

Asia Internet Coalition (AIC) Industry Response to the Exposure draft of a biometric processing code of practice: consultation paper, New Zealand

To
Michael Webster, Privacy Commissioner,
The Office of the Privacy Commissioner (OPC)
New Zealand
biometrics@privacy.org.nz

On behalf of the Asia Internet Coalition (AIC) and its members, I would like to thank the Office of the Privacy Commissioner (OPC) for continuing your thoughtful engagement with stakeholders on this topic. As we described in [our response](#) to OPC's position paper consultation in 2022 and [our response](#) to OPC's discussion paper in 2023, our members share OPC's belief in the benefits of biometric technologies.

We sincerely appreciate OPC's efforts to not preclude these benefits by incorporating stakeholders' feedback into the exposure draft of a biometric code. In particular, we applaud OPC for recognising that, although consent can be a valuable privacy safeguard, a general consent requirement for any and all uses of biometric information would not be an optimal protection. OPC is also right to eschew a blanket prohibition on using biometric information for marketing, focussing instead on other, potentially more intrusive uses of biometric information.

That said, we believe that the code could benefit substantially from further refinements. On balance, its breadth remains a challenge and creates a significant risk of stifling innovation, particularly in artificial intelligence (AI), in New Zealand. Potentially low-risk practices are prohibited without sufficient exceptions, the scopes of several definitions are unclear, and compliance with many provisions as written would be difficult, if not impossible.

Further details on our recommendations are set out below. If OPC adopts them, we firmly believe that OPC will accomplish the goal of protecting New Zealanders, especially underrepresented groups like the Māori, while fostering innovation and economic advancement. We look forward to further dialogue.

1. OPC should focus the code on information used to identify individuals.

As we have described in our previous responses, different uses of biometric information carry different privacy risks. A unique set of risks arises when such information is used to determine someone's real-world identity. In particular, identification enables serious harms like widespread government surveillance, suppression of free speech, and identity theft. It is for precisely this reason that laws and regulations around the world have adopted a common approach - requiring individuals' consent to use their information to uniquely identify them. We still strongly recommend that OPC follow this common approach, not only to produce a code that is properly targeted to the most serious risks, but also to avoid creating significant compliance difficulties for multinational agencies. This is precisely the approach that the Centre for Information

Policy Leadership recommends in [its new report on biometrics](#): “CIPL recommends that laws and regulations with heightened compliance requirements should target biometric systems intended to be used to identify individuals and exclude systems not used for identification purposes.”

If a code is instead applied to many other uses of biometric information, including classification, it must be narrowly tailored. It must recognise what specific risks are most concerning and target any restrictions to those risks. A code that fails to do so will sweep too broadly, disincentivising innovation and economic growth. For the reasons described in the subsequent sections of our response, the draft code is not narrowly and properly tailored. It bluntly prohibits practices that are low-risk, beneficial, and necessary for the functioning of modern technologies, often by relying on definitions whose scopes are unclear and overbroad. It also fails to clarify how certain privacy safeguards could be implemented in practice, creating significant compliance uncertainty.

The easiest and most effective way to remedy these challenges would be to focus the code only on uses of biometric information to identify individuals, which the draft code refers to as “biometric identification” and “biometric verification.” Additionally, the code should expressly exclude images, photographs, videos, and voice recordings that are not specifically and technically processed to enable identification. That said, we recognise OPC’s likely reluctance to do so following the previous rounds of consultation. If OPC maintains its current approach, we strongly encourage OPC to, at the minimum, consider our recommendations in the subsequent sections, intended to refine the contours of the code while still achieving OPC’s objectives.

2. OPC should align the set of prohibited practices with the EU AI Act by removing some prohibitions and adding exceptions.

