technology, such as data quality, bias and transparency (see Part 2 of this paper). In the USAID report [Making AI Work for International Development](), the section entitled 'How people influence the design and use of ML tools' provides helpful detail on how AI models are built and integrated into decision-making processes.

1.3 Levels of complexity: how far advanced are we with AI?

The popular conception of the power of AI is typically far more advanced than is actually the case at present. In understanding how far along we have come, it is helpful to approach AI in terms of in terms of the complexity of things it is currently able to achieve within its three core elements: inputs, processes and outcomes. Each of these components can involve varying levels of complexity. For example, at one end of the spectrum, an algorithm may apply strict rules and weightings to small, structured datasets in order to produce a discrete decision or action. At the other end, an algorithm may be built to process and re-process unlimited quantities of unstructured data, resulting in complex decisions and/or actions.

While the speed at which AI research is developing makes it difficult to clearly delineate the boundaries of our achievements so far, a broad framework of 'low', 'medium' and 'high' levels of AI complexity — such as that published by Nesta (below)[6] — is useful in helping to articulate progress made to date, categorize AI solutions that have been developed, and structure dialogue around broader governance and policy issues.

Others grappling with this issue use the terms 'weak AI' (also referred to as 'narrow' or 'applied' AI) and 'strong AI' (also referred to as 'full AI' or 'artificial general intelligence') to describe low and high complexity AI respectively. Commentators suggest that the vast majority of AI applications that are in use today (including those typically deployed in the development sector) fall into the 'weak' category, while 'strong' AI that truly mimics human reasoning remains an unproven concept, with as yet no tangible examples of systems that have reached this level of sophistication.

Increasing efforts to bridge these two extremes are observable in the emerging body of work on systems that use human reasoning as a framework, but which do not aim to perfectly replicate it. This would include deep learning or artificial neural networks, in which algorithms inspired by the human brain process data in layers (the number of layers indicating the 'depth' of the network). Inputs are ascribed a weighting and passed through one or several 'hidden' processing layers in order to generate an ultimate output. Such systems are able to perform tasks repeatedly in order to learn and continually improve the accuracy of the outputs they produce.

Finally, it is worth noting that many commentators identify an even more advanced level beyond 'strong AI', referred to as 'Artificial Super Intelligence' or 'Superintelligent AI' which describes capability that not only mimics human

COMPLEXITY	CREATION	FUNCTION	OUTCOME
3	Unlimited quantities of unstructured training data such as videos, photos, sound, free text	Dynamic model constantly evolving based on live data	Model used as one part of long and complex decision-producing chain
2	Defined quantities of structured or unstructured training data for a one-time creation of model	Static model created using one-time machine learning process	Model used to cover clearly defined point of decision making process
1	Defined number of structured datasets used for one-time weighing of model	Static model using human inputed rules weighed by machine learning	Model used to cover simple and clearly defined point of decision making process

intelligence, but surpasses it. Whether and when we might reach this level remains unknown.

AI CAN LEARN IN DIFFERENT WAYS

Three key approaches are outlined below. AI solutions may rely on one of these approaches or a combination of them.

Supervised learning: Algorithms are programmed with labelled training data which specifies the 'right' outcomes/answers for the model to replicate in its predictions.

Unsupervised learning: Algorithms independently identify patterns in input data to reveal underlying classifications/ associations/ structure within the data, which can in turn be used to inform decisions.

Reinforcement learning: Algorithms learn how to achieve a goal through trial and error. As it interacts with an environment, a model will receive rewards or penalties depending on the actions it takes, and will aim to maximize the total reward achieved.

The implications of this are significant. As the complexity (and in some cases, opacity) of AI solutions develops, understanding their impact within the development field will similarly require new and more advanced methods for monitoring and evaluating, with an emphasis on identifying those who are benefitting and, crucially, those who are not. Attention should be paid to broader, secondary effects of introducing AI solutions into systems, which might include impacts on employment, human rights, equity, inclusive growth, and many other areas. Developing monitoring and evaluation approaches that respond to the complexities of AI systems may involve, amongst other things:

i. Frameworks for assessing the quality of — and potential biases inherent in — the datasets used by AI systems, the performance of AI models, and the level of confidence in the results they generate;

ii. Robust approaches to interpreting data/analysis generated by AI and how to use it for decision-making, including understanding the potential cost of errors; and

iii. The use of mixed methods including qualitative, participatory approaches in order to validate or invalidate findings.

