



Age Assurance Technology Trial

# PART D

## Age Estimation

*August 2025*



Funded by



**Australian Government**

**Department of Infrastructure, Transport,  
Regional Development, Communications, Sport and the Arts**

Project by



## Findings on Age Estimation

These are our headline findings. In line with the overall findings of the Trial, they relate specifically to the topic of age estimation.

1

**Age estimation can be done** in Australia; is being deployed effectively and across multiple sectors.

2

Most **systems are technically deployable** across standard devices and environments, though edge-case limitations remain.

3

Provider **claims regarding performance were generally substantiated** through independent evaluation; some early-stage systems lacked complete transparency.

4

There is no single approach to **age estimation; must be configured to context.**

5

The age estimation sector in Australia is **dynamic, innovative and responsive to privacy and fairness challenges**; providers are iterating rapidly.

**6**

**Demographic performance consistency is improving;** underrepresentation of Indigenous populations remains a challenge that vendors are beginning to address.

**7**

**Accuracy varies in suboptimal conditions,** highlighting the need for robustness improvements.

**8**

**Vendors are actively mitigating adversarial threats,** including spoofing and injection attacks.

**9**

**Age estimation decisions remained based on real-time,** independently derived evidence – not on self-declared, inferred or parental assertions.

**10**

Providers are aligning with emerging international standards and **demonstrating readiness for certification,** meeting expectations set out in ISO/IEC FDIS 27566-1.

## © Copyright of the Commonwealth of Australia

This document is available for reproduction on royalty-free, perpetual, attributed source, non-commercial rights to reproduce.

This permission allows for unlimited reproduction and distribution, provided that proper credit is given to the original author(s) and source. This grant applies to all formats and media worldwide. For queries about commercial use or the licence in general, please contact the publishers. All rights to materials on links are reserved to the author of those materials.

### Accessibility Statement:

This report was produced in line with the accessibility guidelines found in the Australian Government Style Manual.

### Legal Clearance Statement:

This report includes references to ISO standards through summarisation, referencing and reproduction of definitions only. While the material is not reproduced verbatim, ISO asserts copyright over its publications. For any further use or licensing queries, contact Standards Australia at: [copyright@standards.org.au](mailto:copyright@standards.org.au).

### Published By:

Age Check Certification Scheme  
Unit 321B Broadstone Mill, Broadstone Road  
Stockport, United Kingdom, SK5 7DL

**[www.accscheme.com](http://www.accscheme.com)**

ISBN 978-1-0681646-3-7





# Table of contents

## Introduction and Overview



<b>D.1</b>	Introduction to Part D: Age Estimation	6
<b>D.2</b>	Executive Summary	8
<b>D.3</b>	Who Participated in the Trial of Age Estimation Technology	13

## Context, Standards and Methodology



<b>D.4</b>	What is Age Estimation	16
<b>D.5</b>	Evaluation Approach for Age Estimation Systems	21

## Detailed Analysis of Age Estimation Findings



<b>D.6</b>	Age Estimation Can Be Done	28
<b>D.7</b>	No Substantial Technological Limitations to Age Estimation in Australia	31
<b>D.8</b>	Provider Practice Statements Reflected Technological Maturity and Responsible Policy Application	35
<b>D.9</b>	Importance of Buffer Thresholds and Configuration Management for Age Estimation	42
<b>D.10</b>	Range of Age Estimation Technologies and Performance Characteristics	47
<b>D.11</b>	Vibrant and Evolving Sector: Innovation and Sector-Specific Optimisation	56

<b>D.12</b>	Privacy by Design in Age Estimation: Data Handling, Minimisation and Innovation	61
<b>D.13</b>	Inclusion and Demographic Consistency in Age Estimation	68
<b>D.14</b>	Analysis of Acceptability Characteristics	76
<b>D.15</b>	Training Data Challenges in Age Estimation: Quality, Ethics and Risk Management	79
<b>D.16</b>	Emerging Theoretical Approaches in Age Estimation: Ergonomics, Physiology and Interaction Patterns	85
<b>D.17</b>	Security of Age Estimation Systems: Protecting Biometric Data and Meeting ISO Standards	91
<b>D.18</b>	Threat Mitigation in Age Estimation Systems: Spoofing, Deepfakes and Injection Attacks	97
<b>D.19</b>	Risk of Over-Retention in Age Estimation: Biometric Data and Analytical Residue	104






## Age Assurance Technology Trial

# PART D

## Introduction and Overview

I





## D.1 Introduction to Part D: Age Estimation

**D.1.1** Part D of the Age Assurance Technology Trial focuses specifically on age estimation – a method of determining an individual’s likely age or age range by analysing physical or behavioural characteristics using artificial intelligence or machine learning models. Unlike age verification, which relies on known and validated dates of birth, age estimation applies biometric or statistical techniques (such as facial analysis, voice modelling or motion pattern recognition) to predict age without the need for formal identity documents.

**D.1.2** This section evaluates how age estimation systems perform in the Australian context, including their technical feasibility, statistical accuracy, demographic fairness, privacy protection and resistance to manipulation. The Trial assesses alignment with relevant international standards – particularly ISO/IEC FDIS 27566-1<sup>1</sup>, which provides a functional and privacy framework for age assurance systems and IEEE 2089.1<sup>2</sup>, which outlines performance expectations for demographic consistency.

**D.1.3** In Part D of the report, we present our findings on age estimation systems, including their accuracy across age thresholds, performance across demographic groups, ability to operate in low-friction environments and effectiveness when used alongside other age assurance methods. This analysis supports the development of evidence-based standards, best practices and potential pathways for certification in Australia’s evolving digital safety landscape.

1. All references to ISO/IEC FDIS 27566-1 Standard throughout the suite of reports are referring to ISO/IEC FDIS 27566-1 – Information security, cybersecurity and privacy protection – Age assurance systems – Part 1: Framework.

2. All references to IEEE 2089.1 throughout the suite of reports are referring to IEEE 2089.1-2024 – IEEE Standard for Online Age Verification.





## D.2 Executive Summary

**D.2.1** Age estimation is a method of estimating a user's likely age based on observable characteristics such as facial features, voice or behavioural patterns. Unlike age verification, which relies on official identity documents, age estimation uses statistical models to estimate age without identifying the user. It is increasingly used to enforce age-based access controls in digital and in-person environments where document-based identity is unavailable, inappropriate or unnecessary.

**D.2.2** As part of the Trial, age estimation technologies were evaluated for their accuracy, security, inclusivity, usability and alignment with emerging international standards. The evaluation focused on high-readiness systems (Technology Readiness Level 7 or above) and included, as appropriate, structured technical testing, school-based trials, mystery shopper deployments, practice statement reviews and vendor interviews. Systems were assessed against key benchmarks such as ISO/IEC FDIS 27566-1 (age assurance requirements), IEEE 2089.1 (interoperability and assurance rules), ISO/IEC 27001 (information security) and ISO/IEC 25010 (software quality attributes).

**D.2.3** The Trial confirmed that age estimation can be deployed effectively in the Australian context. Many systems are already live in sectors such as social media, retail and content platforms. Most solutions demonstrated low-friction user experiences, fast estimation times (typically under 20 seconds) and high accuracy outside threshold "buffer zones" (e.g. 13+, 16+, 18+). Some systems achieved mean absolute errors (MAE) of approximately one year in controlled conditions and provided reliable threshold classification when estimated ages exceeded configured buffers. However, it is a fundamental misunderstanding of the capabilities of age estimation to test whether it can implement exactly a specific age-restriction without either accepting there will be a margin of error or applying a buffer age to reduce that margin to an acceptable level, acknowledging that false negatives will then be inevitable and alternative methods will be required to correct them.

**D.2.4** Vendors demonstrated strong alignment with privacy and security expectations, including:

- Temporary biometric processing with no image retention
- On-device or edge estimation architectures
- Secure capture pipelines and encrypted data transmission
- ISO/IEC 27001-certified information security practices
- Presentation attack detection (PAD) and emerging defences against injection and deepfake manipulation

**D.2.5** Inclusivity and demographic fairness were active areas of development. While systems generally performed well across diverse user groups, some showed reduced accuracy for older adults, non-Caucasian users and female-presenting individuals near policy thresholds. Underrepresentation of Indigenous populations in training data remains a challenge, particularly for First Nations Peoples, though vendors acknowledged these gaps and committed to remediation through fairness audits and dataset diversification.

**D.2.6** The Trial also explored innovative and emerging modalities, including gesture-based age classification, voice analysis and motion pattern detection. While these approaches are promising – especially for privacy-sensitive, ambient or child-first environments – they are at earlier readiness levels and require further validation before widespread deployment.

**D.2.7** Critically, while age estimation is highly effective for real-time, contextual age checks, it is not currently suitable for generating verifiable digital credentials (e.g. for use in digital wallets or holder services). Probabilistic age estimates lack the fixed, attestable properties required for credential-based identity systems. However, estimation can support layered or progressive assurance models and serve as a valuable pre-check or fallback when ID-based verification is unavailable or declined.

**D.2.8** Age estimation has emerged as a mature, secure and adaptable tool for enforcing age-based access in a wide range of digital and physical contexts. When configured responsibly and used in proportionate, risk-based scenarios, it supports inclusion, reduces reliance on identity documents and enhances user privacy. Its alignment with evolving international standards – combined with continuous innovation in model accuracy, fairness and spoof resistance – positions age estimation as a key component of modern, privacy-respecting age assurance infrastructure.





## Key Statistics from the Trial on Age Estimation




Figure D.2.1 Key Statistics from the Trial on Age Estimation

**D.2.9** Age estimation is a flexible, low-friction method of age assurance that offers a practical, privacy-preserving way to assess whether an individual is likely to meet a given age threshold – without the need for formal identity documents or declared dates of birth. It provides service providers, regulators and users with a rapid, non-intrusive tool for assessing age eligibility, particularly in lower-risk or high-volume environments such as social media, app stores and content access gateways.

**D.2.10** While it does not offer the binary certainty of verified date-of-birth checks, age estimation can achieve high levels of accuracy, especially when applied to clear thresholds (e.g. under/over 13, 16 or 18). When configured appropriately (e.g. with a buffer age) and supported by transparent confidence scoring, it allows systems to make probability-based age decisions that are contextually appropriate and scalable.

**D.2.11** When deployed using privacy-preserving, bias-aware and standards-aligned practices, age estimation strikes a meaningful balance between:

- Risk-appropriate compliance
- User autonomy and privacy
- Operational scalability and efficiency

**D.2.12** Its adaptability makes it particularly well suited to real-time use cases and it is increasingly being integrated into interoperable ecosystems – such as platforms exploring in-device estimation, in-app gating or signal-based assurance within digital identity frameworks. As confidence in its accuracy and fairness continues to grow, age estimation plays an important role in the broader ecosystem of age assurance methods.



## D.3 Who Participated in the Trial of Age Estimation Technology



Needemand





## Age Assurance Technology Trial



# PART D

## Context, Standards and Methodology





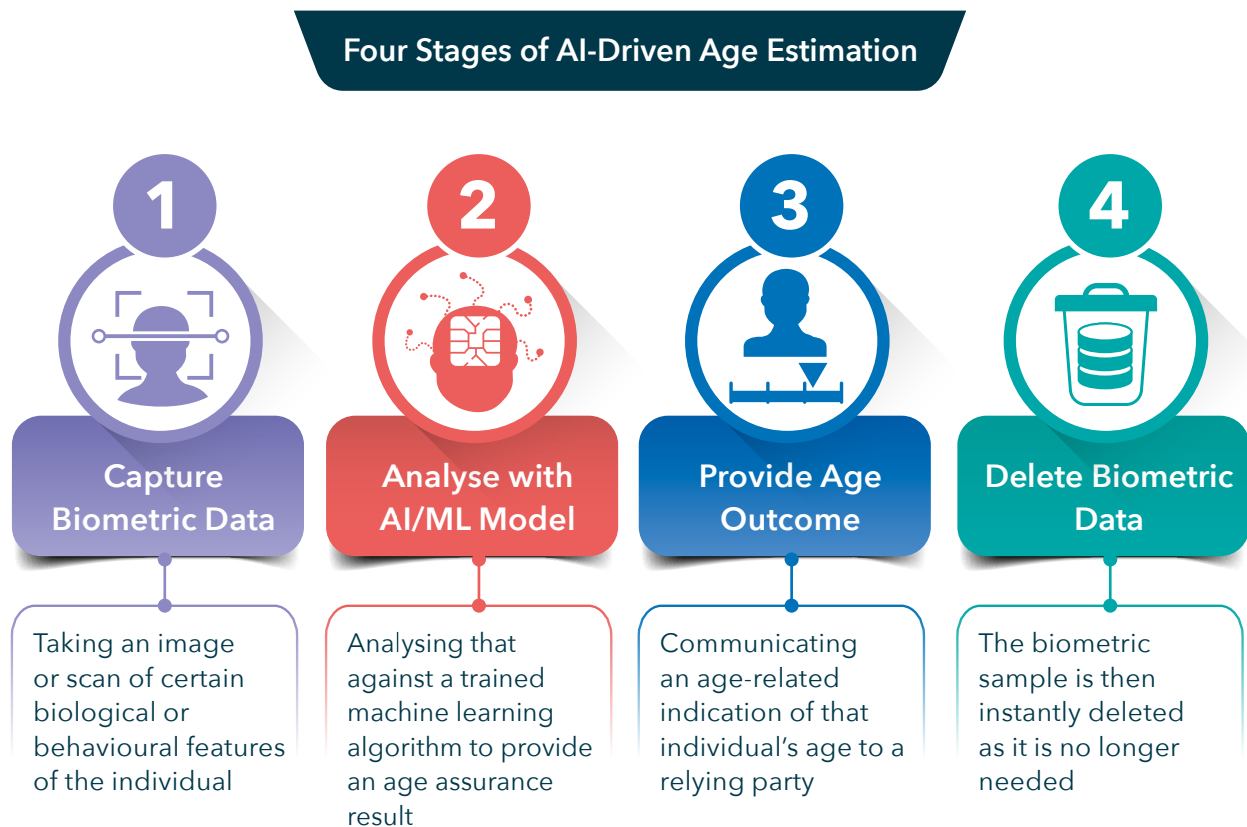
## D.4 What is Age Estimation

**D.4.1** Age estimation is an age assurance method based on analysis of biological or behavioural features of humans that vary with age.

**D.4.2** Age estimation uses artificial intelligence (AI) to deduce a person's likely age based on biometric or behavioural features. Unlike age verification, which checks a user's claimed age against official documents or records, age estimation does not require identity credentials or a known date of birth.

**D.4.3** Instead, it makes an informed analysis – usually about whether someone is likely over or under a certain age threshold (such as 13, 16 or 18) – based on observable traits. Age estimation is typically used where speed, privacy and ease of use are priorities and where exact age or identity is not legally required.

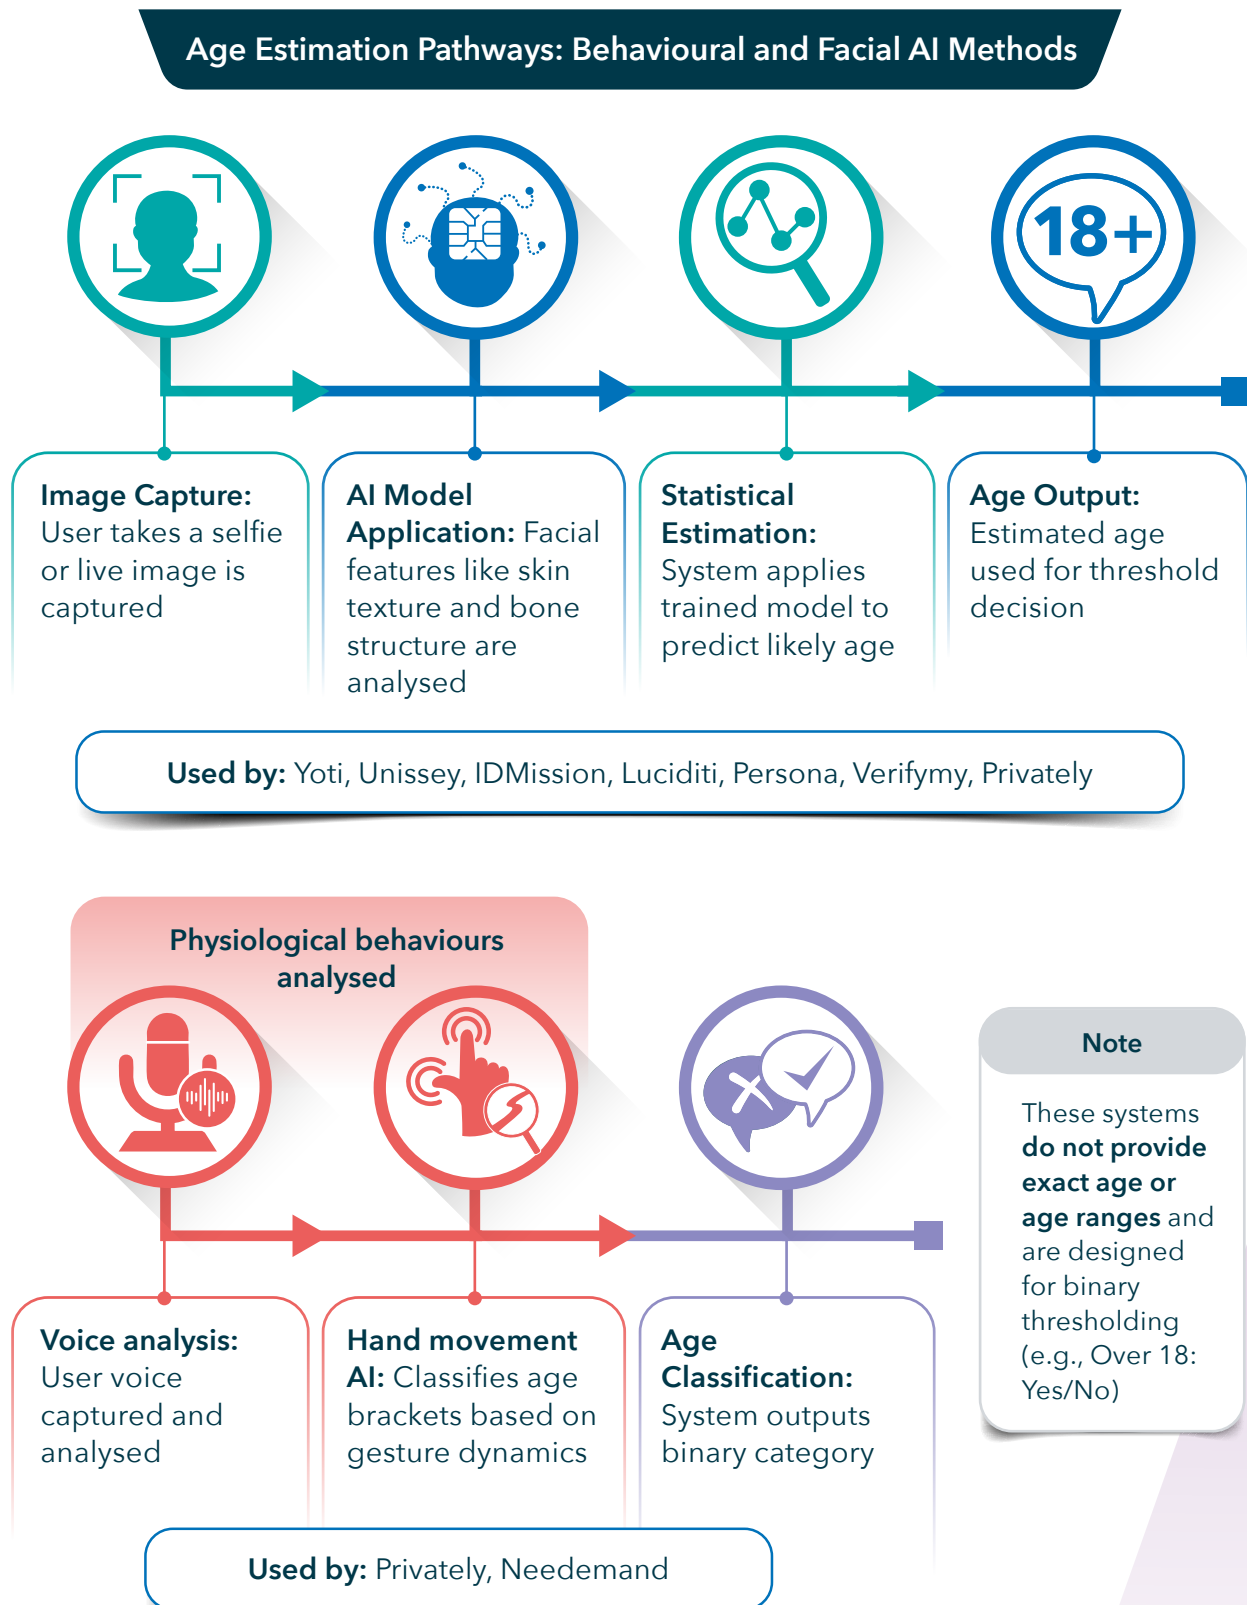
**D.4.4** It involves four stages:



