A GOVERNANCE FRAMEWORK FOR ARTIFICIAL INTELLIGENCE IN QATAR

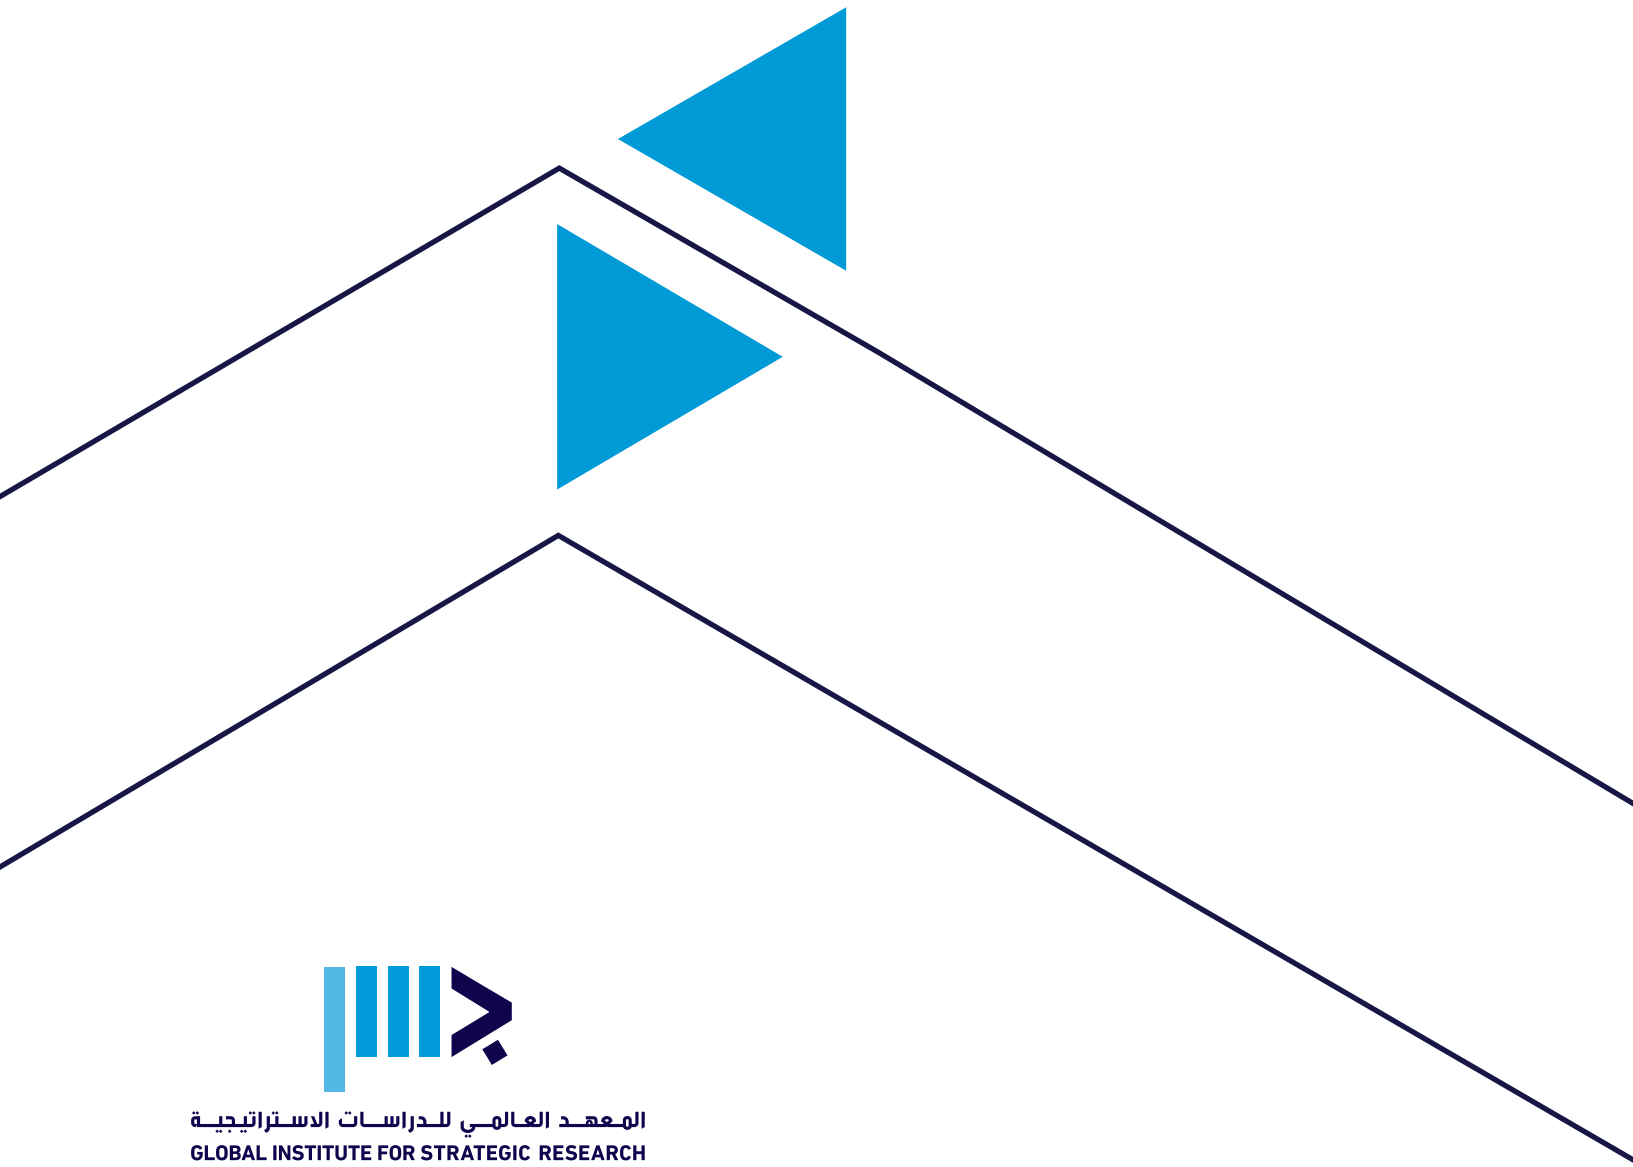

المعهد العالمي للدراسات الاستراتيجية
GLOBAL INSTITUTE FOR STRATEGIC RESEARCH

A GOVERNANCE FRAMEWORK FOR ARTIFICIAL INTELLIGENCE IN QATAR

Ahmed K. Elmagarmid

Lokendra Chauhan

Sanjay Chawla

Nizi Nazar

Ahmed K. Elmagarmid is the Executive Director, Qatar Computing Research Institute Hamad Bin Khalifa University, and an Emeritus Professor of Computer Science at Purdue University (US), and a former chief scientist at Hewlett-Packard HQ in Palo Alto, CA.
He serves on the boards of MEEZA, a leading national/regional IT leader, the Board of Qatar Personalized Health Institute and the Qatar Genome Program, Joint advisory board for Carnegie Mellon University-Qatar, and Sidra Medicine Research Committee. He serves on the Arab League ministerial level committee for AI, and co-chairs the Qatar National AI committee.

He is an accomplished and well-published scientist, authoring six books and more than 180 papers, and has directed several well-funded research programs. In 2019, Dr. Elmagarmid received the ACM SIGMOD Contributions Award. In 2020, he was chosen to serve on the Washington Institute's Council of Advisors and selected to be distinguished fellow. He is an IEEE Life Fellow, an ACM Fellow, and an AAAS Fellow.

Lokendra Chauhan is the Founder of Qen Labs Inc, which builds AI applications that use satellite imagery to measure the progress towards achieving the United Nations' Sustainable Development Goals. Prior to founding Qen Labs Inc., Chauhan worked at the Qatar Computing Research Institute, HBKU and the Qatar National Research Fund.

Sanjay Chawla is the Chief Scientist of Qatar Computing Research Institute, HBKU and leads the Qatar Center for Artificial Intelligence. Before joining QCRI in 2014, he was a Professor in the School of Information Technologies, University of Sydney, Australia and served as the Head of School from 2008-2011. He is known for his work in data mining and machine learning and was the program co-chair of ACM SIGKDD in 2021.