If OPC does not refocus the code on biometric identification and verification, OPC should make substantial edits to which practices it prohibits and the exceptions to these prohibitions. In doing so, we would encourage OPC to take inspiration from the EU AI Act. Following years of negotiations between hundreds of stakeholders, EU lawmakers settled on an approach that we believe comes closer to effectively guarding against potential risks beyond identification without unduly restricting low-risk uses.

Notably, the EU AI Act prohibits only a very narrow set of uses of biometric information: (i) categorising individuals based on characteristics defined as special category data in GDPR Art. 9 (e.g., race, political opinions, sexual orientation); (ii) real-time remote identification in public spaces for law enforcement (unless strictly necessary for certain objectives); and (iii) inferring emotions in workplaces and educational institutions (unless for medical or safety reasons).

Outside of these prohibitions, the Act deems the following uses of biometric information as “high-risk”: (i) remote identification systems; (ii) systems intended to categorise individuals according to sensitive or protected attributes or characteristics; and (iii) systems intended to be used for emotion recognition. The Act applies obligations concerning transparency, data quality, and testing to high-risk systems, many of which are similar to requirements that OPC’s draft code already places on non-prohibited uses of biometric information.

Although some aspects of OPC's code are similar to those of the EU AI Act, there remain critical differences that, on balance, make OPC's code far likelier to unintentionally ensnare low-risk uses of biometric information. Our recommendations below would bring the two schemes closer, ensuring a draft code that offers substantial protections and provides compliance certainty to agencies.

A. OPC should remove from Rule 4(2) the prohibition on using biometric classification to collect information about an individual's inner state or physical state.

Of the prohibitions in Rule 4(2), the prohibition on categorising individuals according to a restricted biometric category most closely aligns with a prohibition in the EU AI Act - that on categorising individuals based on sensitive characteristics, as defined in GDPR. However, Rule 4(2)'s prohibition on collecting information about an individual's inner state or physical state is concerningly broad and has no exact analogue in the EU AI Act. For reasons described more in Section 3 of our response, the terms "inner state" and "physical state" are defined so broadly as to mean virtually anything one could possibly know about a human. For example, "inner state" includes an individual's "intention." But nearly any observation of a human's behaviour could be characterised as an observation of their intention. Imagine a data point collected by a vehicle's acceleration system measuring how far down a driver presses the accelerator pedal. That data point would likely be considered a "behavioural biometric" under the draft code, and, because it represents an individual's "intention" to accelerate the car, it would in turn be considered an inference of someone's "inner state." The draft code, as written, would thus prohibit vehicles from using digital acceleration systems.

This is just one example, but it shows how Rule 4(2)'s broad prohibition, combined with the code's broad definitions, create a wide range of unforeseen and unintended consequences. OPC should remove this overly broad prohibition. Alternatively, OPC could modify it to align exactly with the most analogous prohibition in the EU AI Act; OPC could modify it to be a prohibition on using biometric classification to identify an individual's emotion in workplaces or educational institutions.

B. OPC should add to Rule 4(3) exceptions to the prohibitions in Rule 4(2).

We appreciate OPC's recognition in Rule 4(3) that there are some instances where prohibited practices in Rule 4(2) are in fact low-risk or required for other reasons, warranting an exception. But we would encourage OPC to add the following exceptions to Rule 4(3):

- That the individual has consented to the collection of information;
- That the collection of information is reasonably necessary to provide, maintain, or improve a product or service requested by the individual;
- That the collection of information is for the purpose of protecting against security threats, identity theft, fraud, integrity threats, harassment, or similar threats; and

- That the collection of information is necessary to comply with other legal obligations, processes, or investigations.

These are critical exceptions that are common to almost every other privacy law in the world, and they would help align OPC's code with those laws. For example, the code differs dramatically from other laws—including the EU AI Act—in its prohibition on using biometric classification to collect health information. The common approach around the world is to instead deem health information a kind of sensitive data under comprehensive privacy laws. This often means that its collection is permitted with an individual's consent. But under OPC's code—which lacks a consent exception—collection of such information would be banned, devastating many sectors of New Zealand's economy. For example, there are many wearable health devices, including sports training devices, that arguably collect health information. New Zealand would no longer be a safe jurisdiction in which to build and sell such devices.

