ASIA INTERNET COALITION

Asia Internet Coalition (AIC) Industry Response to the National Guidelines on AI Governance & Ethics for Responsible AI, Malaysia

14 March 2024

To the Ministry of Science, Technology and Innovation (MOSTI)
Government of Malaysia

The Asia Internet Coalition (AIC) and its members express our sincere gratitude to the Ministry of Science, Technology and Innovation (MOSTI) for the opportunity to submit comments on the **National Guidelines on AI Governance & Ethics for Responsible AI** (**AIGE**).

The AIC is an industry association of leading Internet and technology companies. AIC seeks to promote the understanding and resolution of Internet and ICT policy issues in the Asia Pacific region. Our member companies would like to assure MOSTI that they will continue to actively contribute to online safety on digital platforms, products and services in support of the digital economy goals of Malaysia.

We commend MOSTI for steering the model governance framework to holistically address new issues that have emerged. As part of the National Artificial Intelligence Roadmap, the draft National Guidelines on AIGE is an essential step towards the development, deployment and usage of responsible AI in Malaysia. We believe that building international consensus is key, as evidenced by the references made to the European Union’s Artificial Intelligence Act and Singapore’s AI frameworks and many others.

While we support these efforts, we also wish to express our recommendations about some of the requirements proposed in the draft framework. As such, please find attached to this letter detailed comments and recommendations, which we would like MOSTI to consider when preparing the framework.

We are grateful to MOSTI for upholding a transparent, multi-stakeholder approach and further welcome the opportunity to offer our inputs and insights, directly through industry meetings and participating in the official consultations / workshops.

Should you have any questions or need clarification on any of the recommendations, please do not hesitate to contact our Secretariat Mr. Sarthak Luthra at Secretariat@aicasia.org or at +65 8739 1490.

Thank you
Sincerely,

Paine

Jeff Paine
Managing Director
Asia Internet Coalition (AIC)

Detailed Comments and Recommendations

Article/Section	Industry Feedback
General	The document's multiple references to specific organizations within the AI field might be worth reconsidering. Given the evolving public conversation around governance within the sector, it may be more prudent to avoid highlighting individual companies. This would help ensure the focus remains on broader principles. Unless these references are strictly for illustrative purposes, we recommend a more general approach which could help the government's message maintain a neutral and inclusive tone.
Page 20: Section 2.3 Unethical Use of AI and Unacceptable Risk	We seek to clarify how the risk categories will be applied to the development and deployment of AI. Risk assessments should focus on the risk on end applications, and not the underlying general purpose technology. Neither should entire sectors e.g. healthcare or education be categorized as high risk as the risk levels depend on the end application (e.g., the low-risk use of GenAI for administrative activities in the hospital vs the more risky use of GenAI for diagnosing cancer). Different tiers of risk assessments should entail varied accountability and transparency requirements. **<u>We recommend</u>**: We support a risk-based approach to any new regulatory framework, but it is vital to ensure that it is targeted at the right use cases, taking into account the likelihood of harm and not just the severity of harm, as well as consideration of the cost of not using AI in terms of forgone benefits. We recommend scoping the risk of an AI application based on the severity of the harm and likelihood and frequency of its impact because it allows for various combinations of severity/likelihood to qualify as

Article/Section	Industry Feedback
	high-risk. Regulation should include guidance on when the risk classification of a given AI application flips from low or medium to high and reflect that the goal is to mitigate the severity of harm while also reducing its likelihood. An example of such a risk framework can in fact be found in the recent ASEAN Guide on AI Governance and Ethics (see pages 24-25).
Page 23: Section 2.5.1 AI Principles	We understand that while the entire framework is meant to be voluntary, compliance with Section 2.5.1 will be mandatory. We would appreciate clarification on whether the expectation is for organizations to comply with the framework of the 7 principles, or the 7 principles themselves. While this may seem like a nuanced difference, it is important to highlight this as, in reality, responsible AI is far more complex and almost always will involve tradeoffs across these principles. For instance, in order to ensure fairness, an organization may need to collect personal data to conduct fairness assessments or to train an AI model not to be biased against certain personal attributes. However, this could be in conflict with privacy. The organization will therefore have to make tradeoffs between the principles of fairness and privacy, even while complying with the framework of the 7 principles. **We recommend:** We suggest that it be made clear that the expectation is for organizations to comply with the framework of the 7 principles, and not necessarily all the 7 principles. It would be helpful if an example can be provided to illustrate this distinction, similar to the one we've provided
Page 29: Section 2.5.5 Seven (7) AI Principles interacting with Section 2.5.3 Responsible AI in Islamic Perspective and Section 2.6 Rukun Negara, Federal Constitution & AI	In discussing about shared responsibility, we urge caution in invoking economic concepts such as indemnities and insurance at an early phase when (i) the roles and responsibilities of the players in the AI ecosystem are still being worked out and (ii) there has been no extensive study conducted in this space on what are the risks and how they should be apportioned between the various players as well as (iii) the informational asymmetries or externalities that may exist in the AI ecosystem. As such, it would be too premature to discuss indemnities and insurance at such a nascent phase of AI development.
Page 36: Section 3 Guidelines for Stakeholders	Section 3 lays out the various AI stakeholders into three categories: Part A for "end users", Part B for policymakers, and Part C for "developers, designers, technology providers and suppliers". Ideally,

