Submission to the

The Office of the Privacy Commissioner

on the

Biometric Processing Privacy Code Exposure Draft and accompanying Consultation Paper released April 2024

21 May 2024

Prepared by the Biometrics Special Interest Group of Digital Identity NZ (DINZ) with input from individual subject matter experts as well as DINZ member organisation representatives from a mix of large / medium corporates, public service agencies and academia.

Digital Identity New Zealand thanks the Office of the Privacy Commissioner for the opportunity to provide a submission and for the extension of time allowed to enable the participants of our Biometrics Special Interest Group to meet, discuss and curate this response.

As always, we are happy to provide any clarifications in writing, on a call, or in a physical meeting.

Colin Wallis
Executive Director
Digital Identity NZ
E | colin.wallis@digitalidentity.nz

About DINZ

DINZ is a not for profit, membership funded association and a member of the New Zealand Tech Alliance. DINZ is an inclusive organisation bringing together members with a shared passion for the opportunities that digital identity can offer. It supports a sustainable, inclusive and trustworthy digital future for all New Zealanders through its vision - that every New Zealander can easily use their Digital Identity in its mission to empower a unified, trusted and inclusive Digital Identity ecosystem for Aotearoa New Zealand that enhances Kāwanatanga (honourable governance), Rangatiratanga (self-determination & agency) and Ōritetanga (equity & partnerships).

Opening Statement

Position:

- We believe the Code of Practice (CoP) as drafted will potentially worsen privacy outcomes for people and should not proceed.
- We believe there are important considerations that organisations should make regarding privacy when they are adopting biometric technologies. We continue to maintain, as we have from the outset, that Issuing guidelines would be very beneficial, noting that OPC has not demonstrated that it has the necessary expertise to author these guidelines on its own.
- We believe there should be an evidential basis to justify any regulatory changes, especially if the consequence is limiting otherwise legitimate and beneficial activity. Public concerns justify an investigation, not regulation. Empirical evidence that concerns are legitimate problems and that regulation is the **best** solution is necessary for designing good regulatory rules. The OPC's own research reveals no other regulator has responded this way, and all of the case studies cited fail to demonstrate how a CoP would have made any difference.
- We believe biometric classification should be excluded from consideration. Classification is a legitimate and essential activity for any organisation dealing with many people. Biometric technologies have a lot to contribute, including the preservation of anonymity which benefits privacy. The public concerns regarding how businesses use or misuse classifications are not

specific to biometric technologies and are addressed by the Privacy Act privacy principles.

Primary Concerns

1. This CoP (should it proceed) will be unique for OPC. It has been modelled on existing CoPs which are targeted at specific activities by specific agents in regulated industries. This CoP is focused on any activity utilising biometric data or technologies, by anyone. The potential for unintended consequences is incredibly high.

2. The proposed Code is a response to public concerns, rather than actual privacy threat analysis. Most public concerns are well known, and huge efforts have been made in recent years to address them. These concerns are genuine, but some are no longer legitimate as the practices have been superseded. At the same time, there are significant known privacy threats which are not acknowledged or addressed. A CoP, targeted at specific activities by specific agents in regulated industries, could be a valuable addition. However, that is not what is being proposed.

3. There are material discrepancies between the language used in the public and consultation documentation, and the draft CoP. At the simplest level, recognition is not the same thing as identification, an individual is not the same thing as someone, and recordings are not the same thing as biometric data. This mismatch can result in either disappointment from expectations set in the public documents or nullifying of legal obligations contained in the CoP.

4. In other Codes of Practice there is an obligation to review the Code sometime in the future (e.g., 2 years). There is no such provision in this Code, and there should be.

5. It is notable that there is little consideration for the lifecycle of biometric systems (enrol, learn, etc). That is an oversight. It should include a clause that grants biometric service providers access to data that may include PII for the purposes of improving the performance of the system. [reference Rule 11, for example. Rule 12 affects Overseas Biometric providers who NZ Agencies rely on - but does not answer this question]

Unintended Consequence Examples

- Obstructing Board level investment decision making. This is already evident from the direct experience of our members. Many organisations are waiting for clarity on the position the OPC will take on the CoP before even considering investing in solutions that include facial recognition. Innovation gets stifled when risk assessment at the governance level is hard. The muddled rules, definitions, and exceptions will make risk assessment too difficult. The irony is that biometrics have the potential to preserve anonymity, so impedding investment decisions will prevent businesses from making privacy improvements.

- The exceptions necessary to accommodate legitimate and safe uses of biometrics appear to dilute the CoP to the point that some of the Act's privacy principles may be less enforceable.

Question 1: Do you agree with these provisions? Do these rules or considerations adequately respond to concerns about Māori data? Do you have any suggestions for changing them? Have we missed anything?

Answer 1: This is a multi-part question.
CoP Provisions: We do not agree that the provisions in this CoP are appropriate, or will be effective. We believe the provisions of this CoP are misdirected. Rather than regulating specific activities, they attempt to regulate "how" technology and information is used in systems and processes. If the Code was directed at specific activities which people are responsible for, then privacy principles are easily applied. Accountability for conduct and consequences is straightforward. The technologies, systems designs, process flows, and nature of data are constantly changing. Regulation is law governing the actions of people, not the design of solutions. For example, policy and procedures in a business can be updated within 6 months, whereas system-level changes can take years.

Māori Data: We are unsure if the considerations in the Code adequately respond to concerns about Māori data in relation to biometrics. We do recognise that Māori concerns regarding Māori data are not confined solely to biometrics, and so any specific rules in the Code should be respectful of and consistent with other regulatory protection.

Question 2: If you are Māori, do you agree with the way we are proposing to protect your biometric information?

