


The AI for Social Good Summit.

This Summit Is Jointly Developed By



Content

Policy Brief 1

Four Abilities for Governments to Leverage AI for Social Good 1

Policy Brief 2

Seven Challenges to Govern AI 8

Policy Brief 1

FOUR ABILITIES FOR GOVERNMENTS TO LEVERAGE AI FOR SOCIAL GOOD

The AI for Social Good Summit offered a virtual platform over four weeks in November 2020 for government officials, academics, industry and NGO experts to discuss and explore how artificial intelligence (AI) technology can be leveraged most effectively for the good of society in Asia. Artificial intelligence (AI) offers a myriad of technological solutions to today's complex problems and can help us "build back better" as we recover from the impacts of the COVID-19 pandemic. The optimism surrounding the transformative potential of AI, however, has been tempered by concerns regarding possible negative impacts that need to be addressed in time. This Brief provides a strategic summary of policy insights from three expert panels on building an effective enabling environment for AI for social good at the national level.

What is at stake?

With artificial intelligence poised to become as widespread as the internet, its impact in Asia will be widespread. Stories about artificial intelligence (AI) have often focused on its potential impact on the job market. This captures only a small part of its potential: from medical diagnostics for people who lack access to doctors to energy-efficient smart cities, artificial intelligence offers a significant opportunity to improve quality of life in Asia.

Countries that are effective in establishing enabling policies and environments for artificial intelligence that both protect against the risks of artificial intelligence and leverage it for social and environmental good will have the opportunity to make considerable leaps when it comes to achieving the Sustainable Development Goals (SDGs). These could include achievements in providing universal healthcare, ensuring a livable planet, and decent work opportunities for all.

Countries that do not create this enabling environment risk not only forgoing the potential upsides of artificial intelligence, but could bear the brunt of its destructive and destabilizing effects: from weaponized misinformation, to escalating inequalities arising from unequal opportunities, to the rapid displacement of entire industries and job classes.

The four abilities

Building an enabling environment that effectively leverages AI for social good requires that governments develop four abilities:

1. The ability for agile regulation.
2. The ability to negotiate interest groups and ethical considerations.
3. The ability to leverage the private sector for social and environmental good.
4. The ability to build and retain local technical know-how.

1. The ability for agile regulation.

Why it is essential:

- Traditional governance approaches to regulation – setting rules on which technologies can be used and how, plus expecting them to last several years, come short when it comes to AI. Traditional regulatory approaches can alternately fail to protect people from new and

potentially harmful AI applications, and stifle innovation that has the potential to benefit society.

- AI applications are rapidly changing and taking forms that may not have been obvious months earlier. It is not just governments that are unaware of the direction of evolution of artificial intelligence, as many businesses face challenges in predicting how their own applications might be used.
- The behaviour of AI may change over time. As AI systems are continuously learning and adapting, the way they function one week may differ from the way they function the next week.
- In the absence of clear government policy, technology companies are making their own decisions about the balance of risk and reward for citizens. Some may have entirely foregone deploying socially beneficial technologies to avoid exposure to legal risk.

How policymakers can approach this:

- Focus on regulating the acceptable outcomes produced by AI rather than specific technologies and applications. Regulators might oversee an AI system that rates creditworthiness by checking for unwarranted bias, rather than by trying to understand the specific mechanisms through which recommendations are produced. This provides both policymakers and AI creators with space to innovate and adapt.
- Create spaces for trial and error that enable the piloting of AI solutions on a limited basis and under the observation of policymakers, such as regulatory sandboxes, to understand how AI applications behave in the real world while limiting the potential downsides. Once applications are better understood, regulation can be refined and solutions can be scaled up beyond the sandbox.
- Form working groups that enable close collaboration between policymakers, regulators, AI creators, and AI users around risky applications, such as self-driving cars or medical treatment. These working groups hasten the speed of evolution of AI applications and help them get to “safe” more quickly.
- Create governance frameworks that allow policymakers to “peer into” AI models by understanding how they were constructed. “Transparency notes” can document the strengths and weaknesses of AI systems and what types of data were used to construct a model.
- Develop the technology literacy of policymakers involved in regulating AI applications.