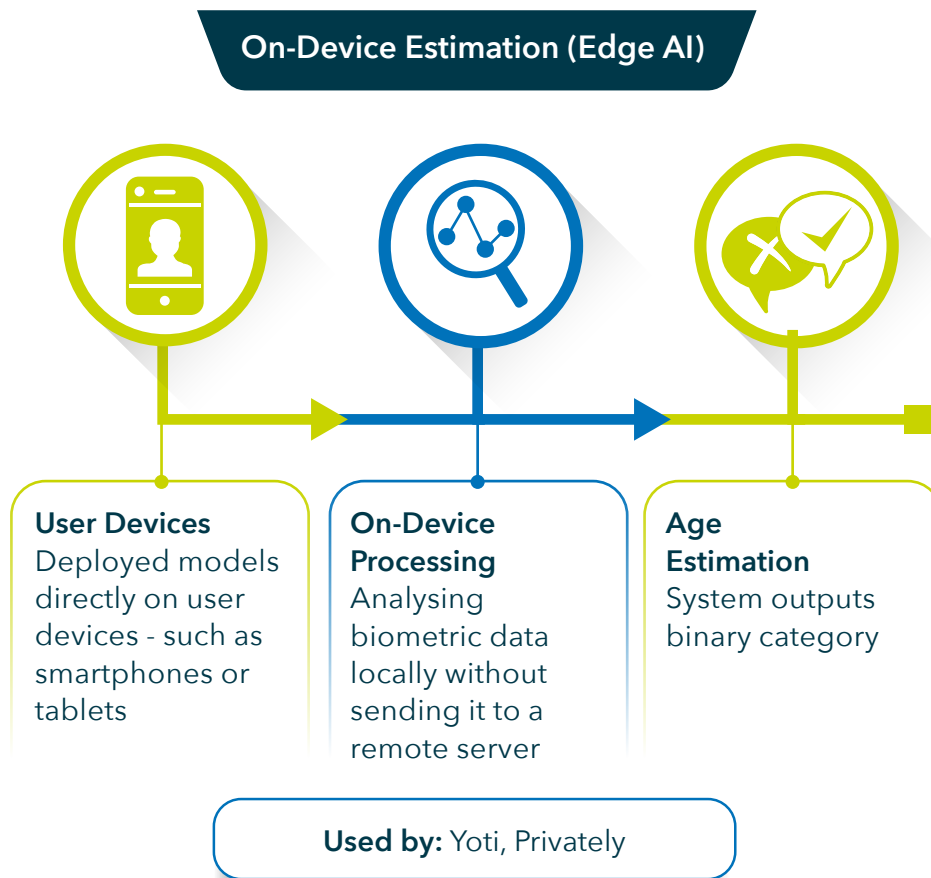
**Figure D.4.1** Four Stages of AI-Driven Age Estimation

## Examples of age estimation approaches in the Trial

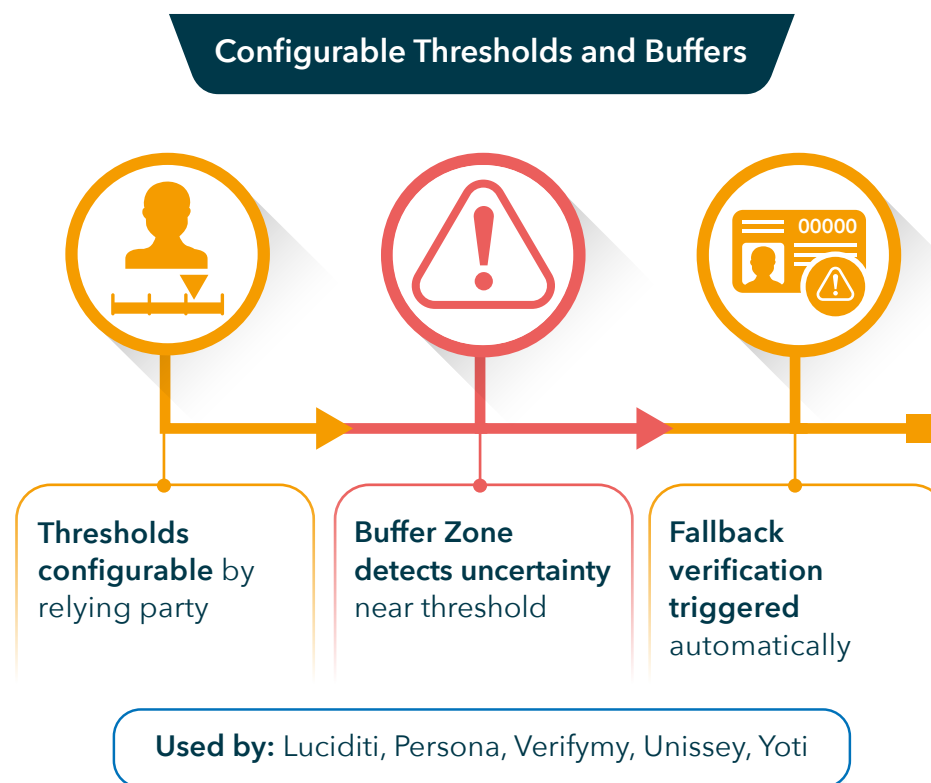
**D.4.5** During the Trial, the following types of age estimation approaches were observed across vendor systems:



**Figure D.4.2** Age Estimation Pathways: Behaviour and Facial AI Methods

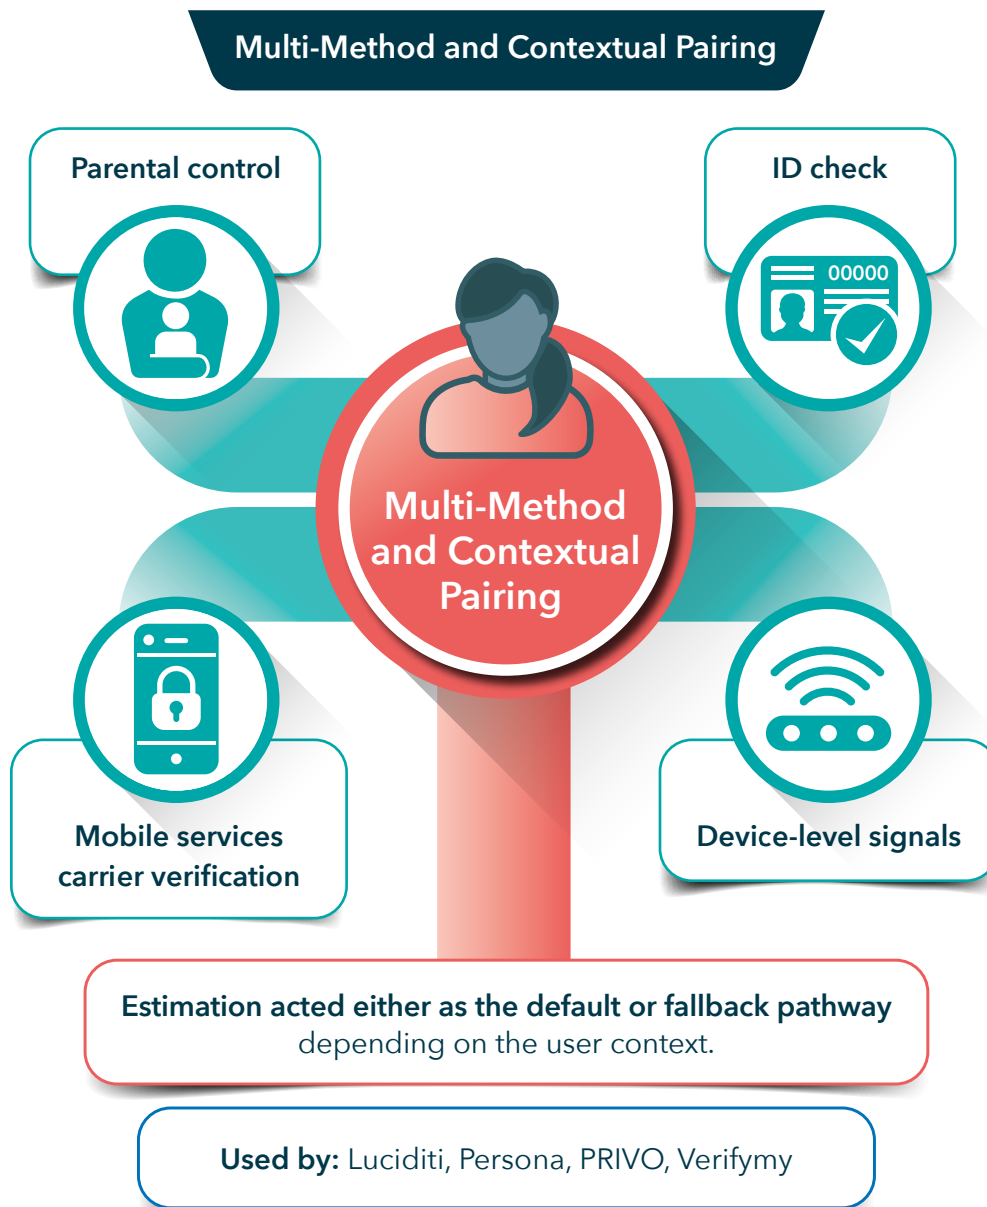


**Figure D.4.3** On-Device Estimation (Edge AI)



**Figure D.4.4** Configurable Thresholds and Buffers





**Figure D.4.5** Multi-Method and Contextual Pairing

**D.4.6** The above infographic illustrates how some vendors in the Trial adopted a multi-method and contextual pairing approach to age assurance. Rather than relying on a single method, these systems combined multiple inputs – such as parental controls, ID checks, device-level signals and mobile carrier verification – to improve accuracy and adaptability. Age estimation played a flexible role, acting either as the primary method or as a fallback mechanism depending on the user’s context and data availability. This layered strategy was particularly useful in real-world scenarios where users presented varying levels of data or identification. By combining signals, providers were able to offer more robust and context-sensitive age assurance solutions.

## What Age Estimation Is - and Is Not

### What is Age Estimation?



Uses physical or behavioural features (face, gestures) to estimate likely age



Provides a probability-based classification (e.g., "likely over 16")



Can be used anonymously without linking to identity

#### No ID required

Often involves no retention of personal data



#### Key Benefit

Age estimation offers a frictionless, privacy-conscious way to implement age-based access controls - especially in online environments where formal ID is unavailable or intrusive.

#### Key Limitation

Because it produces probabilistic results, age estimation may be unsuitable for high-stakes or legally sensitive contexts where verified, deterministic proof of age is required.

### Age Estimation is not the same as



**Verification of a known date of birth** (e.g., from a passport or ID document)



**Use of behavioural patterns**, transaction history or metadata (this is age inference)



**Identity recognition or account matching** (e.g., facial recognition)








**Parental assertion**, user self-declaration or login-based age gates

**Figure D.4.6** What Age Estimation Is - and Is Not

## D.5 Evaluation Approach for Age Estimation Systems

### | Core methodology

**D.5.1** Age estimation systems were tested in line with the overall Trial framework, drawing on:

International Standards	
 <b>ISO/IEC FDIS 27566-1</b>	Framework for age assurance systems
 <b>IEEE 2089.1</b>	Standard for online age verification
 <b>ISO/IEC 25010 and 25040</b>	Software quality models and evaluation processes
 <b>ISO/IEC 29119</b>	Software testing
 <b>ISO/IEC 30107</b>	Biometric presentation attack detection



## | International Standards for Age Estimation Methods

**D.5.2** ISO/IEC FDIS 27566-1, the international standard for age assurance systems, recognises age estimation as a distinct method of determining a person's likely age without using formal identity documents.

**D.5.3** Key provisions relevant to age estimation include:

ISO/IEC FDIS 27566-1	Criteria
<b>Input data and accuracy</b> (Clause 5.2 and 6.3)	<p>Age estimation should be based on relevant, high-quality input data.</p> <p>Systems should express the uncertainty of each estimate (e.g. confidence scores, error margins).</p> <p>Output should be used in a threshold form (e.g. "Is this person likely over 13?") rather than as a precise age.</p>
<b>Bias and demographic fairness</b> (Clause 6.3.3)	<p>Age estimation systems should be evaluated for performance across diverse demographic groups, including age, sex and ethnic background.</p> <p>Developers should mitigate risks of systematic under- or over-estimation for particular populations.</p>
<b>Privacy and minimal retention</b> (Clause 7.2 and 7.3)	<p>Personal data (such as facial images) should be minimised and not retained unless necessary for legitimate purposes.</p> <p>The standard encourages on-device or privacy-preserving models that don't transmit identifiable data to central servers.</p>

ISO/IEC FDIS 27566-1	Criteria
<b>Confidence and use appropriateness</b> (Clause 5.6)	<p>Age estimation should only be used in contexts where probabilistic results are acceptable – not where exact age is legally or contractually required.</p> <p>Relying parties must consider confidence thresholds and potential error rates when making age related eligibility decisions.</p>
<b>Explainability and transparency</b> (Clause 11.2)	<p>Systems should be transparent about how estimates are made and offer information about the model's performance characteristics.</p> <p>Users (and relying parties) should be informed if an estimated age is being used to make decisions about access or eligibility.</p>