Answer 2: Digital Identity New Zealand is committed to being tiriti-honouring by giving mana to Te Tiriti o Waitangi and being an effective treaty partner with Tāngata Whenua. However, in answering this question, we are unable to provide a Māori perspective.

Question 3: Do you agree that the Code should focus on automated processing of biometric information?

Answer 3: No, we do not agree.

What matters are the consequences, regardless of whether some or all tasks in a process are automated. Creating different treatments for manual, automated, and hybrid systems makes the CoP inconsistent and complex. It is legitimate that OPC focus on automation of processes, as they can dramatically scale the impact of privacy infringements. This is true for any automated processes handling any personal data (not just biometrics). However, at the individual level, an infringement is the same whether it was automated or not. When automating, or outsourcing, a business actually has a greater duty of care at the management level to ensure privacy principles are adhered to. This has nothing to do with biometrics though, it's the same for all processes.

Regarding automation:

We suggest hybrid systems need to be in scope as well, and not considering purely manual processes seems contrary to the intention of the Act itself that is concerned with the **impact** of the use of certain types of information: if the impact arises in a 'hybrid' or purely manual scenario it should be caught by the Act / guidance / Code.

Consideration should be given both to the manual processes associated with either making identity decisions or assessing the outputs of automated biometric matching. Insufficient rigour and lacking quality assurance in the human decision-making realm can similarly introduce systemic bias / noise in the identity management process.

The extent of new risks presented by biometrics (e.g., risks that are not present with use of other information or other service delivery models, like purely-biographical identity management, or human-only processes) should also be explored.

Referencing the Consultation Paper, we maintain that

- it is misguided to assume that automated biometric processing has a higher risk profile than human processing without empirical evidence, and

- humans also rely on complex and opaque cognitive processes to make judgements or infer information about humans.

Regarding the definition of *processing* as per the proposed Code:

There are dozens of processing steps covered by the ISO definition that are not explicitly addressed by the OPC definition. Those processing steps generate something other than an artefact defined under biometric result: signal detection, segmentation, biometric feature extraction, quality assessment, biometric model generation, biometric template generation, comparison, biometric comparison decision, compression, decompression, etc. The guidance / Code should follow the established ISO definitions meticulously and not cherry-pick aspects of the definition, or operators will be left with ambiguity as to whether various operational steps count as processing and if the end result of a particular step is to be disclosed to people requesting their own data.

Question 4: Do you agree with the definitions of physiological and behavioural biometrics? Can you think of any types of biometric information that aren't captured within these definitions that should be? Or any types that we should exclude?

Answer 4:

Biometrics are measurements and analysis. Records are not biometrics, they are records. For example, an image is a record whereas a biometric analysis of that image produces measurements/results. The image itself is a record but is not a biometric record. The word "record" should be dropped or clarified to be consistent with international treatment.

The footnotes provide an alternative definition in the Code. This is confusing — which definition should we apply?

There already are (and foreseeably will be many more) biometric modalities that blur the distinction between physiological and behavioural biometrics as per the current definitions. Inner ear acoustics for example is based on the acoustic characteristics of the ear canal and identifies an individual based on the echo sound characteristics from this three-dimensional chamber. The individual doesn't consciously "perform or respond", so it is not a behavioural biometric, but the signal used for identification / verification is also not based on the "physical appearance" of the inner ear, so it is not a physiological biometric. Probably a better distinction can be made with terms like static vs dynamic biometrics, the latter of which relies on time as a dimension of the biometric sample. However, even the static vs dynamic distinction is insufficient in going into issues such as whether gender is a physiological, physical, or behavioural attribute.

At the same time, it is a statistical fact that at a large-enough scale most activities that get logged become sufficiently (often even uniquely) identifying. For example, NZ's now-abandoned Covid tracer app could in theory uniquely identify individuals based on movements and network linkages. Similarly, current telecommunications agencies can trace the movement of mobile phones with the help of the physical tower locations the mobile phone pinged across. Since neither of these two modalities are based on the physical appearance of the human body, currently these would count as behavioural biometrics, but for the above mentioned reasons, dynamic biometrics is probably a less misleading term.

Question 5: Do you agree with the definition of biometric information and the types of biometrics it includes (samples, templates, results)?

Answer 5: No, we don't in their current form.

The ISO definition of *biometric sample* (37.03.21) is "analogue or digital representation of biometric characteristics (37.01.02) prior to biometric feature extraction (37.05.04)". While the ISO definition rightfully includes analogue records, if OPC's intended scope of the Code is automated based processes only, then how do these suppose to apply to analogue records? It is of note also that while the definition itself does not, Footnote 17 in the Code (accidentally) includes in the definition everything that can act as inputs to biometric processing. Given that processing covers potentially dozens of steps, the OPC's definition of sample now also covers templates, while the corresponding ISO definition explicitly says samples are the artefacts *prior to* biometric feature extraction, so templates and "results" are not covered by it.

The starting point should be that only data that has undergone automated processing to identify a specific individual, should be considered 'biometric information' for the purposes of this proposed Code. That is, a physical or digital photograph, a video or audio recording, publicly available information, or data generated therefrom; or information collected, used, or stored for health care treatment, payment, or operations should not be considered 'biometric information' in a privacy Code.

The definition of *biometric result* is both too wide and too narrow at the same time, as it misses fundamental concepts like a comparison score, comparison decision, and a candidate list, all of which are only interpretable in the context of a threshold or other decision policy.

We maintain that likely not all information subsumed under *biometric results* is disclosable to individuals requesting their own data without infringing on other people's privacy (namely, that of the probe and candidates in the same search).

The exclusions of *biological material* from *biometric information* also leads to inconsistencies: besides the fact that both bone and blood are tissues, veins and arteries are organs, which would make the Code not apply to any form of vascular biometrics, such as hand and palm veins, or retina based biometrics.