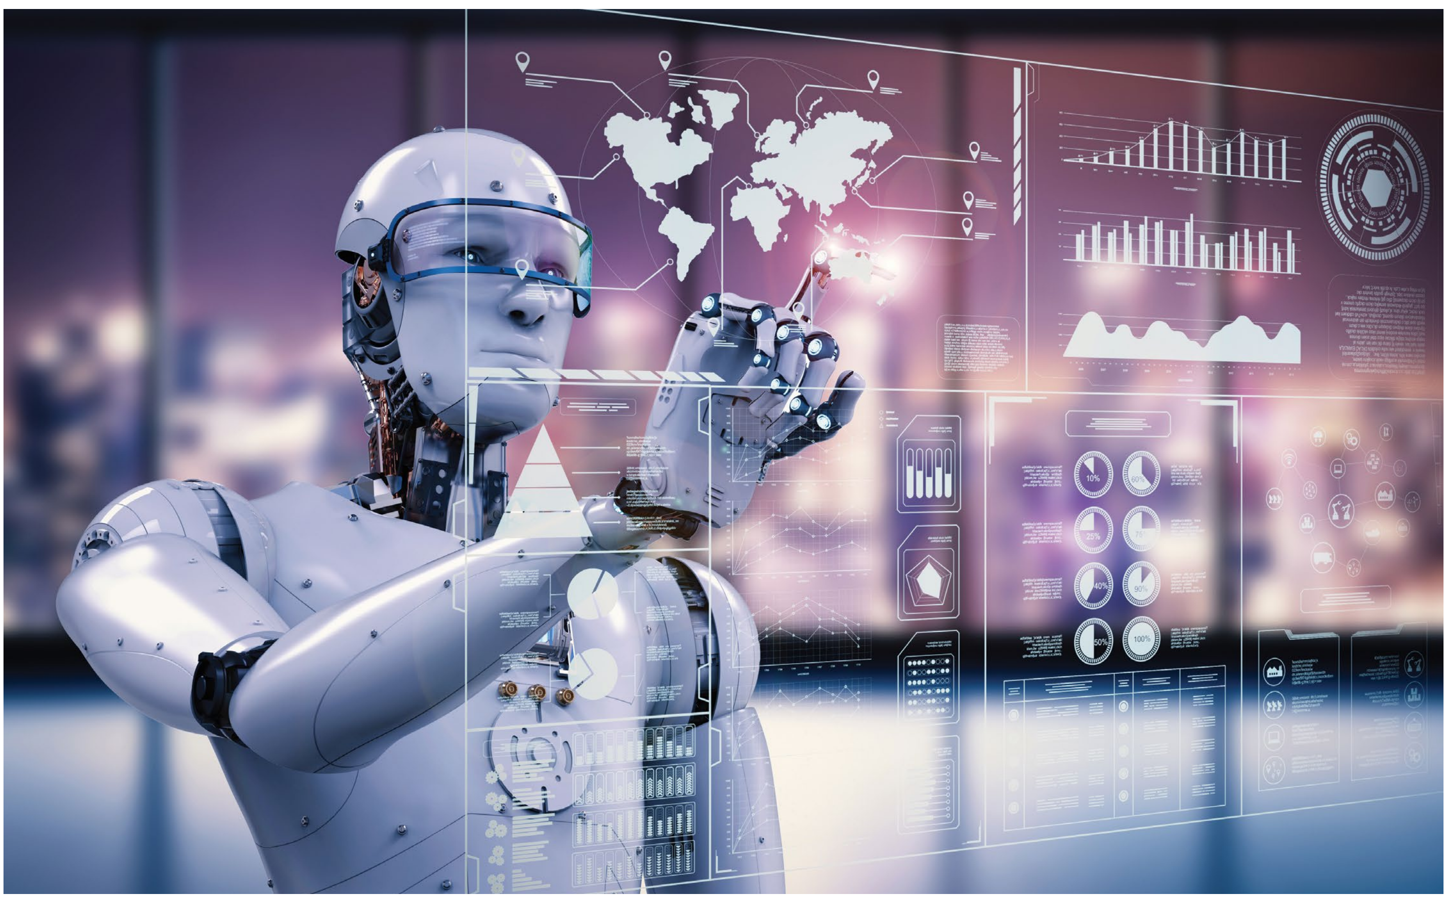

1.4 Current and emerging AI capabilities

The table below draws on publications by McKinsey[7] and Deloitte[8] to summarize some specific AI / machine learning capabilities, their levels of maturity, and their limitations.

	Current capabilities	Emerging capabilities	Limitations
Computer vision	■ Identify and classify image and video data of people/faces, objects, and written characters ■ E.g. Analyzing medical scans for signs of disease	■ Accurately recognize emotions	AI has an innate inability to make 'moral' or common-sense judgments. Other potential limitations which may present significant risks unless explicitly addressed through careful design and programming include: ■ Reliance on extremely large training data sets (and associated challenges of time, cost and expertise needed to responsibly collect, store, manage and interpret such large volumes of data) ■ Need for substantial human input to label / categorize data before it can be used in a model ■ Limited ability to transfer models / learning from one use case to another, or to multitask ■ Inability to detect or mitigate against biased data ■ Inability to explain outcomes / decisions
Content generation	■ Create short-form text content ■ Summarize longer documents ■ E.g. Creation of news articles, personalized emails	■ Generate video and audio content ■ Generate long-form original / creative content	
Natural language processing	■ Identify people/authors through text or speech ■ Translate language ■ Scan bodies of text (documents, websites) for certain information ■ E.g. Smart devices (Alexa, Siri, Google Assistant)	■ Analyze sentiments conveyed through text ■ Understand abstract concepts ■ Understand ambiguous language ■ Sustain coherent dialogue	
Robotics	■ Automatically perform physical functions ■ E.g. Drones collect images to map areas affected by natural disaster	■ Adapt to perform multiple different tasks	
Expert (rules-based) systems	■ Capture and use expert knowledge to provide answers to problems ■ E.g. Provide treatment recommendations for patients	■ Adapt decision-making processes to new information and contexts	
Speech and audio processing	■ Verify individuals' voices ■ Automatically convert speech to text ■ Recognize sounds ■ E.g. Telephone assistance services	■ Recognize emotions ■ Process different languages and accents with accuracy	

While the above is not an exhaustive list, it nonetheless demonstrates the wide variety of functions that AI is able to perform, many of which are already being deployed in a development context (see Section 2.2 for examples). Key drivers behind the introduction of these technologies are often broadly described in terms of their ability to accelerate processes through automation, and/or augment processes by either providing deeper or more accurate information or driving efficiencies that enable humans to dedicate more time to the interpersonal and creative aspects of work which — at least currently — cannot be fulfilled by AI.

Indeed, in the face of widespread hype around the apparent speed of AI development it is vital to ground truth how quickly it is really likely to be adopted at scale within different country contexts, particularly in light of studies highlighting a rapidly widening gap between the 'frontier' companies ready and able to leverage AI and the vast majority of other actors who are still coming to terms with the technology.[9] As one commentator notes:

"Ignoring the perspective of technological laggards can have far-reaching policy implications, especially if techno-boosterism (or alarmism) diverts attention from pressing problems facing education systems and labor markets in the here and now. If governments start allocating more resources to train the high-skilled professional elite of tomorrow, they could foster even deeper inequality today."[10]

In a similar vein, it is important to acknowledge that the applications of AI that will likely add the most value to development programming in the near future will be simple machine learning systems that optimize processes that we are already familiar with. In other words, the greatest opportunities for AI in the development space do not necessarily lie in extremely advanced or complex technologies, and we should not ignore the ways in which more 'mundane' AI solutions might be able to transform development outcomes.

1.5 How are governments responding?

It is only over the last 18-24 months that a number of national governments (predominantly from the Global North) have released AI strategies and/or committed funding for the development of AI.[11] Several, including Australia, Canada and the United Kingdom, have announced significant financial commitments to advance their agendas in this rapidly emerging field. The strategies are varied in scope, with some governments focusing more on research and development, others on policy, legal and regulatory frameworks, some on digital infrastructure and others on education and skills. It is notable that only a few Global South countries have so far developed targeted AI agendas, with India, Kenya and Tunisia among the first to announce their plans to harness this technology, with India expressing particular focus on leveraging AI for social inclusion through its 'AI for All' initiative.

At the international level, there is substantial activity around the topic. Examples include the European Union's AI Alliance; a Memorandum of Understanding between the UAE and India to establish a partnership on AI; and an international study group on inclusive and ethical AI led by Canada and France.