Article/Section	Industry Feedback
	we recommend that these categories be mutually exclusive and collectively exhaustive. Most AI policy frameworks start with at least three key stakeholders: developer, deployer, and user. The developer is an individual or entity that develops an AI model and provides the model for use by the deployer. The deployer is the individual or entity that uses the AI model provided by the developer to offer a service to the user. The user is the recipient of the services offered by the deployer. While these three stakeholders may not necessarily be exhaustive, they are necessary and fairly sufficient to establish the fundamental building blocks of an AI ecosystem. In Section 3, it is unclear which category the concepts of developer, deployer and user should be mapped to. Part A for “end users” sounds like it may refer to the deployer, but at times, it also seems to refer to the user. While Part C is primarily related to “developer”, there are times when it also seems to include deployers. To be clear, a single entity can certainly play more than one role. However it should be clear when its role switches from, say, a developer to a deployer. **We recommend:** We recommend starting Section 3 with an overview of the lifecycle of an AI model (eg, acquiring the AI infrastructure, developing the AI model algorithm, training the AI model, fine-tuning the AI model, deploying the AI model to offer a service, consuming the service that is powered by AI), and the key stakeholders involved in each of these stages. With that, a clear shared responsibility framework can be developed to detail different roles and responsibilities among the key stakeholders, which we recommend to be developers, deployers, and end users. The shared responsibility framework needs to ensure that there is a clear delineation of roles and responsibilities between AI developers/providers and deployers.
Page 43: Part A2.2 Elaboration of Seven AI Principle for End Users	This section states that “AI systems should not discriminate based on race, gender, or religion, and algorithm developers must be cautious of unintentional biases in the data. It is also important to consider the equitable distribution of AI benefits to avoid leaving certain groups without access to its advantages.” These two sentences, while well intended, may contradict each other at times. In the case of underprivileged segments of the population, in order to ensure

Article/Section	Industry Feedback
	equitable distribution, such segments may need to be given greater access than other segments. Some may argue that this is also a form of discrimination. **<u>We recommend</u>**: We suggest that throughout the framework, the use of concepts such as bias and discrimination should be qualified. Some form of bias and discrimination may be necessary to ensure equitable treatment. Therefore, it is unfair or unjustified bias and discrimination that should be avoided.
Page 54: Part B2.0 Responsible AI for Policy Maker on FAIRNESS “This may involve creating guidelines for unbiased data collection, algorithmic decision-making, and regular audits to detect and rectify biases.”	Government interventions should focus on mitigating risks on the output level instead of focusing on the inputs. An excessively interventionist approach on how data is collected, and algorithmic decision-making can conflict with trade secret protections, impede innovation and create security vulnerabilities. We recommend instead that the government set out certain outcomes e.g. to ensure fairness to the extent feasible and to offer avenues for redress/feedback, which the industry should meet, instead of prescribing how exactly they should do so. Audits should be based on certain international standards. We note however, that international and industry standards on audits are still in the midst of development. **<u>We recommend</u>**: We recommend that any independently validated risk assessment or audit be aligned to international, industry-accepted criteria to ensure consistency. Further, independent auditors would need to be professionally qualified and entrusted to only certify organizations that meet the appropriate standards. Regulators would need to balance such audit requirements against the risk of creating security vulnerabilities, exposing trade secrets and confidential information, or hindering innovation or the development of useful applications.
Page 54: Part B2.0 Responsible AI for Policy Maker on RELIABILITY, SAFETY AND CONTROL	Due to the nascent nature of the industry, there is a lack of benchmarks and consensus standards on reliability. Introducing a certification process is akin to licensing every AI application and technology that is developed. This will greatly hinder innovation, add unnecessary friction to the development and deployment of AI, and disproportionately hurt SMEs and start-ups.