In summary the definition is not easy to read, and has inter-dependencies with other definitions. This makes it difficult to understand for non-experts. It also doesn't clearly require that the information to be used to identify an individual. Could a photo, video or audio recording that doesn't identify an individual be captured under the current definition?

We strongly suggest the OPC use and not modify the ISO standard definitions in their proposed guidance, and not invent an informal standard.

Question 6: Do you agree with the exclusion of heartbeat from the definition of behavioural biometrics, or do you think it should be covered by the Code? Why?

Answer 6:

If heartbeat biometrics can be as accurate as fingerprints in identifying a person as inferred here:

https://www.biometricsinstitute.org/types-of-biometrics-heartbeat-key-considerations/ then not having this biometric type in scope seems counterproductive. But in doing so it opens up a proverbial 'can of worms'.

This is actually a good example of how complex this CoP will be to apply. Since a heartbeat can be superficially observed or deeply analysed for a variety of uses, the privacy implications come from the activity rather than the type of data. The same thing can be said for any biometric, which is why Codes of practice normally focus on activities rather than types of data. This question illustrates the point we make perfectly..

We think guidance should refrain from scoping in or out specific technologies as doing so is likely to lead to loopholes and unintended exceptions, while also being less technologically agnostic and less future-proof than the Act in this rapidly-changing emerging technology sector.

With respect to verification and identification, we note that the EU's General Data Protection Regulation (GDPR) Article 9(1) is relevant and we suggest the OPC consider alignment with this. The GDPR Article 9(1) states that the processing of photographs is covered by the definition of biometric data only when processed through a specific technical means allowing the unique identification or authentication of a natural person.

DINZ recommends that only data that has undergone automated processing to identify a specific individual, should be considered 'biometric information' for the purposes of this proposed Code. That is, a physical or digital photograph, a video or audio recording, publicly available information, or data generated therefrom; or information collected, used, or stored for health care treatment, payment, or operations should not be considered 'biometric information' in a privacy Code.

Question 7: Do you agree with the definitions of biometric processing and biometric verification and identification? What would you change and why?

Answer 7: DINZ maintains that any guidance or a Code should not be aimed at a layperson as a target audience, and it must be technically explicit and accurate. Accessibility considerations in a scenario like this are misguided, as they will create the false impression that a naïve operator actually understands the nuances and ramifications of these highly-complex solutions. We maintain that if the OPC's readership does not understand the ISO definitions that the OPC unhelpfully "simplified" beyond usefulness, they should not operate these solutions.

In more detail:

- Many processing steps will generate something else than an artefact currently defined under *biometric result*: signal detection, signal quality assessment, segmentation, biometric feature extraction, biometric model generation, biometric template generation, model updates, comparison, biometric comparison decision, compression, decompression, etc.

- *Biometric verification* decidedly does not mean "the process of seeking to verify the **identity** of an individual". It is only confirming a biometric claim (i.e., does this template look sufficiently similar to this other template to assume that they both originate from the same biometric capture subject) — even in the absence of linked raw sample or any biographic information like name, DoB, driver's licence, or passport number.

- Similarly, biometric identification does not mean or intend to "identify an individual"; it only searches against a biometric enrolment database to find and return the biometric reference identifiers attributable to a single individual.

Question 8: Do you agree with the more technical definitions in the Code (biometric search, query, reference, sample, template and comparison decision)? Are they accurate, too detailed, not detailed enough?

Answer 8:

OPC should be defining activities rather than technical terms. For example an identity service provider using biometric (or other) technologies to validate credentials. Regardless of the technical terms used, it is the business activity and responsibility to care for privacy. Also, technical terms should not be in regulations — they change too frequently and often have multiple interpretations.

More specifically:

- *biometric search*: The ISO definition also contains a biometric candidate list, which necessitates and implies human or other-system intervention / augmentation of the decision-making system.

- *biometric query*: While close enough to the ISO definition, we can confirm that the term *biometric probe* is much more widely used by the industry and operators. For most of the data science community a *query* will be synonymous with a *search*, while *probe* will not have similar misleading connotations.

- *biometric reference*: Without the full scope of the ISO definition, specifically the addition of "attributed to a biometric data subject" this term is not meaningfully distinguishable from the current form of the term *biometric query* (probe).

- *biometric sample*: See under Q5.

- *biometric template*: In reality, these may not be numerical or algorithmic at all, and the definition needs to specify that it is intended to be used as reference in comparisons. Just any digital abstraction created from a sample should not constitute a template, or any digital art depicting faces could technically fall under the definition.

- *comparison decision*: If the definition assumes the existence of a *probable match*, then you have to assume there are human assessors in the loop. We feel that the possible decisions should include further labels, such as *not enough data* and *inconclusive* for verification and identification solutions, or similar along those lines, and possibly many other further labels for classification solutions.

Question 9: Do you agree with our definition of biometric classification i.e. do you agree that a biometrics code should cover these types of biometric classifications? Is it too broad or too narrow? What would you add, amend, or remove and why?

Answer 9:
We do not agree that the use of biometric technology creates a distinctive set of classifications that can be regulated. A business may use biometric data, and other data, to classify people for legitimate purposes. These activities should be, and are, regulated by the Privacy Act and other protective laws. A separate set of classifications and rules when biometric technologies are used is incredibly complex and not warranted.

The use of classification also infringes on the privacy of personal opinions. Inferring the inner state of someone from observing their behaviour is forming an opinion, it is not obtaining factual data. This is why we believe guidelines are valuable to help understand when subjects like observations being factual or opinionated.

The term *physical state* is completely misleading as it is used in the draft Code to refer to *psychological* states, and the distinction between these psychological states and inner states is more complex than OPC allows for.