Nizi Nazar is a research associate at the Qatar Computing Research Institute, HBKU. She received her MSc in Computer Science from Newcastle University in 2020. Her Masters thesis was in machine learning. Her current interest includes evaluation of Large Language Models for Arabic downstream tasks.

Table of Contents

Executive Summary

As AI Systems get rapidly integrated into our daily lives, there is an urgent need to create a governance structure to regulate their deployment and use. The impact of AI on the future of human civilization will potentially be larger than the industrial revolution at the turn of nineteenth century and the information revolution that commenced in the late twentieth century. AI technology is on an exponential trajectory of change and might escape human comprehension and control if safety and human alignment are not built into their design and deployment.

While the private sector in the United States is at the core of AI innovation its practical deployment at industrial scale is being pioneered in China. The European Union has taken the lead in designing a rights-based AI governance framework and is using its commercial influence to shape the development of AI systems.

Qatar which has the modern digital infrastructure and can rapidly embrace AI technology needs to develop an AI governance framework that fits its needs. The paper makes recommendations that are suitable for Qatar.

1. AI Governance should be aligned with the Qatar National Vision of transformation to a knowledge-based economy using the pillars of human, social, economic and environmental development.

2. Adopt a risk-based classification for AI applications to ensure that consumers of AI technology are aware of their associated hazards.

3. Create an AI-free moat for critical infrastructure where all AI-driven decision-making will be mediated through a human in the loop.

4. Commission a national reference AI-stack where the innovative use of AI, data and computational requirements and safety standards can be tested in a controlled but realistic environment.

ABSTRACT

Creating a safe Artificial Intelligence (AI) innovation ecosystem trusted by everyone is key to building an economy of the future. This paper presents the case for building Qatar's AI regulations framework that captures Qatar's unique requirements and reviews three approaches to regulating AI adopted by the EU, the US, and Singapore. The recommendations suggest a path for Qatar to adapt the approaches reviewed for achieving Qatar's twin goals of safeguarding against the risks of AI while becoming an attractive destination for startups and investments in AI.

Keywords and Phrases: Artificial Intelligence, transparency, explainability, privacy, accountability, fairness, non-discrimination, intellectual property rights, innovation, policy, strategy

Introduction

With the introduction of conversational agents like ChatGPT and mainstreaming of Generative AI and Foundational Models, the power of AI systems is now self-evident. In 2021, the National Security Commission on Artificial Intelligence (AI) of the United States of America (USA) concluded that: "AI is going to reorganize the world. America must lead the charge".[1] This and other similar statements from China[2] , Russia[3], and other countries suggest that nations worldwide realize that transforming quickly into an AI-driven economy could provide them with a significant and sustainable competitive advantage.

AI has gained prominence and become pervasive in today's technology-based environments, but the risks of using AI without any constraints is not acceptable and could lead to results that are not human-aligned. Therefore, along with the rise of AI, calls for regulating AI have also increased. It is essential to build an AI innovation ecosystem that everyone trusts. Companies should be able to benefit from AI innovations and operate freely, and people should feel assured and confident when using AI technologies. Therefore, lawmakers and policymakers worldwide are working to find the best strategy to develop AI-related laws and policies.

Another critical point to note is that the pace of technological innovation in AI is very rapid and is expected to remain that way for the coming decades. This means that new risks will continue to emerge, while societal norms and the capacity of governments to enforce rules will also change. Hence, the regulations also need to continue evolving, underscoring the dilemma of policymakers to bring clarity and certainty for everyone while ensuring that the rules are flexible and do not rapidly become obsolete. As most regulators have been unable to decide how to realize the fruits of AI while mitigating the risks, countries have avoided enacting broad-spectrum AI laws.

Reasons for the need to regulate AI

To understand why AI must be regulated, one must understand how AI systems work. The popular conception of AI and its actual capabilities is often inaccurate. Therefore, this section will outline how AI systems work and why they need to be regulated.

What is AI?

While the concept of building systems that can automatically learn was first tested in 1958,[4] the paradigm that emerged to be successful in the second half of the 20th century was centered around rule-based systems. Researchers in academia and industry worked towards building computer systems that could manipulate symbols to perform logic-based tasks, thereby emulating reasoning and intelligence. This meant explicitly representing human knowledge and logical reasoning frameworks in a declarative form of facts and rules. To build such rule-based systems, experts from specific problem domains codified their acquired knowledge into hand-crafted rules that a computer could follow and execute. Tax-filing software or the earlier chess-playing program, are good examples of such rule-based systems.

The underlying assumption of this paradigm is that the world can be understood in terms of structured, logical representations or rules.[5] As someone must first define these rules; these systems effectively use a top-down approach. For clearly defined situations, these rule-based AI systems are exceptional at processing the facts and working through them, assuming that the exhaustive set of rules covering all possibilities has been codified. These white-box AI systems can reason based on the built-in rules; hence, they are intelligent, explainable and transparent for the narrowly defined problem. However, these systems fail when encountering uncertainty, which requires other dimensions of intelligence, such as perceiving observable phenomena, learning from them, and abstracting the learned knowledge for use elsewhere.

The limitations of these top-down rule-based AI systems slowly gave way to bottom-up statistical learning-based technologies, referred to as machine learning and deep learning. The main idea was to start with the simple logical principles or models as the basic units learned from data and then keep increasing the complexity by letting the system form different types of interconnections among those units.[6] The design of deep learning systems is inspired by the way neurons in the brain work.[7] The small computer programs/functions, the basic units, are also called neurons and form multiple connected layers of neural networks. The interconnections formed in the system by these neurons and their relative importance represent the learned rules governing the more significant phenomenon being studied. These interconnections are not predefined but are an emergent behavior of the system or learning. The complexity of these emergent networks makes it impossible to decipher the rules that the system has learned making them unexplainable and of black-box nature.

At the core, these deep learning systems repeatedly do specific data manipulations billions (or even trillions) of times. Each time, a different pair of data and answers (or labels) is used to determine the model performance, which helps calibrate the parameters to perform better every time. This iterative process is called model training, and after going through all the training data, i.e., data + labeled answers, we get a trained AI model that has learned to solve the specific problem at hand when encountering new data, e.g., identifying an object type in images, or recognizing a person's face. Deep learning models can solve most problems if adequate data has been provided. In other words, we show the system what we want it to do, and it learns that in a bottom-up manner.[8]

Every niche problem has nuances and they must be understood to build a narrow and problem-specific AI or deep learning model that works. Implementing these AI solutions is a highly specialized job requiring characterizing the problem domain, creating statistical models to learn the observed phenomenon, and selecting appropriate metrics to evaluate model performance.

The ability of these AI models (specifically deep learning models) to perceive and learn from data about the problem at hand, along with the exponential increase in digitization and falling computational costs, fueled the spectacular rise of AI over the last decade. These models implicitly learn a set of complex rules to predict whether an object seen in the image is similar to a category of the objects identified (i.e., whether it is a dog or a cat or a car, or a specific person) or predict what would be the next item in a sequence, for example, the suggestions that pop up while writing a text message on a smartphone. AI systems can use these learned rules to classify objects, identify patterns and detect trends.

Realizing the capabilities of AI today, organizations are changing the way they think about problems, tasks, and even business models. Wherever enough training data is available, problems, processes, and tasks are reframed so that AI can be used. For example, the book Prediction Machines[9] highlighted that even driving skill can be decomposed into many minor prediction problems that AI can solve, such as what a pedestrian is most likely to do or whether the car in front is about to make an abrupt turn.

The most recent form of AI which underpins conversational agents and large language models is known as Generative AI. As the name suggests, Generative AI uses deep learning to produce content based on the "prediction" it makes. The generated content appears to be created or ideated by a human. For example, by posing tasks as "next symbol prediction", Generative AI can produce a fluent paragraph of content in any language – natural or synthetic. The key technical breakthrough that underpins Generative AI is an algorithmic procedure to capture "context." For example, consider the two sentences:

1. Hamid went for a walk along the river bank

2. Maryah withdrew cash from a bank

There are now computer algorithms (known as transformers) that can automatically learn to distinguish that the usage of "bank" is different in the two sentences. By repeatedly carrying out this procedure over large troves of data (e.g., all digital content on the Internet), Generative AI models acquire the capacity to produce new content. While the technology is cutting-edge, scientists are unable to fully explain why it works. From the perspective of governance, the lack of credible scientific explanation to explain the human-like performance of Generative AI system is a cause of concern because the output of Generative AI does not come with guarantees about its veracity. If not properly regulated, Generative AI has the capacity to overwhelm human generated content and lead to tsunami of misinformation/disinformation.

