PERSONAL DATA PROTECTION COMMISSION SINGAPORE

ADVISORY GUIDELINES ON USE OF PERSONAL DATA IN AI RECOMMENDATION AND DECISION SYSTEMS

Issued 1 Mar 2024

Supported by:

INFOCOMM MEDIA DEVELOPMENT AUTHORITY

In support of:

SG:DIGITAL

TABLE OF CONTENTS

PART I: EXECUTIVE SUMMARY

1.1 The Personal Data Protection Act 2012 (**the "PDPA"**) establishes a general data protection law in Singapore that governs the collection, use and disclosure of an individual's personal data by organisations in a manner that recognises both the right of individuals to protect their personal data and the need for organisations to collect, use and disclose personal data for purposes that a reasonable person would consider appropriate in the circumstances.

1.2 The purpose of the Advisory Guidelines on the Use of Personal Data in AI Recommendation and Decision Systems ("**Guidelines**") is to provide organisations with certainty on when they can use personal data to develop and deploy systems that embed machine learning models ("**AI Systems**"), and give consumers assurance on the use of their personal data in AI Systems, since they are typically used to make autonomous decisions or assist a human decision-maker through recommendations and predictions.

1.3 Generally, organisations can use personal data where there is meaningful consent. Alternatively, organisations can rely on exceptions to consent under the PDPA, e.g., for business improvement or research purposes. The Guidelines set out criteria for how these exceptions can apply. To illustrate, the business improvement exception ("**Business Improvement Exception**") can apply when businesses are developing AI Systems to enhance an existing product or service e.g., an AI System to provide personalised product recommendations for consumers. The research exception ("**Research Exception**") can apply when organisations are conducting commercial research to develop AI Systems that have public benefit, e.g., for precision medicine.

1.4 Third-party developers of bespoke AI Systems ("**Service Providers**") are data intermediaries who have obligations under the PDPA i.e., Protection and Retention Obligations. When developing AI Systems, such Service Providers will handle personal data provided by their client organisations. As required under their Protection Obligation, Service Providers should guard against unauthorised modification of the personal data they are processing. Good practices that Service Providers could undertake include data mapping and labelling, as well as the maintenance of provenance records.

1.5 To assure consumers that their personal data is being used appropriately, the Guidelines encourage organisations to be more transparent. To this end, organisations are encouraged to provide relevant information at the point of data collection so that consumers can give meaningful consent. They are also encouraged to include in their written policies about safeguards and practices they put in place to ensure that AI Systems are trustworthy, especially where the outcome has high impact on consumers.

1.6 The Guidelines also recommend additional resources from the Personal Data Protection Commission's (**the "Commission"**) and Infocomm Media Development Authority (IMDA) for organisations to consider using to facilitate the deployment of trustworthy AI, such as data protection impact assessments and AI Verify.

PART II: LEGAL EFFECT AND SCOPE

2 Application of the PDPA to collection and use of data to design and/or deploy AI Systems

2.1 The PDPA is broad-based legislation that applies to all collection and use of personal data by an organisation, including the collection and/or processing of personal data to develop, test and monitor AI Systems, or as part of their deployment process.

2.2 These Guidelines should be read in conjunction with the Commission's Advisory Guidelines on Key Concepts in the PDPA, Advisory Guidelines on Selected Topics as well as its Guide to Basic Anonymisation. These Guidelines are advisory in nature, are not legally binding on the Commission or on any other party, and do not constitute legal advice. They neither modify nor supplement in any way the legal effect and interpretation of any laws cited, including, but not limited to, the PDPA and any subsidiary legislation issued thereunder. The provisions of the PDPA and any subsidiary legislation will prevail over these Guidelines in the event of any inconsistency. These Guidelines should not be construed to limit or restrict the Commission's administration and enforcement of the PDPA.

3 Scope of the Advisory Guidelines

3.1 These Guidelines are provided for situations where the design and/or deployment of AI Systems involve the use of personal data in scenarios governed by the PDPA. The aim of these Guidelines is to (i) provide certainty by clarifying how the PDPA applies when organisations use personal data to develop and train AI Systems; and (ii) provide consumers assurance by setting out baseline guidance and best practices for organisations on how to be transparent about whether and how their AI Systems use personal data to make recommendations, predictions, or decisions.