It feels *biometric category* as a term is misleading as these are purely demographic categories — why prefix biometric? From the prose in the consultation document it seems that inaccuracy is the main driver for exclusion from fair use — but this will stifle innovation.

Question 10: Do you agree with the intent to exclude some processes from the definition of biometric classification? What do you think of the two exclusions we've proposed (detection of readily apparent expressions and integrated analytical features) and the way they are defined?

Answer 10: No, because if we argue that human judgements are within scope, these need to be in scope of the definition of biometric classification.

Readily apparent expression: the interpretation of these is really far from trivial or low risk. Did that person with Parkinson's nod or is it tremors in the neck muscle? Don't even get us started on the number of sexual assault cases that hinged on the interpretation of whether it was an encouraging smile or not. Audio volume / amplitude changes are also not trivial to interpret in a given noisy context, so whether humans or machines make these judgments about readily-apparent "expressions", a bit more nuance is needed. Arguably the potential to harm exists in this context, so based on that the Act / guidance should cover these use cases.

Question 11: Do you agree that the code should apply to any organisation that starts using biometrics after the code becomes law?

Answer 11:

If something becomes law, any organisation is subject to it. Even if the law is detrimental to a business, we are governed by the rule of law. A better question would be how will the OPC respond if it is made aware that the Code is causing harm, either within scope of privacy or wider socio-economic impact?

Question 12: Do you agree that organisations already using biometrics when the code comes into force should have more time to comply? If you are an organisation that is already doing biometric processing, do you think the additional six-months to bring your activities into alignment with the code is fair?

Answer 12: If the focus was on privacy policies and procedures within an organisation, 6 months to revise existing policies and procedures is not unreasonable, as a target. Those deep into the industry say that large or complex organisations may need more time, and OPC would need to be accommodating of reasonable exemption requests, including after the commencement of the Code.

However the focus is on processes, systems, and data. These are transformative changes to a business, which can take years to plan and implement, at huge expense. Even small businesses would likely need an exemption from OPC. Given the prioritisation, budgeting, procurement, project management and testing obligations of large government agencies, as well as the complex legal and technical negotiations in a multi-vendor situation, 24 months is a more realistic timeframe.

DINZ maintains that guidance and CoPs for governance and management of activities adhere to the privacy principles, rather than operational activities - that focus would be more appropriate.

Question 13: Do you agree with the exclusion for health agencies?

Answer 13: In principle DINZ agrees with the exclusion of health agencies, but as the OPC's own consultation document exemplifies, blatant loopholes like insurance agencies using biometrics to price policies being excluded will need to be resolved some way.

Question 14: Do you agree that health agencies collecting non-patient biometric information should have to comply with the code?

Answer 14: DINZ supports this requirement.

Question 15: Do you agree with the additional requirement that organisations must ensure the biometric processing is proportionate?

Answer 15: No, because while the judgement element in the proportionality test is important, it's unclear what this additional requirement would add/intended to achieve. Proportionality is implicit in only collecting information that is necessary as per IPP 1 in the Act. OPC suggests that agencies should only proceed "if they consider the risks are appropriately minimal or able to be managed", but just being able to demonstrate that the proposed solution operates at a lower privacy risk than alternative solutions (while also serving the business's requirements of volume processing, risk, FTE, timeliness, cost, service friction, etc) should be a good-enough reason for implementation.

This is a subjective provision which needs specific policy statements from OPC to be enforceable. It adds risk to any decision to adopt biometrics, without achieving any privacy goals. Any two people can subjectively reach different conclusions on necessity or proportionality. All comes down to the legal consequence of OPC disagreeing with a business decision. For example, if a business makes a reasonable determination at the time and there is a subsequent issue, will the OPC accept the business made a reasonable determination after the fact?

Question 16: Do you agree with the six factors listed in rule 1(2) that an organisation must consider when considering proportionality? Would you amend, add, or remove any of these factors and why?

Answer 16: Our main concern here is that it's hard to conceive sector-agnostic definitions of *proportional* collection, *effectiveness*, or *discrimination*.

The requirement that biometric processing should be *effective* makes it look like that effectiveness is binary (something is either effective or not, with nothing in between), and as a result it seemingly cannot accept gains in effectiveness as a legitimate business goal. Arguably, agencies demonstrating improvement over the already-existing solution or alternative solutions should be sufficient.

Regarding *cultural impacts and effects*:

- discrimination, or intentional profile-driven treatment is at the base of all risk assessment and security application

- if discrimination is off the table, so is positive discrimination

- biometric technology is still not sentient: the intention to surveil is the cause of surveillance, tracking and profiling, and not the size or type of biometrics holdings. This is a concern about the misuse or abuse of technology, and not about the technology or the magnitude of the collected data itself.

Additionally, noting other legal frameworks relevant here, there is the further consideration that all surveillance, tracking, and profiling are all lawful and supportive of national security if done responsibly.

Question 17: Do you agree with our definition of privacy risk? Do you agree with the privacy risks listed? Would you amend, remove, or add to any of these risks?

Answer 17: Apparent *inaccuracy* due to biometrics records "aging out" of use during the legally-mandated retention time is a fact of life for many agencies with long retention times: NZ Police, DIA, Corrections, and INZ are legally obliged to retain information for decades, including face images.

Bias is similarly a fact of life and should be considered in a more nuanced way than the simplistic black-and-white approach that the draft Code displays:

- modern automated systems – even those that transparently display bias typically outperform human resolvers in assessing biometrics when it comes to neutrality, fairness, speed, and consistency
- even a 100% error-rate differential between two demographic groups in identification, verification, or classification can be inconsequential, when put into the context of volume: as an example, 1 error in 1 million male faces vs 2 errors in 1 million female faces may create the impression of a media sensation-worthy, unfair gender difference, while still vastly surpassing known human benchmarks.