Failures of AI

As AI is becoming pervasive in our lives and its usage is proliferating, failures of AI are also getting noticed widely. Therefore, to understand the rationale for regulating AI, knowing how and where AI falls short of its promise or expectations is imperative. Here is a list of such issues:

I) INCREASING TRAINING DATA REQUIREMENTS

AI models are statistical representations or abstractions of the available observed data. This means that the models pick the dominant patterns from the data and extrapolate those patterns to make predictions. The nuances, namely outliers or patterns that are not appropriately represented, get lost in the process. Finding innumerable examples to train such models for various outlier cases and all situations is often prohibitively costly and often impossible. For example, the latest large language models are trained on trillions of word tokens that require massive computational power to process. While the AI models are open-source, accessing and preparing training data (especially in low-resource languages) and building a computational infrastructure requires massive investment. There is a danger that only a few will be able to train such models which will exacerbate the digital divide between the rich and the poor.

II) GENERALIZABILITY

The models assume that the training data covers all existing possibilities; thus, they cannot predict what they have not seen, effectively limiting their robustness when used in the real world. As of now, neither rule-based nor learning systems can abstract or generalize knowledge learned from one domain and apply it to an entirely different domain. Also, neither of these systems currently exploits multiple and diverse varieties of knowledge and data, which would be necessary to describe and react to the whole context of any problem. For example, simultaneously extracting information from unstructured text and tabular data, something that is routine for humans is not error-proof in the latest AI systems.

III) SAFETY AND TRUST

As the power of AI models increase, their lack of safety guarantees is a cause of major concern. AI sy stems need equivalent of "seat belts" to make them safe. For example, in Generative AI, models tend to confidently produce content which is known to be false – a phenomenon known as hallucination. For now, there is no known principled way of preventing Generative AI models from hallucinating. Without strict safety measures overall trust in AI systems will decline. Similarly, there is no uniformly accepted way to watermark AI generated content so it can be attributed to a legitimate source. Achieving a balance between making a watermark detectable and keeping it hidden from unauthorized removal can be challenging. Finally, the most important source of information in AI models is captured in "model weights" which are vulnerable to being manipulated.

IV) BLACK-BOX (OPAQUE) NATURE OF AI MODELS

While the rule-based systems could provide clear explanations with reasons for their decisions and recommendations, the rules learned by today's AI systems based on deep learning are opaque, which makes these systems unexplainable and thus not understandable. This creates a massive barrier to the large-scale adoption of AI, as people and organizations cannot trust what they cannot understand.

V) TRANSPARENCY, ACCOUNTABILITY AND FAIRNESS

Decision-makers are accountable for their actions, which makes it vital for them to understand why and how any AI system is making a prediction or recommending a particular action. As deep learning-based AI systems are opaque black boxes, this lack of transparency makes it hard to pinpoint accountability.

Additionally, this lack of generalizability and explainability, combined with extreme dependence on large amounts of representative data, creates one of the most concerning problems of AI – discrimination or unfair outcomes. For example, lending decisions made by banks using such models[10] or the recidivism-risk scoring model used by the judicial departments in the US are biased against people of color.[11] Another more relevant example for Qatar is that even the large pre-trained language model of OpenAI, GPT-3, displayed extremely high anti-Muslim bias in tasks like sentence completion and analogy suggestion.[12] While newer versions of large language models have introduced filters (guardrails) that prevents generation of content that discriminates against specific communities there are easy "jailbreaks" that can circumvent the filters.

VI) MISSING LOCAL CONTEXT

The biases observed in AI models often reflect the training data used, which can be corrected by using data that better represents the actual demographics of the users – a supposedly easy solution to implement. Still, it doesn't happen. For example, Qatar, a net importer of technology solutions, is dependent on multinational corporations to supply AI systems. Still, almost no commercially available AI system (even in life-critical areas like the medical domain) is trained on data representative of Qatar's demographics. It means the recommendations or predictions of such systems will most likely be unreliable for the local users in Qatar. This could mean missed diagnosis or incorrect treatment recommendations for patients because the AI system was trained on data from population of primarily European ancestry.

VII) MODEL BIAS

All AI models have an "inductive bias." Due to gaps in data, models can learn spurious relationships which are not supported by underlying causal relationships. Due to the opaqueness of models, discovering spurious relationships before they get applied is challenging and can often lead to unfair and discriminatory outcomes.

VIII) USE OF SOFTWARE PIECES FROM THIRD PARTY SOURCES

Just like the majority of software out there, AI systems are often built using a vast array of third-party software libraries and components. Many software systems have become so intricate and dependent on numerous components that it's challenging for any single individual to fully grasp or account for their entire operation. In a study of security risks in deep learning implementations, it was found that Caffe integrates over 130 dependent libraries, while TensorFlow and Torch utilize 97 Python modules and 48 Lua modules, respectively[13]. Such components pose potential security risks to the final systems. For instance, a flaw in the Python library 'numpy' could lead TensorFlow applications reliant on it to malfunction. Additional security gaps could prompt AI platforms to incorrectly identify inputs or even allow external threats to remotely breach a system.

IX) INTELLECTUAL PROPERTY RIGHTS RELATED TO AI MODELS

AI models or algorithms are eligible for copyright protection, and data used to train AI models can be protected too. Still, it is unclear how one would protect a trained AI model on different datasets. Further complications arise due to the common practice in the AI domain of releasing the code in open source. A single AI model can be trained separately on data of people from different places like the US, China, India, and Qatar. Each of these versions would independently hold value based on what they have learned from these different datasets. Still, no class of IP rights could cover the independent ownership of these trained AI models separately.

Another IP issue often cited is who owns or gets inventor rights for the creations made by AI models. Is it the developer or data owner, or should AI systems themselves become owners/inventors?

Therefore, to ensure the fairness of AI systems, AI must become explainable, making it understandable, accountable, and trustworthy. Explainable systems will accelerate the adoption of AI, increasing the pace of innovation and economic growth that nations seek through AI. Considering the potential harm caused to society through the discriminatory results of AI systems and the benefits that will be accelerated by AI adoption through trustworthy and explainable systems, the growing demand to regulate AI to achieve these outcomes are well justified.

Governing AI

Governments around the world recognize the potential and inevitability of the AI revolution. They fully understand that AI provides a set of effective tools to compete in our fast-changing world. Therefore, the purpose of AI policies for nation-states is to maximize the benefits of AI for society while minimizing the risks of potential harm. The primary risks of AI that policymakers around the globe are concerned about include:

i. Lack of safety guarantees accompanying AI systems

ii. Unfair and discriminatory decisions by AI systems

iii. Threats to national security

iv. Increasing unemployment due to AI-enabled automation

v. Privacy-related issues

The next set of concerns that policymakers have are around adapting to an AI-driven world through changes in various laws. These include intellectual property rights for AI, traffic laws for autonomous vehicles, immigration rules to attract AI talent, and to manage cross-border data flows.

Among all these issues, the concerns around safety, unfair and discriminatory predictions and recommendations of AI systems are the most pressing and are therefore generating public pressure to regulate AI. While methods for safe AI are still in early stages of research a consensus is emerging on the principles to address concerns around the fairness of AI systems.[14] This consensus is that fair outcomes can be achieved by ensuring accountability through human control of AI systems and mandatory requirements for transparency and explainability of AI recommendations or decisions.

Different approaches to regulating AI

The only mechanism to mandate and enforce such requirements involves enacting new laws and regulations. After reviewing the policy and legislative initiatives around the globe,[15] the following four primary approaches emerge for regulating AI:

1) **Individual rights-based:** The European Union (EU) self-identifies with this approach involving deliberative stakeholder consultations in the planning stage, leading to legal mandates claiming to protect individual rights. These policies provide extensive details of execution.