2. AI and the sustainable development agenda

2.1 Applications of AI across the Sustainable Development Goals

AI demonstrates significant potential to solve some of the most pressing issues facing society, by automating or augmenting human inputs in order to make processes more efficient or effective. The ability to collect and analyze vast amounts of data rapidly, and generate deeper insights to inform decision-making is already having a transformative effect on the development sector, with the significant advantage of shortening the feedback loop between monitoring and implementation to achieve better results.[12] But while the development and application of AI is gathering pace in the Global South, the field is currently dominated by a small number of countries and technology companies from the Global North. This creates risks around inequitable distribution of benefits, and perpetuating biases and social marginalization. A PwC analysis published in 2017 estimated that 70% of the global economic gains generated through AI by 2030 will accrue to China and North America.[13]

Recent research highlights a number of sectors in which technological innovations using AI are currently demonstrating great prospects for improving development outcomes. In healthcare AI is helping to survey populations and generate predictions based on health data, and is providing expertise to health workers and clinicians to diagnose and treat illnesses. In agriculture, AI is helping farmers to better understand critical conditions such as soil quality, climate and crop health in order to increase productivity and respond to risks that might compromise their harvests. In the education sector, AI is playing an important role in three key areas: learner-facing technologies (such as personalized delivery models and adaptive learning platforms), teacher-facing technologies (including automated assessment tools) and sector-level analyses (for instance, using data from across schools to predict school inspection performance).[14]

In addition, while government lags behind the private sector in terms of integrating emerging technologies, there is significant potential for AI to improve the way that public services are delivered, for instance through automating complex assessment/application procedures, personalizing services, and streamlining citizen engagement/response systems. Finally, many in the Global South are heavily reliant on the informal sector, which is expanding at pace and, given the low barriers to entry (in terms of capital and skills), driving inclusive economic growth in many countries. Digital platforms such as those for mobile money and e-commerce show great promise to boost efficiency and productivity in these informal markets, and indeed provide links with, or routes into, the formal sector.

Below are a few illustrative examples from organizations attempting to use AI with the aim to improve outcomes around the Sustainable Development Goals:

Kimetrica is using **facial recognition** technology to detect malnutrition in children aged 0-5 during humanitarian emergencies. An algorithm to analyze facial curvature and other non-traditional markers is able to estimate a child's body mass index, helping to identify those who need nutrition support much more rapidly than the traditional Mid-Upper Arm Circumference assessment method.

Babylon Health is pioneering AI to make healthcare universally accessible and affordable in Rwanda and the UK. It uses **Natural Language Processing** to transcribe consultations, summarize clinical records and chat with users in a 'human' way. **Machine learning** and **deep learning** techniques are also being used to interpret combinations of symptoms, diseases and risk factors.

Dost's program aims to encourage parents to promote children's cognitive development, socio-emotional skills, and school preparedness. Part of the Dost service includes a parent counseling hotline which uses an AI model to **transcribe and classify voice data** to understand parent requests in real time, **automate issue resolution**, and refer parents to counselors.

Springster has developed a **chatbot**, called Big Sis, to provide **expert advice** in response to girls' questions about sexual health. The platform, which operates through Facebook Messenger, provides opportunities for girls to access information confidentially, in contexts where it is often impossible to discuss sexual health and relationships due to social stigma and/or girls are given unhelpful or incorrect information.

Alto Analytics and the World Economic Forum used **AI-powered image recognition** on photographs of toilets all over the world to estimate the number of people globally impacted by unsafe sanitation conditions.

Tala has developed an innovative approach to credit scoring using **machine learning** on huge numbers of non-traditional data points from mobile devices in order to generate credit scores for the previously unbanked. Their smartphone app allows users to apply for loans and receive instant **decisions**, regardless of their financial history.

Livox is an alternative communication software that enables non-verbal people with disabilities to communicate and learn. **Machine learning** and **natural language processing** enable users with disabilities to communicate up to 20 times faster, and the technology adapts to each user's individual abilities. It is currently implemented in Egypt, Jordan and Brazil.

The Radiant Earth Foundation provides users access to free **geospatial and machine learning tools** which simplify the discovery and use of **satellite, aerial and drone imagery** in order to guide climate-positive practices. For example, one application uses satellite imagery to map areas affected by deforestation and predict sites that are vulnerable to future deforestation.

WeRobotics specializes in building local capacity for using **AI-enabled aerial and underwater robotics**. Their EcoRobotics program supports local communities in a number of developing countries to generate better data in more effective and efficient ways to support with sustainable farming and fisheries practices.

16 PEACE, JUSTICE AND STRONG INSTITUTIONS

BarefootLaw uses social media, websites and SMS to connect lawyers to individuals and small businesses seeking legal advice who are traditionally underserved. **Machine learning** is helping **identify trends** in legal needs (e.g. land conflict incidents increasing in specific areas or at particular times of year) and AI-enabled **response-automation systems** are being developed to provide quick answers to simple legal questions.

17 PARTNERSHIPS FOR THE GOALS

Organizations including GRID3 and Facebook are using satellite imagery and census data combined with machine learning to **map populations, settlements and infrastructure** with unprecedented accuracy. This technology could help improve development outcomes, with more granular information on population distribution enabling, for example, more targeted, efficient disaster response or vaccination distribution.

McKinsey's 2018 discussion paper Applying AI for Social Good mapped use cases for AI against the SDGs. The exercise indicated that certain SDGs appear much better served by emerging technologies than others. The SDGs which appeared to have most support included 'Good health and wellbeing'; 'Peace, justice, and strong institutions'; and 'Quality education'. On the other side of the spectrum, very few of the use cases targeted 'Life below water'; 'Affordable and clean energy'; or 'Clean water and sanitation'.

The examples above demonstrate the vast range of ways in which AI can be applied in development programming. Of course, it is vital to consider the depth of impact they could achieve, and this includes consideration of whether AI can provide the optimal solution to a given development challenge, weighing up the costs and benefits of AI-based options in comparison to more traditional approaches. There is also a risk of 'AI washing', whether claiming AI capability is being used or is core to an approach when it is not, or shoe-horning more complex AI solutions into projects where simpler alternatives would suffice or indeed be more effective. This is already frequently observed in business contexts and will likely also affect the public sector as demand and funding for AI grow more rapidly than the capacity to supply AI solutions and fully understand their impacts.

It is also important to consider how certain sectors that are more resilient to the impacts of AI and automation present significant opportunities for inclusive growth. The Pathways for Prosperity Commission describes how, for instance, service sectors and certain roles within agriculture and manufacturing require skills such as empathy and judgment, which AI is not currently able to supply, and that the relatively low wage costs in developing countries mean that they are well-placed to export these labor-intensive services.[17]

Having outlined above how a number of development programs are leveraging AI, there are also many opportunities for development actors to streamline their internal operations using the same technologies. As just one example, UNDP developed a system to automate their Rapid Integrated Assessment (RIA), which is used to evaluate the extent to which national development priorities are aligned to the SDG targets. Typically, this would have involved manual review of thousands of pages of documentation over the course of several weeks. Instead, a system based on natural language processing was able to complete the process in a matter of days, including identifying alignment between national strategies and SDGs that were not picked up by experts in the manual review process.