The inclusion list doesn't describe risks, it describes scenarios. A risk is a definition of a harm caused by something - the list doesn't identify the harms. For example over collection is a breach of principles, but the actual risk of harm to an individual is not described.

What is missing from the requirements around privacy risk is the explicit consideration of already existing privacy risk in current-state solutions, which are often based on purely biographical identity management, or human-only processes.

Question 18: Do you agree with the definition of benefit? Do you agree that the higher weighting should be given to public and individual benefit (as opposed to the benefit to the organisation)?

Answer 18: No, we don't. Taking the example of Foodstuffs North Island's FRT trial, intent of the system implementation is primarily to weed out offenders convicted of or known for assaulting staff. That is an individual benefit to organisation's staff, but OPC's definition makes a presumption that an individual benefit is the antithesis of an organisation benefit.

Question 19: Do you agree with the requirement for organisations to adopt reasonable and relevant privacy safeguards to mitigate privacy risk?

Answer 19: This requirement is already in the Act. We do not see how this requirement changes that. If the intent is to change what is in the Act, please be more prescriptive.

Question 20: Do you agree with the definition of privacy safeguards? Do you think the list of privacy safeguard covers appropriate safeguards for biometric processing? Would you amend, add, or remove any of these factors and why?

Answer 20:
OPC should not prescribe the safeguards, but rather require businesses to maintain appropriate safeguards. The list is only relevant for a limited set of use cases. For example most are not relevant if someone is using biometric

technologies built into their phone, or online with an Identity Service Provider selected by the user/client. Or where the biometric system performs a task that is outsourced to a third party provider. The safeguards in all these instances are fundamentally different to the safeguards listed.

Also:
The phrase *reasonably practicable* makes these not (easily) enforceable.

Regarding *informed decision / consent*

Requiring consent fundamentally shifts the basis for processing of personal information under the Act and it does not address the risk of individuals potentially losing control of their biometric information. It is not clear who would "consent" on behalf of children, young adults, missing persons, trafficked persons, and persons who need to be enrolled into a watchlist (i.e., they are not on a watchlist yet). It would likely be proportionate and practical to require *certain* sectors or use cases to require informed consent and to provide reasonable alternatives when consent is not provided.

If due diligence has been performed at the outset (per the guidance), then biometrics should not be being collected and used in an automated system without legitimate cause, in which case, voluntary exemption would likely mean an overall worse outcome for the subject (requirement to collect even more personal information, slower service provision, higher likelihood of lower-quality management of them as an entity). There would be a wide range of use-cases in which not having the data collected is not possible (e.g., where cameras cannot reasonably exclude one individual in an environment who does not consent), and mandatory adherence to allowing exception could significantly increase the complexity of administering any system — reducing the benefits and the likelihood the technology will be used.

It is unclear from the definition if post hoc revocation of the authorisation is still possible, but foreseeably could introduce a wide range of complex requirements on agencies to administer identification that the claimant has sufficient rights to request withdrawal, which in of itself may create a circular loop of information being collected.

Regarding *biometric watchlists*

3(3)(b)(i-ii):

3(3).b.i

How is this limited to instances where suspected foul play is suspected and needs 'proof'? And also in light or 3.(4)? It is often neither possible, nor reasonable to attempt to inform individuals of their watchlist status:

- in ongoing investigations it defies the purpose of watchlists to inform watchlist capture subjects

- in biometric-only watchlist an agency simply may not have a known biographic record associated with the biometric capture subject

- a biometric sample might be watchlisted much later than when the capture subject is enrolled in the database

3(3)(b)(iii-iv): Adverse actions taken by a business or agency (like denying credit, not issuing a passport, not allowing access to the controlled-substance room, or entry to a casino) are not within the scope of biometrics, but firmly in business / service-decision territory, which should remain in the control of the business and should not be present in the guidance or Code.

Regarding *system testing*

It is probably prudent to add that the expectation is that the system was not only tested by the vendor providing the solution on their own laboratory test dataset, but was tested independently of the vendor on operational (or operationally relevant) data.

It is unclear if this definition includes the testing or benchmarking of staff involved in these systems.

It is important to note that this requirement would increase compliance costs in terms of additional or contract data scientists and setting up and maintaining a testing environment, which in turn could inhibit beneficial use of this technology.

Regarding *monitoring* and *training*

Training staff is certainly needed, but research evidence shows it is not enough when employing human assessors to compare images. Assessors also need to display a very high level of natural aptitude in this domain, so we recommend aptitude testing (personnel selection), training and ongoing benchmarking staff as the bare minimum requirements here. Research conducted at the University of New South Wales (White et al 2014) found that passport-issuing officers had high error rates (including 14% false acceptance), which was on par with naïve student participants. Passport officers showed no performance advantage over the general population on a standardised face-matching task. In fact, across all tasks, very large individual differences were observed: while average performance of passport staff was poor, some officers performed very accurately, but more importantly, the accuracy was not related to length of experience or training.

After explicitly testing face image comparison training courses on a sample size of 236 practitioners and 152 novices, the same academic workgroup (Towler at al 2019) observed that improvements due to training were small, inconsistent across tests, and training did not produce the qualitative changes associated with examiners' expertise. This lack of improvement is all the more striking in the light of 93% of all trainees believing their performance had improved.