2) **Free market-driven:** This approach aligns with the United States (US) ethos of capitalism and is characteristically ad hoc and uncoordinated. Currently, multiple initiatives across different parts of the US government and industry are working to influence the law-making processes in an issue-based manner at the various levels of government. It is a very organic and supposedly bottom-up approach that assumes the best of the competing ideas will win and reign.

3) **Soft law:** This approach involves coordinated deliberations with non-mandatory guidelines in the beginning. Singapore has adopted this approach of starting by nudging the stakeholders to comply voluntarily. The objective is to learn over time and enact relevant laws as and when deemed appropriate.

4) **Statist:** China's approach towards managing the transformation to an AI-driven society is statist. It does not involve explicit stakeholder deliberations and is driven by the goals defined by the state. This means that rules, whenever they are enacted, can be changed swiftly to make exceptions, or change the course.

This paper attempts to understand the approaches to frame rules for governing AI; the Statist approach does not align with that objective, so it won't be covered in this document.

The developments over the last few years in the US, the EU, and Singapore make the first three approaches clear. The 116th US Congress had a bill titled "Algorithmic Accountability Act" to address fairness concerns by mandating disparate impact assessments for AI systems.[16] The bill did not go anywhere. While many more bills promoting similar ideas and objectives are being discussed, it is not certain what regulations the US will implement. On the other hand, the European Union is close to enacting a comprehensive legislation entitled the "Artificial Intelligence Act,"[17] providing a holistic framework for regulating AI that focuses on a calibrated approach based on the risk of harm involved. Finally, Singapore's Model AI Governance Framework (Model Framework)[18] is a set of established guidelines or soft laws to be followed by private sector companies.

All three of these jurisdictions, the EU, the US, and Singapore, have declared the fairness of AI systems and strengthening AI competitiveness as their primary objectives. Their actions (or lack thereof) suggest that providing businesses with enough clarity and certainty about the regulatory environment around AI is another equally important objective. As enterprises seek predictability of regulations, those jurisdictions must invite the best talent and corporations working on AI.

As Qatar aspires to become a competitive knowledge-based economy that attracts companies, investments, and talent in the AI industry, launching a well-designed governance framework for AI would serve that aspiration well. Singapore's approach of moving fast but taking a cautionary stance on AI is driven by the same objective.

Questions faced by Policymakers

No overarching laws governing AI technologies have been enacted in any country or jurisdiction. This is because policymakers worldwide are still looking to answer many questions through deliberations. The key questions policymakers face include:

1) How to ensure the preservation of individual safety and fundamental rights in the age of AI for their citizens, especially considering the threats like potential codification or institutionalization of discriminatory practices against specific communities ?

2) How to promote AI adoption and innovation in the entire society, especially in the private sector? Determining the appropriate level of regulatory requirements to enforce is necessary to ensure that innovation does not get stifled by unnecessary compliance burdens.

3) How can public authorities and private entities strike a balance between the energy consumption and emission impacts of AI systems and the advantages of AI advancements?

4) What kind of risks does AI-enabled automation entail to the jobs of people? What policies and government interventions can counter such risks and prevent any further increase in inequality?

5) Does AI demand an entirely new set of legal and design principles, assumptions, and systems, or should be treated as an extension of existing systems and the same laws? For example:

   i) AI algorithms have a strong capacity to learn from data and can be trained to discover or invent innovative ideas by themselves. Can an AI legally claim inventorship on the patent filing for such innovations? If an individual does an apprenticeship to learn how to innovate in a domain and then goes ahead to create new inventions, then those inventions are attributed to the individual. What happens when the AI learns about a domain under the supervision of a human but later invents something that the supervisor could not? Does that imply giving AI personhood under the law?

   ii) Do we need a different kind of IP right for AI algorithms that derive their value by being trained on a specific dataset and learning from it versus the untrained AI algorithms that are essentially the same as software and can be protected through copyright/patents?

   iii) If individuals demand to be forgotten by an AI system, is simply deleting their records from the training data enough, as the AI has already learned from that data?

   iv) Do we need a new liability category between personal and corporate liability to ensure those building AI systems remain accountable but don't get harassed unnecessarily?

6) How to ensure that the law keeps evolving with the rapid and exponential growth in AI technology. Enforcing a ban on some activity or technology is more manageable when that activity is very costly or hard to do. Still, technology can change within years, so how does one enforce such bans when the technology is widely available in open source ?

Overview of the US, the EU, and Singapore's Approaches

We reviewed policy papers, legislation, regulatory proposals, and soft law for regulation in three representative jurisdictions (US, EU, and Singapore) that have witnessed significant legislative activity. Here are summaries of these reviews covering the most critical elements of the three different approaches:

EUROPEAN UNION

The European Union has been conducting vibrant policy discussions around AI since 2018; the AI Act released in Apr 2021 and the Digital Services Act (DSA) released in Dec 2020 16 are two highly consequential pieces of legislation representing the primary outcomes of these years-long efforts. The DSA introduced new rules for hosting online content. It now requires digital platforms to provide a "clear and specific statement of reasons" to users whenever any user-contributed content is removed or disabled on the platform. More importantly, it requires advertisers and platforms to show users accompanying explanatory labels with each advertisement to understand why they are seeing it.

The AI Act is the most comprehensive effort to legislate AI governance by expanding these requirements from DSA for AI explainability to increase transparency. The EU has been explicit that they want to position the AI Act as model legislation to be adopted and adapted by other countries. They have intended to take a cautious yet considerate approach to regulating AI, aiming to strike a balance between safeguarding the rights of EU citizens vs. not burdening the innovators with unnecessary compliance requirements. The core idea of the AI Act is to build a foundation of harmonized rules covering the development and use of AI systems, where the rules vary by characteristics and risks posed by AI systems.

The objective of the risk-based calibration of regulations is to ensure that trustworthy technologies get developed in the EU, as that will lead to the uptake of AI systems. People can trust technology only if they are convinced that measures to protect their safety and fundamental rights are effective. At the same time, such measures (regulations) shouldn not be cumbersome and inhibit AI innovation. The EU's approach of risk-based calibration of regulatory requirements is an attempt to resolve this trade-off.

The AI Act requires all AI applications to be classified into four risk categories. It prohibits all AI applications under the "Unacceptable Risks" category. The threshold to define unacceptable risk is when an AI application threatens the fundamental rights of EU citizens. The next category is "High-Risk" AI applications, subject to strict compliance requirements to safeguard people. This category includes AI applications for managing and operating critical infrastructure and essential private and public services like education, finance, and law enforcement. The compliance requirements for high-risk AI systems include: a) establishing a risk management system; b) ensuring appropriate human oversight; c) registration in a publicly accessible EU database; and d) risk mitigation through design development, training, and testing of AI models.

The remaining AI applications, which cover most of the existing products and services, are classified into either "Limited risk" or "Minimal risk." For such uses, the AI Act takes a soft-law approach of encouraging the adoption of codes of conduct. It only requires transparency about their service, i.e., users must be made aware that they are interacting with an AI system and then allowed to choose whether to proceed.

UNITED STATES OF AMERICA

Historically, to encourage new technologies, the US has taken the path of allowing permissionless innovation. However, regulating AI has been gaining momentum in the US because the emerging consensus is that AI is becoming a potent force with no accountability. Therefore, an increasing number of lawmakers, academics, and even industrialists support the idea of putting guardrails for steering this force in the direction of positive social impact. The congressional interest in AI and legislative activity with references to AI is at an all-time high.[19] The forces increasing this momentum include growing public distrust in AI and multiple US states and cities introducing bills or laws regulating AI uses. Internationally, the pressure comes from initiatives like the EU's introduction of the AI Act as model legislation for the

world and claiming a leadership role in ethical AI. The Pope and Vatican have joined the Rome Call for AI Ethics initiative.[20]

The US government is a significant consumer of technology, so it uses its substantial buying power to steer innovation through funding allocations and requirements for government procurement. Another incentive the US uses to drive innovation is budget allocations for research funding. Following this tradition, the US has already committed substantial additional funding for AI research and much larger allocations for government procurement. As part of the annual Defense Budget bill for 2021,[21] the US allocated $6 billion in additional funding for AI research, along with measures for increased federal-level coordination of AI initiatives and to promote explainability and accountability for AI systems. More recently, in June 2021, the U.S. Innovation and Competition Act[22] was signed into law with a $250 billion budget for advancing research and US competitiveness in advanced technologies like AI.