2.2 Where to find more case studies and examples

- The UN's International Telecommunication Union (ITU) launched a Global AI repository following the first AI for Good Global Summit. This database contains AI related projects, research initiatives, think-tanks and organizations that are using AI to accelerate progress towards the SDGs.
- The UN's Global Pulse Initiative works through a network of regional innovation labs and has developed more than 75 data innovation projects and tools for sustainable development.
- USAID's Artificial Intelligence in Global Health report draws on over 200 use cases of AI in healthcare and lists large numbers of companies operating in this field.

Artificial Intelligence in International Development

Navigating the Challenges

Development actors are, like everyone else, both consumers and producers of data. As consumers of data they are principally interested in its value in helping to better understand development problems, the complex ways in which they change over time, and the different responses that might be effective in addressing them. As producers, they are primarily focused on generating data that will help track the progress of their initiatives and their ultimate impact, as well as to the production of data for the public good that can help inform better decision-making around development problems.

A third emerging role is for development actors to function as data stewards or support national actors such as governments with data stewardship, understood as the process of helping to (re)imagine roles and responsibilities to steer the use of data and the application of the insights it can generate in addressing society's biggest questions and challenges. Data stewardship also includes safeguarding data privacy both now and in the future, and working to reduce the power asymmetry between the individuals that produce the data and the organizations that take advantage of it. This role is likely to become more important in the future, particularly when recognizing that AI deployment is creating a need for new jobs / responsibilities to help generate all of the data that it continuously demands.

In practice, development actors face a range of barriers that constrain their ability to fulfill these three roles within the course of their work. These include high cost of /low capacity for data collection and analysis; poor quality and reliability of data leading to a short-shelf life; and lower than desired uptake or usage of data to inform decision-making (both internally and externally). To a development actor therefore, the idea of AI is highly attractive in the sense of:

i. helping them harvest **more of the data they need** at a much lower cost;

ii. leveraging big data computing power to reveal **key trends, insights, risks and opportunities** across increasingly complex development problems; and

iii. enabling more **efficient intervention design**, decision-making, evaluation and measurement of impact – including through harnessing AI for predictive analysis and modelling. This includes reducing transaction costs through AI predictions that make service delivery affordable, as in the case of assessing the credit worthiness of smallholder farmers for a crop insurance product, where AI-enabled yield predictions could be used to make these a viable business transaction.

However, AI is unfortunately not a standardized 'plug and play' tool you can simply apply at different points of the program cycle, and the same reasons that make it attractive to development actors also render it a potential minefield when considering their commitment to 'Do No Harm'. **Articulating a clear, value-driven strategy for how AI should (and shouldn't) be deployed has therefore become a priority for any development actor looking to safely and responsibly leverage AI for development impact.** However, this is about so much more than creating a Code of Ethics or set of Guiding Principles, which, while an important first step among most organizations seeking to inform their engagement with AI, have been shown by various studies to lack any legal accountability and often remain disconnected from or too simplistic for the kind of complex decisions that AI presents in practice. Rather, it is about the more direct capacity building of development practitioners to individually own and be accountable for AI development and deployment, and the parallel empowerment of their partners and clients in the Global South to take greater control over their own data and develop, own and maintain solutions at the local level.

Challenges in deploying AI in development

1. Availability, accessibility and quality of data

AI is only as good as the data that it is based on — and **in the development world, challenges around the accessibility, quality, depth, diversity and volume of data are particularly common, especially in fragile / conflict-affected contexts.** While the pool of available data has taken a quantum leap forward in recent decades thanks to the rapid penetration of mobile telephony and technological infrastructure across the Global South, it still suffers from significant issues. For example, fewer than half of all African countries completed more than a single national survey on incomes and wealth between 2000 and 2010,[18] while the adoption rates of electronic medical records (EMRs) in LMICs is estimated to be less than 40%.[19] When combined with the problems of unreliability and fragmentation that undermine the data that has been collected, we are in a situation where, according to Kofi Annan, the development community is effectively "flying blind".[20]

In addition, as the development community and countries in the Global South more broadly have typically been slower to embrace digitized forms of data collection, much of the data that is available is not in a format readily accessible for AI deployment. This increases the **dependence of development actors on data predominantly collected by those in the private sector such as telecom companies, social media platforms and search engines which may be less relevant, inaccessible or exacerbate existing inequalities and marginalization.** For example, the existing gender digital divide exacerbates the sexist data crisis. Even then, the 'big data' that these companies generate usually need to be backed up / correlated with the 'thick data' of statistics collected through instruments such as household surveys and national censuses to avoid creating a distorted picture — which comes back to the original problem of availability. This is before taking into account the enormous data requirements of some of the more advanced forms of AI such as deep learning, where millions (rather than thousands) of datasets are needed. Finally, the scaling up of AI-enabled tools across LMICs is further complicated by (a) differing and uncertain regulatory and policy environments – not only between countries, but also across regions and states within countries; and (b) the fact that most AI applications currently available have been designed for deployment among countries in the Global North.

PROMISING INITIATIVES

(a) **Combining multiple datasets.** In Sudan, UNDP is working with partners including German academic institutions on real-time insights on poverty at the household level. Official poverty statistics are being used to ground-truth data analysis of night-time lighting, mobile phones and electricity consumption. Policymakers, and in this case UN agencies, thus get a real-time picture of when and how poverty levels per household change and are able to provide services in a more targeted manner.

(b) **Improving the relevance and utility of existing development data.** In 2018, the OECD-DAC concluded that the development community "does not yet have the right dashboard in place to monitor progress" against the SDGs, despite collecting large amounts of data through member reporting. Since then, the OECD has started to explore how they can use AI to monitor progress in SDG financing across their members to estimate how much aid targets each SDG and how the adoption of the 2030 Agenda has changed the behavior of donors. In this way, the OECD intends to use AI to make better use of the large volumes of complex information it collects, and respond faster to changes in the policy environment.

(c) **Leveraging innovative data proxies.** In an environment of scarce data resources, finding potential proxies is critical and requires thinking outside of the box. For example, UNHCR's Project Jetson ran into difficulty getting hold of the right data points its AI algorithm needed to accurately predict the movements of displaced people in Somalia. Working closely with the refugee community themselves, they identified market prices of goats as a suitable proxy to enhance their dataset, recognizing that a drop in prices typically reflected large numbers of goats being put up for sale to help fund imminent travel.