Towler A, Kemp RI, Burton AM, Dunn JD, Wayne T, Moreton R, et al. (2019) *Do professional facial image comparison training courses work?* PLoS ONE 14(2): e0211037. https://doi.org/10.1371/journal.pone.0211037

White D, Kemp RI, Jenkins R, Matheson M, Burton AM (2014) *Passport Officers' Errors in Face Matching*. PLoS ONE 9(8): e103510. https://doi.org/10.1371/journal.pone.0103510

Monitoring and reporting are necessary, but it may be reasonable to guide that agencies should not implement technological solutions that they cannot reasonably monitor, tune, and maintain over time proportionately to their use case. It is relevant that there are existing avenues for recourse for unfair practices (e.g. consumer protection laws, complaints to the Ombudsman and industry bodies). However, even if accuracy can be meaningfully operationally defined (alluded to here as *false positives* and *negatives*), specific requirements may be overburdensome and result in denying use of efficiency and privacy enhancing technologies to smaller agencies that cannot shoulder the costs required to quantitatively assess the concept of accuracy. This may mean the agencies are unable to access the benefits of biometric technology.

Question 21: Do you agree with the additional notification matters? Can you think of any other matters that an organisation should be transparent about?

Answer 21: "A list of any policies, protocols, and procedures that apply to the organisation's use and disclosure of biometric information" might not be feasible as they are often subject to common law restrictions such as confidentiality, or to intellectual property restrictions. It is also unclear how this requirement will apply to service providers who verify an individual's identity on behalf of another service provider.

Question 22: Do you agree with the requirement for organisations to have a conspicuous notice? Do you agree with the definition of conspicuous notice?;

Question 23: Do you agree with the matters that need to be on the conspicuous notice? Are there any items that you think should be added to the conspicuous notice? Or removed?;

Question 24: Do you agree with the requirement for agencies to have an accessible notice? Do you agree with the definition of accessible notice?

Answers 22, 23, 24: Noting that there is already an obligation under the Act to give notice, DINZ is concerned with the necessity of and the meaningful differences between *accessible, noticeable*, and *conspicuous* notices (the distinction between "readily accessible" and "readily noticeable" being a case in point regarding clarity), especially in the light of the requirement and safety mechanism of obtaining informed consent.

Question 25: Do you agree that some exceptions should be removed to strengthen the notification obligations? Would you remove, keep or add some exceptions, and if so, which ones?

Answer 25: We are uncertain of the definition of *recent* in the phrase *recent previous occasion*.

Question 26:

See below, batched together with Q28 and Q29.

Question 27: Because health agencies are excluded from scope, insurance agencies providing health insurance won't be subject to this processing limit on inferring health information (although they'll still have to comply with the HIPC). Do you think this is problematic or a gap in the code's coverage? Are you aware of any other regulation that puts rules in place for insurance agencies that would regulate this?

Answer 27: There is an issue with health information being collected by technology companies, who are not subject to HIPC, as they are not health agencies. All our smart-watch/health trackers are collecting health information that is not managed by health agencies.

Question 26: Do you agree with the fair processing limit on using biometrics to detect or attempt to detect health information?; **Question 28**: Do you agree with the fair processing limit on using biometrics to infer or attempt to infer emotions, personality or mental state?; **Question 29:** Do you agree with the fair processing limit on using biometrics to detect physical state generally? Do you agree with the exception for detecting physical state if necessary to comply with a health or safety standard? Or do you think this use should also be restricted? Is the exception drafted too broadly or too narrowly?

Answers 26, 28, 29: No. The policy case and evidence base for limiting individuals' freedom to contract and consent to use of their biometric information for any purpose is not clear. Particularly where the basis of that consent (and other transparency requirements) is (and should already be actively) regulated. There are particular uses that should only (generally) be permissible based on consent, but any purpose expressly ruled out would be predictably based on OPC's current state understanding of not-necessarily current state use, ignoring current state technological advancements, which would cripple technological innovation and progress of these same contexts unnecessarily, effectively barring them from becoming better and more useful.

OPC's hopeful proviso "If the technology developed and proved appropriately accurate, non-discriminatory, and societally permissible then the Code could be modified" suffers from a circularity problem: if a technology is banned by current-state legislation, no investor or research group will try to develop them knowing that even if they are able to show technical improvements (in accuracy or bias), by then societal attitudes generated by the Code itself will likely not support or trust the technology. Even if the Code and the publicity around it somehow did not shape public sentiment and attitudes, investors also understand that changing legislation takes multiple years, which would further discourage the same investors from trying to market products in NZ years after they sank money into developing them.

The current shape of the Code neglects the crucial keyword of wellbeing from their references to health and safety in relation to people. We argue that wellbeing should be included, as inferring emotions, personality, and mental state are crucial elements to managing the wellbeing of people (staff, visitors, customers, etc) under an employer's duty of care.

As for the terminology, we feel it is misleading to use the word *physical* to mean a certain subclass of *psychological* states covering alertness, fatigue, focus, and attention. We understand the connection: slower than average heartbeat and oxygen levels, physical-physiological processes, can create the psychological perception of tiredness or loss of focus. Even if a fair processing limit were to be placed on using biometrics to detect or infer fatigue and focus related properties, we agree with the exception to facilitate compliance with health and safety standards.

Question 30: If you are an employer or employee, what do you think about this exception? Can you see beneficial or problematic cases for monitoring physical state (attention, fatigue) for health and safety reasons in your workplace?

Answer 30:

While we agree with the general statement in the consultation document that "human emotion; it is incredibly complex and varies across cultures, contexts, and individuals", it is important to note that most of these classification models can in fact be further trained and tuned to an individual in a particular context. For example, we would not expect a fatigue classifier to be applied straight out of the box on any staff member in any use case in any country, having been trained on an arbitrary vendor sample, but getting first further trained on Bob for an amount of time in the particular use case, specifically so it is able to evaluate that Bob is now exhibiting attributes associated with tiredness.