Article/Section	Industry Feedback
“Develop certification standards for AI systems to ensure their reliability. Policymakers can mandate adherence to these standards and establish a certification process to assess the resilience of AI technologies in various conditions.”	**We recommend:** We suggest certification requirements be removed and increase focus on governance and testing framework instead.
Page 54: Part B2.0 Responsible AI for Policy Maker on PRIVACY AND SECURITY “This involves setting standards for obtaining informed consent, ensuring data security, and defining the permissible uses of personal information.”	Large language models are trained primarily on publicly available information, which could include personal information, on the open web. Obtaining informed consent for every single website on the internet is practically impossible. **We recommend:** Delete the requirement on obtaining informed consent. This recommendation should also apply throughout the framework, such as on Pages 29 and 43 as well.
Page 54: Part B2.0 Responsible AI for Policy Maker on TRANSPARENCY “This includes requirements for organizations to provide clear explanations of how their AI algorithms work, disclose data sources, and communicate the impact of AI decisions on individuals or groups.”	Meaningful transparency can help to build users' trust. However, requiring the explanation of how AI algorithms work in technical detail is not meaningful to the end user who often does not have the technical expertise to understand it. Requiring the disclosure of data sources also could be in tension with trade secrets, and is practically impossible for large language models that are trained on the open web. **We recommend:** We recommend that transparency requirements be based on these four general principles: 1. A developer/provider of an underlying foundation model should provide documentation outlining how the model is intended to be used, known inappropriate uses, known risks, and recommendations for deployers and users to manage risk. 2. The organization deploying an AI application should be solely

Article/Section	Industry Feedback
	responsible for any disclosure and documentation requirements about the AI application because it is best positioned to identify potential uses of a particular application and mitigate against misuse. 3. Whenever AI is playing a substantive role in decision-making (such as the allocation of government services or healthcare), that fact should be easily discoverable. 4. Disclosures should be presented in clear, salient language to be meaningful to a wide audience and should provide an overview of the key tasks with which the AI is being deployed to assist, within the context of the application being offered.
Page 56: Part B2.2 Open-Source Data Sharing Open-source AI allows for easier identification of biases and malfunction (algorithmic auditing), allowing for rectification before and after deployment. Hence, open AI systems promote transparency, accountability and therefore help to discover and mitigate biases and other risks.	There appears to be a conflation and confusion between open datasets and open models. **We recommend:** Point 2 should be removed as it relates to open models rather than open datasets.
Page 57: Part B2.3 AI Regulations	There should be a more holistic ecosystem approach in which regulation is not the only tool to mitigate risks of AI technologies. There are different venues to ensure responsible AI development and deployment, including co-regulation, public education, standards of procurement for government.
Page 58: Part B2.3. Initiatives on the formulation of AI ACT to support Responsible AI	There are some factual inaccuracies in the slide. For instance, the US has not enacted nor is anywhere close to enacting AI regulation. It is instead pursuing AI governance via mechanisms such as the White House Voluntary Commitments and the WH Executive Order on AI.

Article/Section	Industry Feedback
	<u>We recommend:</u> Updating the draft with the latest developments.
Pages 57-60: Part B2.3 AI Regulations	We applaud the approach to build on existing regulations and laws where relevant, instead of rushing to "regulate AI". We recommend considering existing rules that are applicable to AI applications. For example, in most jurisdictions, discrimination in lending is clearly defined in existing law, whether or not loan decisions are made by a human loan officer or an algorithm. Where existing discrimination laws provide clear guidelines and accountability mechanisms, new rules may be unnecessary. Governments should look first to existing regulatory experts, frameworks, and instruments that may encompass AI applications. Such sectoral experts typically will be well-positioned to assess context-specific uses and effects of AI and to determine whether and how best to regulate them, although sometimes additional resources may be required, including internal technical AI expert capacity. For instance, health-focused agencies are best positioned to evaluate the use of AI in medical devices and energy regulators are best positioned to evaluate the use of AI in energy production and distribution. It will also be useful to have consistency in oversight and the expectations for human and machine actors performing the same task unless there are justifiable grounds for difference.
Page 60: Part B2.3 AI Regulations	The recommendation that "It is best to have different and separate guidelines on AI for government, private sector and industry and the general public" is somewhat confusing. While we agree that "These groups have different level of expectations, responsibilities and duties towards responsible AI", there needs to be consistency in how AI is governed across the economy. **<u>We recommend:</u>** We recommend that there be a single set of guidelines for AI that is sector-agnostic to ensure that there is an overarching national policy framework. Sectoral regulators can then apply this national policy framework to their individual sectors, since the implementation of these guidelines almost always depends on the use case, where sectoral expertise is necessary. Our recommended approach seems to be aligned with the approach set out on Page 51, so it might be helpful to