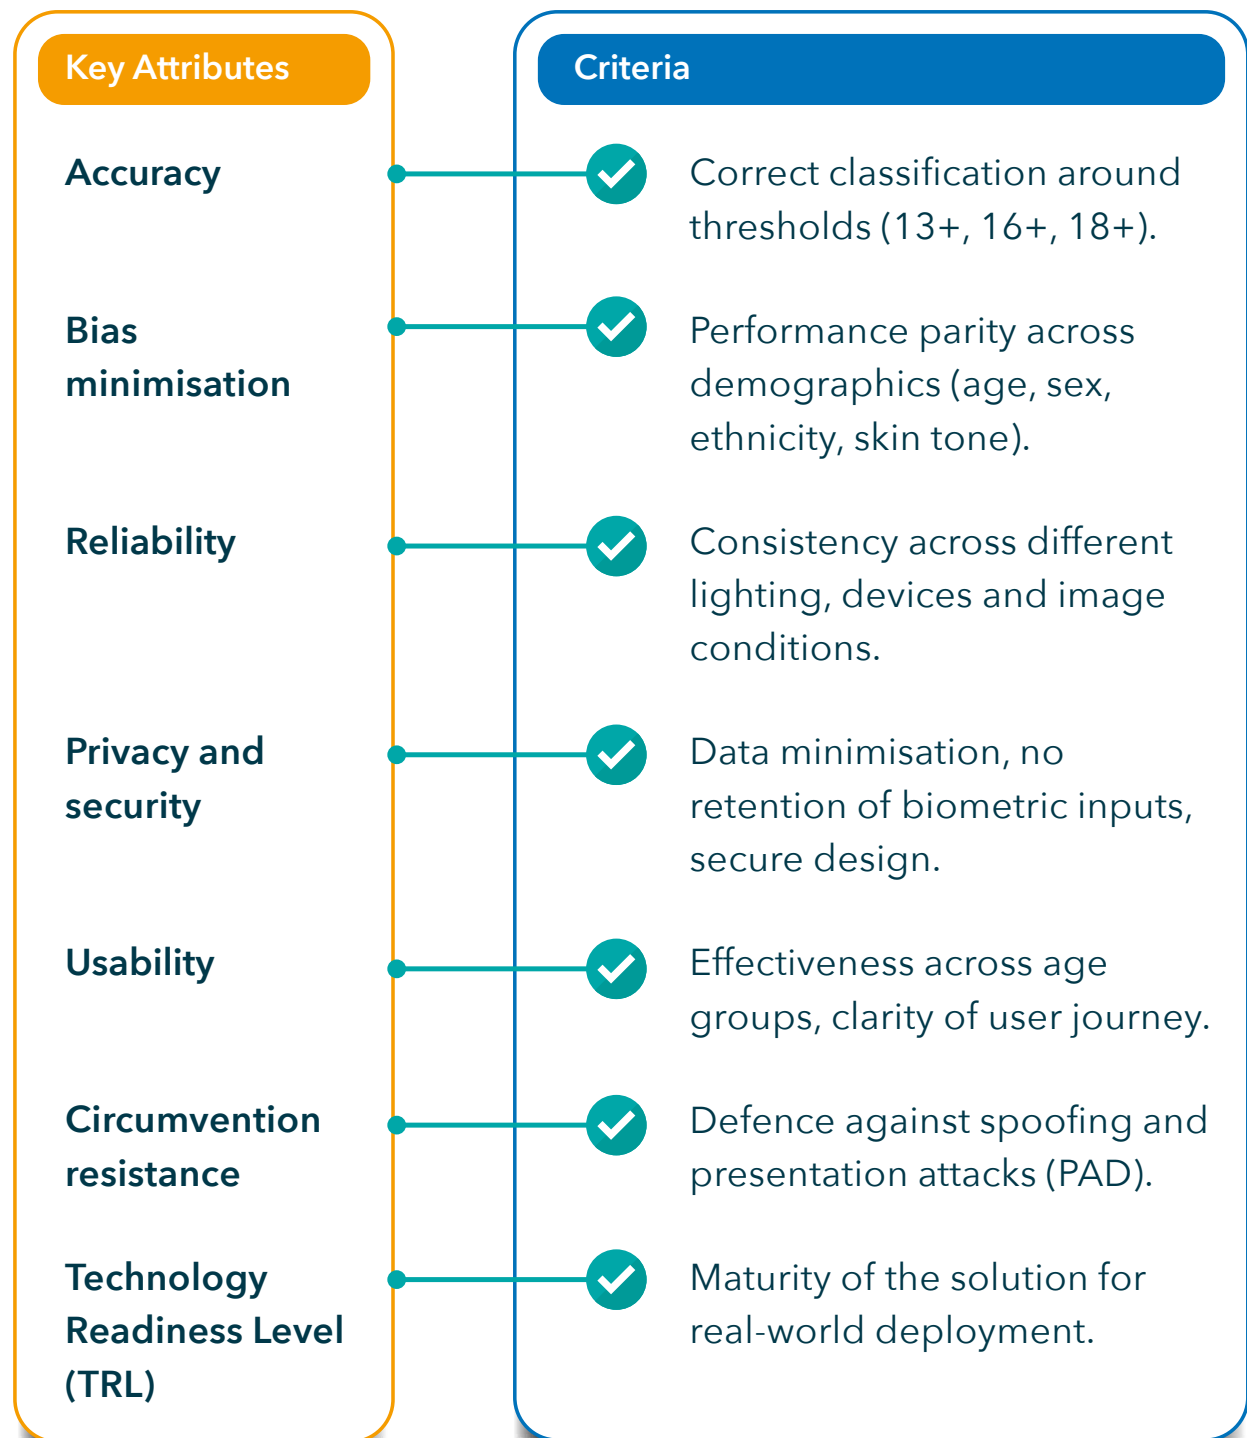
## | Summary

**D.5.4** ISO/IEC FDIS 27566-1 treats age estimation as a legitimate, low-friction method of age assurance – provided it is used transparently, evaluated for fairness, applied in appropriate contexts and designed to minimise data risk. It does not recommend age estimation as a standalone method for high-risk or compliance-heavy scenarios, but rather as a supportive tool in age-restricted environments, especially where speed and convenience are critical.

## Evaluation criteria

**D.5.5** Age estimation systems were assessed against key attributes including:

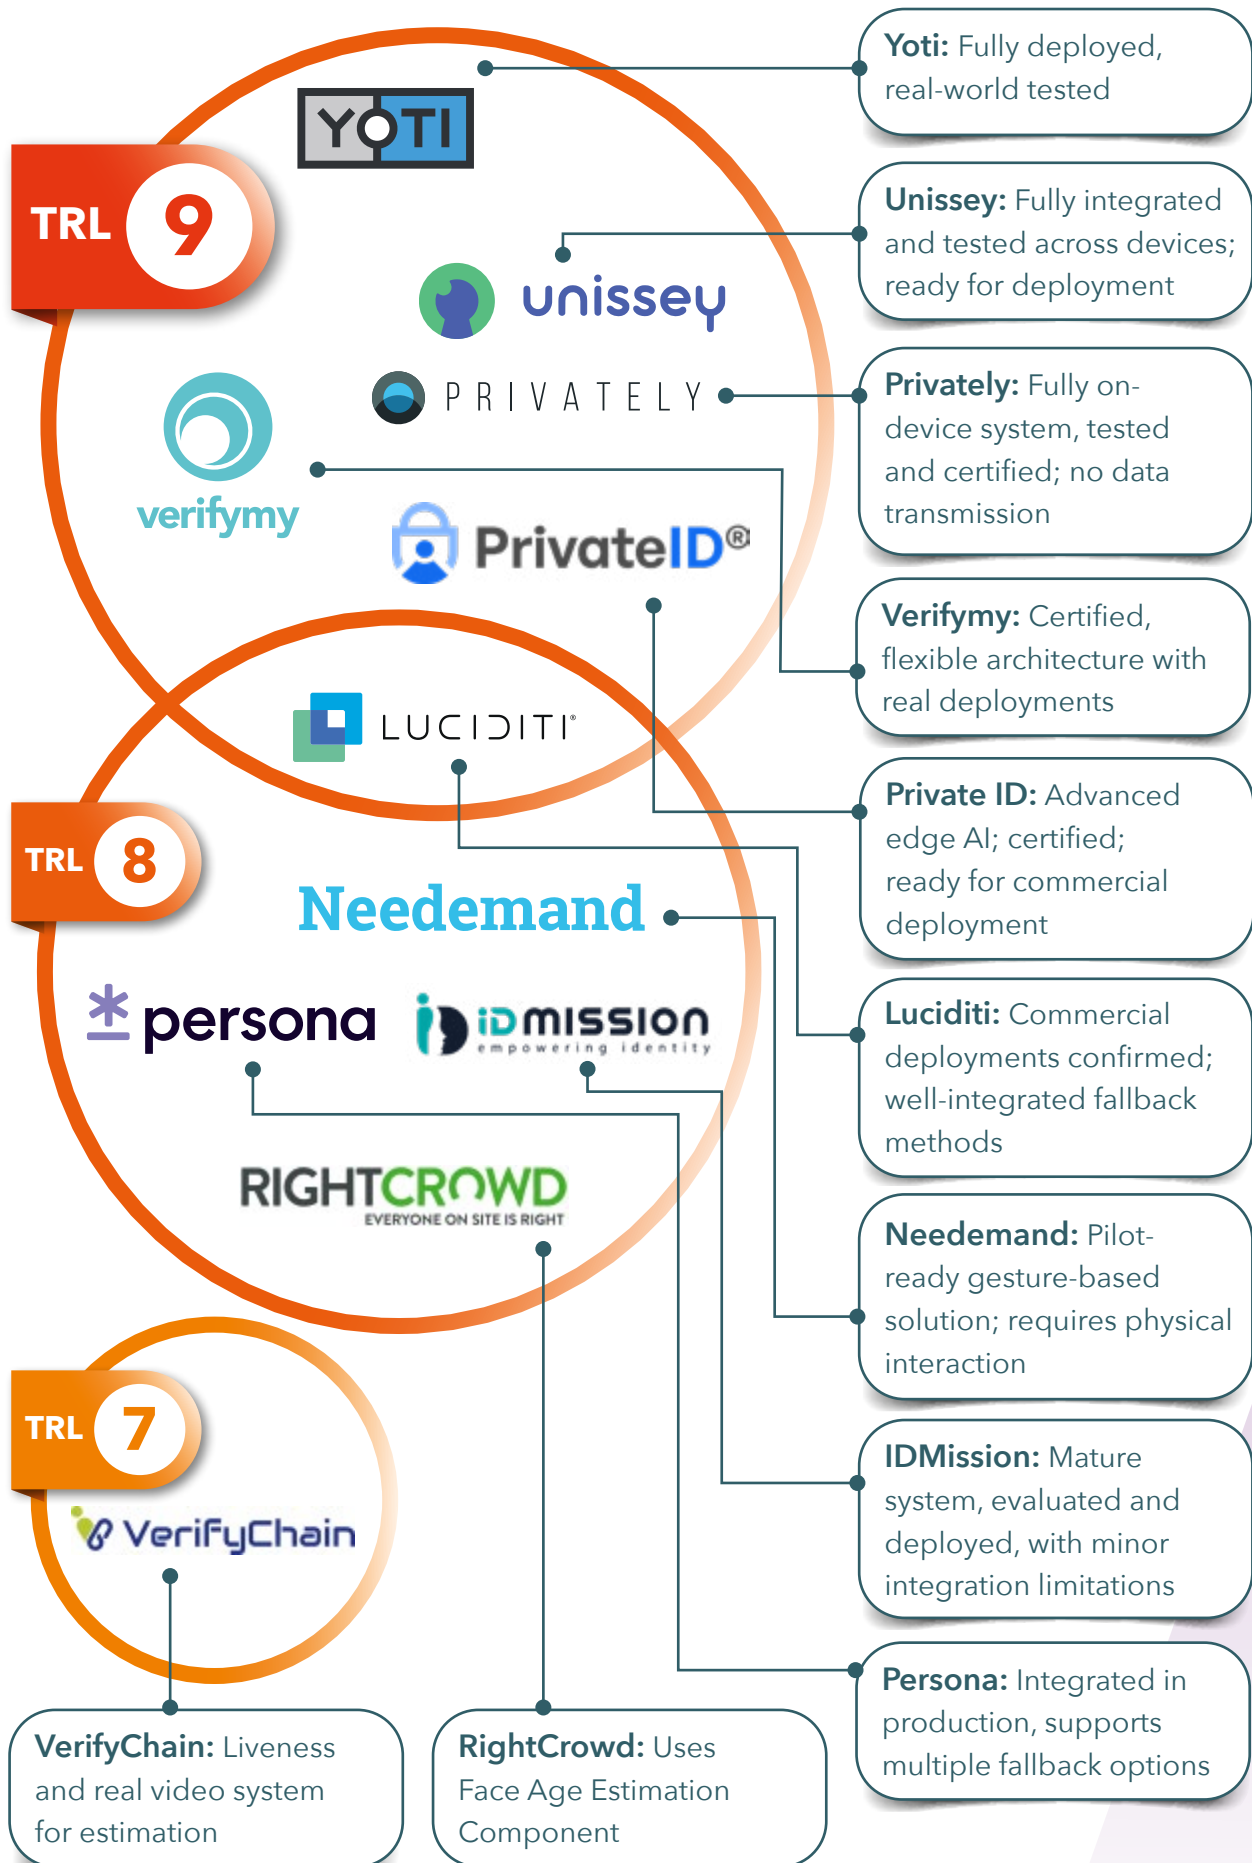
### Key Evaluation Criteria for Age Estimation Systems



**Figure D.5.1** Key Evaluation Criteria for Age Estimation Systems



## Technology readiness assessment for age estimation systems





## Age Assurance Technology Trial



# PART D

## Detailed Analysis of Age Estimation Findings



## D.6 Age Estimation Can Be Done

### | Summary finding

**D.6.1** Age estimation can be implemented in Australia and is already being used or trialled in real-world services. Advances in AI-based facial analysis – along with emerging non-facial methods such as gesture-based classification – make it a practical tool for scalable, low-friction age checks across online and offline contexts. Privacy-preserving techniques (e.g. on-device, zero data retention) support responsible deployment in line with international standards.

### | Detailed analysis

**D.6.2** Age estimation is a practical and scalable method of age assurance, particularly effective for threshold-based decisions (e.g. determining whether a user is likely over 13, 16 or 18). It is increasingly being integrated into consumer-facing applications, both in live deployments and in controlled field trials.

**D.6.3** While most systems tested in the Trial were based on facial image analysis, at least one used gesture-based movement classification. These approaches offer effective alternatives where users may not wish to share images of their face. Trial participants demonstrated maturity in system design and operational readiness, with deployments seen in digital services and emerging offline use cases (e.g. kiosks and in-store systems).

**D.6.4** This aligns with Clause 4.3.3 of ISO/IEC FDIS 27566-1, which recognises age estimation as a valid assurance method when based on high-quality, relevant input data – such as facial images or behavioural signals – and when used transparently and appropriately in the context of risk and confidence.

## | Functional characteristics (ISO/IEC FDIS 27566-1 reference)

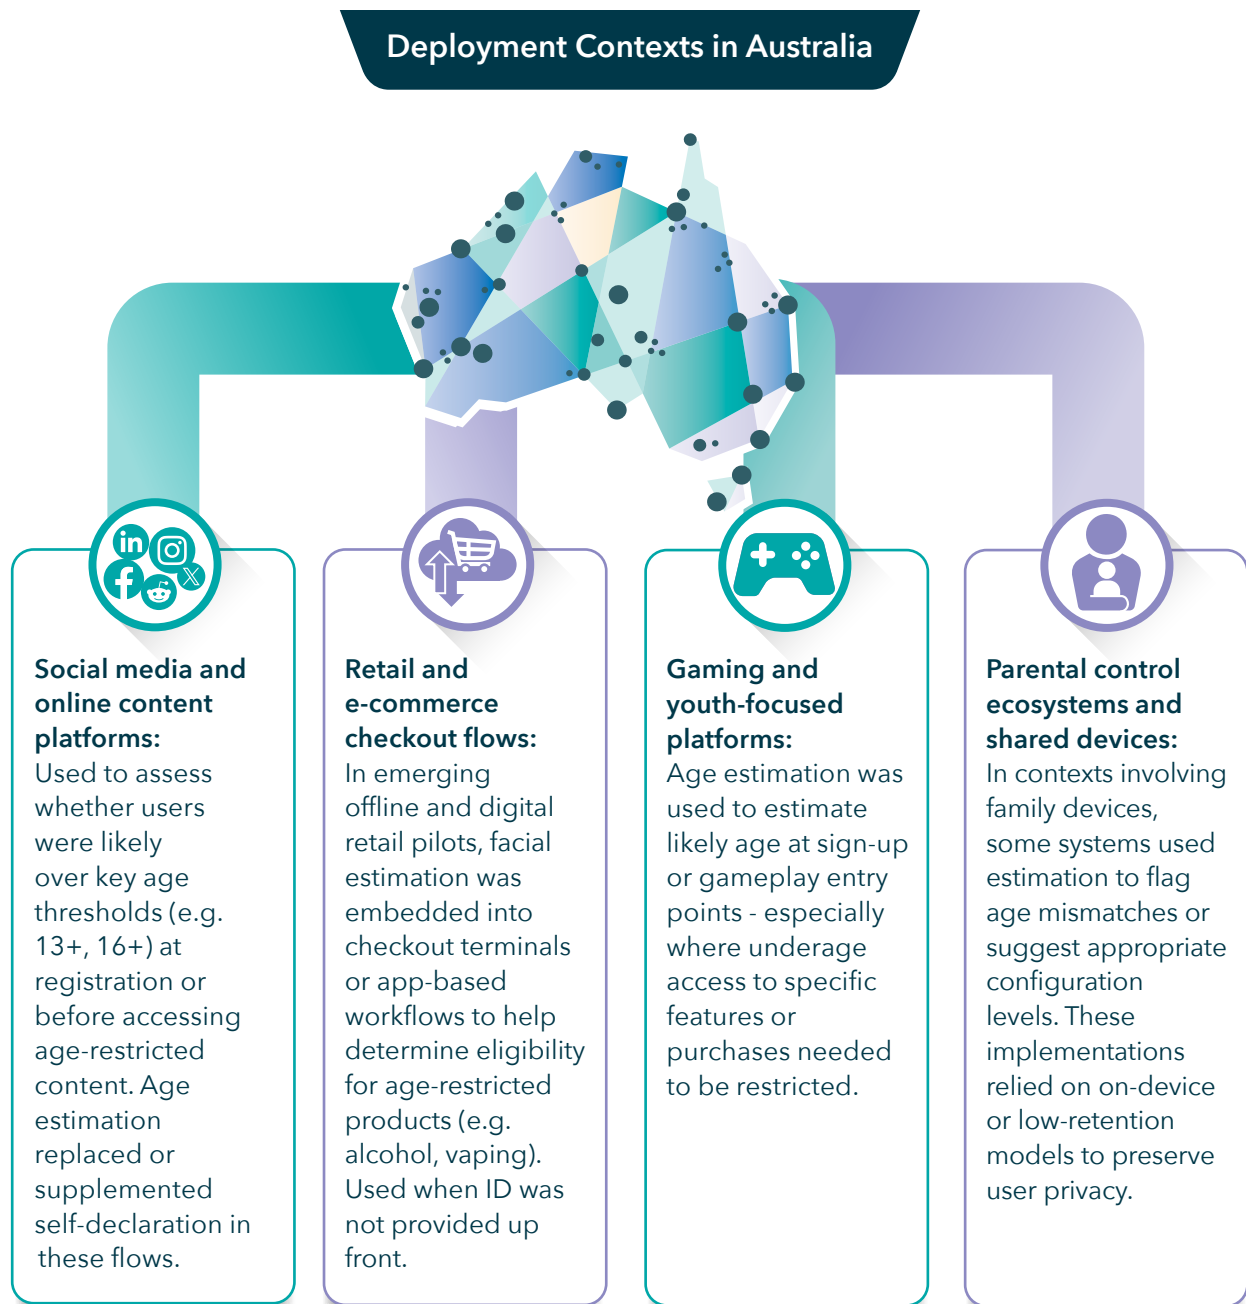
**D.6.5** Age estimation technologies exhibit the core functional characteristics described in ISO/IEC FDIS 27566-1:

ISO/IEC FDIS 27566-1	Criteria
<b>Friction Minimisation</b> (Clause 6.6)	Age estimation allows services to check age with minimal interruption to the user experience. It does not require account creation, document upload or third-party identity verification.
<b>Selective Disclosure</b> (Clause 7.3.2)	The output of estimation systems typically involves binary or thresholded results (e.g. "Likely over 16: Yes/No") and avoids revealing full or inferred ages unnecessarily.
<b>Appropriateness to Risk</b> (Clause 4.3.3)	Age estimation is particularly suitable for moderate-risk contexts where identity is not required but age-restricted access should be enforced – such as social platforms, online games or media platforms.
<b>Adaptability and Context Awareness</b> (Clause 5.6)	Providers demonstrated the ability to tailor estimation thresholds, confidence levels and system responses to the specific operational or regulatory environment of the relying party.



## Deployment contexts in Australia

**D.6.6** During the Trial, age estimation systems were observed in live or pilot deployments across several operational contexts. These deployments focused on threshold-based age checks using facial analysis or gesture-based methods, with minimal user friction and no requirement for identity disclosure:



**Figure D.6.1** Deployment Contexts in Australia

## D.7 No Substantial Technological Limitations to Age Estimation in Australia

### | Summary finding

**D.7.1** Age estimation technologies in Australia face no substantial technological limitations to implementation. Providers and relying parties have shown thoughtful, responsible deployment practices, with strong attention to privacy, data protection and security. The alignment of technology and policy enables effective, privacy-preserving age estimation that supports age-based requirements without requiring full identity disclosure.

