Internet Society
Singapore Chapter
(in rejuvenation)

Internet Society Singapore Chapter response to PDPC's public consultation on proposed advisory guidelines on use of personal data in AI recommendation and decision systems

August 31, 2023

About us

The Internet Society Singapore Chapter is registered in Singapore under the Societies Act, operating under a chapter agreement with the Internet Society. The chapter's object is to advance and promote the use of the Internet and its associated technologies and applications — both as an end in itself, and as a means of enabling organisations, professions, and individuals worldwide to more effectively collaborate, cooperate, and innovate in their respective fields and interests in Singapore.

We are still working towards normal operation as a chapter after a period of inactivity. Consequently we are required to make clear that we are "in rejuvenation" at present.

The Internet Society is a global nonprofit organisation empowering people to keep the Internet a force for good: open, globally connected, secure, and trustworthy. The society is a nonprofit corporation formed under the laws of the District of Columbia and headquartered in Virginia, USA.

Enquiries

Enquiries about this submission may be directed to Ankur Gupta via isocsg@gmail.com.

About this submission

Creation

This submission was produced on the basis of an in-person session with chapter members and a small number of interested members of the public. Most participants had professional, commercial, or academic interests in applying AI to personal data.

Motivation

The Internet Society has for many years included security and trustworthiness in its mission[1], both for the network itself and for its applications. In 2017 we identified AI specifically as one of the key Drivers of Change that represent challenges for the future Internet[2]. The immense surge in interest and activity around AI in the last 12 months has borne this out with remarkable force. Giving each individual control over their personal data, its collection, and its use in the context of AI applications is a fundamental norm towards fostering public confidence that choice is not undermined in usage of apps which have AI running under the hood. These apps would serve as gateways which affect more profoundly than ever before, how users access and benefit from the Internet. Similarly, motivating effective risk management by organisations applying AI to personal data with respect to the risks to the organisation itself, to individuals, and to society at large, materially improves security and trustworthiness of the Internet and its applications.

The Singapore Chapter has been an active civil society actor fostering community building, fostering awareness building, and acting as a node for thought leadership on varied issues around regulation, standards and policies impacting the Internet, framed in local terms.

[1] https://www.internetsociety.org/mission/
[2] https://www.internetsociety.org/resources/doc/2017/global-internet-report-2017/

Summary of major points

The proposed guidelines are generally excellent and we support their publication and use.

We do note that the guidelines don't make their intended audience(s) clear, and that there is significant room to better help non-expert audiences to understand their obligations, assuming that they are in fact an intended audience.

We are concerned that the level of disclosure advocated for organisations providing AI goods and services may be insufficient for reliance upon by non-expert customer organisations, particularly SMEs.

Although it is outside the current scope of the PDPA, we note that there is no formal right to contest unfair or unjust decisions made on the basis of the automated processing of personal data generally, let alone where AI has been applied.

Comments

Our comments are grouped largely by document section.

3 Scope of the Advisory Guidelines

Audience

It would be helpful for the guidelines to make explicit the intended audience.

Not specifying the intended audience complicates choices about what to include or exclude. For example:

- A large organisation with an existing personal data compliance team, perhaps with obligations in multiple jurisdictions, perhaps with a Chief Privacy Officer appointed, will typically study regulator-issued guidance to:
    - (a) ensure that a particular regulator's priorities have not been overlooked; and
    - (b) understand how that regulator will view particular issues that the guidance addresses.

  An organisation of this type will generally look directly to legislation and regulation to identify and understand their legal obligations.
- There are hundreds of AI-focussed startups in Singapore. At least in their early stages they don't have substantial compliance teams, nor the resources to contract extensive independent advice. Organisations of this type will look to regulator guidance to understand what their obligations are in the first place, in preference to trying to understand the entirety of all legislation and regulation that affects them.

These are clearly very different uses of the guidelines, by very different audiences.

The AI ecosystem is complex and evolving. Relevant stakeholders include: AI platform companies, app developers, app services operators, platform developers, platform operators, protocol developers, network operators, retailers and resellers of every size, policymakers and regulators, and users (these include individuals, businesses and governments). It would be desirable to make clear who the intended audience(s) is/are, and to scope the guidelines accordingly.

Harms

It would be helpful for the guidelines to spell out in the introduction what the assumed likely harms are.

While the approach taken is quite properly to help the reader understand what their obligations are and where to go to look for details, it is likely to be helpful to spell out up front what the assumed likely harms are, particularly for audiences that don't have established personal data compliance teams. During our discussion session it became clear that even people who work in contexts which PDPA affects haven't grasped its full extent beyond the fundamental obligation to keep confidential data confidential. A considerable amount of time

went into spelling out what harms typical model training on personal data might give rise to. We'd suggest that this indicates the existence of a specific, important gap that the guidelines could fill.

This might best be addressed by providing a set of example AI processing operations, and their likely harms. The point is not to provide an exhaustive treatise, but enough to help orient non-experts.

4 Business Improvement Exception and Research Exception

Boundary between the exceptions

During our discussion session it became clear that understanding the boundary between these exceptions was both poorly understood and a matter of some importance, at least amongst those working in startup-adjacent research environments. This would again appear to be a relevant gap that the guidelines could reasonably fill. Completely restating the Act would be pointless, but providing concrete examples of what might fall under either exception, or even both.

4.1

> a) The Business Improvement Exception is relevant when the organisation has developed a product or has an existing product that it is enhancing.