Article/Section	Industry Feedback
	clarify the potential confusion.
Page 70: Part B3.6 Create Your Own Responsible Checklist	We recommend developing a more flexible approach focusing on the end goals/outcomes to allow a thriving environment for innovation.
Page 72: Part B4.1 Significant Roles of Independent Advisory Agency	This section recommends an independent advisory agency to "Establish and enforce regulations related to AI principles." However this centralized approach seems to contradict the sectoral approach discussed on Pages 51 and 60. **<u>We recommend:</u>** We recommend that this be clarified, preferably aligned with our aforementioned recommended approach of a national policy framework, but enforced by sectoral regulators.
Page 77: Part C2.0 Responsible Principle for Sector Players	There needs further alignment on responsible principles for sector players and the guidelines for policymakers. For instance, point #3 PRIVACY AND SECURITY does not involve informed consent mechanism for sector players while principles for policy makers require informed consent. **<u>We recommend:</u>** Align recommendations for policymakers with the proposed application of the principles for sector players in Section C.
Page 78: Part C2.2 Responsible AI Algorithm Development On Transparency: "Providing explanations and justifications for the decisions made by AI	Like any system, including human-based processes, AI systems are not perfect. They do, however, offer the opportunity to dramatically improve on current human-based decision making. Thus, the operational benchmark for AI systems should not be perfection, but instead the performance of comparable current processes (if existing) or an available human-powered alternative. There is a real risk that innovative uses of AI could be precluded by demanding that AI systems meet a standard that far exceeds that required of non-AI approaches. Sometimes this may be deliberate due

Article/Section	Industry Feedback
algorithms. This includes making the decision-making process transparent and understandable to users.”	to artificial protectionist constraints, but more often it is likely to be due to a lack of understanding about hidden flaws in existing non-AI decisions, and people’s natural tendency to be more forgiving of mistakes made by a human vs a machine. There should be parity in terms of expectations between AI and non-AI approaches, unless there is a clear justification put forward as to why it should differ for a particular use case and context. **We recommend:** We recommend deletion of this section.
Page 78: Part C2.2 Responsible AI Algorithm Development On User Feedback: “Incorporating user feedback into the algorithm development process. This includes actively seeking input from users and incorporating their perspectives and needs into the design of the AI system.”	User feedback is helpful after the app is deployed, which can then be taken into consideration for the improvement of the app. It is highly unlikely that user feedback would be useful in the algorithm development process itself. **We recommend:** We recommend deletion of this section.
Page 81: Part C2.4 Data Sharing “Include information related to the lineage of datasets used, model training and selection process and expected behaviour of the AI solution to help organizations that do not develop AI models in-house to be able to deploy AI models, and the AI systems should	Where relevant, the sharing of open data sets can be helpful to encourage innovation and broader AI ecosystem development. However, industry should not be compelled to do this at the expense of trade secret protection and the security of models. For instance, publicly sharing how the model was trained could open up new risks of bad actors circumventing and attacking the AI system. **We recommend:** We recommend deletion of this section.

Article/Section	Industry Feedback
still be set up and comply with AI Governance guidelines.”	
Page 87: Part C 3.3 Embedding Life Cycle in Human-Centred AI	The 4 phases of an AI life cycle also references specific tools and metrics associated with each phase. These references may appear like a recommendation or a restrictive list of tools and metrics to be applied. **We recommend:** We recommend that the specific tools and metrics be removed or clear verbiage included to indicate it being an example or is provided for reference only.