These funding allocations mandate that the US government source ethically and responsibly developed AI technology, which makes these requirements significant steps towards achieving accountability, safety, and fairness for AI. Additionally, the American AI Initiative, titled "Maintaining American Leadership in Artificial Intelligence (No. 13,859)," was launched through a Presidential Executive Order (EO) in 2019.[23] Its guiding principles seek "scientific diversity, economic competitiveness, and national security," the protection of civil liberties, and "foster public trust in AI technologies." While the American AI Initiative called for national-level coordination of AI-related efforts, it also advocated that regulators must protect AI innovation and growth by not imposing high or 'precautionary' standards.

In August 2019, the National Institute of Standards and Technology (NIST) released its plan for Federal Engagement in Developing Technical Standards.[24] Like the American AI Initiative, it is another overarching effort that outlines the following nine areas for the development of AI standards: (i) concepts and terminology; (ii) data and knowledge; (iii) human interactions; (iv) metrics; (v) networking; (vi) performance testing and reporting methodology; (vii) safety; (viii) risk management; and (ix) trustworthiness. In 2022, the Office of Science and Technology (OSTP) release a non-binding document "Blueprint for an AI Bill of Rights" which introduces five principles (i) Safe and Effective Systems, (ii) Algorithmic Discrimination Protections, (iii) Data Privacy, (iv) Notice and Explanation and (v) Human Alternatives, Consideration and Fallback.

While all these developments and growing public distrust in AI make it seem that the regulations on AI in the US are inevitable, there is no clear theme or exact details available yet, because a comprehensive bill like the EU's AI Act does not exist. The conventional US technology policymaking narrative assumes regulatory actions for protecting consumers' rights to be at odds with innovation and growth. Hence, the US prefers to give free rein to innovators, and if any policy prohibitions are enacted, they are implemented when abuses of such freedom are observed. This ad hoc approach of policy making has been criticized by stakeholders, academics, and industry actors for being uncoordinated, vague, non-inclusive, and relatively 'laissez-faire.'[25]

SINGAPORE

Singapore is among the most business-friendly jurisdictions in the world and attracts significant investments. As Singapore expects itself to be a more substantial consumer of AI than a producer, they need to ensure that AI deployments in the country and the whole ecosystem can be trusted. At the same time, they see an opportunity to position themselves as a destination where various stakeholders

of AI businesses can come and work together in a business-friendly environment that also provides a predictable regulatory regime for AI.

The Singaporean government adopted a human-centric risk and accountability-driven approach like the EU. They adopted the voluntary governance approach to not unnecessarily burden the companies. Singapore's Personal Data Protection Commission (PDPC) was responsible for the Model AI Governance Framework (Model Framework). It is a voluntary, sector-agnostic, and ready-to-use tool that rests on the following two guiding principles:

1) Decisions made by AI should be explainable, transparent, and fair

2) AI systems should be human-centric.

The Model Framework is expected to keep evolving with technological and public priorities developments and remain sector-agnostic; it leaves room for sector-specific guidelines to be added later. To help companies get started, it outlines four key areas of consideration that help in voluntary adherence to the Model Framework:

1. **Internal governance structures and measures** to monitor and manage AI risks through processes and training for the responsible use of AI.

2. **Determining the Level of Human involvement in AI-augmented Decision-making** to minimize the risk of harm to individuals.

3. **Operations management** for minimizing bias in data and models using a risk-based approach to implement correcting measures like data representativeness and model explainability.

4. **Stakeholder interaction and communication** ensure that AI policies are articulated and communicated well so that users can easily understand and provide feedback.

The Model Framework does not mandate algorithmic audits but proposes that companies engage third-party experts to conduct AI audits if the sectoral regulators request. It offers to conduct a cost-benefit analysis and risk assessment before AI audits. The third-party AI auditors must comply with the appropriate non-disclosure and discretion requirements, as algorithms have commercial value as intellectual property.

As Singapore and Qatar are quite similar in their size, objectives, and governance system, Singapore's soft-law approach is a good direction to pursue for Qatar as well. The best approach would be to pick elements from the EU approach but go further than Singapore regarding details, requirements, and providing regulatory certainty on AI.

Qatar's Context

The leadership of Qatar recognized early on that they needed to use the country's natural resource advantage as a hedge against future uncertainties and started working towards diversification of the Qatari economy. With a vision to gradually transform itself into a knowledge-based economy and society, Qatar National Vision 2030 (QNV 2030)[26] was adopted by the nation in November 2008. It provided a blueprint of action plans for the next few decades.

Qatar's Investments in building an AI-driven economy

While the QNV 2030 was being drafted, Qatar made significant investments in building vital higher education, research, and innovation ecosystems in parallel. The Education City campus was set up to host the satellite campuses of internationally renowned universities like Carnegie Mellon University, Northwestern University, Texas A&M University, Virginia Commonwealth University, and Weill Cornell Medical College. These investments were accompanied by establishing three national research institutes with state-of-the-facilities, the Qatar National Research Fund, and a special economic zone functioning as a science and technology park with infrastructure that could attract top technology companies to open their R&D offices in Qatar.

Qatar also has its home-grown universities, Qatar University and Hamad Bin Khalifa University, which have strong research programs in AI and other disciplines. It is a result of these strategic and continued investments that the World Economic Forum (WEF) recognized Qatar as the leading nation among MENA countries in two categories of vital importance for building an AI-based economy: 1) Quality education systems; and 2) Ease of finding skilled employees.[27]

Beyond these investments in AI research, which focus on making Qatar a producer of AI technologies, the country expects to spend significant amounts and primarily be a consumer of AI in various sectors. To develop preparedness for the consumption of AI in all walks of life, Qatar's Ministry of Transport and Communication (MoTC) launched a QR 6 billion (~US $1.65 billion) funding initiative, TASMU, with the stated objective of enabling AI technologies for generating up to QR 40 billion (~US $11 billion) in economic activity.[28] In addition, MoTC has also released guidelines for the protection of personal data of individuals, along with detailed data governance guidelines, and controls on data flow for the organizations that impact AI deployments.

To transition from a hydrocarbon to a knowledge-based economy, Qatar has invested heavily in building a physical and a digital infrastructure. However disproportionate investments in infrastructure have resulted in dramatic fall in Qatar's overall productivity[29]. To reverse the trend in declining productivity, Qatar will need to invest in projects that leverage on the built infrastructure. AI can be a driver to catalyze a reverse in Qatar's overall productivity and an AI governance framework should be designed accordingly.

Developments in the Neighborhood

In the immediate neighborhood of Qatar, Saudi Arabia and the United Arab Emirates (UAE) have declared similar initiatives to increase the adoption of AI and attract AI companies. Saudi Arabia surprised the world by granting citizenship to the robot Sofia[30] in October 2017. In that same month, UAE became the first country with a cabinet Minister for AI and created their AI Council.[31] Saudi Arabia's move is essential from the legal perspective as it makes AI a subject of law[32] and equal to a person. Following these announcements, in 2019 UAE officially adopted its national strategy for AI[33], and Saudi Arabia established a new agency called the Saudi Data and Artificial Intelligence Authority (SDAIA) "to drive and own the national data and AI agenda to help achieve Vision 2030's goals and Kingdom's highest potential".[34]

All the countries in the Middle East's Gulf region are primarily consumers of AI-enabled products and services with the large public sector; the impact of AI is expected to be felt most in the governments. That is why at the strategic level, the main objective is effective governance within the country through AI-enabled government services, education, and a robust national security apparatus.

The UAE AI Council also forged a strategic bilateral partnership with India with open engagement across borders, fostering innovation ecosystem development through AI regulatory sandboxes to facilitate startups in both jurisdictions and collaborate on futuristic research and development.[35] Within the UAE, the Abu Dhabi Department of Health has adopted a "Policy on the Use of Artificial Intelligence in the Healthcare Sector" as a soft-law guideline for using AI in Healthcare.[36] It specifies the responsibilities of various stakeholders and risks of AI systems in healthcare that need to be addressed quickly, especially for those issues that can't be tackled like other laws of the country.