3.2 These Guidelines are organised according to the typical stages of AI System implementation as follows:

Section	Stage of AI System Implementation	Topics
Part III	**Development, testing and monitoring:** Using personal data for training and testing the AI System, as well as monitoring the performance of AI Systems post deployment.	• Consent • Business Improvement and Research Exceptions • Implementing data protection measures • Anonymisation
Part IV	**Deployment:** Collecting and using personal data in deployed AI Systems (**"business to consumer" or B2C**).	• Notification and Consent • Accountability
Part V	**Procurement:** Service Providers for bespoke AI Systems developed using personal data in organisations' possession (**"business to business" or B2B**).	• Notification and Consent • Accountability

PART III: USING PERSONAL DATA IN AI SYSTEM DEVELOPMENT, TESTING AND MONITORING

4 Business Improvement Exception and Research Exception

4.1 Organisations may occupy the role of an AI developer by developing AI models in-house, or engaging Service Providers to develop bespoke AI applications using personal data in the organisations' possession. The following sections cover the PDPA obligations that AI developers should pay attention to.

4.2 Besides seeking consent to use personal data to train an AI System, organisations who are AI developers may wish to consider relying on the Business Improvement or Research Exceptions.

a) The Business Improvement Exception[1] is relevant when the organisation has developed a product or has an existing product that it is enhancing. It is also relevant when an AI System is intended to improve operational efficiency by supporting decision-making, or to offer more or new personalised products and/or services such as through offering recommendations to users. The Business Improvement Exception caters for sharing with **related companies**

[1] See Part 5 of the First Schedule and Division 2 under Part 2 of the Second Schedule to the PDPA.

within a group of companies, as wells interdepartmental sharing within a company[2].

b) The Research Exception[3] is relevant when the organisation is conducting commercial research to advance the science and engineering without a product development roadmap. It also caters for sharing data between **unrelated companies** for jointly conducted commercial research to develop new AI Systems.

5 Application of the Business Improvement Exception

5.1 The Business Improvement Exception enables organisations to use, without consent, personal data that they had collected in accordance with the PDPA, where such use falls within the scope of the following relevant purposes[4]:

a) Improving, enhancing existing goods and services or developing new goods or services;

b) Improving, enhancing existing methods or processes or developing new methods or processes for business operations in relation to the organisations' goods and services;

c) Learning or understanding the behaviour and preferences of individuals (including groups of individuals segmented by profile); or

d) Identifying goods and services that may be suitable for individuals (including groups of individuals segmented by profile) or personalising or customising any such goods or services for individuals.

5.2 In addition, organisations will need to ensure the following[5]:

a) The business improvement purposes cannot reasonably be achieved without using the personal data in an individually identifiable form; and

[2] Organisations should note that where related companies are transferring data to each other, these companies should be bound by a contract or agreement or binding corporate rules requiring the recipient of personal data to implement and maintain appropriate safeguards for that data.
[3] See Division 3 under Part 2 of the Second Schedule to the PDPA.
[4] Defined in para 1(2) under Part 5 of the First Schedule to the PDPA.
[5] See Part 5 of the First Schedule and Division 2 under Part 2 of the Second Schedule to the PDPA.

b) The organisation's use of personal data for business improvement purpose(s) is that which a reasonable person would consider appropriate in the circumstances[6].

5.3 Relevant considerations for organisations on whether to rely on the Business Improvement Exception to justify the use of personal data for the development, testing and monitoring of AI Systems are set out below:

a) Whether using personal data for this purpose contributes towards improving the effectiveness or quality of the AI Systems and their output;

b) Whether it is technically possible and/or cost-effective to use other means to develop, test or monitor the AI Systems without using personal data (e.g., it may not be cost-effective where the personal data constitutes a small part of a wider data set containing non-personal data and it would require disproportionate effort to anonymise it);

c) Common industry practices or standards on how to develop, test and monitor such AI Systems; and/or

d) Whether such use will contribute to the effectiveness or improved quality of new product features and functionalities that help organisations innovate, improve competitiveness, become more efficient/effective, and enhance consumer choice, experience, and usability.