2. Capacity to engage with or use AI

Lack of technological infrastructure remains a significant barrier to the effective development and deployment of AI solutions in the Global South and in rural areas especially. Worldwide, there are around four billion people still without Internet access, which further hinders the ability of these populations to engage with — let alone design or take control of — AI technologies. More generally, those in the development sector also lag far behind their private sector cousins in terms of the expertise and resources they need to build and deploy data science capabilities. **Globally, there is a shortage of knowledge and skills in the field of AI, as well as a concentration of AI expertise in the hands of a select few (white males based in countries in the Global North) which is limiting its spread across more geographically and ethnically diverse groups.**

An additional problem noted by DataKind is that **most civil society organizations and governments still struggle to identify problems that are appropriate for AI-enabled interventions**, or determine what that intervention should actually do.[21] Similarly, **few data scientists realize what their skillsets can achieve when applied to development challenges**, with the vast majority of the world's growing data scientist population on career paths that typically steer them towards the private sector where demand and compensation are high. Partnerships between development actors are also timebound and often focused on specific projects rather than long-term capacity building.

All of this means that having some form of in-house capacity for data science is quickly becoming a critical priority for development actors who see it as a platform for both greater operational efficiency and continued relevance and efficacy within a data-driven international community. Moreover, it becomes essential to be able to leverage the potential of 'transfer learning', where AI algorithms and models from the private sector such as facial recognition are repurposed within development programs. **Those operating in this sector must consider the funding needed to train experts at the nexus of data science and international development.** In addition, it will be important to leverage investment around reducing the dependency on data from a select few private sector companies, for instance supporting the collection and production of training data to feed machine learning systems.

PROMISING INITIATIVES

(a) **Deepening data science capacity within the development sector.** In January 2019, The Rockefeller Foundation and the Mastercard Center for Inclusive Growth announced a joint, $50-million investment over five years to build the field of data science for social impact through a transformational model for collaborative philanthropy. In addition to funding, this partnership is based around a shared commitment to mobilize additional partners, resources, and networks to accelerate the use of data science by empowering non-profit, civic, and government organizations with the tools, expertise, and knowledge they need to help solve the world's most pressing challenges.

(b) **Strengthening local capacity in the Global South.** In 2019, Google opened an AI Research Lab in Accra, Ghana – the first of its kind on the continent. This is part of a broader wave of AI-related initiatives in Africa that includes the launch of an African Master's for Machine Intelligence degree, and Deep Learning Indaba, the annual meeting of the African machine learning community.

(c) **Building individual accountability.** At Nesta, experts in AI are also arguing for a targeted capacity building approach to increasing data science capacity at an individual accountability level within governments, proposing that they "ditch the codes of ethics" and choose the "harder but potentially more effective path of educating, guiding and then trusting in the professionalism of public sector staff". As part of this approach, they have developed a set of 10 questions behind the proposed use of algorithmic decision making that if staff are unable to answer render the deployment of AI 'unacceptable'.

(d) **Broadening the discipline.** The AI Now Institute, an interdisciplinary research center dedicated to understanding the social implications of AI, is also calling for educators to expand AI courses beyond computer science and engineering and create a new blend of data scientist specifically trained to integrate knowledge and expertise from social and human sciences.

3. Data ownership, privacy and security

When AI makes the news, it is usually in a negative light, ranging from exposure of the biases in AI-enabled HR recruitment platforms through to high-profile global scandals exposing the illegal manipulation of the personal data of millions of people by companies such as Facebook and Cambridge Analytica. However, these issues appear not to have dampened enthusiasm among both public and private sector actors to hand over more and more of their work to AI and automated decision systems (ADS).[22] In fact, the pace at which personal data is now being harvested from individuals and manipulated to serve different political and market-driven agendas around the world far outruns the parallel efforts to tighten data security and privacy. **Development actors and their partners working in sectors that involve large amounts of individually-identifiable data (such as health and education) are at particular risk of undermining privacy and ownership, even (or perhaps especially) when their use of this data is for apparently benign or public good purposes.**

For example, in 2018, a major corporate player in the education sector committed to improving learning outcomes for all, chose to insert "social-psychological interventions" into one of its commercial learning software programs to test how 9,000 students would respond. Without the consent or knowledge of students, teachers or parents, the company tracked whether students who received "growth-mindset" messages while learning performed better than students who did not.[23] Similarly, one municipal government in Denmark has been experimenting with a system that uses AI algorithms to identify children at risk of abuse, allowing authorities to target the flagged families for early intervention that could ultimately result in forced removals of children from their parents.[24] Meanwhile, an AI-powered voice recognition system in the UK designed to detect immigration fraud ended up cancelling thousands of visas and deporting people in error.[25] These kinds of intervention, regardless of how benevolent their overall aims may be, raise ethical and privacy concerns and underscore the importance of development actors and their partners ensuring due transparency in any solution that seeks to understand or influence an individual's behavior.

While enabling the obvious improvements arising from the rapid processing and accessibility of data, digitalization also brings a number of its own risks, not least of which is the increased vulnerability of organizations to cybercrime and malicious attacks that lead to costly data breaches. Hacking is cheaper and easier than it has ever been, while data protection software and systems are growing ever more complex and expensive for institutions. **With regulation largely unable to keep pace with (or enforce) the increasing challenge of data protection, it is more important than ever for individual organizations to fully understand and act upon these considerations. This includes thinking through in a transparent way what data is absolutely necessary to collect and whether the level of accuracy that using AI to handle that data justifies its deployment in the first place.** Very few AI models ever reach the perfection of having a 100% accuracy in their predictions, so the question of what level (50%? 75%?) is morally appropriate / efficient enough to warrant using the AI (vs. not offering the service at all, or to far fewer people) is a key one that all actors must grapple with.

PROMISING INITIATIVES

(a) **Improving the digital literacy of citizens.** The DQ Institute in Singapore is an international think-tank dedicated to setting global standards for digital intelligence education and policies. In the last few years it has curated a set of global standards and tools to guide citizens in developing their 'DQ' (Digital Intelligence Quotient), including an adaptable framework aligned with the SDGs that covers 24 skills related to digital literacy/ readiness.

(b) **Helping governments assess their level of data maturity.** DLT, a US-based company working to accelerate the deployment of new technologies within the public sector, has developed a simple tool for governments to assess their data maturity across various areas of data analytics, management, personnel, systems and governance.

(c) **Giving back data control to the individual.** Tim Berners-Lee, the inventor of the world-wide-web, is currently leading a movement to restore the power and agency of individuals in controlling their data through the open source project Solid, which is designed to give every user a choice about the storage of and access to their data.