National AI Strategy for Qatar

The Qatar Center for Artificial Intelligence (QCAI) was launched in 2018 by the nation's leading AI research institute, the Qatar Computing Research Institute (QCRI). After its launch, QCAI proposed a blueprint for the nation's AI strategy. The Ministry of Transportation and Communication of Qatar adopted that document and released it as the National AI Strategy for Qatar in Oct 2019.[37] This strategy identified the following six pillars for Qatar to seize the opportunity and transform itself into an AI-based knowledge economy:

1. **Talent**: Build a nurturing and attractive ecosystem for developing and upgrading local talent in AI and attract international researchers and entrepreneurs to come to Qatar and grow.

2. **Data access**: Data is the fundamental strategic resource underpinning AI technology. Qatar should develop a policy on data generation and access that will balance the trade-off between ensuring privacy and respecting cultural norms while allowing its use for creating innovative and customized solutions for its citizens. It should also initiate and lead multilateral diplomatic efforts for global data-sharing for inclusive growth in AI.

3. **AI-Augmented Jobs**: AI will transform the current state of employment worldwide. Both blue-collar and white-collar jobs will be affected. The government of Qatar needs to appoint a task force to study how the workforce in Qatar will be impacted by the emergence of AI technology and develop a plan for leveraging these changes to realize QNV 2030.

4. **Wealth Creation**: Qatar should position itself as the most attractive destination to build and grow AI businesses. It should invest in people and cloud infrastructure and take the lead in international issues surrounding AI, including policy creation, standards, and protocol setting to ensure a level playing field for all stakeholders.

5. **Transforming Qatar into an AI + X nation**: The AI-enabled future will lead to a "winner takes all" paradigm. Qatar needs to make research investments in areas of strategic importance and where the country enjoys a "natural advantage," including Arabic content businesses, oil & gas, healthcare, national security, transportation, and health.

6. **Thought Leadership in AI Ethics and Governance**: Develop an "AI Ethics and Governance" framework that is rooted in the local context and aligns with international norms. As AI expands into all aspects of life and society, its impact on policy and ethical issues related to governance, law & order, health, and warfare require thoughtful debate and leadership.

To ensure nationwide coordinated actions for implementing the AI Strategy, the government of Qatar established the Artificial Intelligence Committee of Qatar within the Ministry of Transport and Communications. The objective of this inter-ministerial committee is to supervise the coordination across ministries and government agencies and establish implementation mechanisms for the AI strategy.[38]

Finally, to build upon the National AI Strategy, QCAI released a whitepaper entitled "Impact of Artificial Intelligence on Jobs in Qatar"[39], which is focused on the pillar of AI-Augmented Jobs. The paper's objective was to inform Qatar's leadership and assist them in making informed decisions to implement the National AI Strategy for Qatar.

The current whitepaper follows the aforementioned whitepaper from QCAI. The objective of this paper is to review the legislative efforts to regulate AI abroad and recommend a path forward for Qatar to achieve its goals through AI. It will aid in developing an agile and evidence-based approach to the regulation of AI within Qatar.

Recommendations for Qatar

To realize the Qatar National Vision (QNV) 2030 of transforming Qatar into a knowledge-driven economy, the adoption of AI at all levels within the society and economy is of paramount interest. This requires taking an ecosystem approach which also ensures that the people of Qatar trust AI technologies in all aspects of their life. Adequate safeguards for consumer interests are the starting point for people to be confident in adopting AI. At the same time, ensuring businesses and innovators have a conducive ecosystem to innovate is vital for making AI a competitive advantage for Qatar. Therefore, the objectives for AI regulations in Qatar are to 1) ensure offerings from corporations do not overreach into the rights of Qatar's citizenry and 2) make Qatar an attractive jurisdiction for AI companies around the globe.

Considering the twin goals for AI governance in Qatar, our recommendations for building a nurturing AI ecosystem in Qatar also follow two themes:

i) To adopt policy ideas from other jurisdictions that align with emerging consensus around regulating AI to safeguard fundamental human values, like fairness and social justice.

ii) To select and take a leadership role on policy ideas that are foundational from the intellectual property and the emerging AI businesses perspective but have not yet gained the required momentum because of the internal political considerations of the leading jurisdictions for AI, such as the US and the EU.

Harvard University's Berkman Klein Center has analyzed 36 prominent AI governance frameworks and published their findings.[40] They found that there was a growing consensus around eight key thematic trends in AI governance: 1) privacy; 2) accountability; 3) safety and security; 4) transparency and explainability; 5) fairness and non-discrimination; 6) human control of technology; 7) professional responsibility, and 8) the promotion of human values.

Thematic ideas for Qatar's AI Governance Framework

Our review of the legislation introduced in the US and the EU for governing AI confirms Harvard University's findings about the emerging consensus on regulating AI to ensure privacy, accountability, transparency, fairness, and non-discrimination. Therefore, the case for adopting critical elements from those pieces of legislation is strong. Singapore's approach to striking a balance between promoting AI innovation and safeguarding consumer rights makes the most sense, considering the many similarities between Singapore and Qatar. Combining these two approaches, we make the following recommendations for Qatar to build an AI ethics and governance framework:

1. **Aligning AI ethics within Qatar's value system:** Qatar should develop its AI ethics principles consistent with Qatari social, cultural, and religious norms. These principles can then be translated

into policies, rules, and mandatory legal requirements by adopting and adapting relevant clauses from the EU's AI Act, Singapore's Model Framework, or other relevant documents.

2. **Soft-law approach at the national level:** In the short term, Qatar would benefit from adopting a soft-law approach of providing guidelines to the providers of AI solutions. This would communicate what regulations to expect as more evidence of their usefulness is gathered in the coming years. It would also provide flexibility to drop or update requirements based on what is realistic, enforceable, and valuable.

3. **Risk-based calibration of regulations:** Qatar should adopt the EU's risk-based classification of AI applications and adapt corresponding reporting requirements to Qatar's context. The transparency requirements determined in this way should be recommended for compliance voluntarily. Including impact assessments and records, such as report cards for AI models[41] and datasets[42 & 43] would increase transparency.

4. **Regulatory Sandbox:** For testing experimental policy ideas before enactment into law, the recommendation is to implement those in controlled settings. Qatar has Free Zones with independent jurisdiction separate from the mainland (something most countries do not have). Therefore, testing bold legal ideas by enforcing them first in a Free Zone is possible. The successful implementation of such ideas would provide Qatar with a competitive edge in attracting AI startups and external investments. The following ideas should be considered for such experiments:

   i) A new type of IP right for trained AI models: As trained AI models have commercial value but do not fit perfectly into either copyright or patent frameworks, they require a new class of IP rights. The ability to secure such rights and its commercialization potential should attract many AI businesses to Qatar. Such rights do not exist currently and would take considerable time to be included in the leading IP jurisdictions, which presents Qatar with an opportunity.

   ii) AI audits / Impact assessments: AI policy discussions around the globe end up including the concept of AI audits or impact assessments. Various templates and ideas exist for how those should be conducted, what requirements to include, and in what sequence to fulfill those to ensure minimum compliance while achieving fairness and transparency objectives. A controlled testing ground for implementing curated ideas for such assessments to learn fast and iteratively from experience can catapult Qatar into a leadership position in implementing this fundamental requirement of AI policies.

   iii) A reference AI-stack: Qatar should build a reference AI-stack to test the development and evolution of AI. A modern reference stack should comprise three layers: computing and storage infrastructure; pre-trained foundational model; application portfolio. By working with a select group of partners, different aspects of AI-stack can be investigated in controlled but realistic settings giving valuable feedback to fine-tune AI regulation and create best practices. For example, protocols that ensure safety of AI models by carrying out extensive testing, security of model weights and watermarking of AI generated content could be part of the reference stack.

Figure 1: Qatar should build a reference AI-stack which can be used to test the use of AI in realistic settings and providing feedback to fine-tune regulations.