5.4 The following are examples of purposes where the Business Improvement Exception could be relevant to AI System development:

a) Recommendation engines in social media services that offer users content more aligned to their browsing history;

b) Job assignment systems that automatically assign jobs to platform workers;

c) Internal HR systems used to recommend potential job candidates by providing a first cut in matching candidates to job vacancies; or

d) Use of AI Systems to provide new product features and functionalities to improve competitiveness of products and services.

[6] See para 1(2) under Part 5 of the First Schedule to the PDPA, and para 1(2) under Division 2 under Part 2 of the Second Schedule to the PDPA.

5.5 In addition, organisations may wish to consider relying on the Business Improvement Exception to use personal data to test AI Systems or for bias assessments. The following paragraphs discuss how the exception can apply, as well as additional considerations for organisations' noting.

When using personal data to test AI Systems

5.6 Organisations could rely on the Business Improvement Exception to use personal data to test AI Systems, taking into consideration the requirements as set out in paragraphs 5.1 to 5.3 above. Organisations may need to use personal data to test an AI System to improve or assess model performance e.g., to assess its accuracy in a live environment with personal data; ensure that de-biasing of the model is effective; or to check if privacy enhancing measures have compromised the accuracy of the AI System.

5.7 Organisations are to take note that different standards for securing and protecting the datasets apply (see paragraphs 7.1 to 7.7 below), depending on the type of data used.

When using personal data for bias assessment

5.8 The Business Improvement Exception could apply to the use of personal data for bias assessments. The Commission understands that personal data may need to be used to check if protected characteristics, such as race or religion, are well represented in datasets or to assess the bias of the training dataset, so that adjustments may be made for de-biasing during AI System development.

5.9 In considering whether the Business Improvement Exception applies, organisations should consider:

a) Whether using personal data for this purpose is relevant for the effectiveness or improved quality of the AI Systems and its output;

b) Whether it is technically possible and cost-effective to use other means to debias models without using personal data; and/or

c) Common industry practices or standards on how to debias datasets used for AI Systems.

5.10 The Commission understands that generally, industry best practice is to use personal data to debias datasets used for model training.

6 Application of the Research Exception

6.1 The Research Exception is intended to allow organisations to conduct broader research and development that may not have immediate application to their products, services, business operations or market.

6.2 Organisations may use personal data for a research purpose[7], subject to the following conditions:

a) The research purposes cannot reasonably be accomplished unless the personal data is provided in an individually identifiable form;

b) There is a clear public benefit to using the personal data for the research purpose;

c) The results of the research will not be used to make any decision that affects the individual; and

d) If results of the research are published, the organisation must publish the results in a form that does not identify the individual.

6.3 Organisations may rely on the Research Exception to disclose personal data for a research purpose, including disclosure to another company for joint research and development of new AI Systems. In such a scenario, apart from the conditions in paragraphs 6.2(a) to (d), organisations will also need to assess whether it will be impracticable to seek the consent of the individual for such disclosure[8].

6.4 Relevant considerations for organisations on whether to rely on the Research Exception to justify the use of personal data in research on AI Systems that would satisfy paragraphs 6.2(a) and 6.2(b) include:

a) How and to what extent developing such an AI System will improve understanding and development of science and engineering;

b) Potential of application of the AI System to increase innovation in products or services that benefit society by improving the quality of life;

c) The use of personal data helps develop more effective methods to improve quality or performance of the AI System; and/or

[7] Pursuant to Division 3 under Part 2 of the Second Schedule to the PDPA.

[8] Required in para 1(b) of Division 2 of the Second Schedule. Please refer to the Advisory Guidelines on Key Concepts in the PDPA for more details on how to make an assessment on whether it would be "impracticable" for the organisation to seek consent of the individual.

d) Developing industry practices or standards for the development and deployment of AI Systems.

7 Data Protection Considerations when using Personal Data

7.1 When developing AI Systems, organisations should practise data minimisation as good practice. Using only personal data containing attributes required to train and improve the AI System will reduce unnecessary data protection and cyber threat risks to the AI System. Organisations should also limit the volume of personal data necessary to train the AI System and base this on relevant time periods and any other relevant filter e.g., market/customer segment, attributes, etc. Organisations may wish to refer to the Commission's Guide to Data Protection Practices for ICT systems for further guidance.