### | What is meant by no technological limitations

**D.7.2** The evaluation found that age estimation technologies are technically and operationally feasible in Australia, with no substantial technological limitations to their implementation. The participating providers demonstrated a high level of system maturity, compatibility with mainstream consumer devices and adaptability to diverse use cases. These systems are well-positioned to meet the requirements of relying parties who seek non-intrusive, privacy-preserving mechanisms to determine whether a user likely meets an age threshold.



## Vendor Case Study



PRIVATELY

Website

[privately.eu](https://privately.eu)

Privately offers a lightweight, on-device facial age estimation system designed for privacy-by-design deployments, especially in youth-focused contexts such as education and family settings.

## Three Key Facts

1

Age estimation is performed entirely on the user's device, with no images transmitted or stored externally.

2

Estimation times were typically under 5 seconds, even on entry-level Android devices.

3

Functional tests confirmed successful operation across a range of conditions (lighting, device types and demographic inputs).

## Strengths

Strong privacy by design: zero biometric retention, on-device only architecture, no cloud contact - aligns with ISO/IEC FDIS 27566-1 expectations.

Flexible integration options with local fallback mechanisms and threshold tuning make Privately adaptable across different use cases and risk levels.

## Practice Statement

[ageassurance.com.au/v/prv/#PS](https://ageassurance.com.au/v/prv/#PS)

## Privacy Policy

[ageassurance.com.au/v/prv/#PP](https://ageassurance.com.au/v/prv/#PP)

## Technology Trial Test Report

[ageassurance.com.au/v/prv/#TR](https://ageassurance.com.au/v/prv/#TR)

## Technology Trial Interview

[ageassurance.com.au/v/prv/#VI](https://ageassurance.com.au/v/prv/#VI)

## Summary of Results

Privately's edge-based solution is a strong example of a privacy-preserving, operationally ready technology that eliminates the need for identity documents or server infrastructure - demonstrating no significant technical limitations to deployment in Australian contexts.

## | Rapid evolution of AI capabilities

**D.7.3** Age estimation relies heavily on machine learning and artificial intelligence (AI) – particularly in the form of facial analysis models trained to infer age based on patterns in human facial features. One of the most notable developments observed during the Trial was the rapid improvement in estimation accuracy and model generalisability, reflecting ongoing advances in:


- Training data quality and diversity.
- Neural network optimisation, including use of smaller, faster models deployable on mobile devices.
- Bias reduction techniques for fairer performance across age, gender and ethnicity.
- On-device estimation and federated learning approaches that preserve user privacy.

**D.7.4** Providers also demonstrated that they were regularly improving their models by using feedback and updated versions to improve accuracy and reduce errors – particularly at critical policy thresholds (e.g. age 13 or 18). These improvements are increasingly measurable in terms of mean absolute error (MAE), false positive/negative rates and confidence scoring tuned to context.



## | Yoti's published data: Average MAE across all ages

**D.7.5** Yoti's published, verifiable data, showing the average MAE across all age groups (typically 6–70 years) from 2019 to 2025:



**Figure D.7.1** Yoti MAE Reduction Over Time (2019-2025)

**D.7.6** These verified data points substantiate a clear downward trend in Yoti's MAE – from 3.22 yrs in 2019 to 2.4 yrs in 2025 and now approaching ~1 yr at key age thresholds such as around the age 18+ gateway in 2025 – demonstrating definitive, incremental improvements grounded in published findings.

## | No substantial limitations to deployment

**D.7.7** Throughout the Trial, age estimation technologies were shown to be compatible with Australian infrastructure, devices and user expectations. Systems operated effectively:

- Across desktop and mobile environments.
- Using standard cameras and device sensors.
- In diverse real-world lighting and image conditions.



## D.8 Provider Practice Statements Reflected Technological Maturity and Responsible Policy Application

**D.8.1** Our analysis of practice statements provided by age estimation technology providers with a system of TRL 7 or above accurately and fairly reflected the technological capabilities of their products, processes or services. The statements demonstrated how age analysis was being undertaken and how the provider's internal policy decisions were being applied to the output of the systems.

### | Accurate reflection of technological capabilities

**D.8.2** The content of these practice statements accurately reflected the observable capabilities of the systems under evaluation. Providers were generally clear and consistent in describing:

- The underlying method of age analysis (e.g. facial estimation via convolutional neural networks).
- The sources of training data, including how diversity and representativeness were addressed.
- Estimation processes, such as threshold application, confidence scoring and error handling.
- Measures taken to prevent spoofing, tampering or adversarial input manipulation.

**D.8.3** These declarations aligned closely with real-world system behaviour observed in the evaluation, including classification performance, responsiveness and consistency across a range of demographic groups and operating conditions.

Provider	Practice Statement Summary	Privacy Policy Summary	Consistent?
<b>Yoti</b>	Uses facial age estimation; deletes images after inference; explains methods transparently.	No biometric storage; GD-PR-aligned; anonymous age checks possible.	✓
<b>Privately</b>	Edge-AI based estimation; no data leaves device; privacy-first design.	No data collection or storage; operates fully on-device.	✓
<b>Luciditi</b>	Estimation with fallback to ID/banking; deletes biometric input post-processing.	Images deleted after use; identity not retained; UK/EU data compliance.	✓
<b>IDMission</b>	AI-based facial analysis with encrypted server-side processing.	No data stored on user device; SOC 2 and GDPR compliant.	✓
<b>Verifymy</b>	Flexible age estimation and verification; modular workflow options.	Data retention configurable; compliant with data protection norms.	✓
<b>GBG</b>	Combines facial estimation with ID/document checks; flexible workflows for verification fallback.	Offers configurable retention; aligns with GDPR and Australian Privacy Principles.	✓
<b>Needemand</b>	Gesture-based classification; does not use biometric or identity data; fast, lightweight process.	No PII or metadata stored; gesture tokens are non-identifiable.	✓
<b>Persona</b>	Uses facial estimation with optional ID fallback; governed model updates for bias reduction.	Minimal retention; supports GDPR-compliant data control and opt-out.	✓
<b>VerifyChain</b>	Blockchain-based validation of age tokens; decentralized model enhances auditability.	User-consent driven logging; verifiable yet anonymised token exchange.	✓
<b>Unissey</b>	Neural network-based facial estimation; real-time inference with no ID requirement.	Biometric data processed transiently and deleted post-use; privacy compliant.	✓

## | Application of internal policy to system outputs

**D.8.4** Importantly, the statements also demonstrated how internal policy decisions shaped the use of age estimation outputs – highlighting the maturity of governance practices within TRL 7+ providers.

Examples included:

- Decisions about what confidence thresholds should be applied for different use cases (e.g. higher certainty for 18+ content, lower friction for 13+ access).
- Whether estimation results should be used standalone or combined with other inputs (e.g. device signals or user-declared age).
- How to handle uncertain or borderline results, such as invoking fallback mechanisms<sup>3</sup>, providing advisory warnings or referring users to human review.
- Commitments to not retaining facial images or biometric templates, beyond the time needed for estimation processes to be completed

**D.8.5** These policy applications demonstrate thoughtful implementation aligned with ISO/IEC FDIS 27566-1, which encourages systems to be both functionally effective and contextually appropriate (see Clause 5.1: Appropriateness to Risk; Clause 6.3: Expression of Confidence and Uncertainty).

3. A more comprehensive look at this can be found in Part F, the report on Successive Validation.

Provider	Confidence Thresholds	Standalone or combined	Handling Uncertainty	No Retention Commitment
<b>Yoti</b>	Tuneable by use case (13+, 13-17, 16+, 18+)	Used standalone or with user-declared age	Uses buffers; can refer to verification	Images deleted after estimation
<b>Privately</b>	Adjustable threshold logic at edge	Edge-only, can integrate device signals	Local fallback triggers; optional parental guidance	No storage or server contact
<b>Luciditi</b>	Distinct thresholds for each gate (13/16/18)	Combined with ID or banking in fallback	Triggers ID verification for borderline cases	Deletes biometric inputs post-use
<b>IDMission</b>	Adjustable threshold values by client	Typically used standalone	Refers to manual review or system message	No biometric data stored on device
<b>Verifymy</b>	Configurable per policy gate	Often combined with ID or mobile carrier	May escalate to verification	Supports zero-retention mode
<b>GBG</b>	Customisable by service context	Usually part of verification bundle	Prompts further checks or warnings	Optional minimisation policy
<b>Needemand</b>	Binary threshold for pass/fail logic	Standalone gesture-based method	Advisory message or retry prompt	No biometric or gesture data retained
<b>Persona</b>	Client defines based on context	Supports hybrid workflows	Step-up verification or advisory fallback	Minimal retention, user opt-out supported
<b>VerifyChain</b>	Smart contracts enforce rules	Combined with third-party attestations	Blockchain referral/warning flow	Commitment to anonymisation and non-retention
<b>Unissey</b>	Tuneable thresholds per age gate	Primarily standalone	Refers to human review or verification	Transient biometric processing only

## Vendor Case Study



Website

[privateid.com](https://privateid.com)

PrivateID uses signal-based facial AI estimation within parent-managed workflows. Optimised for under-13 use; COPPA-aligned design with integrated identity and age checks.

## Three Key Facts

1

Local-only processing with zero data retention.

2

Privacy-by-design for under-13 and COPPA contexts.

3

Multimodal signal fusion for higher adaptability.

Practice Statement

[ageassurance.com.au/v/pid/#PS](https://ageassurance.com.au/v/pid/#PS)

Technology Trial Test Report

[ageassurance.com.au/v/pid/#TR](https://ageassurance.com.au/v/pid/#TR)

Privacy Policy

[ageassurance.com.au/v/pid/#PP](https://ageassurance.com.au/v/pid/#PP)

Technology Trial Interview

[ageassurance.com.au/v/pid/#VI](https://ageassurance.com.au/v/pid/#VI)

## Summary of Results

No exact MAE disclosed; but verified homomorphic protocols and strong privacy alignment.



## | What opportunities are there for improvement of age estimation practices

**D.8.6** While the submitted statements from providers were generally strong and demonstrated well-developed, standards-aware systems, the evaluation identified several key areas for ongoing improvement in age estimation practice.

### *Configuration management and integration guidance*

**D.8.7** Providers often lacked detailed guidance on how to configure or integrate estimation tools into relying party systems. This includes:

- Distinguishing default versus configurable thresholds.
- Clarifying fallback behaviour (e.g. when to escalate to verification).
- Preventing misuse or misinterpretation of confidence scores.
- Supporting low-tech relying parties (e.g. SMEs) with integration toolkits.

**D.8.8 Risk:** Inappropriate application of buffer zones or thresholds can reduce fairness and accuracy, especially if end-users interpret probabilistic outputs as deterministic.

### *Buffer zone calibration*

**D.8.9** While many providers allowed for buffer zones around key thresholds (e.g. 18+), few offered clear guidance on:

- The optimal range of buffers for different risk levels.
- When fallback to ID or advisory messages is appropriate.
- How to tune thresholds to balance inclusion vs. over blocking.

**D.8.10 Clause Alignment:** ISO/IEC FDIS 27566-1 Clause 6.3 (Confidence and Uncertainty Expression).

## *Demographic consistency and fairness*

**D.8.11** Though practice statements often referenced bias testing, demographic performance disparities were still observed – particularly:

- Increased false positives for users with darker skin tones or aged 16–20.
- Variability in model output by gender presentation or lighting conditions.
- Limited use of fairness benchmarks or mitigation strategies.

**D.8.12 Opportunity:** Providers should expand fairness testing datasets and disclose disaggregated performance metrics across demographic subgroups.

## *Explainability and user transparency*

**D.8.13** Only a few providers offered meaningful transparency to users or relying parties regarding how age estimates were derived, especially:

- What features were considered (e.g. facial landmarks, behavioural cues).
- How confidence scores should be interpreted.
- Whether the system “learns” or adapts over time.

**D.8.14 Clause Alignment:** ISO/IEC FDIS 27566-1 Clause 9.3 and 10.2 (Transparency and User Communication).

## D.9 Importance of Buffer Thresholds and Configuration Management for Age Estimation

**D.9.1** Opportunities were identified to improve how systems are configured by relying parties, particularly regarding buffer age thresholds. For instance, a party might discard results near a restriction point (e.g. 18–21) and instead use other methods. Since the risk of false positives tends to zero outside this buffer zone, this provides confidence in eligibility decisions.

**D.9.2** Age estimation systems – especially those enforcing thresholds like “over 13,” “over 16,” or “over 18” – use probabilistic models that estimate likely age but cannot offer certainty at the individual level. This makes threshold configuration and clarity essential when deployed by relying parties.

### | What Is a buffer threshold?

**D.9.3** A buffer threshold is a configurable margin around the age restriction that accounts for the inherent uncertainty in age estimation.

**D.9.4** For example:

- A service requiring users to be 18+ may accept those estimated at 21+, reject those estimated at under 15, and apply extra checks for users estimated between 15–21.

**D.9.5** This “grey zone” helps reduce false positives (e.g. underage users gaining access) while ensuring a smooth experience for clearly eligible users.



Table showing Average Buffer Threshold Performance by Age Gate (Across All Vendors):

Age Gate	False Positives Tend to Zero (Above Threshold)	False Negatives Tend to Zero (Below Threshold)
13+	Age 16	Age 11-12
16+	Age 19-20	Age 12-13
18+	Age 21-22	Age 13-14

### *Interpretation and Rationale*

- **False Positives Tend to Zero:** These are cases where underage users are incorrectly passed. Across all gates, the rate of false positives sharply drops as the estimated age moves upward and away from the threshold. By ages 20-22 (for the 18+ gate), false positive rates approach zero.
- **False Negatives Tend to Zero:** These are cases where eligible users are incorrectly blocked. As the actual age falls below the threshold by 2-3 years, the system almost never misclassifies those users as eligible. For example, 11-12-year-olds are reliably rejected by the 13+ gate.

### *Why This Matters*

These patterns reinforce the concept of a buffer threshold:

- Around 2-3 years on either side of the age gate lies a "grey zone" where system uncertainty is higher.
- Outside this zone, system accuracy stabilises and error rates become negligible - supporting high-confidence decisions for accept or deny outcomes.
- In statistical terms, the true rate never actually gets to 'zero', hence the term 'tends to zero'.

## | Why buffer thresholds matter

**D.9.6** Age estimation models tend to be most uncertain at or near the threshold (e.g. distinguishing between a 17.8-year-old and an 18.2-year-old), while accuracy improves as you move away from it. As such, the likelihood of a false decision tends to zero outside a reasonable buffer.

**D.9.7** This means:

- Users confidently above the threshold (e.g. estimated 21+ when the cutoff is 18) can be passed with high assurance.
- Users estimated to be near the threshold can be diverted to supplementary assurance methods, such as ID-based verification or parental consent.
- Users confidently below the threshold can be denied access outright.

**D.9.8** This approach aligns with ISO/IEC FDIS 27566-1 Clause 6.4, which calls for systems to:

*“Express uncertainty associated with the input data and associated thresholds and adapt outputs to the confidence level of the system.”*

**D.9.9** The implementation of buffer thresholds also plays a crucial role in enabling successive validation, a layered approach to age assurance where one method (e.g. age estimation) is followed by a second, more definitive method (e.g. document-based verification) when uncertainty is high. In practice, users whose estimated age falls within a designated buffer zone can be seamlessly referred to a secondary verification process, reducing the risk of incorrect access decisions without introducing unnecessary friction for the wider user base. By integrating buffer thresholds as part of a successive validation model, relying parties can apply risk-based decision-making while respecting user privacy and ensuring regulatory compliance.



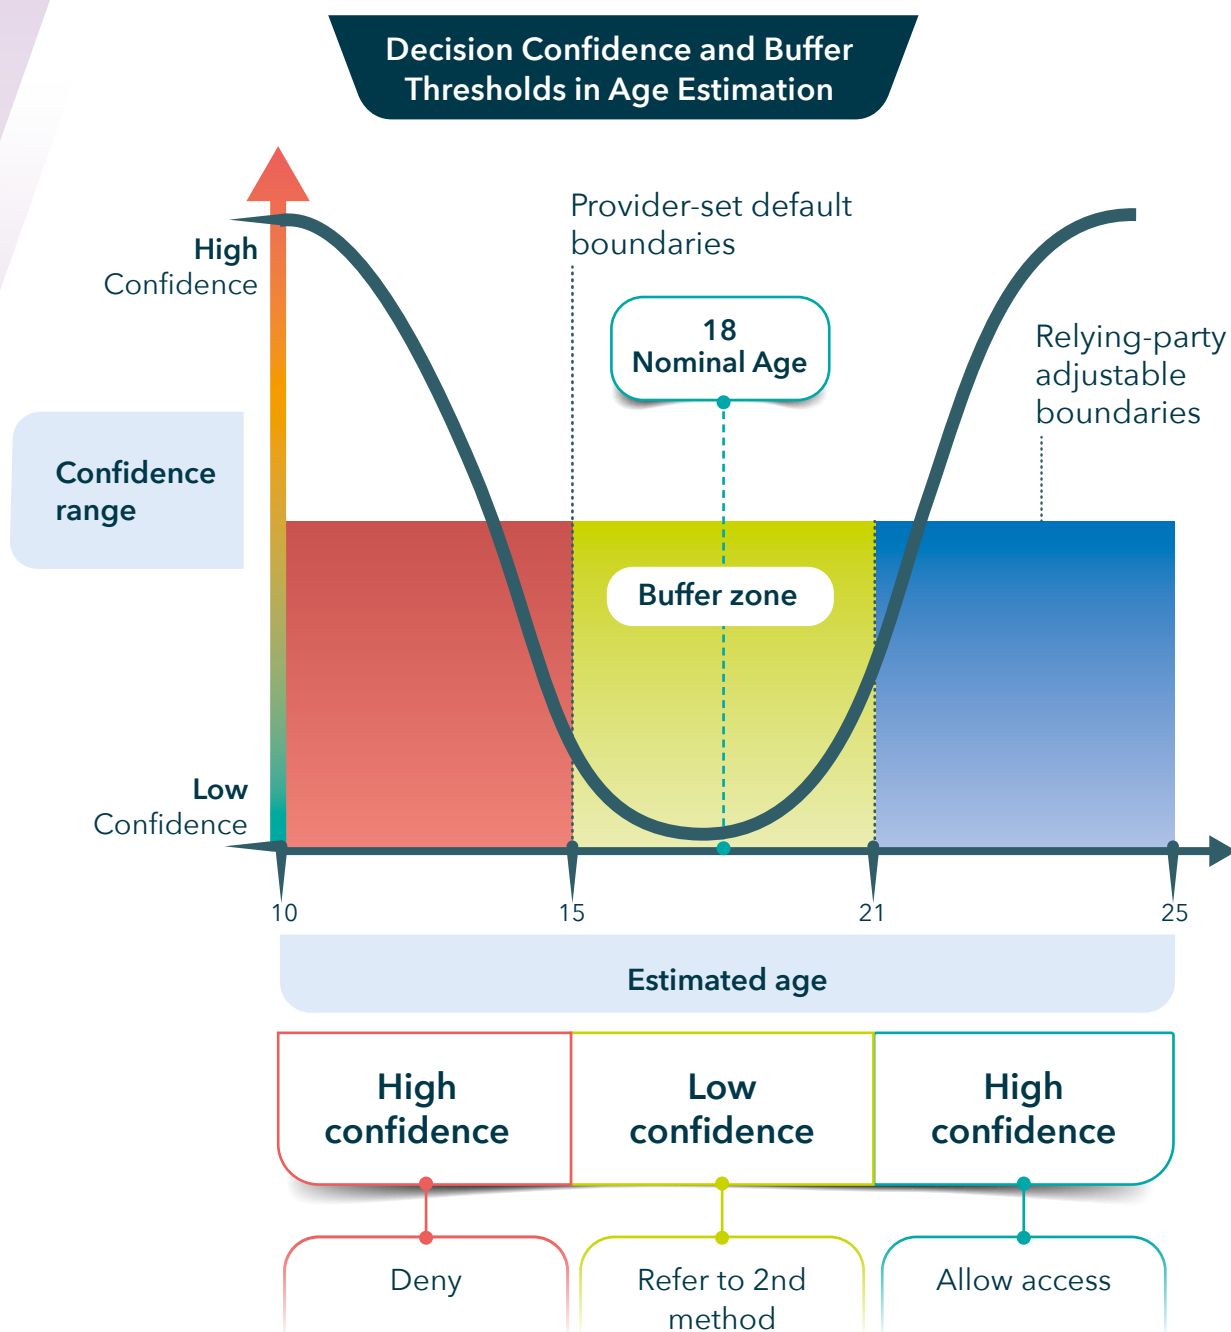
## | Configuration management challenges

**D.9.10** During the Trial, we observed that some relying parties had unclear or inconsistent implementation policies around buffer threshold use. In some cases:

- Relying parties were unsure how to configure or adjust threshold settings for their risk appetite.
- Documentation or provider guidance on how estimation scores should be interpreted was not sufficiently specific.
- Systems defaulted to binary Yes/No responses without explaining confidence margins or uncertainty zones.