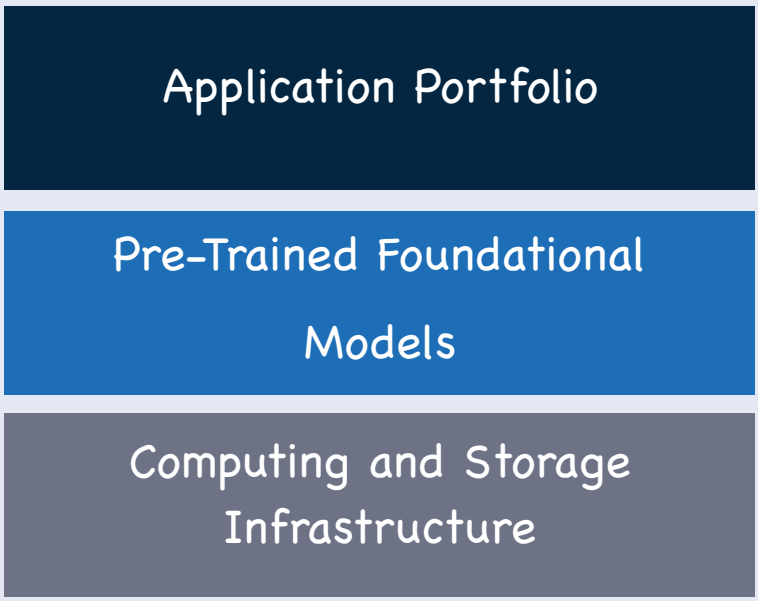

5. **Clear Guidelines, Checklists, and Support:** Clear and detailed guidelines on how to comply with the soft-law provisions and the AI regulations in sandboxed free zones would make it simple and easy for companies to comply. The information included should contain examples from each sector and publicly available tests to benchmark and report the bias in the models and data. This purpose is essential to engage with the prime stakeholders, including industry and public interest groups. They have a deep understanding of the domain and are also impacted. Additional government support services to help organizations comply with these new rules for the first few years would make enforceability easier.

6. **Qatar as a Living Lab for AI Innovations:** Even with a small population of 2.5 million people, Qatar is home to people from almost all world countries. Qatar should leverage this diversity to position the country as a living lab for AI innovations. It will involve making multiple kinds of datasets available, fast approval processes for human subject experiments, clear AI policies, and the vigorous enforcement of rules within the country for AI innovations. It is rare for companies to be able to access datasets with so much demographic diversity. Qatar is uniquely positioned to capture this opportunity to become a launchpad to take AI innovations globally. A predictable, supporting legal environment for AI businesses in the Free Zones with access to a living lab is a tremendous go-to-market strategy for startups, making it a compelling pitch for Qatar to attract AI startups from around the globe.

7. **Data Free Flow with Trust:** Recognizing AI's promise to leverage data and growing privacy concerns, most countries introduced data residency requirements to control data flow. While this distrust is well-founded, a lack of data sharing curbs the potential of AI, especially for smaller or emerging countries that do not have as much data to train AI systems. Following up on the recommendations from the National AI Strategy for Qatar, the country would benefit significantly by promoting data

sharing among emerging economies through mechanisms for "data free flow with trust." The World Economic Forum already recognizes this opportunity[44 & 45] and has been promoting cooperation agreements for the seamless functioning of AI-driven economies of the future. The ideal solution would be a combination of technological and legal or diplomatic approaches. QCRI enjoys the position of a thought leader in the database domain; therefore, building a technical and legal solution to assist Qatar's diplomatic efforts to achieve this would be a high-impact initiative.

8. **International legal accords:** National security in the age of AI is one of the most pressing concerns for all governments. The use of AI in cyberattacks and spreading misinformation targeted for creating disruptions or interfering in elections abroad are genuine threats. Article 5 of the EU's AI Act prohibits using "subliminal techniques" to "distort a person's behavior" that can cause harm in general. To penalize violators, there are huge fines imposed on large companies. However, when it is a willful violation by a small company from another jurisdiction, it is not clear how the EU will be able to enforce its law. This is a problem that all jurisdictions will face. As disconnecting from the world is not an option, the only way forward is to have international legal accords that enforce AI laws across borders. This will demand a vast diplomatic effort, but it will be a worthy investment, as taking the lead in these discussions would help shape the terms of cooperation and power-sharing.

9. **AI-free moat around critical national infrastructure:** Unlike classic engineering systems (e.g. airline or communication systems) that are ultimately based on immutable laws of physics, AI systems are data-driven and there are no know provable safety guarantees about their performance in real world settings. Thus, it is imperative that Qatar's critical infrastructure including the electricity and water grid are isolated from directly interfacing with decision-making AI systems. All interaction with critical infrastructure should be mediated through an expert human in-the-loop.

Endnotes

1 National Security Commission on Artificial Intelligence (NSCAI), NSCAI Final Report, Retrieved July 30, 2021 from https://www.nscai.gov/2021-final-report/

2 Graham Webster, Rogier Creemers, Paul Triolo, and Elsa Kania (2017). Full Translation: China's 'New Generation Artificial Intelligence Development Plan'. Retrieved on July 31, 2021 from https://www.newamerica.org/cybersecurity-initiative/digichina/blog/full-translation-chinas-new-generation-artificial-intelligence-development-plan-2017/

3 James Vincent (2017). Putin says the nation that leads in AI 'will be the ruler of the world'. Retrieved on July 31, 2021 from https://www.theverge.com/2017/9/4/16251226/russia-ai-putin-rule-the-world

4 Frank Rosenblatt (1958). The perceptron: A probabilistic model for information storage and organization in the brain. Psychological Review, 65(6), 386–408. https://citeseerx.ist.psu.edu/viewdoc/download?doi=10.1.1.335.3398&rep=rep1&type=pdf

5 Symbolic Artificial Intelligence, Wikipedia. Retrieved on July 10, 2021 from https://en.wikipedia.org/wiki/Symbolic_artificial_intelligence

6 Smolensky, P. Connectionist AI, symbolic AI, and the brain. Artificial Intelligence Review 1, 95–109 (1987). https://doi.org/10.1007/BF00130011

7 F. Rosenblatt, Perceptrons and the Theory of Brain Mechanics (Cornell Aeronautical Lab Inc., Buffalo, NY, 1961), vol. VG-1196-G, p. 621.

8 Patrick H. Winston (1990), Introduction to "Logical vs Analogical or Symbolic vs Connectionist or Neat vs. Scruffy" by Marvin Minsky, Artificial Intelligence at MIT, Expanding Frontiers Vol.1, MIT Press, 1990. https://web.media.mit.edu/~minsky/papers/SymbolicVs.Connectionist.html

9 Agrawal, A, J Gans, and A Goldfarb (2018b), Prediction Machines: The Simple Economics of Artificial Intelligence, Harvard Business School Press

10 Lisa Rice (2019), National Fair Housing Alliance Missing Credit: How the U.S. Credit System Restricts Access to Consumers of Color. Testimony to U.S. House Committee on Financial Services. Retrieved on July 10, 2021 from https://financialservices.house.gov/uploadedfiles/hhrg-116-ba00-wstate-ricel-20190226.pdf

11 Alexandra Chouldechova (2017). "Fair prediction with disparate impact: A study of bias in recidivism prediction instruments." Big Data, Jun 2017, pages 153-163. https://arxiv.org/pdf/1703.00056.pdf

12 Abid, A., Farooqi, M. & Zou, J. Large language models associate Muslims with violence. Nature Machine Intelligence 3, 461–463 (2021). https://doi.org/10.1038/s42256-021-00359-2

13 Xiao, Q., Li, K., Zhang, D., & Xu, W. (2018, May). Security risks in deep learning implementations. In 2018 IEEE Security and privacy workshops (SPW) (pp. 123-128). IEEE.