7.2 Organisations are reminded that when designing, training, testing, or monitoring AI Systems using personal data, appropriate technical, process and/or legal controls for data protection should be included. Where possible, organisations are encouraged to pseudonymise or de-identify the personal data used as a basic control.

7.3 If pseudonymisation is not possible and raw personal data has to be used e.g., facial images, organisations are reminded of their Protection Obligation under the PDPA[9]. Particular attention should be paid to the data security and protection measures around the development environment and organisations are encouraged to conduct a Data Protection Impact Assessment[10]. Standards for data protection in the development environment should be similar to the standards needed for systems handling personal data. Organisations may wish to refer to the Commission's Guide to Data Protection Practices for ICT systems, as well as the Guide on Responsible Use of Biometric Data in Security Applications where biometric data is used.

7.4 When deciding what kind of controls for data protection should be implemented, companies should consider:

a) The types of disclosure/theft risks that the personal data would be subject to; and

b) The sensitivity and volume of the personal data used[11].

[9] See S. 24 of the PDPA.

[10] Organisations may wish to refer to the Commission's Guide to Data Protection Impact Assessments for the key principles and considerations for organisations on when to conduct a Data Protection Impact Assessment.

[11] See paragraphs 17.2 to 17.4 of the Advisory Guidelines on Key Concepts in the PDPA.

7.5 Generally, privacy controls for internal use need not be as extensive as intra-group or cross-company sharing. However, companies are encouraged to assess the risks and implement appropriate legal, technical and process controls for such personal data use.

7.6 Separately, whether AI Systems are built in-house, externally, or using a combination of both, they will have security risks/points of weakness that can be exploited for privacy attacks to obtain information on the training data used e.g., model inversion attacks. Organisations should take a privacy-by-design approach and assess the risk of such privacy attacks as well as seek to mitigate such risks where possible within the AI System.

7.7 In addition, as per their **Accountability Obligation** under the PDPA[12], organisations must ensure that their policies regarding the use of personal data in their organisations to develop AI Systems are updated and practices are established. For example, establishing policies relating to when model training should be conducted using anonymised or pseudonymised data, and when it is permissible to use identifiable personal data, e.g., when model performance is degraded or for bias testing.

Using anonymised data

7.8 Organisations are encouraged to anonymise their datasets as far as possible instead of using personal data. While anonymised data is not subject to the PDPA, organisations are reminded that such data still bears risks of re-identification and disclosure, and appropriate legal, technical and process controls should be instituted when using or disclosing such data. Organisations can refer to the Commission's Guide to Basic Anonymisation[13] for further guidance.

7.9 The Commission recognises that there are trade-offs with using anonymised data when developing or training AI Systems, such as model accuracy, repeatability, or reproducibility of results. It may be preferable to use personal data. Organisations should carefully weigh the pros and cons of using both types of data, and clearly document internally the reasons for choosing to use personal data over anonymised data. Organisations should employ appropriate corporate governance methods to make such decisions, including consulting relevant stakeholders and having such decisions made at an appropriately senior management level.

7.10 In terms of what would be considered effectively anonymised data for developing or training an AI System, organisations can seek to anonymise the dataset only to the

[12] See S. 12 (a) and (c) of the PDPA.

[13] Organisations may wish to refer to the section on how to anonymise personal data (see pgs 13 – 25).

extent that there is no serious possibility of reidentification. Organisations should refer to Chapter 3 of the Commission's Advisory Guidelines for Selected Topics for the Commission's criteria on what constitutes anonymised data outside the scope of the PDPA.

7.11 As to whether anonymisation is sufficiently robust to reduce the risk of reidentification, this would include considerations such as:

a) Whether the process of chosen anonymisation method is reversible;

b) The extent of disclosure of the dataset and its intended recipients (e.g., internal closed-group sharing vs. cross-company sharing);

c) Whether a motivated individual can likely find means to re-identify the anonymised dataset using either publicly available information or information the organisation already has in its possession; and

d) The extent of controls the organisation has put in place, including within the AI System, to prevent re-identification of the anonymised data[14].

7.12 The Commission is aware that identifiability and anonymisation exists on a spectrum and is inherently context specific. It is likely that where the risk of re-identification is lower e.g., due to limited circulation or extensive controls over the anonymised data to prevent re-identification, the extent of anonymisation could be lesser to preserve the utility of the dataset.