4. Bias, discrimination and inequity

The integrity of an AI application is wholly dependent on training data that it learns from, such that it will replicate and perpetuate any biases that may exist within that data. Controlling for this is extremely difficult, not only because of the vast and often unstructured datasets that are used to train AI models, but also because the **biases within data are often unconscious and therefore unrecognizable to the person / process that collected them**. Moreover, even presuming the unlikely instance of an 'unbiased' dataset, the AI may itself generate new biases through its own learning processes which it would then incorporate into the outcomes and decisions it produces. **Even when biases are obvious, there may be little incentive to address these.** For example, studies have shown that more men than women use digital devices, so the data being used to inform population-based AI tools therefore tell us more about men's preferences than women's. Similarly, it is widely recognized that the African languages are significantly under-represented online, with huge platforms like Twitter and Google AdSense not supporting any African languages because – put simply – providers are typically only interested in languages that can make them profit. This then deepens the challenge of finding sufficient and appropriate African language data on which to train AI, perpetuating their exclusion.

Given the current scale of this 'diversity disaster'[26] within AI, the task for anyone using AI in development becomes more one of *identifying* and *mitigating* bias rather than eliminating it. Many recent efforts to do this have focused on finding technological fixes that try to encode 'fairness' within the mathematical algorithms that AI uses – something that Facebook's Mark Zuckerberg recently proposed as the solution to the company's problems in the areas of censorship, fairness and content moderation.[27] Google and IBM also have dedicated cohorts of staff researching ways to mitigate the transfer of human bias into AI decision-making. While this is likely to be an ongoing process, some are argue that "most uses of AI are likely to represent an improvement"[28] given the 180 different kinds of human bias[29] already identified. However, as others have cautioned, without a framework that accounts for social and political contexts and histories, these formulas for fairness may serve to paper over deeper problems in ways that ultimately increase harm or ignore justice.

PROMISING INITIATIVES

(a) **Tackling AI discrimination through human rights frameworks.** In May 2018, Amnesty International and Access Now launched The Toronto Declaration on Protecting the Right to Equality and Non-Discrimination in Machine Learning Systems and is encouraging organizations to endorse this.

(b) **Minimizing the risk of discrimination through best practice standards.** In March 2018, The World Economic Forum released a White Paper co-created by the Global Future Council on Human Rights that describes eight practical approaches to preventing discriminatory outcomes in machine learning, each of which are supported by guiding questions for agencies and an appendix of key actions they can take when faced with incomplete or biased data.

(c) **Grassroot community building around AI.** In the last few years, a range of communities have organically emerged to try and increase participation of marginalized populations currently under-represented in the AI community. These include both global groups such as Black in AI, Women in Machine Learning and Black Girls Code, as well as very localized initiatives such as AI Saturdays in Lagos, Nigeria. At the same time, evidence is emerging of how the research being done by AI-focused institutions in Africa is helping to overcome the data-related challenges faced by the continent by shrinking the size of datasets needed to accurately train algorithms.[30] Similarly, the lack of internet infrastructure in some areas is already propelling cutting edge technology that gets rid of its reliance on cloud services, making AI programs run faster and more securely for sensitive data, like personal information. Finally, the recent MIT Solve challenge on Work of the Future surfaced a range of solutions for how those most impacted by AI can turn it to their advantage in building new livelihoods.

5. Governance, accountability and transparency

Despite the first experiments with AI taking place almost 70 years ago, efforts to define clear governance, accountability and transparency structures around its responsible use have been very slow to materialize. The most immediate evidence of action in this space has been the increasing number of 'AI Codes of Ethics' developed at the level of individual companies, countries and regions. For example, Google – one of the biggest users and developers of AI – published its seven principles for the use of AI in June 2018, while the European Commission even more recently released their [Ethics Guidelines for Trustworthy AI] in April 2019, with the idea of 'trustworthy' being broken into 'lawful', 'ethical' and 'robust'. Many others are now following suit.

However, while the debate around the ethics of AI is important in surfacing many of the potential risks of this new technology in different contexts, it has been argued that focusing on developing ethical guidelines is largely unnecessary given (a) the existence of human rights frameworks it can tap into and which are legally enforceable (as opposed to mere 'guidelines'); and (b) it can lead those companies / agencies that have instituted a code of ethics into a false sense of security that they have somehow 'covered' or 'protected' themselves from potential AI challenges. Instead, it is argued that **organizations that are truly interested in coming to terms with the consequences (good and bad) of AI need to invest more heavily in two internal areas of capacity: first, developing their knowledge and expertise of data science and what it might offer them in terms of both augmenting their current operations and anticipating challenges in deploying AI; and second, building their scenario mapping capacities** in order to continually stimulate dialogue around how they would practically respond to different circumstances in this rapidly changing environment.

Others have argued that a focus on ethics misses one of the fundamental issues at the heart of AI – the question of power. Today, the vast majority of the world's data is in the hands of just nine companies,[31] and while approaches such as algorithm auditing and transparency mechanisms may be helpful, unless these data monopolies can be disrupted these efforts have been equated with "turning out the lights to solve climate change".[32] Tim Berners-Lee, creator of the World Wide Web, has observed: *"for all the good we've achieved, the web has evolved into an engine of inequity and division; swayed by powerful forces who use it for their own agendas."*[33]

To this end, experts are now suggesting a **shift in the focus of the global debate from 'AI for Good'** (which sits squarely within the ethical frame) to **'Data as a public good**' (which explores the more far-reaching questions of both how to democratize AI as more of an open-source technology, and how to help populations take greater control over the data they are generating every day). The ultimate objective of this agenda would be to work towards the creation of an inclusive business model for AI that works for everyone, and that governments, populations and stakeholders at any level could unite around to promote appropriate AI deployment.

The reality here is that AI currently occupies a black box area of regulation in which efforts to govern its rapid proliferation across countries fall well short of actual impact. Part of the reason for this is that **global or regional codes of ethics are typically too general / high level to be meaningfully enforced, while those at the individual institution level are often so specific and risk-averse as to stifle the potential for innovation and experimentation with AI that are the foundation of learning.** Most significantly, these efforts tend to focus above all on 'top-down' approaches designed to regulate the deployment of AI, rather than on 'bottom up' interventions to incentivize and empower individuals and their communities to take greater control of their own data and manage their own relationship with AI.

The latter point around building technological self-determination - particularly among communities in the Global South where government regulatory structures are typically weaker – is perhaps one area where development actors are well positioned to play a role. However, like all others in this space, they will have to contend with the almost inherent resistance to transparency that AI brings with it by virtue of its algorithmic complexity, which can quickly extend well beyond human capacities to understand or interrogate. Put simply, it is not always easy or even possible to fully trace the calculations that a complex application of AI makes to understand why it has arrived at a particular decision. These efforts at 'algorithm auditing' to better understand and track the component parts of an AI system and the full supply chain on which it relies are important and valuable, but well beyond the capacity of most players which undermines their ability to self-regulate more broadly.