C. OPC should remove restrictions on web scraping from Rules 2 and 11.

Rule 2(3) states that biometric samples obtained via web scraping are not publicly available information under Rule 2(2)(d) and, thus, cannot be collected per Rule 2(1). Further, Rule 11(1)(d)(i) prohibits disclosing biometric samples collected via web scraping to other agencies or persons.

Put simply, these restrictions would cripple New Zealand's AI industry. Cutting-edge AI models, including generative AI models, are trained on information scraped from the web. Given the broad definitions of "behavioural biometric" and "physiological biometric" as incorporated in the definition of "biometric sample," these restrictions may amount to a ban on using photos and videos scraped from the web to train AI models.

Publicly available data is necessary to train generative AI models, as there is currently no reasonable or feasible alternative source from which to obtain the volume, breadth, variety and nuances of language necessary to build a generative AI model. To put this into context, for the size of the vocabulary of a large language model (LLM) to be 50,000-250,000 words, which is what is needed for minimal fluency, the LLM would need to be trained on billions to trillions of tokens.

This is recognized by other global regulators. For example, in the public consultation of the UK Information Commissioner Office (ICO) on AI, ICO notes, "The ICO's understanding is that currently, most generative AI training is only possible using the volume of data obtained through large-scale scraping. Even though future technological developments may provide novel solutions and alternatives, currently there is little evidence that generative AI could be developed with smaller, proprietary databases."

Blocking developers from using publicly available New Zealand data would materially restrict the ability of developers in New Zealand and worldwide to train generative AI models using New Zealand data and potentially inhibit the development of AI capabilities in the country. The sheer quantities of data required to train the LLM simply would not be available to developers for New Zealand data. The likely impact would be that LLMs would not be available or reliably reflect the nuances of New Zealand's distinct culture. A

lack of representation of different demographics or linguistic communities in training data could also lead to patterns of bias in data output.

Generative AI development is likely to have substantial societal benefits. Developments in AI are pushing the frontiers of scientific discovery – from medicine to climate technology. Generative AI will be central to addressing global challenges and has the potential to revolutionise sectors like healthcare, education, and agriculture that can deliver enormous societal benefits. The ability of New Zealand to capitalise on these benefits hinges on New Zealand companies being able to participate in the global generative AI ecosystem, without the hindrance of local regulatory restrictions that global players are not subject to.

Further, given the practical effect of Rule 2(3) is that biometric information must be directly collected from individuals, the OPC's code will restrict the ability to develop AI systems to a small handful of technology companies who have sufficient first-party data (i.e., data collected directly from individuals) to train AI models. This would effectively mean that New Zealand companies could not participate or compete in the generative AI development space.

To avoid letting New Zealand lag behind in today's wave of AI innovation, OPC should remove Rules 2(3) and 11(1)(d)(i).

3. OPC should clarify several definitions whose scopes are currently unclear and potentially overly broad.

As described above, several of the code's most concerning implications result in part from unclear and overly broad definitions. We would encourage OPC to modify the definitions below.

A. "Biometric information"

OPC should exclude from the definition of "biometric information" photos, videos, audio recordings, and data generated from any of the preceding kinds of information, unless such information is used for biometric identification or biometric verification. This is a widespread exemption in global laws, which recognises that risks arise not from photos, videos, or audio recordings themselves, but from how they are further processed and used.

B. "Behavioural biometric"

OPC should remove from the definition of "behavioural biometric" "the individual's pattern of using any digital device." Nearly any interaction with a digital device could be considered information about a "pattern" of usage. For example, music recording software that collects information from a digital piano—a digital device—could be said to be collecting a pianist's behavioural biometrics. And, because these behavioural biometrics indicate the pianist's intention to play certain notes—an inner state—the music recording software would be prohibited under the code.