PART IV: DEPLOYMENT – COLLECTION AND USE OF PERSONAL DATA IN AI SYSTEMS

8 PDPA Applies to Collection and Use of Personal Data in AI Systems

8.1 This section deals with how the PDPA applies when organisations deploy AI Systems in their products or services that collect and use personal data to provide new functionalities or enhance product features. Organisations should be mindful of the following PDPA obligations: **Consent** and **Notification** as well as **Accountability**.

9 Consent and Notification Obligations

9.1 Unless deemed consent or exceptions to the Consent Obligation apply, e.g., Legitimate Interests Exception, pursuant to Section 13 of the PDPA, consent will be

[14] Organisations may also wish to refer to the Commission's Guide to Basic Anonymisation, which provides further guidance on computing and assessing the risk of reidentification.

required for the collection and use of personal data to provide recommendations, predictions, or decisions. This is referred to as the **Consent Obligation**.

9.2 The **Consent Obligation** is complemented by the **Notification Obligation,** which requires that users be notified of the purpose of the collection and intended use of their personal data when seeking their consent. Section 20 of the PDPA sets out organisations' obligations to inform individuals of the purposes for which their personal data is collected, used, and disclosed. Among other things, Section 20(1) requires an organisation to inform the individual of:

a) The purposes for the collection, use and disclosure of their personal data, on or before collecting the personal data; or

b) Any purpose for the use or disclosure of personal data which has not been informed under sub-paragraph(a) above before such use or disclosure of personal data for that purpose.

9.3 As set out in the Advisory Guidelines on Key Concepts in the PDPA, consent should be meaningful, and **Notification** requires giving individuals information about *the types of personal data* that will be collected and processed and the *purpose for the processing*, e.g., to recommend books, songs, or movies.

9.4 The *raison d'etre* for the **Consent** and **Notification Obligations** is to enable individuals to provide *meaningful consent*. Organisations should place themselves in the shoes of consumers and craft notifications that will enable individuals to understand how personal data will be processed to achieve the intended purpose. Notifications need not be overly technical or detailed and should be proportionate to the risks of each use-case, e.g., taking into account potential harm to the individual and the level of autonomy of the AI System.

9.5 Organisations are encouraged to provide information on the following, to the extent practicable, in crafting notifications:

a) The function of their product that requires collection and processing of personal data (e.g., recommendation of movies);

b) A general description of types of personal data that will be collected and processed (e.g., movie viewing history);

c) Explain how the processing of personal data collected is relevant to the product feature (e.g., analysis of users' viewing history to make movie recommendations); and

d) Identify specific features of personal data that are more likely to influence the product feature (e.g., whether movie was viewed completely, viewed multiple times, etc).

9.6 The provision of such information could be through notification pop-ups or included in more detailed written policies that are publicly accessible or made available to end users on request. Organisations should decide the mode of providing such information, based on their own assessment of how this supports their business objectives vis-à-vis user experience.

> **Example:** A bank uses AI to assist in credit scoring when assessing whether to approve applications for credit cards. It prepared a policy document entitled "Bank's Credit Assessment Policy Statement" which provides information about what personal data it collects from applicants and how they are processed by AI when the bank assesses applications. The policy document is provided to applicants who request for the information.

> **Example:** An organisation provides personalised recommendations for content to an individual on its online social media platform. To provide information to individuals as to why specific content is shown to them, the organisation has provided a pop up containing a link to a page to explain why this content is shown and ranked highly on the content feed for the user. The page includes information on why that content is shown, what information has the largest influence over the order of posts in the user's content feed, such as past interactions or membership in specific groups on the platform etc.

9.7 It may also be useful to consider "layering" information. This means displaying the most relevant information more prominently and providing more details elsewhere. For example, notification pop-ups could provide a link to publicly accessible privacy policies; additionally, privacy policies may be structured to have details organised in expanding sections or separate tabs. The Commission recognises that industry is also developing disclosure best practices, such as model cards and system cards[15]. Information necessary to meet the **Consent and Notification Obligations** may also be provided through such model and/or system cards, if the organisation adopts this practice or assesses it to be useful.