2. The ability to negotiate interest groups and ethical considerations.

Why it is essential:

- The progress of promising and socially beneficial AI applications can be blocked by vested interests or a poor understanding of the trade-offs between privacy and social impact.
- In some countries, doctors have held back the progress of promising medical applications, while in others, taxi drivers have fought against the progress of ridesharing applications. Reactionary politics around technologies that are poorly understood can have destabilizing effects and could corner policy makers into decisions that lock out the social and environmental benefits of AI.
- The willingness of people and organizations to share data — the lifeblood of AI — and their willingness to adopt AI applications is a matter of both trust and citizens' preferences for privacy. Calibrating the balance between privacy and the benefits of AI for citizens in a particular country is valuable because it may defer from the privacy trade-offs made by businesses outside the country and outside Asia.
- Policymakers play an important role in resolving conflicts between differing interests and fostering trust between stakeholders that need to share data, and have the responsibility of protecting the rights of their citizens.

How policymakers can approach this:

- Explore creating “just transitions” to support people whose jobs may be disrupted by the advent of AI and ensure they are not left behind. This can include retraining, job transition support, and job guarantees, as well as social safety nets.
- Pursue a human-centred approach to AI and data governance. This means framing stakeholder conversations on the potential benefits and risks to people, rather than in terms of the technology. For example, in the case of diagnostics: how can we quickly get information to an anxious patient about whether they need surgery in days rather than weeks? How do we ensure that overworked radiologists get the time they need to perform effective analyses? Using human-centred questions to guide the development of policy helps stakeholders take into consideration trade-offs for the benefit of people. This stands in contrast to more one-sided questions, such as whether and how data are made available. A human-centred approach is also about ensuring that vulnerable groups are not left behind

and do not buy into things that will hurt them in the future.

- Create multi stakeholder processes that build trust between stakeholders and help reconcile ad hoc challenges (such as whether machines can review medical data or under what conditions ridesharing services can operate) to resolve conflicts in favour of citizens, especially those most vulnerable. In addition to creating progress around important issues, these processes can help people understand better AI.
- Provide people with data self-determination — the ability to decide how their data are being used, and also with the ability to make informed determinations by understanding the potential outcomes of sharing data. In the past, discussions about privacy have tended to be too one-sided (“are you willing to share your personal information?”) without helping people understand how sharing their data might create benefits for themselves or others (“this can help us significantly bring down wait times for scan results around potentially life-threatening illnesses”).
- Create data trusts that centralise, anonymise and render accessible sensitive and valuable data that might not otherwise be shared. Where the data is managed “in trust” by a third party on behalf of the people who originated the data, while making that data accessible to important AI applications.

3. The ability to leverage the private sector for social and environmental good.

Why it is essential:

- While effective regulation of AI can help protect citizens from the downsides of AI, it does not ensure that the country benefits from the transformative effects of AI for society, particularly those furthest behind.
- The majority of these transformative applications will have to come from the private sector. If policymakers do not succeed in creating an enabling environment for AI that tilts businesses toward creating social and environmental good, AI applications will focus exclusively on the most commercially viable and easiest-to-reach applications and segments of the population.
- Experts speaking at the AI for Social Good Summit believed that governments were not using the right mix of policy interventions to create an effective enabling environment for AI, let alone an effective enabling environment for AI for social good.

How policymakers can approach this:

- Match supply-side investment with demand-side enablement by leveraging the government's position as a major buyer of services and technology and its influence on other buyers. In many countries, investments in supply-side approaches such as science parks, innovation challenges, STI research, and even investment incentives for AI technology are not yielding the expected benefits. Making use of the government's role as a regulator is often not enough to leverage AI for social good. It must also leverage its roles as a market facilitator and market player. An effective way to grow the AI ecosystem is by encouraging greater use of AI, which leads to a greater speed of improvement in AI. Governments can act as market players and amplify demand for AI applications by working with priority sectors to identify specific use cases for AI and by nudging potential customers both within and outside government to adopt local AI applications.
- Shift supply-side investments towards AI applications by civil society organisations pursuing technology solutions ("civic tech"). Engaging and enabling civil society organisations around AI for social good applications builds alternatives to the private sector for building AI for social good applications.