There is substantial legislation governing employment relations, conditions, and workplaces. This includes specific enhancements to protection of human, cultural, and privacy rights given to people in NZ. Conditions of employment, company policies and procedures, staff handbooks, and security protocols are all crafted to ensure businesses meet their legal obligations and duty of care to their employees. There are thousands of roles where being physically able to perform duties is critical, and where biometrics could play a valuable role. Driving vehicles, operating machinery, working at height, carrying loads, etc, etc. There are thousands of roles where cultural sensitivity is important. You have to be able to classify people in order to respect certain aspects of their identity. For example dressing appropriately for a monocultural community, being able to speak the native language of a customer, being culturally aware to understand gestures and body language, etc. Biometrics is a tool that can assist with all these things.

The diversity of workplaces and roles, and how we interact with each other, makes any prescriptive provisions for biometric classification in the Code impractical. The existing Act principles and conduct provisions provide a practical foundation to regulate privacy. The other laws governing workplaces and workforces address the specific needs to regulate those scenarios. These classification provisions in the Code make everything more complex and will likely be impractical to comply with.

Question 31: Do you agree with the fair processing limit on using biometrics to place people in categories that are protected under the HRA? Are there any categories we've missed that are intrusive? Can you think of any beneficial uses for placing people into these categories?

Answer 31:

The Human Rights Act is aimed at preventing **unlawful** discrimination. There are legitimate reasons to discriminate, which are essential to helping those who are disadvantaged. Limiting access to information impedes both legitimate and unlawful discrimination. Our human traits are an important aspect of our personal identities. Biometrics are a useful tool in ensuring we know each other. If your business cares for the elderly or the very young, provides women-specific personal services, has a customer base that is faith-focused or mono-cultural, etc, then biometrics may enable powerfully-beneficial use-cases that do not infringe privacy.

The HRA lists over 25 exceptions for various use-cases. This Code suggests all discrimination is harmful, which is not the goal of the HRA.

DINZ contends that individuals, for whatever reason, might *want* to know where they place on various continua along these categories. Similarly, already-existing unequal-treatment biases (such as those known to affect Māori and Pasifika in healthcare) cannot be fought or rectified without the helpful application of positive discrimination). Arguably 21(1)(h) of the HRA could block classifying for intoxication levels, which wouldn't be exempted by the general exceptions.

Question 32: Do you agree with the exception for age-estimation? Do you agree with the way we've drafted the age-estimation exception – can only use it if necessary to comply with lawful obligation to apply an access limit or meet a duty of care?

Answer 32: We agree with the exception for age-estimation.

Question 33: Do you agree with providing the standard 'serious threat' and 'research' exceptions to the fair processing limits? Do you agree that the research exception should be strengthened by adding written authorisation requirement and ethical oversight and approval requirements?

Answer 33: In principle we agree with the serious threat and research exceptions, but the OPC should note that while universities and research institutes will have well-established ethics committees and institutional review boards, it is unclear who would (could) provide the ethical oversight in other agencies and businesses. Additionally, a maintenance of the law exception should be included here for cases like child and migrant exploitation or trafficking.

The individual consent requirement actually makes biometrics more privacy invasive, as many research and statistical activities do not require the identification of the individual. By specifying this just for biometrics, does that mean other forms of personal information require less privacy protection?

In summary, all organisations have governing legislation that require them to act in good faith and have effective governance in place. They strengthen the Privacy Act by attributing personal accountability of those in leadership positions. Prescribing duties in the Code will be far less effective than legislation already in place that is specifically designed to make people accountable.

Question 34: Do you agree with the exception to the fair processing limits for assisting an individual with accessibility? Do you agree with our definition of accessibility?

Answer 34: We agree with the exception to the fair processing limits for assisting individuals with accessibility.

However, the definition of *accessibility* currently does not consider that even in the absence of a disability assisting with accessibility might be desirable. For example, non-native speakers of a target language are not disabled, but would benefit from an audio signal processing algorithm that bins them into native(-like) vs non-native speaker categories in order to facilitate access to translators, interpreters, or more readable documents. Similarly, people wearing different types of footwear might pose different tripping, slipping and catching hazards on an escalator or construction site.

Question 35: Do you think there needs to be other exceptions to the fair processing limits? What exceptions would you suggest and why are they needed?

Answer 35: We believe the fair processing limits in the Code are so problematic that the entire section should be removed. That would remove the need for more exceptions, such as the 27 industry-specific exceptions contained in the HRA.

We have no further exceptions to suggest.

Question 36: Do you agree that the collection exception should be changed so the threshold is higher for relying on it?

Answer 36: We do not have a strong stance on this.

There are a range of activities such as web scraping, data harvesting, personal device monitoring, etc, that have major privacy implications. OPC should be examining these activities as potentially benefiting from a CoP. **Most importantly**, none of these are specific to biometric technology or biometric information, and the egregious acts typically do not involve biometrics.

Question 37: Do you agree that agencies shouldn't be able to rely on this exception to collect biometric information by web scraping? What do you think of our definition of web scraping? Does it cover what we intend to capture?

Answer 37: A useful definition of web scraping is likely to be difficult, but usually there would be a *de minimis* argument. Currently, it is not clear why the unfair collection means principle is not sufficient to address unfair web scraping.

Does the definition capture mere photos, video and audio? If so, prohibiting a web crawler from downloading any photos of people from the internet would have a huge impact for AI training, where the photos are not being used to create identification databases. If the concern is creation of databases without consent, the prohibition can be more narrowly drawn.

Unless the biometric analysis on such data is intended to identify an individual, the risk of harm or surveillance is relatively low and shouldn't be prohibited outright. There are positive use cases that rely on such data that don't require identifying individuals. For example, testing and reducing bias in biometric systems requires large, diverse samples of data to ensure representation.