[15] For model cards, see "The value of a shared understanding of AI models" <https://modelcards.withgoogle.com/about>; for system cards, see "System Cards, a new resource for understanding how AI systems work" <https://ai.facebook.com/blog/system-cards-a-new-resource-for-understanding-how-ai-systems-work/>.

Example: An organisation provides a video streaming service. It informs users that its service uses AI to provide recommendations. Through its notification pop-up, it informs users that it collects and analyses users' declaration of topics of interest, browsing activities and media consumption data to recommend videos that users may be interested in. Users are provided the option to consent or decline the use of this feature. The notification pop-up contains a link to its privacy policy, which contains a section that provides information about what declared topics of interest, browsing activity and media consumption data are collected and analysed. This includes the topic classification of videos that users watch, duration and proportion of the video that is played, how many times the video is played, whether the video is watched in a preview window or in actual size, etc. The organisation also explains that the topics of videos that users watch in full are most likely to influence future recommendations.

Example: A social media platform provides an AI system card to its users to explain how its AI System uses user activity data to generate recommendations for its content feed. The system card contains a step-by-step walk through on how the AI System gathers user activity data and broadly processes it in its AI System with other parameters to generate personalised output for a content feed.

9.8 Notwithstanding the above, the Commission recognises that organisations may need to protect commercially sensitive and/or proprietary information, as well as the security of AI Systems. Where organisations assess that it is necessary to limit or omit detail and, if appropriate, provide a more general explanation instead, it is good practice for these decisions to be justified and documented clearly internally.

Legitimate Interests Exception

9.9 "Legitimate Interests" generally refer to any lawful interests of an organisation or other person (including other organisations). Paragraphs 2 to 10 under Part 3 of the First Schedule to the PDPA relate to specific purposes that would be considered "Legitimate Interests", e.g., evaluative purposes; for managing or terminating an employment relationship. To rely on this exception, organisations must assess and ensure that the legitimate interests outweigh any adverse effect.

9.10 An example of a Legitimate Interest for processing personal data without consent would be the use of personal data as input in an AI System for the purposes of detecting or preventing illegal activities.

9.11 Organisations may wish to refer to the Commission's Advisory Guideline on Key Concepts in the PDPA for guidance on how to make an adverse effect assessment. Organisations who rely on this exception must make it known to individuals that they are relying on this exception to collect and use personal data.

10 The Accountability Obligation

10.1 The **Accountability Obligation** refers to how an organisation discharges its responsibility for personal data which it has collected or obtained for processing, or which it has control over. Sections 11 and 12 of the PDPA detail the actions to be carried out by organisations in fulfilment of this obligation[16].

10.2 Among other things, Section 12 of the PDPA requires organisations to develop policies and practices to meet its obligations under the PDPA. Written policies and documentation of processes enable organisations to show that their internal governance and supervision structures as well as operational practices ensure the responsible use of personal data. Such use should either in line with purposes that individuals have been notified of and consented to or for legitimate purposes that a reasonable person would consider appropriate in the circumstances[17].

10.3 Organisations that make use of AI Systems should be transparent and include in their written policies relevant practices and safeguards to achieve fairness and reasonableness[18]. The level of detail to be provided should be proportionate to the risks in each use-case, e.g., taking into account potential harm to the individual and the level of autonomy of the AI System.

10.4 Section 12(d) requires organisations to make information about such policies and practices available to individuals upon request. As the *raison d'etre* for such external communications with consumers is to help build trust with data subjects by demonstrating accountability in compliance with the PDPA, organisations should consider pre-emptively making such written policies available through their website, and not only upon request. Organisations should also consider making policies available in the form of short policy that is simple, clear, and concise.

---

[16] Please refer to the the Advisory Guidelines on Key Concepts in the PDPA for more details on the accountability obligation under the PDPA.

[17] See s. 17 of the PDPA.

[18] See *RE HSBC* [2021] SGPDPC 3 – where it was found that HSBC met its Accountability and Disclosure Obligations by providing information on how it has used personal data and AI technology to conduct credit facility assessments.

10.5 Written policies can house more detailed information that organisations ought to provide to obtain meaningful consent[19]. Where organisations have relied on exceptions to consent, e.g., Business Improvement and Research Exceptions, written policies can also provide information about the practices and safeguards that were adopted to protect the interests of individuals[20]. Developing industry best practices, such as model cards and system cards, can also form part of an organisation's written policies.