**D.9.11** Improved configuration management – including clear documentation, policy guidance and dashboard-level controls – would support more consistent and defensible use of age estimation in live service environments.





**Figure D.9.1** Decision Confidence and Buffer Thresholds in Age Estimation

## D.10 Range of Age Estimation Technologies and Performance Characteristics

**D.10.1** The Trial conducted structured performance testing of facial age estimation systems using a combined cohort of school students and mystery shoppers. The majority of samples were drawn from school-based testing, ensuring statistical strength in the 13–19 age range and enabling robust evaluation of threshold-based classification performance.

**D.10.2** This section summarises how well the tested systems could determine whether users were above or below the 13+ and 16+ age thresholds. Key metrics included:

- True Positive Rate (TPR): Correctly accepting eligible users.
- False Negative Rate (FNR): Incorrectly rejecting eligible users.
- False Positive Rate (FPR): Incorrectly accepting ineligible users.
- Mean Absolute Error (MAE): Average difference between estimated and true age.

## | Analysis of performance characteristics

**D.10.3** This section measures system behaviour under realistic use conditions, with focus on accuracy, error rates, robustness and fairness.

**D.10.4** The landscape of age estimation technologies is rapidly advancing, with mature facial estimation models now offering high-precision results and newer techniques showing promise for broader adoption. Systems that align with the performance principles set out in ISO/IEC FDIS 27566-1 and have undergone statistically grounded evaluation, such as in the Trial, are increasingly ready for deployment in live, regulated environments. As the technology continues to improve, so too does its capacity to serve a wide range of age assurance scenarios – accurately, ethically and with minimal user intrusion.

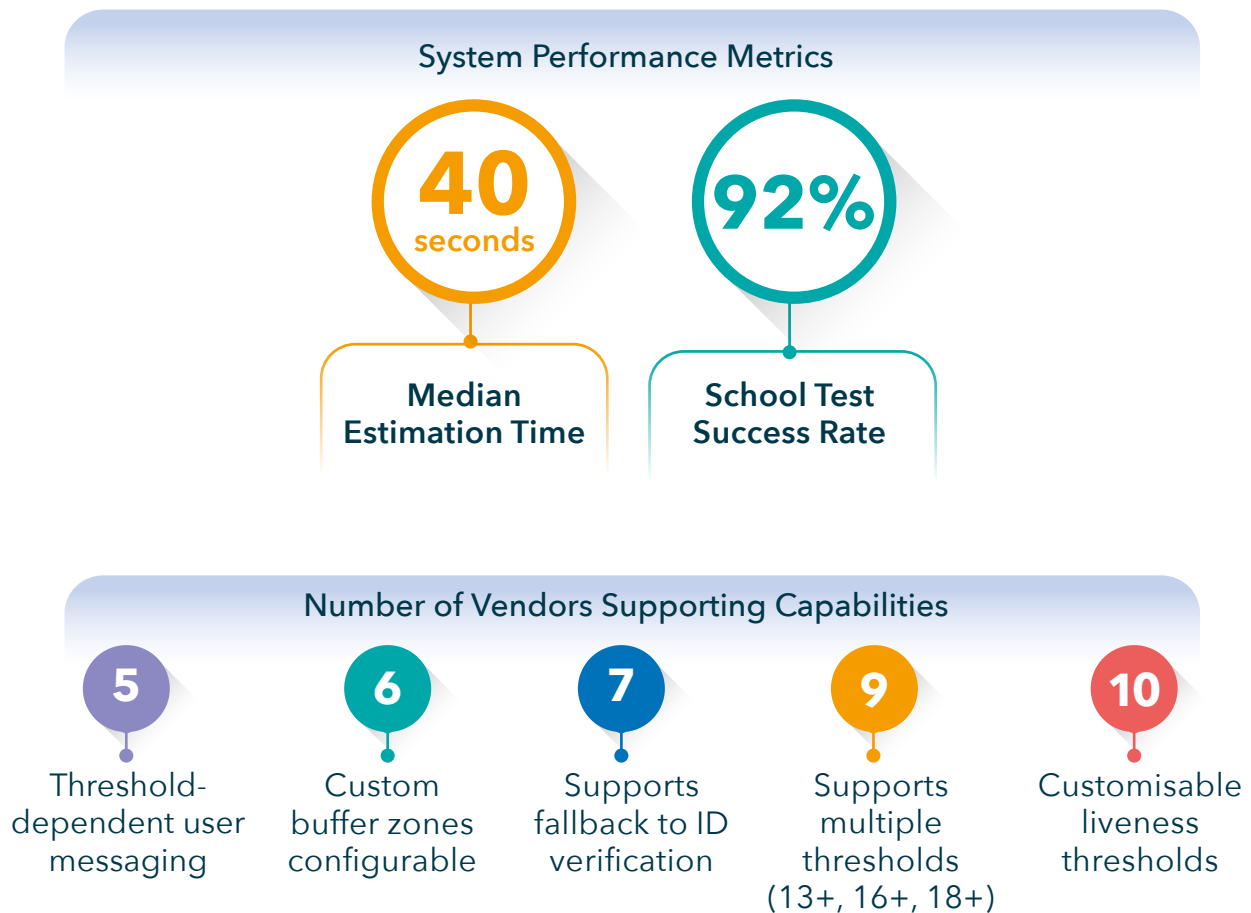
## | Evaluation approach

**D.10.5** The Trial followed a robust testing framework, drawing on the methodology outlined in the Evaluation Proposal Report. Age estimation systems were tested using regression-based performance metrics, such as:

- **Mean Absolute Error (MAE)** – the average difference between estimated and actual age.
- **Threshold Accuracy** – classification accuracy at key age points (13, 16, 18, 21).
- **False Acceptance / Rejection Rates** – misclassification rates relevant to policy enforcement.
- **Statistical Sampling** – subgroup analysis by age, gender and skin tone, based on a labelled dataset with known ground-truth values.

**D.10.6** Systems were also evaluated for model explainability, handling of uncertainty near threshold boundaries and their robustness to real-world variations in lighting, image resolution and facial occlusion.

## System Friction, Usability Data and Capabilities



**Figure D.10.1** System Friction, Usability Data and Capabilities

### | 13+ Age threshold

**D.10.7** For the 13+ gate, systems consistently underperformed at the nominal ages of 13, 14 and 15. False rejection rates (FNR) at these ages ranged from 22% (age 13) to 6% (age 15), meaning a significant number of eligible users were blocked.

**D.10.8** Across all systems at the 13+ age gate, the aggregate TPR exceeded 95% only for users aged 16 and above. This implies that to meet a 95% accuracy threshold for certification, relying parties would need to set a conservative buffer, granting access only when the system estimates the user to be at least 16.



**TABLE - Performance at 13+ Age Gate by Nominal Age**
**7**  
Providers  
Tested

Does not Meet Certification Threshold ←							→ Exceeds Certification Threshold			
True Age	10	11	12	13	14	15	16	17	18	19+
Sample	183	226	1315	1485	1596	453	923	1072	2359	5380
MAE	3.39	2.69	2.57	2.46	2.61	2.72	2.5	2.33	2.29	2.52
StDev	3.97	3.03	2.65	2.48	2.69	2.76	2.7	2.46	2.5	2.36
FPR	44.81 %	46.9 %	62.89 %							
TNR	55.19 %	53.1 %	37.11 %							
TPR				76.43 %	89.66 %	94.04 %	98.7 %	98.69 %	99.79 %	99.63 %
FNR				23.57 %	10.34 %	5.96 %	1.3 %	1.31 %	0.21 %	0.37 %
MoE	0.58	0.4	0.14	0.13	0.13	0.25	0.17	0.15	0.1	0.06

**Samples/Vendor <20 Excluded (Small Samples)**
**Margin of Error based on a 95% Confidence Interval**
**Figure D.10.2 Performance at 13+ Age Gate by Nominal Age (7 providers)**

## | 16+ Age threshold

**D.10.9** At the 16+ gate, systems again fell short of 95% accuracy at the target age. For users aged 16 and 17, false rejection rates remained above acceptable levels (8.5% and 2.6%, respectively).

**D.10.10** TPR consistently exceeded 95% only at age 19 or higher, suggesting that certification-compliant implementations would require relying parties to set a threshold at or above age 19 to ensure reliability.

**D.10.11** These results highlight the need for careful configuration of access thresholds. While facial age estimation technologies have improved, natural error margins and demographic variability require systems to apply age buffers – often granting access only when the estimated age exceeds the legal threshold by 2-3 years.



**TABLE - Performance at 16+ Age Gate by Nominal Age**
**6**  
 Providers  
 Tested

	Does not Meet Certification Threshold ←						→ Exceeds Certification Threshold			
True Age	<13	13	14	15	16	17	18	19	20	21+
Sample	1438	1245	1329	374	770	901	1987	2084	1174	1321
MAE	2.99	2.76	2.91	3.02	2.7	2.41	2.33	2.16	2.27	2.93
StDev	3.01	2.58	2.84	2.92	2.88	2.63	2.68	2.43	2.24	2.56
FPR	25.24 %	38.31 %	59.37 %	73.26 %						
TNR	74.76 %	61.69 %	40.63 %	26.74 %						
TPR					86.1 %	91.34 %	97.58 %	93.71 %	96 %	98.56 %
FNR					13.9 %	8.66 %	2.42 %	6.29 %	4 %	1.44 %
MoE	0.16	0.14	0.15	0.3	0.2	0.17	0.12	0.1	0.13	0.14

**Samples/Vendor <20  
Excluded (Small Samples)**

**Margin of Error based on a  
95% Confidence Interval**

**Note**

- 2 Providers were excluded from this sample due to consistent poor performance at this age gate unfairly skewing the results.
- 2 Providers performed sub-optimally at the nominal ages over 16 which would require further investigation on certification.

**Figure D.10.3 Performance at 16+ Age Gate by Nominal Age**

## Vendor Case Study



Website

rigr.ai

Rigr AI uses AI-driven facial age estimation with privacy-preserving, on-device or edge-enabled architecture to deliver real-time age assurance without storing biometric data, supporting diverse, low-friction digital contexts.

## Three Key Facts

1

Real-time facial age estimation using lightweight AI models optimised for edge devices with no biometric data retention.

2

Binary age decisions (e.g. "Over 18: Yes/No") and integrates configurable thresholds aligned with international standards.

3

Strong performance in live trials, including minimal latency, fast estimation, and compatibility with standard consumer devices.

## Strengths

Privacy-centric by design: no personal data is stored; biometric inputs are processed and deleted locally in milliseconds post-estimation.

System is configurable with buffer zones, allowing relying parties to manage uncertainty while maintaining high user experience quality.

## Practice Statement

[ageassurance.com.au/v/rgr/#PS](https://ageassurance.com.au/v/rgr/#PS)

## Privacy Policy

[ageassurance.com.au/v/rgr/#PP](https://ageassurance.com.au/v/rgr/#PP)

## Technology Trial Test Report

[ageassurance.com.au/v/rgr/#TR](https://ageassurance.com.au/v/rgr/#TR)

## Technology Trial Interview

[ageassurance.com.au/v/rgr/#VI](https://ageassurance.com.au/v/rgr/#VI)

## Summary of Results

Rigr AI showed strong technical readiness, privacy alignment and usability. It delivered reliable threshold-based age decisions with fast, on-device processing and low false-positive rates in real-world scenarios.



## | Use case expansion through greater granularity

**D.10.12** A key observation from the Trial is that increased age estimation accuracy enables new use cases beyond binary “child or adult” classification. Systems capable of estimating age within 1-2 years allow relying parties to enforce differentiated policies across a range of thresholds, including:

- **13+** for online privacy requirements (e.g. COPPA<sup>4</sup>-style regulations)
- **16+** for targeted advertising or social networking access
- **18+** for adult content or legal contracts
- **60+** for age-restricted benefits or senior-specific content

**D.10.13** This level of granularity positions age estimation as a versatile and context-aware age assurance tool, adaptable to evolving regulatory and service design needs.

4. *The Children’s Online Privacy Protection Act, or COPPA, is a law that was passed by the United States Congress in 1998 with the aim of protecting the privacy and personally identifying information of children under the age of 13 who use online services.*



## Vendor Case Study



Website

verifymy.io

Verifymy provides flexible AV solutions integrated with digital wallets, document verification and cross-jurisdictional datasets. It supports selective disclosure and privacy-first age checks, delivering binary outcomes (e.g., "Over 18: Yes") via APIs and reusable credentials for platforms such as gambling, e-commerce and education.

## Three Key Facts

1

Interoperable design with APIs/ SDKs integrates across platforms; supports real-time and asynchronous use cases effectively.

2

Passed tests for 13 +/18+ users; strong fallback triggers like document upload or carrier validation when uncertain.

3

Configurable data retention, including zero-retention; outputs as threshold signals, not exact ages, per ISO/IEC FDIS 27566-1 guidance.

## Strengths

Successfully passed manual testing for key user types, including those at the 13+ and 18+ age thresholds.

Verified strong fallback pathways when confidence was low - triggering document upload or mobile carrier validation.

## Practice Statement

[ageassurance.com.au/v/vmy/#PS](https://ageassurance.com.au/v/vmy/#PS)

## Privacy Policy

[ageassurance.com.au/v/vmy/#PP](https://ageassurance.com.au/v/vmy/#PP)

## Technology Trial Test Report

[ageassurance.com.au/v/vmy/#TR](https://ageassurance.com.au/v/vmy/#TR)

## Technology Trial Interview

[ageassurance.com.au/v/vmy/#VI](https://ageassurance.com.au/v/vmy/#VI)

## Summary of Results

Verifymy demonstrated operational flexibility, strong fallback mechanisms and successful deployment across digital service contexts, confirming no substantial technological barriers for scalable use in Australia.

## D.11 Vibrant and Evolving Sector: Innovation and Sector-Specific Optimisation

**D.11.1** The Trial observed a dynamic, innovative age estimation sector, with providers demonstrating continuous research, model iteration and sector-specific optimisation. Most participating systems had been developed or tuned to align with the needs of the commercial, regulatory and operational domains they serve - ranging from youth-facing social platforms to high-assurance retail and content environments.

**D.11.2** Across the sector, providers showed evidence of:

- Ongoing model improvements, including reductions in error margins.
- Iterative design changes informed by real-world use and feedback.
- A strong focus on reducing user friction and optimising mobile workflows.
- Increasing granularity to support a wider range of age thresholds.



## | Sector-specific optimisation

**D.11.3** Structured interviews and practice statement reviews revealed that providers had taken distinct design paths depending on their target domain:

- **Social Media Platforms:** Prioritise real-time, in-app estimation embedded at friction-sensitive moments (e.g. sign-up, profile edits) to trigger age-related interventions without disrupting user experience.
- **Age-Restricted Content Providers:** Focus on high-accuracy estimation at upper age thresholds (18+, 21+), with fallback to verification if confidence is low.
- **Parental Control Environments:** Use estimation as a first-line filter to flag likely under-13 users before engagement or data collection occurs.

**D.11.4** These variations confirm that age estimation is not a one-size-fits-all solution – instead, systems are being customised to suit different levels of risk, compliance obligations and user flow expectations.



While not all platform implementations were tested in the Trial, several major platforms contributed documentation and interviews that illustrate real-world deployment of age estimation:

### **Meta (including Instagram)**

Meta integrates facial age estimation as a background process – typically triggered when users attempt to change their date of birth. This design uses estimation as a silent, risk-triggered control, aligning with the principles of proportionality and privacy as outlined in ISO/IEC FDIS 27566-1. The system is optimised for mobile and low-friction use, enabling targeted age checks without interrupting broader platform use.

### **Snap**

Snap uses facial age estimation at key user entry points, such as during account creation or before accessing age-sensitive filters. Embedded within the Snapchat app, the system is designed for fast, first-pass screening and supports immediate fallback if required. This lightweight integration supports high churn and prioritises user experience for youth-centric audiences.