14 Fjeld, Jessica and Achten, Nele and Hilligoss, Hannah and Nagy, Adam and Srikumar, Madhulika, Principled Artificial Intelligence: Mapping Consensus in Ethical and Rights-Based Approaches to Principles for AI (January 15, 2020). Berkman Klein Center Research Publication No. 2020-1 http://dx.doi.org/10.2139/ssrn.3518482

15 Lokendra Chauhan (2021). Geospatial AI/ML (GeoAI) Applications and Policies – A Global Perspective. Policy Report. World Geospatial Industry Council. Retrieved from https://wgicouncil.org/report-geoai-geospatial-ai-ml-applications-policies-global

16 Algorithmic Accountability Act of 2019, H.R.2231, 116th US Congress https://www.congress.gov/bill/116th-congress/house-bill/2231/text

17 Artificial Intelligence Act of 2021, COM/2021/206 Final, European Commission. Retrieved from https://digital-strategy.ec.europa.eu/en/library/proposal-regulation-laying-down-harmonised-rules-artificial-intelligence-artificial-intelligence

18 Singapore's Model AI Governance Framework – 2nd edition (2020), Personal Data Protection Commission. Retreived on July 15, 2021 from https://www.pdpc.gov.sg/-/media/files/pdpc/pdf-files/resource-for-organisation/ai/sgmodelaigovframework2.pdf

19 AI Index 2021 Report: Page 171. Retrieved on July 23, 2021 from https://aiindex.stanford.edu/wp-content/uploads/2021/03/2021-AI-Index-Report_Master.pdf

20 Rome Call For AI Ethics. Retreived on July 31, 2021 from https://www.romecall.org/

21 Jon Harper, Federal AI Spending to Top $6 Billion (2021), National Defense Magazine. Retrieved on July 31, 2021 from https://www.nationaldefensemagazine.org/articles/2021/2/10/federal-ai-spending-to-top-$6-billion

22 Jake Harrigton and Riley McCabe, What the U.S. Innovation and Competition Act Gets Right and What It Gets Wrong (Jul 2021), CSIS. Retrieved on July 31, 2021 from https://www.csis.org/analysis/what-us-innovation-and-competition-act-gets-right-and-what-it-gets-wrong

23 Whitehouse, American AI Initiative (2019). Retrieved on July 31, 2021 from https://trumpwhitehouse.archives.gov/ai/

24 National Institute for Standards and Technology (NIST), U.S. LEADERSHIP IN AI: A Plan for Federal Engagement in Developing Technical Standards and Related Tools

(2019). Retrieved on July 31, 2021 from https://www.nist.gov/system/files/documents/2019/08/10/ai_standards_fedengagement_plan_9aug2019.pdf

[25] James Vincent, White House encourages hands-off approach to AI regulation (Jan 2020), The Verge. Retrieved on July 31, 2021 from https://www.theverge.com/2020/1/7/21054653/america-us-ai-regulation-principles-federal-agencies-ostp-principles

[26] Qatar National Vision 2030. Retrieved from https://www.gco.gov.qa/wp-content/uploads/2016/09/GCO-QNV-English.pdf

[27] World Economic Forum, The Future of Jobs and Skills in the Middle East and North Africa. Retrieved on July 26, 2021 from http://www3.weforum.org/docs/WEF_EGW_FOJ_MENA.pdf

[28] TASMU Smart Qatar Program. Retrieved from https://www.tasmu.gov.qa/en

[29] https://www.researchgate.net/publication/348740808_Falling_productivity_in_Qatar_and_the_challenge_it_poses_to_the_aim_of_diversification_issues_and_implications_Irfan_Aleem

[30] Saudi Arabia Is First Country In The World To Grant A Robot Citizenship (2017), Center for International Communication, Saudi Arabia. Retrieved on July 31, 2021 from https://cic.org.sa/2017/10/saudi-arabia-is-first-country-in-the-world-to-grant-a-robot-citizenship/

[31] UAE Minister of State for Artificial Intelligence Announced (2018). Retrieved on July 31, 2021 from https://uaecabinet.ae/en/details/cabinet-members/his-excellency-omar-bin-sultan-al-olama

[32] Saudi Arabia Is First Country In The World To Grant A Robot Citizenship (2017), Center for International Communication, Saudi Arabia. Retrieved on July 31, 2021 from https://cic.org.sa/2017/10/saudi-arabia-is-first-country-in-the-world-to-grant-a-robot-citizenship/

[33] UAE National AI Strategy. Retrieved on July 31, 2021 from https://ai.gov.ae/wp-content/uploads/2021/07/UAE-National-Strategy-for-Artificial-Intelligence-2031.pdf

[34] Saudi Data and AI Authority. Retrieved on July 31, 2021 from https://sdaia.gov.sa/?Lang=en&page=SectionAbout

[35] UAE, India to partner on artificial intelligence (2018), Gulf Business. Retrieved on July 31, 2021 from https://gulfbusiness.com/uae-india-sign-mou-artificial-intelligence-aim-generate-20bn-benefits/

[36] Policy on "Use of Artificial Intelligence (AI) in the Healthcare" (2018), Department of Health Abu Dhabi. Retrieved from https://www.doh.gov.ae/-/media/E9C1470A575146B18015DEBE57E47F8D.ashx

[37] National AI Strategy for Qatar. Retrieved from https://www.motc.gov.qa/sites/default/files/national_ai_strategy_-_english_0.pdf

[38] Hukoomi Qatar, Cabinet Approves Draft Decision Setting up Artificial Intelligence Committee. Retrieved on July 20, 2021 from https://hukoomi.gov.qa/en/news/cabinet-approves-draft-decision-setting-up-artificial-intelligence-committee

[39] Qatar Center for Artificial Intelligence, Impact of Artificial Intelligence on Jobs in Qatar. Retrieved on July 23, 2021 from https://qcai-blog.qcri.org/index.php/2020/11/20/impact-of-ai-on-qatars-job-market/

[40] Fjeld, Jessica and Achten, Nele and Hilligoss, Hannah and Nagy, Adam and Srikumar, Madhulika, Principled Artificial Intelligence: Mapping Consensus in Ethical and Rights-Based Approaches to Principles for AI (January 15, 2020). Berkman Klein Center Research Publication No. 2020-1 http://dx.doi.org/10.2139/ssrn.3518482

[41] Mitchell, Margaret & Wu, Simone & Zaldivar, Andrew & Barnes, Parker & Vasserman, Lucy & Hutchinson, Ben & Spitzer, Elena & Raji, Inioluwa & Gebru, Timnit. (2019). Model Cards for Model Reporting. 220-229. 10.1145/3287560.3287596.

[42] Saravanan Thirumuruganathan, Mayuresh Kunjir, Mourad Ouzzani, and Sanjay Chawla. 2021. Automated Annotations for AI Data and Model Transparency. ACM J. Data Inform. Quality, 11 pages.

[43] Saleiro, Pedro & Kuester, Benedict & Stevens, Abby & Anisfeld, Ari & Hinkson, Loren & London, Jesse & Ghani, Rayid. (2018). Aequitas: A Bias and Fairness Audit Toolkit.

[44] World Economic Forum, Rebuilding Trust and Governance: Towards Data Free Flow with Trust (2021). Retrieved on Aug 7, 2021 from https://www.weforum.org/whitepapers/rebuilding-trust-and-governance-towards-data-free-flow-with-trust-dfft

[45] World Economic Forum, Data Free Flow with Trust (DFFT): Paths towards Free and Trusted Data Flows (2020). Retrieved on Aug 7, 2021 https://www.weforum.org/whitepapers/data-free-flow-with-trust-dfft-paths-towards-free-and-trusted-data-flows

DECLARATION OF COMPETING INTEREST

The authors declare that they have no known competing financial interests or personal relationships that could have appeared to influence the work reported in this paper.

The Global Institute for Strategic Research is an independent and inter-disciplinary think tank concerned with the global challenges of governance, advancement – defined as progress, prosperity, and development – and peace (GAP). **GISR** seeks to document trends, understand processes, unpack power relations, analyze root causes, and explore policy solutions. Our work combines structural analysis with the power of imagination to go beyond the conventional extrapolation of current developments and trends. **GISR's** research and ability to convene global policy dialogues amongst stakeholders offer policymakers in Qatar (and the region) the foresight necessary for making better-informed decisions, based on having considered future eventualities, scenarios, and outcomes.

Founded by Qatar Foundation's Hamad Bin Khalifa University and partnered with a network of leading local, regional, and global collaborators, **GISR** is international in its orientation while aiming to make a meaningful impact on global policy from an Arab and regional point of view.