4. The ability to build and retain local technical know-how.

Why it is essential:

- AI superpower countries are built on a critical mass of technical talent that has been trained, attracted to the country, and retained. Building a local AI ecosystem that serves local needs and potentials will require building up talent as well as retaining it.

How policymakers can approach this:

- Target and attract the diaspora to return to the country by showcasing economic opportunities and building their confidence in having a fair chance at being successful. An asset many countries in Asia have is a diaspora who have been trained in AI at leading universities and who have worked with leading AI firms. The AI ecosystems in China were seeded by foreign-trained nationals who returned attracted by greater opportunities and stayed because of their confidence in the country's meritocracy.
- Shift university incentive systems to place greater emphasis on taking products to market. Building an AI startup ecosystem requires entrepreneurial scientists who are looking to find expression beyond publication, by forming businesses, building products, and taking them to

market. Universities whose incentive and promotion systems are deeply vested in publications will have a limited impact on the AI ecosystem.

- Invest in modern learning approaches, such as short technical programmes. The primacy of universities as a primary source for education is in decline. New types of technical training programmes that provide modern skills in a matter of weeks or months are more accessible and can have a greater impact on talent availability in the ecosystem.

Taken together, these four abilities — for agile regulation, to negotiate interest groups and ethical considerations, to leverage the private sector for social and environmental good, and to build and retain technical know-how — provide a blueprint to leverage AI for social benefit.

SEVEN CHALLENGES TO GOVERN AI

The AI for Social Good Summit offered a virtual platform over four weeks in November 2020 for government officials, academics, industry, and NGO experts to discuss and explore how artificial intelligence (AI) technology can be leveraged most effectively for the good of society in Asia. Artificial intelligence (AI) offers a myriad of technological solutions to today's complex problems and can help us "build back better" as we recover from the impacts of the COVID-19 pandemic. The optimism surrounding the transformative potential of AI, however, has been tempered by concerns regarding possible negative impacts that need to be addressed in time. This brief provides a strategic summary of policy insights from three expert panels on AI governance and accountability for social and environmental benefit.

What's different about governing AI?

Artificial intelligence differs from other technologies that policymakers might regulate. AI systems are adaptive and evolving, meaning their behaviour is neither fixed nor deterministic. AI is a discipline that endows artificial structures with intelligence, meaning that AI systems gain the ability to operate autonomously. These special elements, in combination with AI's ever-growing scope of application, make regulating AI uniquely challenging.

The seven challenges

Effectively leveraging the transformative potential of AI for social good requires that governments master seven challenges:

1. Understanding the long-term implications of fast-moving technologies and effectively calibrating risk.
2. Preventing bias and unfair outcomes produced by AI systems that are not and likely cannot be made transparent.
3. Getting private sector stakeholders around the table when the government and private sector have very different motives and time horizons.
4. Having a lead agency and guarding against its policy blindspots.
5. Creating processes that can answer both moral and policy questions.
6. Addressing the conflict between making data private and limiting AI growth vs. making data open and potentially creating privacy challenges.
7. Taking advantage of emerging AI technologies in sectors that have limited readiness and few capable partners.

1. Understanding the long-term implications of fast-moving technologies and effectively calibrating risk.

The challenge:

- In creating AI-related policies, policymakers face three challenges related to uncertainty. First, AI technology and its applications are rapidly evolving. Second, today's AI applications can have long-term implications that require foresight. Third, policy decisions need to strike an adequate trade-off between potential benefits and potential risks.

How policymakers can approach this:

- Develop foresight capabilities. Policymakers must be in a position to anticipate different ways technologies might be used before the use cases arise and build consensus on how to respond. This requires developing capability in the disciplines of foresight, speculative design, and future scenario planning to govern AI use cases.
- Consider a broader range of users and stakeholders. The people who design the technology are often not the people who use it nor are they the people who regulate it. Engaging a broader range of users and stakeholders (users in rural areas, users with disabilities, and users from minority groups in addition to developers, regulators, entrepreneurs, and researchers) can help in identifying unexpected or unintended risks. Assembling diverse

voices around the table minimizes blind spots when it comes to the impacts of AI.