## | Innovation and continuous improvement

**D.11.5** Across the sector, the Trial observed a sustained commitment to technical advancement, model refinement and user-focused design. Most participating providers had:

- Developed multiple generations of age estimation models, with measurable accuracy improvements validated through independent benchmarking, including participation in initiatives such as NIST FATE<sup>5</sup>, academic research or certification schemes (e.g. ACCS).
- Created proprietary datasets and domain-specific training pipelines, improving performance across diverse demographic groups and operational settings.
- Invested in dedicated user experience (UX) and accessibility efforts, recognising that reducing friction is essential to adoption and compliance.

**D.11.6** Several vendors also reported active research into emerging input modalities, such as gesture-based classification, behavioural inference and exploratory work on voice-based estimation. While these methods are still in early stages of maturity, they reflect the sector's innovation culture and experimental mindset.

5. NIST FATE stands for Face Analysis Technology Evaluation (FATE), one half of their two track Face Recognition Vendor Test (FRVT) program. The FATE tests provide information on the capabilities of algorithms to inform developers, end users, standards processes and policy and decision makers.



## | Growing use and deployment

**D.11.7** Practice statements and interviews confirm that age estimation is no longer experimental. It is now actively deployed at scale:

- Meta (including Instagram) uses facial age estimation globally to assess users attempting to modify their date of birth – embedding estimation as a background safeguard to enhance age-appropriate design.
- Snap integrates age estimation into its mobile app workflows, applying it during account creation and content gating to deliver fast, low-friction age checks.
- E-commerce and media services use age estimation as a first-layer filter, enabling automated gating before triggering identity verification or parental consent processes.

**D.11.8** These real-world use cases demonstrate that facial age estimation is technically viable, operationally mature and increasingly relied upon to support regulatory compliance with age-based access and privacy rules.

**D.11.9** The age estimation sector – both in Australia and internationally – is mature, adaptable and driven by ongoing innovation. Providers have shown:

- A long-term investment in technical excellence, fairness and bias mitigation.
- Continuous improvements in accuracy, configurability and explainability.

**D.11.10** Sector-specific tailoring to address the distinct needs of platforms, retail and youth safety contexts.

## D.12 Privacy by Design in Age Estimation: Data Handling, Minimisation and Innovation

**D.12.1** The Trial found strong evidence that participating age estimation providers applied robust privacy-preserving practices and had clear internal policies governing the use, handling and disposal of biometric data. Across the sector, we observed:

- Immediate deletion of facial images post-estimation or use of on-device processing where biometric data never left the user's device.
- Separation of data uses, distinguishing between:
  - Real-time operational activity
  - Offline model training
  - Internal benchmarking
  - External third-party validation
- Minimal data retention, with most providers storing only hashed, non-identifiable transaction codes for audit or system diagnostics.

**D.12.2** Some vendors also demonstrated innovation in using non-identifiable physiological inputs – such as gesture-based systems that do not involve facial images or traditional biometrics. These systems, while still emerging in terms of precision, reflect a growing emphasis on privacy-first design and have been independently validated for binary age classification (e.g., child/adult).

## | Alignment with responsible practice and privacy by design

**D.12.3** Participating providers exhibited a consistent commitment to privacy, data protection and ethical design. In line with Clause 7.1 of ISO/IEC FDIS 27566-1, age estimation systems were built and deployed with:

- Minimal data retention, often discarding facial images after estimation
- Client-side or local processing, reducing exposure to centralised risk
- Clear user consent pathways, with transparent explanation of the process
- No requirement for identity disclosure, unlike verification methods

**D.12.4** This reflects a broader alignment with Australia's privacy principles and global data protection expectations, making age estimation an attractive option for low-risk or privacy-sensitive environments.



Provider	Minimal Data Retention	Client-side Processing	Consent Pathways	No ID Required	ISO 27566-1 Alignment
<b>Yoti</b>	Facial images not stored	On-device supported	Clear and transparent	Estimation only mode	Strong
<b>Privately</b>	No transmission; no retention	Fully local (edge AI)	Built into local app flow	Always document-free	Strong
<b>Luciditi</b>	Biometric data deleted post-inference	Partial; fallback steps server-based	Explicit app-based consent	In estimation mode only	Strong
<b>IDMission</b>	Encrypted; no device storage	Server-side inference	Consent governed by client config	Estimation works stand-alone	Moderate
<b>Verifymy</b>	Configurable (incl. zero-retention)	Mostly cloud-based	UX-integrated	Estimation option available	Strong
<b>Needemand</b>	No PII or biometric data stored	Gesture data processed locally	Minimal UI; local language shown	Gesture-only model	Strong
<b>Persona</b>	Minimal data retained	Mixed (client-server)	Consent and opt-out options	Estimation possible	Moderate
<b>Unissey</b>	Deleted after inference	Client capture, server process	Transparent display at capture	Yes	Strong
<b>Private Identity</b>	No biometric data transmitted	On-device only	Consent via interface	Always ID-free	Strong
<b>Privo</b>	Parental consent structure	Varies (parent-managed)	Explicit opt-in via parent	Not ID-free for parent	Partial
<b>Rigr AI</b>	Federated, no central retention	Edge inference	Unknown	No direct ID link	Emerging

## Vendor Case Study



Website

[privo.com](https://privo.com)

PRIVO uses facial and voice biometrics for age estimation via homomorphic encryption. Supports local and cloud-based processing with no template storage. Uses face, voice and fingerprint for biometric estimation, with full privacy via local authentication.

## Three Key Facts

1

No biometric template stored; end-to-end FHE used.

2

Designed for online/ offline use with realtime responses.

3

High-trust applications with strong local binding.

Practice Statement

[ageassurance.com.au/v/pvo/#PS](https://ageassurance.com.au/v/pvo/#PS)

Technology Trial Test Report

[ageassurance.com.au/v/pvo/#TR](https://ageassurance.com.au/v/pvo/#TR)

Privacy Policy

[ageassurance.com.au/v/pvo/#PP](https://ageassurance.com.au/v/pvo/#PP)

Technology Trial Interview

[ageassurance.com.au/v/pvo/#VI](https://ageassurance.com.au/v/pvo/#VI)

## Summary of Results

No full MAE disclosed; positioning reflects future-proof architecture rather than production maturity.



## | Emerging non-facial analysis methods

**D.12.5** While facial estimation was the most mature method under test, the Trial also observed early-stage development of alternative techniques designed to avoid the use of facial or uniquely identifying features:

- **Gesture analysis:** One provider demonstrated a binary classifier using hand movement patterns (no face captured), independently validated for high child/adult accuracy.
- **Body pose or voice-based signals:** Other research presented systems capable of classifying likely minors based on pitch analysis or skeletal movement – not yet trial-tested, but indicative of sector-wide innovation.

**D.12.6** These developments reflect an increasing regulatory and public demand for non-intrusive, zero-retention age estimation methods. While still in early deployment phases, they highlight the sector's responsiveness to concerns about biometric overreach and offer promising paths forward for privacy-maximised age assurance.

## Vendor Case Study

# Needemand

Website

[needemand.com](https://needemand.com)

Needemand's BorderAge solution offers a compelling example of a non-facial, privacy-preserving age estimation modality. Unlike most Trial participants, Needemand does not rely on facial analysis, voice or any biometric traits traditionally associated with identity. Instead, it uses hand gesture dynamics, captured via a device's camera, to whether a user is likely an adult or a child.

## Three Key Facts

1

Uses hand geometry only, with no facial or identifying data required.

2

Preferred by many users during mystery shopping for its noninvasive, document-free approach.

3

Demonstrated strong underage detection in in controlled school trials, with low false positives for users below Age Gate 18.

## Strengths

Fast and lightweight, Needemand's system is ideal for constrained devices or physical environments like kiosks or mobile retail setups. Enables inclusive access by avoiding biases linked to facial recognition; useful in contexts where users decline image-based processing.

### Practice Statement

[ageassurance.com.au/v/nee/#PS](https://ageassurance.com.au/v/nee/#PS)

### Privacy Policy

[ageassurance.com.au/v/nee/#PP](https://ageassurance.com.au/v/nee/#PP)

### Technology Trial Test Report

[ageassurance.com.au/v/nee/#TR](https://ageassurance.com.au/v/nee/#TR)

### Technology Trial Interview

[ageassurance.com.au/v/nee/#VI](https://ageassurance.com.au/v/nee/#VI)

## Summary of Results


Needemand's gesture-based system proved lightweight, privacy-preserving and effective for simple age gates. It offers an innovative alternative to facial AI, suitable for specific contexts with low biometric tolerance.

## | Contextual privacy protections

**D.12.7** A consistent finding across provider statements and system designs was that privacy safeguards were calibrated to the deployment context, in alignment with the risk-based approach taken by ISO/IEC FDIS 27566-1:

- Low-risk environments (e.g., content intended for children or app filtering) favoured on-device inference, zero retention and no fallback to identity.
- Higher-risk contexts (e.g., gambling, adult content) used confidence thresholds and fallback mechanisms, such as optional identity verification for borderline cases.

**D.12.8** This risk-aligned approach reflects a sector that is actively implementing privacy-by-design principles at scale. The use of temporary biometric data, modular fallback logic and non-identifying input methods (e.g., gesture-based estimation) demonstrates that age estimation can be deployed proportionately and ethically, even in anonymous user flows.



## D.13 Inclusion and Demographic Consistency in Age Estimation

**D.13.1** Age estimation provides a document-free option for assessing age eligibility, which can be particularly useful in situations where formal ID is unavailable, a digital identity has not been established or users prefer to remain anonymous. This approach may help reduce access barriers for groups such as First Nations and Torres Strait Islander Peoples, newly arrived migrants and younger users, who are statistically less likely to hold official identity credentials. However, its effectiveness depends on the accuracy of the underlying technology and its suitability across diverse populations.

**D.13.2** These characteristics make age estimation particularly relevant to the accessibility and equity goals outlined in ISO/IEC FDIS 27566-1, including:

Title and Clause	Requirements
<b>Accessibility and Inclusion</b> (Clause 9.2)	Systems must accommodate users without ID, accounts or stable connectivity and avoid bias.
<b>Friction Minimisation</b> (Clause 9.3)	Age assurance methods must not impose disproportionate barriers or delays for certain groups.

## | Known limitations in demographic coverage

**D.13.3** Despite these strengths, the Trial identified a recurring concern: the underrepresentation of First Nations and other non-majority ethnic groups in the training data used by facial age estimation systems. Several vendors acknowledged this issue and recognised that:

- Common AI training sets are often Western-centric, leading to reduced precision for users with underrepresented skin tones, facial features or age-expression patterns.
- Environmental factors – such as poor lighting in rural areas or low-bandwidth mobile capture – can further degrade performance for underserved communities.

**D.13.4** Although the Trial found no consistent or statistically significant evidence of adverse performance for First Nations users, it also noted that sample sizes in subgroup testing were inadequate and further validation is warranted.

## | Provider responses and sector trends

**D.13.5** In response to this challenge, participating vendors described ongoing efforts to improve fairness, including:

- Fairness audits and disaggregated performance monitoring.
- Synthetic data augmentation to improve representation of under-sampled demographics.
- Consent-based data gathering in collaboration with local communities.
- Exploration of non-facial modalities (e.g. hand gesture classification) to reduce reliance on biometrics that may carry demographic bias.



**D.13.6** These actions reflect a growing sector-wide awareness of inclusion risks and a willingness to align with the ethical expectations of diverse user groups. As the field of age estimation continues to evolve, efforts to improve demographic consistency will be crucial to ensuring equitable, privacy-preserving age assurance at scale.

### **| Acceptable demographic variance thresholds**

**D.13.7** While the IEEE 2089.1 standard includes a column referencing a 1% threshold for outcome error parity, its interpretation is currently ambiguous. Specifically, the standard does not clearly define whether the 1% limit refers to a percentage of the total range, prediction error, population or another basis. This lack of clarity makes it difficult to implement the threshold in a transparent, reproducible way. We have provided this feedback to the IEEE standards development team who have confirmed this will be addressed as part of the current corrigendum review of its standardised levels of age assurance.

**D.13.8** For the purposes of this Trial, we have adopted the Four-Fifths Rule as an alternative and well-established measure of fairness for outcome error analysis. This standard was originally established by the U.S. Equal Employment Opportunity Commission (EEOC) for evaluating adverse impact. Although not formally legislated in Australia, the rule has been referenced in Australian Human Rights Commission (AHRC<sup>6</sup>) publications as a relevant international benchmark for fairness evaluation in algorithmic decision-making. Under this framework, outcome error parity is deemed acceptable so long as no group experiences more than a 20% disparity relative to others, which provides a transparent and practical standard for assessing fairness across demographic groups.


6. *Human Rights and Technology Discussion Paper - December 2019*

[humanrights.gov.au/sites/default/files/document/publicationtechrights\\_2019\\_discussionpaper\\_0.pdf](https://humanrights.gov.au/sites/default/files/document/publicationtechrights_2019_discussionpaper_0.pdf)

## | Outcome error parity by skin tone

**D.13.9** We evaluated the performance fairness of participants in the Trial providing age estimation solutions across individuals with varying skin tones.


**D.13.10** All test subjects were categorised against the Fitzpatrick Scale. The Fitzpatrick scale is a widely used classification system that categorises human skin tones into six types based on their response to ultraviolet (UV) exposure. Originally developed in dermatology, it ranges from Type I (very fair skin that always burns and never tans) to Type VI (very dark skin that never burns and tans very easily). The scale provides a standardised framework for assessing skin pigmentation and is commonly used in medical research, cosmetics and increasingly in evaluating fairness in AI systems, where skin tone may impact model performance.



**Figure D.13.1** Skin Tone Shades

**D.13.11** Skin tone is a known source of bias in age estimation systems. As light reflects off a subject's face and into the camera sensor, lighter skin tones tend to return more detail-rich signals because they reflect more light; darker skin absorbs more light, reducing contrast and lower data availability for the algorithm to analyse. This intrinsic disparity, combined with historical underrepresentation in datasets, can lower accuracy for darker-skinned users. Addressing this requires more inclusive data and improved image processing across lighting conditions.


**D.13.12** We conducted a Mean Absolute Error (MAE) parity analysis. This assessment measures whether individuals from different demographic groups experience materially different prediction errors. Fairness in regression contexts such as this can be evaluated by comparing group-wise MAE ratios to the average MAE.



**Figure D.13.2** Mean Absolute Error (MAE) Parity Analysis

## | Metric and methodology


**D.13.13** Adopting the **Four-Fifths Rule** (or **80% rule**) as a threshold of acceptable disparity, we calculated the MAE for each Fitzpatrick skin tone group (I-VI), expressing it as a ratio relative to the overall average MAE across all groups. The results are as follows:



**Figure D.13.3** Skin Tone Group MAE Ratio (vs Average MAE)

**D.13.14** All skin tone groups fall within the 0.80 to 1.25 range, which satisfies the Four-Fifths Rule. This indicates no significant evidence of adverse impact or systemic disparity in error rates across skin tone groups enabling us to state that they are broadly consistent across demographic groups.

- **Fairness threshold:** Under the Four-Fifths Rule, any group whose outcomes fall below 0.80 (i.e., experiencing 20% or more error disparity) would trigger a fairness concern. No group in our data meets this condition.



- **Relative fairness:** Although groups V-VI have slightly higher error ratio (1.11), this remains within the accepted 1.25 upper bound. Conversely, groups I-II and III-IV show slightly lower-than-average errors but still exceed the 0.80 threshold.
- **Regulatory relevance:** While Australia does not formally legislate this rule, its use by AI vendors and recognition in AHRC submissions supports it as a reasonable fairness benchmark in Australian deployment contexts.

**D.13.15** Based on the Four-Fifths Rule, the Trial participants demonstrate acceptable outcome error parity across all evaluated skin tone groups. Although the distribution of errors suggests some variation, no group is subject to disproportionate error and the model therefore meets a minimum fairness standard as understood under international norms referenced by Australian human rights bodies.

**D.13.16** This section focuses on error parity across skin tone groups, with all results falling within accepted MAE bounds. Further refinements may be warranted by examining classification behaviour, such as false positive and true positive rates across skin tone groups and age gates, to strengthen fairness assurance.



## Vendor Case Study



Website

[withpersona.com](https://withpersona.com)

Facial age estimation with fallback to ID verification. Includes audit-backed fairness metrics and governed update process; privacy-preserving design with opt-out controls.

Practice Statement

[ageassurance.com.au/v/per/#PS](https://ageassurance.com.au/v/per/#PS)

Technology Trial Test Report

[ageassurance.com.au/v/per/#TR](https://ageassurance.com.au/v/per/#TR)

Privacy Policy

[ageassurance.com.au/v/per/#PP](https://ageassurance.com.au/v/per/#PP)

Technology Trial Interview

[ageassurance.com.au/v/per/#VI](https://ageassurance.com.au/v/per/#VI)**Summary of Results**

Supports user opt-out and clear user interface, strong performance with First Nations youth in school trials, demographic audit trails supported. MAE ranged from 0.86 to 3.31 across different cohorts; strong at 13+ thresholds. Accuracy varied more widely for under-13 users.

## D.14 Analysis of Acceptability Characteristics

**D.14.1** Acceptability in the context of ISO/IEC FDIS 27566-1 refers to whether an age assurance tool is understandable, inclusive, fair and trustworthy from the user's perspective. While age estimation systems are typically backend or low friction, the Trial examined elements relevant to user perception, control and fairness.

Attribute and Standard Reference	Indicator	Evidence
<b>Accessibility</b> (IEEE 2089.1 §1.4(f), ISO/IEC 25010)	Supports assistive tech, multiple languages	User interface tests, WCAG <sup>7</sup> conformance
<b>Transparency</b> (IEEE 2089.1 9.3.b.9)	Clear age assurance explanations for users	User guidance, helpdesk feedback
<b>Bias Minimisation</b> (IEEE 2089.1 Annex A.3, ISO/IEC 25010)	Fair performance across demographic groups	Error parity reports
<b>User Control and Remedy</b> (IEEE 2089.1 §10.3.i)	Right to challenge incorrect estimations	Support logs, escalation flows

7. WCAG conformance refers to Web Content Accessibility Guidelines conformance. WCAG is a set of international standards developed by the World Wide Web Consortium (W3C) to make web content more accessible to people with disabilities.

### D.14.2 *Key insights from the Trial*

- Vendors with on-device or non-persistent systems (e.g. Yoti, Privately, Needdemand) scored highest on user privacy and low friction, which are strongly correlated with acceptability in youth-facing platforms.
- Fallback paths (e.g. escalation to ID or retry) were used to mitigate frustration in borderline cases – an important usability safeguard.
- Fairness testing showed acceptable parity across age and gender, with some known challenges in accuracy for underrepresented skin tones or facial features (see D.13.9).
- Transparency mechanisms varied: only a few providers included clear user-facing messages explaining that estimation – not verification – was being used.