C. *“Readily apparent expression”*

OPC should modify this definition to more clearly align with what we believe is its intent - ensuring that “inner state” does not inadvertently ensnare any mere observation of an individual’s appearance or behaviour. We appreciate this intent, but, by listing only a few specific behaviours—expression, gesture, movement, and the level or pitch of a voice—the definition leaves significant gaps. For example, it does not include an observation that someone is sweating, which is just as visually recordable as an expression or a gesture. We would suggest that the definition be modified to mean “any aspect of an individual’s appearance or behaviour that can be recorded visually or aurally without biometric processing.”

D. *“Physical state”*

OPC should modify this definition to exclude the detection of a readily apparent expression, just as is excluded from the definition of “inner state.” The purpose of a biometrics code should not be to place restrictions on mundane observations that humans make every day about other humans simply because they are made using technology instead of eyes or ears. We believe this is the motivation for excluding the detection of readily apparent expressions from “inner state,” and the same logic should apply to “physical state.” This is especially true because the states listed in the definition of “physical state”—fatigue, alertness, or attention level—are all states that are readily visually observable.

4. OPC should provide additional clarity on how to implement certain privacy safeguards.

Our final recommendations relate to two privacy safeguards contemplated by the code: notices and an assessment of proportionality.

A. *OPC should clarify the relationship between a conspicuous notice, an accessible notice, and a privacy policy in Rule 3(2).*

The draft code’s provisions concerning notices in Rule 3(2) are unclear, would create significant compliance difficulties, and would confuse individuals. Rule 3(2) states that both the accessible notice and the conspicuous must contain the same information from Rules 3(1)(a), 3(1)(b), and 3(1)(h). But the definition of “accessible notice” requires it to be presented independently of an agency’s privacy policy, and the definition of “conspicuous notice” requires it to be presented independently of both the accessible notice and the privacy policy. Further, many agencies’ privacy policies already disclose the information contemplated under Rules 3(1)(a) and 3(1)(b) by virtue of adhering to the Privacy Act’s IPP 3(1)(a) and 3(1)(b). The net result of this web of related requirements is that individuals would be faced with three notices—a privacy policy, an accessible notice, and a conspicuous notice—that all present the same information about the

collection of biometric information and the specific purposes for which it is being used. This would not only confuse consumers about what each notice is meant to disclose, but also make them less likely to read the notices, entirely defeating the purpose of notices.

To remedy this, OPC should remove the requirements for an accessible and conspicuous notice, but modify Rules 3(1)(a) and 3(1)(b) to state that an individual must be made aware of those pieces of information before their biometric information is collected or at the point of collection. This would still achieve the goal of the code's current requirements—ensuring that individuals are adequately informed before or during data collection—without confusing them with three long and redundant notices.

B. OPC should clarify how to determine proportionality under Rules 1 and 10.

In our 2023 response, we advised against requiring an assessment of proportionality, arguing that proportionality could carry very different meanings in different contexts, making compliance difficult. We appreciate that OPC recognised this nuance in Rules 1(1)(d) and 10(1)(b) by looking to an agency's belief "on reasonable grounds" considering "the particular circumstances."

That said, we believe there is room for further clarity. In particular, Rule 1(2) states that, when conducting a proportionality assessment under Rule 1(1)(d), an agency should take into account, among other things, the degree of privacy risk and whether the benefit of achieving a lawful purpose outweighs the privacy risk. We would suggest that OPC add a subrule to Rule 1 stating that, if the biometric processing is reasonably necessary to provide, maintain, or improve a product or service requested by an individual, then the benefit will be presumed to outweigh the privacy risk. This clarification would ensure that agencies have leeway to conduct biometric processing that individuals fully expect to occur—and desire—when seeking out products whose functionality is dependent on that biometric processing.