Despite this challenge, **some jurisdictions are already taking concrete steps to legally mandate transparency in the decision-making of automated systems.** For example, the European Union now requires, by law, that machine-made decisions be explainable. At our present level of AI complexity, this is still largely feasible and yet as one commentator notes:

"As deep learning progresses, algorithmic processes will only become more incomprehensible to human beings, who will be relegated to merely relying on the outcomes of these processes, without having meaningful access to the data or its processing that these algorithmic systems rely upon to produce specific outcomes."[34]

Yet while the nascent field of algorithmic auditing may ultimately be constrained by challenges of technical complexity, the ability of those deploying AI to transparently explain their own intentions and decision-making processes against established best practices is not. The process in which algorithms are themselves developed is an important window of opportunity here, with many calling for more pro-active algorithmic governance structures at the front end that ensure meaningful democratic participation and regulation in their design.

PROMISING INITIATIVES

(a) **Embedding ethical AI principles in policy and practice.** There are a growing number of examples of countries moving beyond a code of ethics for AI. These range from tools for systematizing Algorithmic Impact Assessment within the responsibilities of government workers developed by the AI Now Institute, and recently taken on by the Canadian Federal Government; to Japan's decision to embed their AI strategy directly within its Society 5.0 social development framework. The growing Map of AI Governance that Nesta has created will also continue to be a useful tool in understanding different trends and approaches influencing effective AI governance around the world.

(b) **Growing AI leadership in the Global South.** Given the critical importance of context in the deployment of AI, many commentators are now calling for the establishment of "AI Centers of Excellence" in the Global South in order to ensure governments, companies, NGOs and communities in these countries are not forced to simply 'inherit' AI technologies and policies designed and built elsewhere, but can rather take the lead in creating and testing solutions that are locally-owned and regulated. In January 2019 the Harvard Business Review published a step-by-step article on How to Set Up an AI Center of Excellence, while the IDRC recently announced plans to launch a network of excellence in sub-Saharan Africa, the first such regional network to connect AI researchers with social scientists, ethicists, development actors, policymakers and sources of funding.

ENDNOTES

[1] Paul, A, Jolley, C and Anthony, A (2018) 'Reflecting the Past, Shaping the Future: Making AI Work for International Development', *USAID*. Download the report: https://www.usaid.gov/digital-development/machine-learning/AI-ML-in-development

[2] PwC (2017) 'The rise of robotics and AI'. View the infographic: http://usblogs.pwc.com/emerging-technology/rise-robotics-ai-infographic/

[3] Press, G (2016) 'A very short history of artificial intelligence (AI)', *Forbes*. Read the article: https://www.forbes.com/sites/louiscolumbus/2017/12/24/53-of-companies-are-adopting-big-data-analytics/#11aa857d39a1

[4] Columbus, L (2017) '53% Of Companies Are Adopting Big Data Analytics', *Forbes*. Read the article: https://www.forbes.com/sites/gilpress/2016/12/30/a-very-short-history-of-artificial-intelligence-ai/#4fa9c8cf6fba

[5] Diagram produced based on visuals produced by Geospatial World (https://www.geospatialworld.net/blogs/difference-between-ai%EF%BB%BF-machine-learning-and-deep-learning/) and Oracle (https://blogs.oracle.com/bigdata/difference-ai-machine-learning-deep-learning).

[6] Copeland, E (2019) 'Does the public sector really need a code of AI ethics?', *Nesta*. Read the blog: https://www.nesta.org.uk/blog/does-public-sector-really-need-code-ai-ethics/

[7] Chui, M et al (2018) 'Notes from the AI Frontier: Applying AI for Social Good', *McKinsey*. Read the report: https://www.mckinsey.com/~/media/McKinsey/Featured%20Insights/Artificial%20Intelligence/Applying%20artificial%20intelligence%20for%20social%20good/MGI-Applying-AI-for-social-good-Discussion-paper-Dec-2018.ashx

[8] Eggers, W, Schatsky, D and Viechnicki, P (2017) 'AI-augmented government: Using cognitive technologies to redesign public sector work', *Deloitte*. Read the article: https://www2.deloitte.com/insights/us/en/focus/cognitive-technologies/artificial-intelligence-government.html

[9] Andrews, D, Crisculo, C and Gal, P (2015), 'Frontier Firms, Technology Diffusion and Public Policy'. Download the report: https://www.oecd.org/eco/growth/Frontier-Firms-Technology-Diffusion-and-Public-Policy-Micro-Evidence-from-OECD-Countries.pdf

[10] Campanella, E (2019) 'The Digital Revolution's Silent Majority', *Asia Times*. Read the article: https://www.asiatimes.com/2019/04/opinion/the-digital-revolutions-silent-majority/

[11] Dutton, T (2018) 'An Overview of National AI Strategies', *Medium*. Read the article: https://medium.com/politics-ai/an-overview-of-national-ai-strategies-2a70ec6edfd

[12] Cohen, J and Kharas, H (2018) 'Using big data and artificial intelligence to accelerate global development', *Brookings Institution*. Read the article: https://www.brookings.edu/research/using-big-data-and-artificial-intelligence-to-accelerate-global-development/

[13] PwC (2017) 'Sizing the prize: What's the real value of AI for your business and how can you capitalise?' Read the report: https://www.pwc.com/gx/en/issues/analytics/assets/pwc-ai-analysis-sizing-the-prize-report.pdf

[14] Baker, T, Smith, L and Anissa, N (2019) 'Educ-AI-tion rebooted? Exploring the future of artificial intelligence in schools and colleges', *Nesta*. Read the article: https://www.nesta.org.uk/report/education-rebooted/

[15] Ng'weno, A and Porteous, D (2018) 'Let's Be Real: The Informal Sector and the Gig Economy are the Future, and the Present, of Work in Africa', *Center for Global Development*. Read the article: https://www.cgdev.org/publication/lets-be-real-informal-sector-and-gig-economy-are-future-and-present-work-africa

[16] Note that these are illustrative examples and their inclusion in this paper does not necessarily endorse them as effective / impactful innovations in the international development field.