### | Improving acceptability in future deployments

**D.14.3** To enhance public trust and uptake of age estimation tools, relying parties could:

- Offer user-visible messaging that explains how age estimation works and what happens if a user is blocked.
- Include fallback or retry options when confidence is low, especially for users near key thresholds (e.g. 13, 16, 18).
- Ensure estimation systems are tested across diverse populations, with published fairness audit results.
- Adopt document-free default modes, reserving identity linkage only for high-risk exceptions.

## Vendor Case Study



Website

[yoti.com](https://yoti.com)

Yoti performs high-accuracy facial estimation either on-device or at the edge. Offers rapid, anonymous checks with no data storage. Certified for privacy and security compliance. Functional testing confirmed interoperability, robustness and privacy aspects of the systems.

## Three Key Facts

1

Uses facial age estimation; deletes images after inference; explains methods transparently.

2

No biometric storage; GDPR-aligned; anonymous age checks possible.

3

Uses anti-spoofing SDK with protection against static and video injection; penetration tested.

## Strengths

Automated lab testing showed high accuracy for Age Gates 13 and 16, with True Positive Rates consistently above 94% from age 13 upwards and Mean Absolute Error (MAE) values under 2 years for ages 13-20. 80% of users reported being either satisfied or very satisfied with the experience.

## Practice Statement

[ageassurance.com.au/v/yot/#PS](https://ageassurance.com.au/v/yot/#PS)

## Privacy Policy

[ageassurance.com.au/v/yot/#PP](https://ageassurance.com.au/v/yot/#PP)

## Technology Trial Test Report

[ageassurance.com.au/v/yot/#TR](https://ageassurance.com.au/v/yot/#TR)

## Technology Trial Interview

[ageassurance.com.au/v/yot/#VI](https://ageassurance.com.au/v/yot/#VI)

## Summary of Results

There were very few usability issues arising from Mystery Shopper trials. While response time is not as fast as some systems tested, but users reported that the system met or exceeded their expectations in terms of task time. The majority of users had no issues in completing the evaluation task and reported the system to be either easy or very easy to use.

**D.14.4** Age estimation has strong potential to deliver inclusive, low-friction age assurance – especially in document-scarce populations. To achieve high acceptability, systems would need to be:

- Understandable
- Fair
- Reversible when wrong; and
- Designed to respect user autonomy and privacy

**D.14.5** When configured accordingly, age estimation aligns well with the acceptability expectations of ISO/IEC FDIS 27566-1 and IEEE 2089.1.

## D.15 Training Data Challenges in Age Estimation: Quality, Ethics and Risk Management


**D.15.1** The Trial identified a critical opportunity for improvement in the sector: the availability, quality and ethical sourcing of ground-truth training data used in age estimation systems.

**D.15.2** Age estimation models – particularly those using facial analysis – require biological or behavioural data linked to known dates of birth to train and validate accurate predictions. These “ground-truth” datasets are essential for improving precision and fairness, but are challenging to obtain, especially when minors are involved.



## | Observed age estimation performance by age group

**D.15.3** The Trial's accuracy data demonstrates that performance improves with subject age, which is partially attributable to the amount and diversity of training data available at each age band.



**Figure D.15.1** Age Estimation Performance

## Challenges in ground-truth dataset collection

### Four Key Challenges of Data Sources



#### Privacy and consent

Minors cannot legally consent; parental or guardian actions are required, increasing complexity.



#### Demographic skew

Many publicly available datasets disproportionately reflect lighter-skinned, Western & adult faces.



#### Temporal validity

Datasets with outdated or incorrectly labelled ages reduce model precision.









#### Context bias

Posed or high-resolution photos (e.g. celebrities) do not reflect everyday user conditions like webcams or mobile selfies.

**Figure D.15.2** Four Key Challenges of Data Sources

## | Ethical vs. dubious data sources

**D.15.4** It is important to consider the ethics of data sources.

Ethical/Responsible Sources		Dubious or High-Risk Sources	
	<b>Volunteer data with consent</b> (e.g. Yoti <sup>8</sup> , Microsoft <sup>9</sup> programs)		Scraped web or social media images without consent
	<b>IRB-approved academic datasets</b> (e.g. FG-NET <sup>8</sup> , Morph-II <sup>9</sup> )		IMDB <sup>10</sup> , WIKI <sup>11</sup> and others with estimated or unclear age labels
	<b>Synthetic data and AI augmentation</b>		Unauthorised reuse of identity verification images

8. <https://www.yoti.com/>

9. [https://www.microsoft.com/en-us/microsoft-365/products-apps-services?msocid=26daf\\_a6f2ff16ffb3aa2efc92ee46ea2](https://www.microsoft.com/en-us/microsoft-365/products-apps-services?msocid=26daf_a6f2ff16ffb3aa2efc92ee46ea2)

10. FG-Net is a dataset for age estimation and face recognition across ages. It is composed of a total of 1,002 images of 82 people with age range from 0 to 69 and an age gap up to 45 years.

11. The MORPH-II dataset is one of the largest longitudinal morphological face database available to the public. It is comprised of mug shots taken of arrested persons over a period of five years and has been used for advancement in facial recognition in many settings.

12. IMDb, historically known as the Internet Movie Database, is an online database of information related to films, television series, podcasts, home videos, video games, and streaming content online. It was founded in 1990 and since 1998, has been owned by Amazon.

13. Wikipedia is a free online encyclopedia written and maintained by a community of volunteers, known as 'Wikipedians'. It was founded in 2001, exists in over 340 languages and is the world's eighth most visited website.

## | Standards-based risk management

**D.15.5** While ISO/IEC FDIS 27566-1 does not dictate how training data should be collected, it imposes clear privacy, fairness and transparency requirements:

ISO/IEC FDIS 27566-1	Criteria
<b>Privacy and Data Protection</b> (Clause 7)	Only data required for the stated purpose should be collected or retained.
<b>Fairness and Accessibility</b> (Clause 6.3.3)	Systems should demonstrate performance consistency across diverse populations.
<b>Confidence Expression</b> (Clause 7.2)	Where uncertainty exists (including from data limitations), the system should expose this in its decision logic.



**D.15.6** Long-term, solutions such as federated learning, shared data stewardship or privacy-preserving training techniques may help the sector scale ethically.



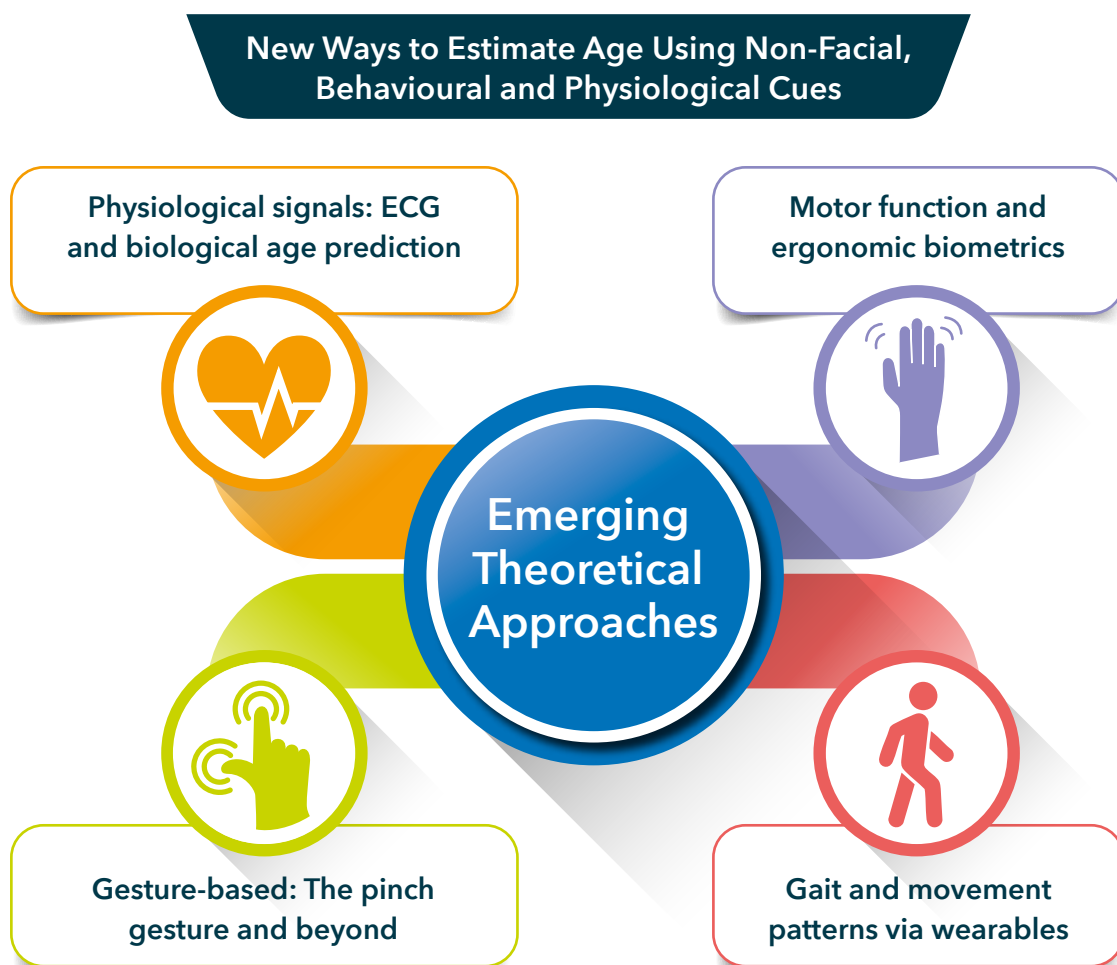
**Figure D.15.3** Compliant Training Data Pipeline for Age Estimation



## D.16 Emerging Theoretical Approaches in Age Estimation: Ergonomics, Physiology and Interaction Patterns

**D.16.1** While the Trial focused on commercially viable systems – primarily based on facial analysis – there is a growing body of academic research exploring new ways to estimate age using non-facial, behavioural and physiological cues. These methods remain largely theoretical or in early-stage development (TRL<3), but they show promise for future age estimation systems that are less intrusive, more ambient and potentially more privacy-preserving.

**D.16.2** These experimental techniques span multiple disciplines including biomedical engineering, human-computer interaction, ergonomics and behavioural science and could eventually enable age estimation embedded in-app, in-game or in ambient environments.



**Figure D.16.1** New Ways to Estimate Age Using Non-Facial, Behavioural and Physiological Cues

## Physiological signals: ECG and biological age prediction



- ECG (electrocardiogram) data varies with age due to cardiovascular changes.
- Deep learning models can infer chronological or biological age from features like heart rate variability and waveform morphology.

### Examples:

- Attia et al. (2019)<sup>14</sup>: MAE of ~6 years using 12-lead ECG data.
- Zhen et al. (2022): Used ECG to estimate “biological age” and health risks.

## Motor function and ergonomic biometrics



- Age influences motor control and grip strength, which can be detected via touch-based or mobile interfaces.
- Age-related differences have been observed in gesture fluidity, tap speed and force modulation.

### Examples:

- Seidler et al. (2010): Precision grip and motor decline with age.
- Voelcker-Rehage et al. (2007): Age affects ergonomic movement control.

14. All references to academic studies and research on pages 86-87 can be found in Part D's Bibliography section within the Part K Report.

## | Gesture-based age estimation



- Common gestures like pinch, swipe and zoom show measurable variation by age group – particularly amongst children.
- Gesture dynamics (speed, accuracy, pressure) offer potential non-biometric age signals.

### Examples:

- Weber et al. (2011): Gesture timing differences by age.
- Frank et al. (2013): Touch dynamics used for age-influenced authentication.

## | Gait and movement via wearables



- Body motion and posture change predictably with age.
- Wearables (e.g. wrist or ankle IMUs<sup>15</sup>) can capture kinematic data used to estimate age or gender.

### Example:

- Jamil et al. (2020): Used deep learning on wearable sensor data to estimate age.

These approaches could integrate with augmented/virtual reality, fitness or gaming systems for low-friction age filtering. These emerging modalities show potential for next-generation age estimation systems that are:

- Non-facial and non-identifying
- Ambient and frictionless
- Suitable for low-data or privacy-sensitive environments

15. An IMU refers to an Inertial Measurement Unit (IMU) that is attached to or embedded near the ankle to capture motion-related data.

## | Age estimation and digital credentials: Limitations for verified use in holder services

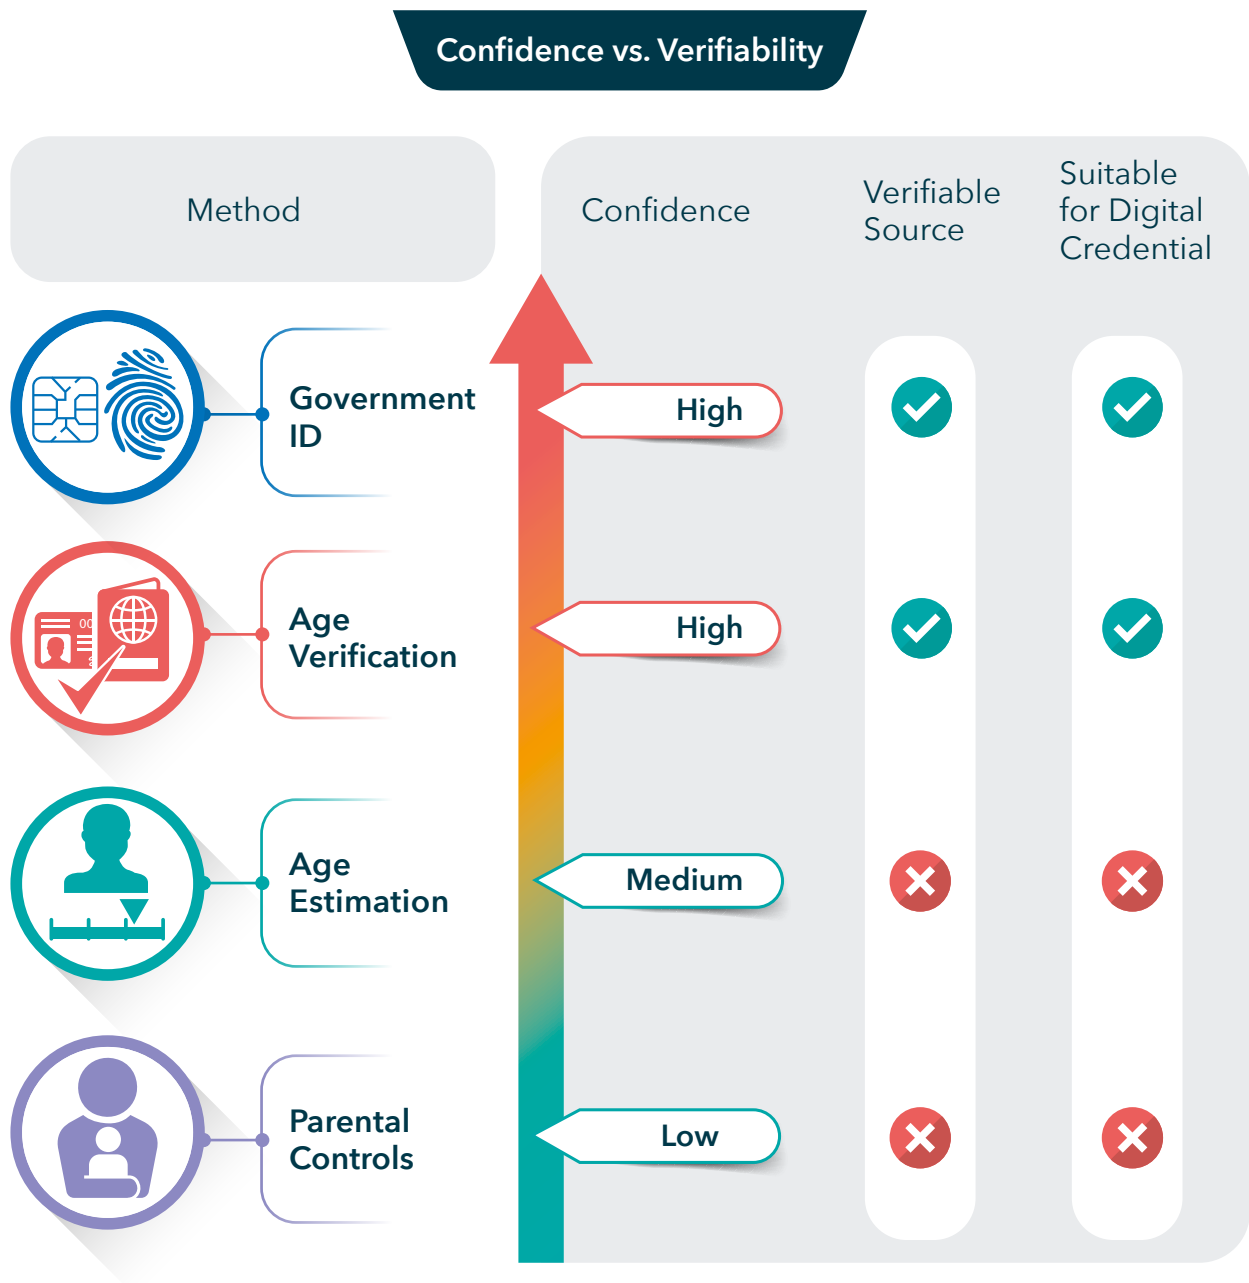
**D.16.3** On the basis of the evaluation results, age estimation methods in isolation may not always be sufficiently dependable to serve as verified credentials in digital wallets or holder services – particularly for regulatory thresholds at ages 13, 16, or 18.

**D.16.4** Digital credentials that state “Over 13” or “Over 18” offer value in systems that require interoperability and privacy, such as online platforms, e-commerce, and government services, however they are not currently sufficient on their own to support the creation of verified age credentials within digital wallets or holder services for use in regulatory or compliance-critical contexts (e.g. thresholds at 13, 16 or 18).

This limitation is not necessarily a reflection of the underlying model accuracy, but rather of the challenges associated with treating a probabilistic, context-sensitive output (such as “estimated age 18+”) as a reusable, persistent credential that carries the same assurance level as a verified date of birth or government-issued document.

**D.16.5** Although it is technically possible to store an age estimation result in a digital wallet—for example, as a token stating the user was previously assessed as “18+” – these tokens currently lack a standardised, interoperable structure. Essential metadata, such as confidence thresholds, buffer zones, fallback methods and time of assessment, are not consistently encoded or available for cross-system interpretation.

**D.16.6** Without such standardisation, estimated age tokens cannot yet support robust enforcement, audit or escalation across services – particularly in compliance-critical settings. Until supporting frameworks mature, age estimation is better suited to real-time decision-making scenarios, where context and confidence thresholds can be assessed dynamically at the moment of use.