10.6 Written policies also play an important function in education and confidence-building, which are necessary ingredients for building consumer trust and confidence. Policies could therefore include behind-the-scenes measures taken to ensure that the personal data is used in a safe and trusted manner within the AI System, such as:

a) Measures taken to achieve fairness and reasonableness for recommendations, predictions, and decisions for the benefit of consumers during model development and testing stages. These can include measures relating to bias assessment, ensuring quality of training data or other data governance measures, or the repeatability/reproducibility of results using personal data.

b) Safeguards and technical measures taken to protect personal data. These can include measures to protect personal data during model development and testing (e.g., pseudonymisation and data minimisation), or steps to ensure personal data is protected in the AI System via ensuring the security of such systems before and after they are deployed.

c) For outcomes that have a higher impact on the individual, organisations may wish to consider whether it is useful to provide information on how proper accountability mechanisms and human agency and oversight have been implemented. It may also be useful to provide information on safety and/or robustness of the AI System i.e., how the AI System will operate when encountering adversarial or unexpected input.

10.7 Information on the above-mentioned measures is not always required. Organisations using personal data for model development and testing, and in deployed AI Systems, should consider adopting measures that a reasonable person

[19] This is discussed in the preceding section on the layering of information when meeting the **Consent and Notification Obligations**.

[20] It is not strictly necessary to disclose reliance on exceptions to the consent requirement unless disclosure is expressly required, e.g., when relying on the Legitimate Interests Exception.

would consider appropriate in the circumstances. Having done so, organisations are encouraged to consider providing sufficient information about such measures to build consumer trust and confidence.

10.8 Organisations are generally encouraged to provide more information on data quality and governance measures taken during AI System development. This is only if such information is deemed relevant and doing so does not compromise security, safety, or commercial confidentiality. Information that organisations can consider including are:

a) Steps taken to ensure the quality of personal data in the training dataset (e.g., how representative it is of the market and how recently it was compiled) to improve model accuracy and performance;

b) Whether model development was conducted using pseudonymised data, and if not, what organisation, process or technical safeguards were adopted to restrict access to personal data to developers and/or testers who had access;

c) Whether it was necessary to use personal data when conducting bias assessment to check if protected characteristics, such as race or religion, are well represented in the training dataset or to assess the bias of the training dataset;

d) If personal data was used, what process or technical safeguards were adopted to secure the testing environment and to limit access to testers; and

e) Whether data minimisation was practised at all stages of model and/or AI System development and testing.

Additional resources

10.9 Organisations may wish to refer to the Model AI Governance Framework for further suggestions on managing stakeholder interaction (see in particular Section 3, pages 53 – 55). Organisations may also find the guiding questions and examples on stakeholder interaction provided in Section 5 of the Implementation and Self-Assessment Guide for Organisations helpful.

10.10 Organisations can consider using technical tools such as AI Verify to validate the performance of AI Systems. Information from the testing report can be used to support information that organisations wish to include into their notifications or written policies. For example:

a) Results of explainability testing can be used to identify the data features that are most likely to influence the recommendation, prediction, or decision.

b) Results of fairness testing can be used to illustrate differences in model outcomes across demographic groups to show that there has not been unreasonable discrimination or bias in the use of personal data by an AI System. This can also be supported by process checks for repeatability/reproducibility.

c) Process checks for security can support an organisation's statement in their notification that they have taken steps to ensure that personal data used in an AI System is protected.

d) AI Verify also includes process checks that organisations may find useful to validate any claims in their notifications that they have included on accountability/human oversight and safety of the AI System. Robustness testing may also be useful if organisations intend to provide information on the robustness of the AI System in their notification.

Where possible, improvements should be introduced.

10.11 It is good practice for organisations to develop processes to regularly review the quality of the information provided, as well as the effectiveness of its notifications, policies, and practices for their intended audience.

10.12 Organisations are also encouraged to perform impact assessments, particularly data protection impact assessments, where these are deemed to be useful. These can help support organisations in their efforts to identify and mitigate data protection risks in an AI System. Organisations may wish to refer to the Commission's Guide on Data Protection Impact Assessments for more guidance on this area.