[17] Pathways for Prosperity Commission (2018) 'Charting Pathways for Inclusive Growth: From Paralysis to Preparation'. Download the report: https://pathwayscommission.bsg.ox.ac.uk/charting-pathways-report

[18] Lobell, D (2019) 'A Year in the Making…', *Medium*. Read the blog: https://medium.com/atlasai/a-year-in-the-making-71a7e0323b1e

[19] USAID (2019) 'AI in Global Health: Defining a Collective Path Forward', p18. Download the report: https://www.usaid.gov/cii/ai-in-global-health

[20] Annan, K (2018) 'Data can help end malnutrition across Africa', N*ature – International Journal of Science,* Vo.555, 1 March 2018. Download the article: https://www.ncbi.nlm.nih.gov/pubmed/29493625

[21] Porway, J (2019) 'How do we ensure 'data for good' means data for all?'. Read the blog: https://www.rockefellerfoundation.org/blog/ensure-data-for-all-consider-three-principles/

[22] See the AI Now Institute website for examples of automated decision systems in operation in New York City. https://ainowinstitute.org/nycadschart.pdf

[23] AI Now Institute (2018) 'AI Now Report 2018'. Dounload the report: https://ainowinstitute.org/AI_Now_2018_Report.pdf

[24] Mchangama, J and Hin-Yan, L (2018) 'The Welfare state is committing suicide by Artificial Intelligence', *Foreign Policy*. Read the blog: https://foreignpolicy.com/2018/12/25/the-welfare-state-is-committing-suicide-by-artificial-intelligence/

[25] AI Now Institute (2018) 'AI Now Report 2018'. Download the report: https://ainowinstitute.org/AI_Now_2018_Report.pdf

[26] Paul, K (2019) ''Disastrous' lack of diversity in AI industry perpetuates bias, study finds', *The Guardian*. Read the article: https://www.theguardian.com/technology/2019/apr/16/artificial-intelligence-lack-diversity-new-york-university-study?linkId=66248341

[27] AI Now Institute (2018) 'AI Now Report 2018'. Download the report: https://ainowinstitute.org/AI_Now_2018_Report.pdf

[28] Copeland, E (2019) 'Does the public sector really need a code of AI ethics?', *Nesta*. Read the blog: https://www.nesta.org.uk/blog/does-public-sector-really-need-code-ai-ethics/

[29] Wikipedia, 'List of cognitive biases'. Read the article: https://en.wikipedia.org/wiki/List_of_cognitive_biases

[30] Snow, J (2019) 'How Africa is Seizing an AI Opportunity', *Fast Company*. Read the article: https://www.fastcompany.com/90308114/how-africa-is-seizing-an-ai-opportunity

[31] Webb, A (2019) 'The Big Nine: How the Tech Titans and their Thinking Machines could Warp Humanity'.

[32] Tanya O'Carroll, Director of Amnesty Tech at Amnesty International, presentation at the 2019 Skoll World Forum, Oxford 9-12 April 2019.

[33] Berners-Lee, T (2018) 'One Small Step for the Web…', *Medium*. Read the blog: https://medium.com/@timberners_lee/one-small-step-for-the-web-87f92217d085

[34] Mchangama, J and Hin-Yan, L (2018) 'The Welfare state is committing suicide by Artificial Intelligence', *Foreign Policy*. Read the blog: https://foreignpolicy.com/2018/12/25/the-welfare-state-is-committing-suicide-by-artificial-intelligence/

Additional resources

In addition to those linked directly in the text, the following resources were consulted in the development of this paper:

Anyoha, R (2017) 'The History of Artificial Intelligence', *Harvard University*. Read the blog: http://sitn.hms.harvard.edu/flash/2017/history-artificial-intelligence/

Cheney, C (2018) 'The promise and pitfalls of artificial intelligence for global development', *Devex*. Read the article: https://www.devex.com/news/the-promise-and-pitfalls-of-artificial-intelligence-for-global-development-91881

Genc, O (2018) 'Notes on Artificial Intelligence, Machine Learning and Deep Learning for curious people', *Towards Data Science*. Read the article: https://towardsdatascience.com/notes-on-artificial-intelligence-ai-machine-learning-ml-and-deep-learning-dl-for-56e51a2071c2

Hammond, K (2015) 'What is artificial intelligence?', *Computerworld*. Read the article: https://www.computerworld.com/article/2906336/what-is-artificial-intelligence.html

Iriondo, R (2018) 'Differences Between AI and Machine Learning, and Why it Matters', *Medium*. Read the article: https://medium.com/datadriveninvestor/differences-between-ai-and-machine-learning-and-why-it-matters-1255b182fc6

Marr, B (2018) 'The Key Definitions Of Artificial Intelligence (AI) That Explain Its Importance', *Forbes*. Read the article: https://www.forbes.com/sites/bernardmarr/2018/02/14/the-key-definitions-of-artificial-intelligence-ai-that-explain-its-importance/#132633c44f5d

Raicea, R (2017) 'Want to know how Deep Learning works? Here's a quick guide for everyone'. Read the article: https://www.freecodecamp.org/news/want-to-know-how-deep-learning-works-heres-a-quick-guide-for-everyone-1aedeca88076/

Smith, M (2019) 'Excitement, concern, and hope for AI in the Global South', *International Development Research Centre*. Read the article: https://www.idrc.ca/en/resources/perspectives/excitement-concern-and-hope-ai-global-south

The International Development Innovation Alliance has collectively created a number of resources tackling different aspects of development innovation. Download these reports, and access other useful resources and insights at idiainnovation.org.

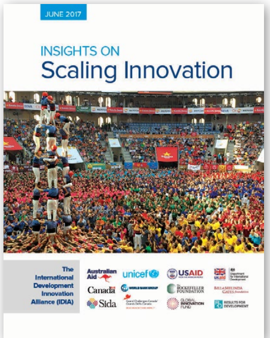

Insights on Scaling Innovation
This paper presents a high-level architecture comprising six scaling stages, eight good practices, and a matrix of influencing factors to help guide funders through the long and complex process of scaling innovation.

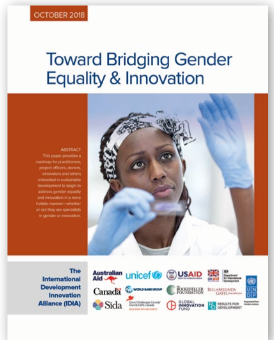

Toward Bridging Gender Equality and Innovation
This paper provides a roadmap for practitioners, donors, innovators and others interested in sustainable development to begin to address gender equality and innovation in a more holistic manner — whether or not they are specialists in gender or innovation.

Scaling Innovation: Good Practice Guides for Funders
This document explores the eight Good Practices identified in Insights on Scaling Innovation in more detail, and provides funders with further guidance on tools and knowledge products that can help them start to operationalize these good practices within the context of their own agencies.

Development Innovation Principles in Practice
This resource looks at how the eight Whistler principles adopted by the G7 Development Ministers are brought to life across a range of sectors and geographies, drawing from a shared repository of over 60 innovation stories contributed by IDIA member agencies. Questions for reflection, resources and tools for practitioners looking to integrate the principles into their own practice are also included

Insights on Measuring the Impact of Innovation
The companion to Insights on Scaling Innovation looks at challenges around measuring the impact of innovation, and presents an approach highlighting key impact domains and indicators. It also includes a case study on projecting the future impact of innovation created by Grand Challenges Canada and Results for Development.

The International Development Innovation Alliance (IDIA)

Artificial Intelligence in International Development