1 December 2023

Mr Sean Coley
Committee Manager
Integrity and Oversight Committee
Parliament House
East Melbourne VIC 3002
By email: ioc@parliament.vic.gov.au

Inquiry into the Freedom of Information Act 1982 (Vic)

About us

The **UNSW AI Institute** (UNSW.ai) is the flagship research institute at UNSW focused on artificial intelligence (AI), data science (DS) and machine learning (ML). It proudly supports the endeavours of more than 300 UNSW academics and over 50 research groups, labs, and centres. The Institute's researchers possess a remarkable track record in AI research and development, spanning across multiple faculties such as Engineering, Science, Business, Law, Medicine, Arts, Design & Architecture, and UNSW Canberra. The AI Institute serves as a frontrunner in advancing the field. It is driven by various objectives, including the facilitation of interdisciplinary collaborations in teaching and research, active engagement in public dialogue on AI, and the promotion of research commercialisation.

The **UNSW Allens Hub for Technology, Law and Innovation** ('UNSW Allens Hub') is an independent community of scholars based at UNSW Sydney. As a partnership between Allens and UNSW Law and Justice, the Hub aims to add depth to research on the diverse interactions among technology, law, and society. The partnership enriches academic and policy debates and drives considered reform of law and practice through engagement with the legal profession, the judiciary, government, industry, civil society and the broader community. More information about the UNSW Allens Hub can be found at http://www.allenshub.unsw.edu.au/.

About this Submission

We are grateful for the opportunity to make this submission in response to invitations addressed to Professor Toby Walsh and Professor Lyria Bennett Moses. Our submission reflects our views as researchers; they are not an institutional position. This submission can be made public. As neither of us is an expert on Victorian freedom of information law per se, we focus on only some of the issues raised, in particular:

- Informal release (drawing attention to research on this topic done at UNSW for the NSW Information and Privacy Commission);
- Changing data practices and artificial intelligence and the impact on the coherence of concepts in legislation.

Responses to selected issues

2. Mechanisms for proactive and informal release of information, including the effectiveness of information publication schemes

UNSW produced a report for the NSW Information and Privacy Commission on informal release in NSW. The report is available here, and may be useful to consider in the context of the analogous Victorian mechanism. The NSW IPC's response is here.

3. Efficient and timely mechanisms for persons to access their own personal and health information

Our discussion on terminology in question 6 below is also relevant to this question. In particular, there are circumstances in which information held about a person will be diffuse and not in a single "document". In such circumstances, there are important questions to consider as to whether an agency ought to be required to collate such data for the purposes of freedom of information.

Another challenge to consider here is the issue of inferred information. Agencies will hold not only information initially collected about a person but will also draw inferences from that data. Such inferences may be held in a "document", but they may also be more diffuse, for example in machine learning models. Consider, for example, a large language model such as that used by ChatGPT. This has a concept of "Lyria Bennett Moses" and "Toby Walsh" and is able to make statements about each of us when queried (some of which is true, or close to true). However, this information is not stored in a "document" and is difficult to disentangle from the larger model. As government increasingly relies on deep learning, large language models and other kinds of artificial intelligence, the question of what is required to allow individuals to access "their own personal information" will become more difficult to answer because government will have access to extensive, often inferred, information about individuals that is not held in a "document".

4. The information management practices and procedures required across government to facilitate access to information

This issue is closely related to the issue above. In other words, what should government be required to do in a context where it holds information about an individual but needs to engage in some process (database search, chatbot prompts, etc) to extract that information?

5. Opportunities to increase the disclosure of information relating to government services using technology

Some comments made in the UNSW Allens Hub submission to the Commonwealth on their 'Safe and Responsible AI in Australia' paper might also be apposite here:

> Transparency is a concept that many people are agitating for, but the crucial questions are *what* is rendered transparent, *to whom*, *how* and in which contexts. A driver of an automated vehicle does not need a continuous output from an automated vehicle explaining the logic behind a particular automated decision to steer slightly left to stay in a lane.