- Adopt a prototyping mindset. Policy should enable the practice of testing and iterating AI applications. If a product can be tested for a small subset of users or in a controlled pilot test where the users are closely monitored, governments can get early signals about the ways in which a product might have unintended harms. Regulatory sandboxes, mentioned in the Policy Insights Brief ‘Four Abilities for Governments to Leverage AI for Social Good’, are a space where technologies can be piloted within a limited jurisdiction. They enable AI creators and policymakers to see what the impacts of AI applications and AI policy are in real life. This setup allows regulators to cooperate with innovators in the private sector, observing how an AI application might work “in situ”, and provides evidence about the safety of AI applications.

2. Preventing bias and unfair outcomes produced by AI systems that are not and likely cannot be made transparent.

The challenge:

- Incomplete datasets, bias of developers, and other social biases are not simply a matter of effectiveness but also one of justice. AI systems can become a victim of “bad teaching” through datasets that are biased, leading to the perpetuation of unjust decisions. Impoverished datasets create bias because data may not represent all the users for whom an AI solution is intended. For example, one country used COVID-19 symptom data from an automated phone system but had to account for the fact that only citizens who had the skills to navigate an interactive voice response system (IVR) had their data recorded, leading to a blind spot in data about how and where COVID-19 is progressing in the country.
- Addressing the trade-off between the *sensitivity* of an AI system and its *specificity*. AI systems that are more sensitive (i.e. are more likely to make a correct determination) tend to be more specific (i.e. work only in a specific set of situations), and vice versa. For example, in healthcare, a highly sensitive system that produces a very accurate disease indication may only work for a narrow segment of people.

How policymakers can approach this:

- Implement continuous regulatory checks and system tests against criteria that define desirable and undesirable outcomes from an AI system, including biased outcomes.
- Subject important AI systems to independent peer-review to identify and address conscious or unconscious biases. Carefully evaluate AI systems and ensure that there are no known sources of bias, namely, data representation, classification, teaching accuracy, and goals and so on through testing. Create a process that brings together a diverse set of actors who will

challenge each other about their biases and ensure that different points of view are equally articulated and represented.

3. Getting private sector stakeholders around the table when the government and private sector have very different motives and time horizons.

The challenge:

- Given most AI applications are produced by the private sector, the government needs to find ways to bring the private sector to the table to leverage AI for social good. However, a mismatch in timeframes and interests can derail multi-stakeholder processes.
- While governments often take an interest in AI from the long-term perspective such as inclusion or economic development, businesses have a very different goal set focused on commercialisation, profit and loss in the short term.

How policymakers can approach this:

- Provide clear goals and short time frames for AI policy so that it can be made relevant to business. While governments are often interested in plans that span five years or more, these time periods exceed the planning horizon for most businesses. AI initiatives need to be broken down into shorter time frames and made meaningful to the broader context of the business to enable their buy-in.
- Use complementary forms of regulation. External regulation (laws and other types of formal mechanisms external to an organisation) can establish a “floor” or set of minimal behaviours for compliance. Moving beyond this floor to a greater social and environmental responsibility for protecting people who are impacted by AI technologies also requires internal or self-regulation, such as voluntary standards or internal ethics programmes, amongst businesses. Self-regulation also plays an important role in regulating emerging technologies in their initial stages, before external regulation has had the chance to catch up.

4. Having a lead policy formulation agency and guarding against its policy blindspots.

The challenge:

- Government agencies have generally not been successful in self-organising to create holistic AI policies. A lead policy formulation agency is usually needed to ensure a coordinated approach.

- The regulatory body leading the development of AI policy will shape the interests represented by that AI policy. An AI policy led by a Ministry of Trade or Economy will represent a very different set of interests from an AI policy led by the Ministry of Science and Technology — each presents an incomplete approach to governing AI.

How policymakers can approach this:

- Consider establishing leading the development of AI policy from the Prime Minister's or President's office. Alternatively, a comprehensive committee, rather than a specific government agency should lead the organisation of an AI policy, particularly in its early stages of development.

5. Creating processes that can answer both moral and policy questions.

The challenge:

- Governments sometimes allow gaps in AI regulation because governing AI is a complex topic. This complexity is partly due to the fact that attempts to regulate AI frequently surface moral issues. Sometimes these moral issues require a balancing act. For example, different groups of people may have different levels of tolerance for taking on risk.
- This challenge is further complicated by the fact that, while citizen engagement may be welcomed, topics involving AI governance may be difficult to understand without expertise.