We note that the current proposed definition of *web scraping* does not capture 1) an agency or individuals *manually* extracting large scale collections, or 2) any kind of extraction happening on *non-public* online sources, both of which are easily exploitable by malicious actors, despite presumably not being the intended outcome by OPC.

The *publicly available* information in the definition limits the freedom of operators of digital platforms to decide how information on their platform will be made available, and their users to elect to use such services. If an individual consents to their information being made available on a service that permits web scraping, it is not clear why the OPC / the Act / a Code should prohibit this.

Further, the basis for belief that people "would not reasonably expect their information to be used in this way" is not clear. There is a view that the typical user of such services is aware of the risks of posting information about their real identity online because it could obviously be used by any agency or actor to identify them (sometimes this is the point of posting biometric information in the form of photographs), and thus potentially to impersonate them. We question whether there is a reasonable expectation of privacy in information posted to public-facing websites, though this will always be a matter of fact and degree (including relevant terms of service). What is likely to be clear is that providing the information to a service provider does involve giving up some control, usually in return for the benefit(s) the service offers.

Question 38: Do you agree that an organisation should have to tell the individual what form of biometric information they hold about them?

Answer 38 & 39: No, we do not. As already noted in our response to Question 5, we maintain that the proposed definition of *biometric information* is an overreach, needs to use the ISO definition and just confirming what type of *biometric sample* (if any) is held about the individual should suffice.

Notifying an individual that an agency holds a *biometric template* as well is likely meaningless to the requestor, and disclosing a by-default encrypted, unreadable, unviewable template or feature set is even less meaningful to the requestor, while also possibly harming intellectual property law and (at scale) potentially creating a reverse engineering opportunity for bad actors.

Particularly concerning is that since *biometric result* is subsumed under *biometric information*, 6(1)(c) would oblige agencies to disclose similarity scores and identification / verification decisions with regards to other natural persons, whether the requestor was the source of the probe sample or a candidate sample. Furthermore, when "results" are produced in the context of national or international data sharing agreements, data sharing agreements would need to be reviewed to seek partners' agreement on disclosing the results of searches against *their* database.

Question 39: Do you have ideas for other ways rule 6 could be modified to give a person more oversight of what information is held by the organisation?

Answered in #38

Question 40: Do you agree with the intent of this modification? Do you agree with how this provision is drafted?

Answer 40: We agree with the intent of this modification and the way it is drafted.

Question 41: Do you agree that rule 12 should require the organisation to make sure the overseas jurisdictions they're sending to have protections that reflect the heightened protections in the biometrics code, rather than the general Privacy Act?

Answer 41: No, we do not agree. No NZ agency or business has cross-jurisdictional power to mandate other countries we already share data with or have data sharing obligations towards to create or obey a similar Code.

Recurring themes

DINZ has made comments regarding these matters in previous consultation rounds and still stands by them.

1. The Code is substantially less flexible and less technology-agnostic than the Privacy Act principles.

2. Human / manual processes should be within scope of interpretation (be it the Act, guidance or the proposed Code)

3. The OPC is not a subject matter expert in the activities discussed here, and any investigation or decision by the OPC should include consultation with experts in the field.

4. The Goldilocks Problem: In previous consultation rounds it is alleged that some technologies are not accurate enough, and then in subsequent ones that these technologies are too accurate. What is the Goldilocks point for accuracy? And which end of the spectrum is driving the fair use principles?

5. DINZ still asserts that the next step should be issuance of Guidance, and only if that was not successful then the issuance of a CoP. Currently the language of the

Code is often ambiguous, arbitrarily redefining known ISO definitions, exempted, or otherwise unenforceable.

6. Terminology: Terms are introduced that either modify or depart from internationally accepted terms, in some cases without reference to the original source. That an internationally well accepted term could be used to mean something different in NZ does not reflect well on any stakeholders. It will become incredibly confusing for any organisation attempting to comply, and for a market as small as NZ we suspect many international organisations will simply not bother.

Late Contribution

A contribution from a member that represents two Australian based biometrics vendors supplying IRD, MSD, MBIE and the BNZ arrived too late to be incorporated into the main response. While the responses to some of the questions below are already addressed in our responses, given the size of the installations and as a result making it easier for New Zealander to authenticate themselves in phone conversations with contact centre agents, we have included it here to hopefully help OPC to appreciate types of queries it will need to address in order for industry to maintain operations.

1. Rule 13 seems out of context. Is it put in here to fix an issue in the parent Act this Code modifies / extends?
2. With regards Rule 2 (1) Vendor would need clarification on whether it can use information that has previously been collected for other purposes (e.g. use existing call recordings that were obtained from the individual).
3. With Regards Rule 3 (1).(d) What happens if the vendor holds the data for the agency in VBaaS?(Voice Biometrics as a Service) Does this Code apply to as-a-service operators?
   (1).(l) How is this different from 3.(1).(e)?
   (5).(a).i What classifies an “offence”? Does the Privacy Act define this?

4. With regards to Rule 5 (b) should “person” also include “service providers”?; Clarify this clause about providing information so a vendor can access it for tuning purposes.
5. With Regards to Rule 6: Clarification should be given to what information should be included / provided. One presumes handing over the Voice print is unlikely to be a satisfactory outcome for the person requesting access? (e.g,. They will not be able to do anything with it?)
6. With Regards Rule 7: (1) and (2) are standard PPI / data governance clauses; biometrics operate differently. The mechanics of correction are very different from updating a CRM record. (3) is hard to imagine to be practicable. (4) and (5) would work in individual cases but not in batch processing.
7. With Regards to Rule 10: This focuses on the sample; but what about the metadata that is associated with the sample? (3).a How can we defend that Forensics is an extension of Identity & Verification (ID&V)?