Rather, they want to know that the car has been evaluated (overall) as safe. On the other hand, the public should be able to find out the logic behind government systems that make decisions affecting them, the nature and quality of training data used, the testing and evaluation of systems that has been conducted (and the results of such), the assumptions on which a system relies, and so forth. Mandating uniform transparency requirements across sectors and contexts would not be helpful in almost all cases. An exception is the proposal (across sectors and contexts) to prohibit misleading uses of AI and automated systems. People should have a right to know when they are interacting with a machine rather than a human (unless they voluntarily relinquish that right for a specific activity, for example in the context of AI research). Similarly, there should be transparency about the involvement of AI in content-generation, so that (for example), an AI-generated image is labelled as such rather than represented as a human artwork.

One way in which governments can provide signalling as to best practice, would be to include the model cards[1] (where applicable) used by the Commonwealth in its use of AI. Similarly, a requirement in public sector procurement that model cards are a mandatory part of supply of g AI products and services would assist with transparency.

A model card is a human-readable document that provides critical information about a machine learning model. It is used to help people understand how the model works, its limitations, and its potential biases.

Model cards usually include the following minimum information:

(a) **Model name and version**: This information helps to identify the model and to track its development over time;

(b) **Model type**: This information describes the type of machine learning model, such as a neural network, large language model, decision tree, or support vector machine;

(c) **Model inputs and outputs**: This information describes the types of data that the model can take as input and the types of data that it produces as output;

(d) **Model training data**: This information describes the data that was used to train the model. This information can be used to assess the model's performance on different types of data;

(e) **Model evaluation metrics**: This information describes how the model was evaluated. This information can be used to assess the model's performance on different tasks; and

(f) **Known limitations and biases**: This information describes any known limitations or biases in the model. This information can be used to help users interpret the model's results and to make informed decisions about its use.

[1] Margaret Mitchell et al, 'Model Cards for Model Reporting' (2019) Proceedings on the Conference for Fairness, Accountability, and Transparency https://dl.acm.org/doi/10.1145/3287560.3287596.

6. The purposes and principles of access to information and whether the Act meets those purposes and principles, including: a. the object of the Act as set out in section 3; b. the definition of document in section 5; and c. the operation of exemptions and exceptions in Part III and Part IV;

There is a need for more consistency around how terms such as 'document' are used in legislation, not only in relation to this Act in Victoria, but in relation to legislation across Australia. Given the nature of electronic data, there is a need for greater clarity about what are 'data', 'information', 'document', 'record' etc and what distinctions need to be drawn (if any) across these terms. One important question in the context of freedom of information laws is the circumstances in which government should be required to provide not only existing files but other collective forms of information, such as responses to particular database queries. These are not in a single, existing "document" per se, but can often be collated with minimal effort. Provided the effort required is not onerous, there is little reason to condition transparency on the pre-existence of information in a collated form. While this issue is important in the context of freedom of information laws, it is not unique to them, and a better set of definitions (of a variety of data-related terms) would help government agencies and other organisations navigate legislation in the context of modern data practices.[2]

On exceptions in Part IV, the exception in s 33 (Document affecting personal privacy) raises issues at the intersection of freedom of information and privacy law. The federal *Privacy Act 1988* is currently under review and it would be useful to consider harmonisation of state laws at the appropriate time. It is worth, however, mentioning two issues related specifically to developments in artificial intelligence:

1. The issue of inferred information (not collected directly from an individual) and whether there ought to be restrictions on the circumstances in which it is generated and used in line with data protection principles; and
2. The reliance on de-identification techniques as a basis for assuming that a document does not affect personal privacy. The increased availability of related datasets and the improvements in re-identification techniques mean that whether a document affects personal privacy is likely to change over time. There is a need to carefully consider the scope of the exception in this light.

Yours sincerely,

Lyria Bennett Moses

Toby Walsh

[2] On the need for terminology to be clearer in the context of modern data practices more generally, see Lyria Bennett Moses, 'Who Owns Information? Law Enforcement Information Sharing as a Case Study in Conceptual Confusion' (2020) 43(2) University of New South Wales Law Journal 615.