PART V: PROCUREMENT OF AI SYSTEMS – BEST PRACTICES FOR HOW SERVICE PROVIDERS MAY SUPPORT ORGANISATIONS IMPLEMENTING AI SYSTEMS

11 Business to Business Provision of AI solutions

11.1 This section is relevant for Service Providers (e.g., systems integrators) who are engaged by organisations to provide professional services for the development and deployment of bespoke or fully customisable AI Systems. It is not relevant to organisations that develop AI Systems in-house or who retail commercial off-the-shelf solutions that make use of AI for their product features and functions.

11.2 Where Service Providers, as part of developing bespoke or fully customisable AI Systems, process personal data on behalf of their customers, they take on the role

of data intermediaries and have to comply with applicable obligations under the PDPA[21]. It is good practice for such Service Providers to adopt the following practices:

a) At pre-processing stage, use techniques such as data mapping and labelling to keep track of data that was used to form the training dataset.
b) Maintain a provenance record to document the lineage of the training data that identifies the source of training data and tracks how it has been transformed during data preparation.

11.3 The above measures will support data intermediaries in assessing whether there has been unauthorised access and modification of the training data sets in their possession[22]. They will also provide the deploying organisation with necessary information to assess whether such unauthorised access or modification is a notifiable data breach and the scope of impact of any modification of training data sets on the AI System. This is particularly important given the direct impact training data generally has on AI Systems. Further, data mapping and labelling will be useful in helping data intermediaries identify whether there is sensitive personal data in their possession and calibrate security and data protection measures accordingly.

11.4 In addition, Service Providers developing such bespoke or fully customisable AI Systems are encouraged to support organisations in meeting their **Notification**, **Consent** and **Accountability Obligations**. This is because some customers may rely on the technical expertise of these service providers to meet their own obligations under the PDPA. Service Providers may be asked to provide technical clarification or consultation on the adequacy and accuracy of information in policy documents developed by organisations for their customers.

11.5 The subsequent paragraphs contain best practices on how such Service Providers may support these organisations.

Step 1: Understand the information that customers are likely to require based on their needs and impact on users

11.6 Service Providers are encouraged to be familiar with the types of information described in paragraphs 9.5 and 10.6 – 10.8 above that contribute towards meeting their customers' **Consent**, **Notification** and **Accountability** Obligations. To do so, Service Providers will have to pay attention to the context and impact the AI System will have on individuals. Information that is likely to be relevant should be identified,

[21] Data intermediaries are subject to the Protection Obligation and Retention obligation under S. 24 of and S. 25 of the PDPA respectively.

[22] Data intermediaries have a duty under S. 26C of the PDPA to report data breaches of data they are processing to the organization they are processing the data on behalf of.

and these Service Providers are encouraged to engage their customers on what will be helpful for them.

Step 2: Design your system to ensure that you can obtain relevant information

11.7 As part of implementing privacy-by-design, Service Providers are encouraged to build in processes when designing bespoke or customisable AI System that facilitate the extraction of information relevant to meeting their customers' PDPA obligations. This will enable Service Providers to better support customers who may require their assistance in developing policy documents or notifications. These include:

Process	Description of best practice
Providing information for user organisation's internal use	• Translate the operation of the AI System into easily understandable language for operators so that they can understand how the outcome is arrived at. • This can be done through textual explanation, visual aids such as video, graphs or table, or a combination of all or some of these. • Use technical tools such as AI Verify to aid in the development of various types of explanations. Please refer to paragraph 10.10 above.
Supporting user organisations to understand the acquired AI System	• Where needed e.g., where the AI System is more complex, sufficient instructions, training, or information for human decision-makers involved in the AI-assisted decision-making process to ensure that they have adequate and appropriate knowledge of how to use the AI System. • Where the AI System is intended for autonomous decision-making, providing training or a clear explanation to ensure that user organisations deploying the AI System properly understand how the AI System operates.

11.8 While Service Providers can support organisations in achieving their **Consent** and **Notification** Obligations as well as **Accountability** Obligation, the Commission reiterates that organisations bear the primary responsibility for ensuring that the AI System they have chosen to use can meet their obligations under the PDPA.

END OF DOCUMENT