How policymakers can approach this:

- Separate the "moralising" stage of policy from the "regulating" stage. The moralising stage in the development of a policy focuses on ethics, values, and outcomes. While stakeholders and organisations may have their own ethical frameworks, there is often a great deal of overlap and space for agreement at this stage of dialogue. More often, differences emerge between stakeholders at the regulation stage. In cases where it is harder to develop consensus around moral questions, it is useful to separate moralising and regulating conversations and focus on the moralising stage before attempting to move to policy making.
- Anticipate the full diversity of people who may use AI products to ensure that AI applications do not disadvantage minorities, underprivileged groups, or others "at the margins".

6. Addressing the conflict between making data private and limiting AI growth vs. making data open and potentially creating privacy challenges.

The challenge:

- Privacy and data availability are often seen as opposing ends of a spectrum: stronger privacy protections impede the growth of AI by limiting the availability of data. This is not necessarily the case.

How policymakers can approach this:

- Encourage “data loops” that provide feedback to users on how their data are being used. Data self-determination, mentioned in our previous policy insights brief titled ‘Four Abilities For Governments To Leverage AI For Social Good’, allows users to choose how their data is being used. These approaches allow users to retain the choice to privacy without putting large amounts of data out of reach. These approaches often require an agency letting people know how their data is being used.

7. Taking advantage of emerging AI technologies in sectors that have limited readiness and few capable partners.

The challenge:

- AI provides technological solutions that can help develop a variety of economic sectors, from agriculture, to health, energy, and transport. However, often only a few sectors have the necessary intermediary organisations and local expertise (e.g. software developers, data providers) to support the development of AI applications. Without intermediaries and expertise who can serve as partners, it may not be possible for the government to develop and implement an AI strategy beyond a few key sectors.

How policymakers can approach this:

- Screen industries that are candidates for AI-enabled growth policy for the necessary enabling organisations before building an AI strategy. Look for organisations that have experience providing or enabling AI solutions for the industry, from businesses to civil society organisations, associations, and universities, and that have a disposition to collaborating for the growth of AI solutions in the industry.
- Organisations that can support the development and implementation of AI applications are essential to being able to implement AI policies that provide industry support.
- Invest in the development of intermediary organisations in industries where AI-enablement is strategically important for the country as an early building block for building AI adoption in an industry.

Each of the seven challenges to govern AI plays a crucial role in determining how successful a country is in leveraging AI to improve quality of life and deal with critical challenges. Across the Asia-Pacific

region, policymakers are already adopting strategies and tactics allowing them to master these challenges. The successful mastery of these challenges is not just important for AI policy but holds the key to developing next-generation policymakers ready to tackle the complex and fast-moving challenges of the 21st century.

About the AI for Social Good Summit

The AI for Social Good Summit offered a virtual platform over four weeks in November 2020 for government officials, academics, industry and NGO experts to discuss and explore how AI technology can be leveraged most effectively for the good of society in Asia. Two highly interactive Policy Insight Briefings, aimed at senior government officials from the Asia-Pacific region making decisions on developing capacities and guiding the governance of AI, shared and further refined key insights from the discussion panels and the AI for Social Good research. See all summit session recordings, speaker details, and Policy Insight Briefing reports at the [AI for Social Good Summit web page](#).

The Summit is part of a broader project that aims to develop insights about how to cultivate an ecosystem that will foster and enhance AI for Social Good and maximise the technology's potential in Asia-Pacific, through collaborative effort from expertise across the region. The United Nations ESCAP, APRU, and Google partnered in 2018 to bridge the gap between the growing AI research ecosystem and limited research into AI's potential to positively transform economies and societies. [A publication](#) was produced bringing together thought pieces and research on developing an enabling environment and a governance framework. A *Project Advisory Board* of multi-stakeholder experts from Japan, Hong Kong, Australia, Indonesia, India, Korea, and Thailand provided advice and input throughout the development of the project.

For enquiries:

christina.schoenleber@apru.org

Visit apru.org/aiforsocialgood for more details

Jointly Developed By  APRU  ESCAP 

In Partnership With  THE ROCKEFELLER FOUNDATION