This looks like a mistake. Presumably "is developing" rather than "has developed". Development would almost certainly include working with personal data, meaning that the obligations of PDPA come into effect at the beginning of the process, not the end.

5 Application of the Business Improvement Exception

5.3

> b) The organisations use of personal data for business improvement purpose is that which a reasonable person would consider appropriate in the circumstances.

Editorial:

- "purpose" should be "purposes" (plural)
- "organisations" (plural) should be "organisation's" (singular possessive).

5.8

> protected characteristics, such as race or religion,

10.8

> c) … protected characteristics, such as race or religion,

What is meant by “protected characteristics” here?

The term is not defined in these Guidelines, nor in the Act. There are different protected characteristics relevant in different contexts, however it would be helpful to spell out the sets of characteristics that are relevant in [selected] specific contexts in Singapore, along with pointers to legislation etc.

We note for example Singapore’s Constitution Article 12(1) specifies “religion, race, descent and place of birth” with respect to laws, employment, etc., but assume that there are other obligations. There was discussion during the in-person session about the potential for political sensitivity here, but assume that if parliament or a regulator has decided on concrete protective obligations then at least pointing to them would be helpful here.

5.9

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

This appears to be a mistake. We assume that the intention was to make clear that the use of personal data in bias reduction is permissible, but the text as written goes much further and presents it as best practice by itself. We note that the use of personal data for this purpose is not best practice — at least by itself — as it often risks **worsening** bias as it’s not in general possible for developers to get access to personal data that’s representative of all relevant values of parameters that are likely to lead to bias. There’s an entire body of research in this topic, but we’d suggest at least making clear that e.g. transparency about data sources and the use of synthetic data are as important.

Further, although it’s perhaps a terminological quibble, bias **removal** is impossible. Several terms better express the objective. Any of: reduction, moderation, mitigation, or minimisation would be better choices.

7 Data Protection Considerations when using Personal Data

7.10

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

Editorial:

- Delete “process of” as it’s redundant.
- Alternatively: insert “the” before “chosen” to make the text grammatically sound, if ungainly.

9 Consent and Notification Obligations

Please provide many more examples

This is actually feedback about every part of the document, but the presence of examples in section 9 was the motivation for it.

We received feedback that — for non-expert readers — the examples in this section were extremely helpful in understanding the points that were being made. We suggest that in almost every section of the document, providing concrete examples would be helpful for non-expert users to more readily understand what's being conveyed.

11 Business to Business Provision of AI solutions

Degree of disclosure by intermediaries

There was a general sense that this section was weaker than it could be. In particular, SMEs relying on AI services provided by intermediaries and therefore dependent upon them for part of their own PDPA compliance are unlikely to even know what the right questions to ask are. There would appear to be benefit in particularly clear disclosures in this situation. Unfortunately we did not get to the point of discussing specific improvements to this end, we merely note that commercial incentives and limited accountability tend to breed harmful externalities. Guidance to provide clear expectations today — and with a view to potential eventual formalisation in legislation or regulation — would appear desirable.

11.2

> Service providers who are data intermediaries should adopt the following practices:
> …

These obligations only appear in section 11 in the context of the provision of AI goods and services to other organisations, however the underlying obligations also apply to in-house development. These should perhaps be relocated to section 7.

> (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.

Editorial: consider inserting "and model(s)" after "lineage of the training data". Strictly speaking this addition is redundant as the model is transformed training data, however it would appear to be worthwhile to make clear that the traceability/provenance should apply to the derived model(s) as well as to things that look like training data.

11.6

> As part of implementing privacy-by-design, service providers are encouraged to try to build in processes

Editorial: This seems a little weak, or redundant. Suggest removing “try to”. Organisations are already encouraged rather than required, further softening seems unnecessary.

Outside of the current scope of PDPA

Contesting decisions made by or with the aid of AI

A key area of risk in the processing of personal data at all, but far more so with the application of AI, is that decisions will be made and acted upon which are in some sense unfair or unjust. The ability to make reasonably good decisions in complicated areas **very cheaply** will almost certainly motivate automating the execution of those decisions without human involvement, and therefore cause harm in the corner cases that would likely not have arisen had a human officer made and/or executed the decision. This risk has been addressed in the EU as a right to object, and frequently in research under the general heading of contestability.

It may make sense to include good practice in this area under section 10, however accountability is only part of the picture. The means of contesting a decision with a view to having it changed is not mere accountability, and not currently provided for by PDPA. We therefore offer this concern for PDPC’s consideration at the next revision of PDPA, rather than solely as an improvement to the proposed guidelines.

Human in/over the loop approaches

A closely related concern is the desirability of human-in-the-loop or human-over-the-loop approaches to decision-making in preference to purely automated approaches[3]. It’s not immediately clear to us where to place this in the proposed guidelines, but it’s part of existing PDPC guidance and we’d suggest worth drawing attention to somewhere in the guidelines.

[3]e.g. as described in https://www.pdpc.gov.sg/-/media/Files/PDPC/PDF-Files/Resource-for-Organisation/AI/Primer-for-Discussion-Paper-on-AI-and-PD---050618.PDF

Conclusion

We appreciate PDPC's initiative in assembling guidelines of this type to help navigate what would otherwise be a debilitating body of law for hundreds of organisations developing AI goods and services for others, and for many thousands of organisations applying AI. We hope that our submission provides some worthwhile improvements.