**Figure D.16.2** Confidence vs. Verifiability

Government ID – Confidence: High; Verifiable Source: Yes; Suitable for Digital Credentials: Yes.

Age Verification – Confidence: High; Verifiable Source: Yes; Suitable for Digital Credential: Yes.

Age Estimation – Confidence: Medium; Verifiable Source: No; Suitable for Digital Credential: No.

Parental Controls – Confidence: Low; Verifiable Source: No; Suitable for Digital Credential: No.

A vertical arrow on the left visually indicates confidence levels from low at the bottom to high at the top.



## Vendor Case Study



LUCIDITI®

Website

[luciditi.co.uk](https://luciditi.co.uk)

Luciditi uses facial biometrics with fallback to ID and Open Banking. Biometric inputs are deleted post-inference. Offers threshold tuning and inclusive usability features.

## Three Key Facts

1

Combines facial estimation, ID document fallback and Open Banking checks.

2

MAE of 1.29 years for 18-20-year-olds. High fallback usability and multiple assurance pathways.

3

Better clarity needed on exact error performance across all age bands.

## Strengths

Passed all controlled test scenarios, including:

- Poor lighting and low-quality selfies
- Face partially obscured by glasses, hats or scarves
- Liveness challenge using spoof images or pre-recorded selfies

## Practice Statement

[ageassurance.com.au/v/luc/#PS](https://ageassurance.com.au/v/luc/#PS)

## Privacy Policy

[ageassurance.com.au/v/luc/#PP](https://ageassurance.com.au/v/luc/#PP)

## Technology Trial Test Report

[ageassurance.com.au/v/luc/#TR](https://ageassurance.com.au/v/luc/#TR)

## Technology Trial Interview

[ageassurance.com.au/v/luc/#VI](https://ageassurance.com.au/v/luc/#VI)

## Summary of Results

Manual testing confirmed robust interoperability, resilience to input variation, and strong privacy controls. The system deletes biometric data after inference. The solution is assessed at TRL 9, aligned with the vendor's self-assessment and is considered deployment-ready with privacy strengths and performance trade-offs.

## D.17 Security of Age Estimation Systems: Protecting Biometric Data and Meeting ISO Standards

**D.17.1** We found that the age estimation systems were generally secure and consistent with information security standards. Most of the providers were able to demonstrate ISO/IEC 27001:2022 certified information security management and some had other supplementary security protocols (such as SOC 2 or Fintech-level security).

**D.17.2** The Trial found that age estimation providers demonstrated a strong commitment to information security, particularly in how they managed the collection, processing and disposal of biometric data (e.g. facial images or video streams) used to estimate age. Providers consistently implemented systems that were well-engineered, securely deployed and aligned with international information security standards, ensuring that sensitive inputs could be handled without introducing unnecessary risk to users.



## | Biometric security requirements in ISO/IEC FDIS 27566-1

**D.17.3** Age estimation systems, by their nature, often involve the collection and processing of biometric samples – typically facial images used as inputs for machine learning models. These samples, even when not used for identification, are classified as sensitive personal data under most data protection laws and standards.

**D.17.4** ISO/IEC FDIS 27566-1 recognises this risk and makes specific provisions for the secure handling of biometric data, including:

ISO/IEC FDIS 27566-1	Criteria
<b>Privacy and Data Protection</b> (Clause 7.2)	Requires that personal data, especially biometric inputs, are processed with data minimisation and security in mind and only retained for as long as necessary for the age estimation process.
<b>Security</b> (Clause 8.1)	Mandates that biometric data should be protected against unauthorised access, tampering or disclosure and that risks to confidentiality and integrity are appropriately mitigated.
<b>Expression of Confidence and Uncertainty</b> (Clause 6.3)	Requires that age estimation systems clearly represent the confidence level of the estimation, reducing the risk of misuse or overreliance on uncertain outputs.

## | Best practice observed during the Trial

**D.17.5** The age estimation systems evaluated during the Trial adhered to best practices in biometric security, including:

- **Temporary processing:** Most systems were designed to process the facial image on-device or in a secure enclave, after which the image was immediately discarded. This minimised the risk of post-processing exposure.
- **No biometric retention:** In nearly all cases, no biometric sample (e.g. image, video or feature vector) was stored after the estimation process was completed. Instead, systems retained only a non-identifiable transaction code or binary result (e.g. "Likely over 18").
- **End-to-end encryption:** Where processing occurred remotely, biometric inputs were transmitted over secure channels (TLS<sup>16</sup>) and processed in segregated, access-controlled environments.
- **Controlled inference environments:** Some vendors offered age estimation via software development kits that run locally on the user's device, meaning no biometric data ever leaves the device – an approach aligned with both data minimisation and zero-trust architectural principles.
- **Access control and audit:** Internal access to age estimation infrastructure was tightly controlled, often using role-based permissions and accompanied by detailed logging and monitoring to detect anomalous or unauthorised access.

16. TLS refers to Transport Layer Security. It is a security protocol that provides privacy and data integrity for Internet communications.

## | Analysis of security characteristics







**D.17.6** Evaluates the ability of the system to protect data and resist attacks during age estimation operations.

### *Security controls evaluation*

Control Area and References	Description	Evidence	Alignment
<b>Secure Transmission</b> (ISO/IEC 25010 §4.2.5)	Use of TLS 1.3 or higher	Vendor interviews / Network configuration	✓
<b>Data Storage Isolation</b> (IEEE 2089.1 §14.2)	Encrypted and access-controlled	Vendor interviews / System logs, configs	✓
<b>User Control Over Data</b> (IEEE 2089.1 §14.2)	Users can manage retained data	Practice statements / Static review UI, UX evidence	✓



**D.17.7** Some examples of approaches to information security management by age estimation service providers:

Approaches to Information Security Management by Age Estimation Service Providers		
Solution Name	ISO/IEC 27001 Certified	Additional Certifications
 LUCIDITI®	✓	ACCS audited; PASS 5:2023; UK DIATF
 iDmission empowering identity	✓	SOC 2 Type 2
 verifymyage	✓	Certified by A-LIGN, ACCS audited
 VeriDas Age Assurance	✓	SOC 2 Type II, SOC 3, iBeta PAD, ACCS
 agechecked	✓	Internal/external audits; aligning with ISO/IEC FDIS 27566-1, ACCS
 YOTI Age Estimation	✓	Certified liveness detection, GDPR compliant, ACCS

**Figure D.17.1** Approaches to Information Security Management by Age Estimation Service Providers

## Vendor Case Study



Website

[idmission.com](https://idmission.com)

IDmission offers facial AI-based age estimation trained over 10 years on global datasets, with server-side encrypted processing. Liveness detection and spoof resistance tested under various real-world conditions. Demonstrated high precision with <1 year MAE and strong demographic inclusivity, except with indigenous users.

Practice Statement

[ageassurance.com.au/v/idm/#PS](https://ageassurance.com.au/v/idm/#PS)

Technology Trial Test Report

[ageassurance.com.au/v/idm/#TR](https://ageassurance.com.au/v/idm/#TR)

Privacy Policy

[ageassurance.com.au/v/idm/#PP](https://ageassurance.com.au/v/idm/#PP)

Technology Trial Interview

[ageassurance.com.au/v/idm/#VI](https://ageassurance.com.au/v/idm/#VI)

## Summary of Results

Lab and school testing showed high false positives for underage users across all gates, and MAE above 4.5 years, indicating low age estimation precision. False negatives were low for eligible users. While the vendor assessed their ToE at TRL 9, independent evaluation concluded it to be at TRL 8, suggesting further validation is needed for full commercial deployment.



## D.18 Threat Mitigation in Age Estimation Systems: Spoofing, Deepfakes and Injection Attacks

**D.18.1** The age estimation providers were acutely aware of the threat vectors from spoofing and artificial intelligence and have taken steps to mitigate those risks in accordance with ISO/IEC 30107 (Biometric Presentational Attack Detection). This means that it is hard (but not entirely impossible) for a user to present a deepfake age-altered image of themselves (or someone else) to spoof a system.

**D.18.2** We found that providers were increasingly aware of injection attack vectors, whereby the user can bypass the sensor on the device (such as a camera) and inject code or images into the age assurance system workflow. International standards in this respect are developing (see ISO/IEC AWI 25456 – Biometric Data Injection Attack Detection), but providers were able to demonstrate some resilience to this type of attack.

**D.18.3** While no biometric system can be entirely impervious to attack, the systems reviewed during the Trial demonstrated robust resilience to known presentation threats and showed promising awareness of the next generation of adversarial techniques.

## Vendor Case Study



# unissey

Website

[unissey.com/en](https://unissey.com/en)

Unissey provides real-time facial age estimation using AI models deployed on-device. It enables rapid, privacy-preserving age checks with no biometric retention, supporting threshold-based access decisions across digital services.

## Three Key Facts

1

Neural network-based facial estimation to determine if users meet age thresholds without storing or transmitting biometric data.

2

Supports tuneable thresholds and buffer zones, allowing accurate classification across 13+, 16+, and 18+ age gates.

3

Fully deployed, tested and compatible with common devices, supporting real-world use across regulated and commercial environments.

## Strengths

Demonstrated high accuracy and minimal latency, delivering fast decisions with high true positive rates at 18+.

Built-in fallback mechanisms and human review options make Unissey flexible and robust for edge-case handling in sensitive access scenarios.

## Practice Statement

[ageassurance.com.au/v/uni/#PS](https://ageassurance.com.au/v/uni/#PS)

## Privacy Policy

[ageassurance.com.au/v/uni/#PP](https://ageassurance.com.au/v/uni/#PP)

## Technology Trial Test Report

[ageassurance.com.au/v/uni/#TR](https://ageassurance.com.au/v/uni/#TR)

## Technology Trial Interview

[ageassurance.com.au/v/uni/#VI](https://ageassurance.com.au/v/uni/#VI)

## Summary of Results

Unissey demonstrates that age estimation can now be a stand-alone solution in many age assurance scenarios. The systems high throughput, predictable accuracy, ethical processing practices and alignment with ISO/IEC FDIS 27566-1 make it suitable for deployment in live, regulated environments - from online safety tools to age-gated commerce.

## | Presentation attacks and deepfake manipulation

**D.18.4** Age estimation systems that rely on facial analysis are susceptible to presentation attacks, where an attacker attempts to trick the system using:

- A static image (e.g. printed photo)
- A pre-recorded video
- A digitally altered face (e.g. via deepfake generation tools)
- A re-aged or age-morphed image to appear younger or older

**D.18.5** Providers participating in the Trial addressed these risks using techniques compliant with ISO/IEC 30107-3: Biometric Presentation Attack Detection. These included:

- Liveness detection to confirm that a real human face is present (e.g. via micro-movement, blink detection, reflectivity checks)
- Texture and edge analysis to detect printed or synthetic images
- Motion tracking and depth detection to verify spatial consistency
- Rejection of inconsistent facial landmarks or unrealistic age cues

**D.18.6** Some vendors had also begun training their systems to detect age-specific deepfake alterations, where facial features are selectively modified to alter perceived age without triggering detection. This represents a new category of spoofing risk requiring updated detection algorithms.



Vendor	Liveness Detection	ISO/IEC 30107-3 Alignment	Deepfake/Morph Detection	Injection Attack Mitigation	Notes
<b>Yoti</b>	Yes (see Yoti White Paper)	Level 2 iBeta certified	Yoti has injection detection and anti-spoofing. See Yoti White Papers	Secure SDK, local inference	High standard of detection; no image storage
<b>IDMission</b>	Level 2 iBeta certified	ISO 30107-3 certified	Not specifically stated	TLS 1.3+, secure pipeline	Strong infrastructure; audit trails included
<b>Persona</b>	In development (age morphing aware)	Fully aligned	Under exploration	Secure capture and transmission	Deepfake testing ongoing
<b>Verifymy</b>	Secure SDK, local inference	Claimed alignment	Not disclosed	Secure capture modules	Strong user-side validation
<b>Privately</b>	High standard of detection; no image storage	Partial alignment (PAD-ready)	Not required (non-facial model)	Local-only input, no API vulnerability	Gesture-based approach bypasses facial spoofing risks
<b>Unissey</b>	Yes (passive, certified)	Aligned	Exploring adversarial training	TLS-secured stream verification	Strong edge detection; expanding deepfake countermeasures
<b>Needemand</b>	ISO 30107-3 certified	N/A	N/A	Device-restricted input	Spoofing not relevant due to modality
<b>Luciditi</b>	Not specifically stated	Internally benchmarked	Limited disclosure	Custom secure SDK	Additional validation planned post-Trial

## | Injection attacks and sensor bypass

**D.18.7** In addition to presentation attacks, the evaluation considered injection attack vectors, where the attacker bypasses the camera sensor entirely and feeds altered or synthetic media directly into the estimation system via software injection or API tampering.

**D.18.8** These attacks are particularly relevant in web-based or embedded SDK deployments, where:

- A malicious script may intercept or replace a live camera feed
- A pre-processed video file may be passed as if it were a live user stream
- Adversaries may spoof the system input environment, making it appear that input is coming from a legitimate sensor

**D.18.9** Though standards for mitigating injection attacks are still under development, providers demonstrated preliminary controls, including:

- Secure capture modules that tightly bind the image capture process to the device camera
- Checksum and signature validation to detect tampering
- Basic environment validation (e.g. verifying browser permissions, frame timing consistency)

Vendor	Secure Capture Module	Checksum / Signature Validation	Environment Validation (e.g. frame timing, browser check)	Injection Attack Mitigation Summary
Yoti	✓	✓	✓	Uses anti-spoofing SDK with protection against static and video injection; penetration tested.
Persona	✓	✓	✓	Implements secure API design with ISO/IEC 27001 controls; regular vulnerability scanning.
Verifymy	✓	Partial (encrypted tamper-evident)	✓	Secure SDK with liveness + capture binding; incident response procedures in place.
Privately	✓ Yes (on-device only)	Not applicable (no transmission)	✓	Local inference removes injection vector; robust device control.
Unissey	✓	Partial	✓	Active research on adversarial robustness and frame-level consistency.
Needemand	✓ (Gesture-only input)	Not applicable	✓	Modality avoids injection risk altogether (no facial imagery).
Rigr AI	✓ Yes (API deployable)	Partial	No active liveness detection	No spoof or injection defences in base model; relying party must add safeguards.

## | Project DefAI and the future of testing and certification

**D.18.10** The Trial also recognises the relevance of Project DefAI – a collaborative UK-Swiss Collaborative Project – which is advancing research, benchmarking and certification methods for:

- Deepfake presentation attack detection
- Video and image injection testing
- Security robustness assessments specific to age estimation and age assurance contexts

**D.18.11** Project DefAI seeks to develop standardised, repeatable test environments for evaluating biometric security claims and to ensure that AI-based age systems are resilient not just to conventional spoofing, but also to sophisticated adversarial inputs that target probabilistic vulnerabilities in model design.

- This work complements international standards efforts, including:
- ISO/IEC 30107 Series – for presentation attack detection
- ISO/IEC AWI 25456 – a draft standard under development for biometric data injection attack detection

**D.18.12** The outputs of Project DefAI are expected to influence future certification frameworks, providing a much-needed assurance layer for service providers, regulators and users alike.

*[defaiproject.com](https://defaiproject.com)*

**D.18.13** While age estimation systems inherently face threats from deepfake manipulation and injection attacks, the Trial found that participating providers had taken concrete steps to mitigate these vectors. Using standards-aligned liveness detection, secure input binding and tamper-resistant architectures, providers were able to demonstrate resilience in high-risk use cases. Emerging collaborations like Project DefAI and developing standards such as ISO/IEC AWI 25456 are expected to further mature this space, helping the sector establish verifiable, certifiable defences against next-generation attacks.

### **D.19 Risk of Over-Retention in Age Estimation: Biometric Data and Analytical Residue**

**D.19.1** The Trial had previously found some concerning evidence that service providers were over-anticipating the eventual needs of regulators about providing and over collecting and retaining personal information. This is referred to in the Part C Report. The risk here is less than in the case of age verification, but the risk of biometric sample retention remains.



**Cross Reference:** *Part C - Page 106*

**D.19.2** The Trial observed a generally high standard of data minimisation practices among age estimation providers. Unlike age verification – which may involve document retention, audit trails or verified identity logs – age estimation systems are designed to operate with temporary, instant analysis models. Nonetheless, a residual risk remains: that biometric samples (e.g. facial images) or analytical artefacts could be retained unnecessarily, often in anticipation of potential future demands from regulators or investigators.



**D.19.3** This over-anticipation, while less pronounced than in verification workflows, still warrants attention – especially as age estimation becomes more widely deployed in high-sensitivity contexts, such as online platforms serving children or regulated services like gambling.

### | Biometric sample retention risks

**D.19.4** Most systems reviewed during the Trial were explicitly designed to process biometric inputs temporarily – typically facial images used for real-time age estimation – and delete those inputs immediately after processing. In the best cases, this was enforced through:

- On-device inference SDKs that prevent biometric data from being transmitted.
- In-memory processing, followed by immediate disposal.
- Audit logs that excluded input data, recording only estimation outcomes or transaction IDs.

**D.19.5** However, a minority of providers retained:

- Facial images for “system accuracy auditing” or internal performance tuning.
- Hashes or biometric feature vectors, which – though pseudonymised – still present a theoretical re-identification risk.
- Buffer threshold calculation logs, used to store how close an individual was to a decision threshold (e.g. “estimated age: 17.9” in an 18+ service).

**D.19.6** Even where these practices were justified for quality assurance, they may drift into data over-retention unless governed by strict, time-limited policies and technical controls.

## | Analytical residue: Buffer threshold storage

**D.19.7** Age estimation often involves buffer zones around decision thresholds (e.g. 17.5–18.5 for an 18+ gate). While storing the binary outcome ("Pass" or "Fail") is typically sufficient, some systems were found to retain:

- Estimated age scores (e.g. 17.2).
- Confidence values or probability distributions.
- Metadata about facial features used in the decision.

**D.19.8** When combined with biometric pseudonyms or session IDs, this data could act as a quasi-profile, potentially exposing individuals to behavioural tracking or retrospective scrutiny.

**D.19.9** The risk of over-collection and over-retention in age estimation systems is lower than in age verification, but not negligible. Biometric inputs and analytical outputs must be treated as privacy-sensitive by default, with strong adherence to ISO/IEC FDIS 27566-1 principles of data minimisation, proportionality and explicit purpose limitation.

**D.19.10** Providers are advised to:

- Ensure biometric samples are deleted immediately post-inference.
- Avoid retaining detailed buffer-age calculations unless strictly necessary.
- Establish clear policies on quality assurance vs operational retention.
- Resist speculative data retention for hypothetical regulatory scenarios.

**D.19.11** Doing so will help maintain public trust and ensure compliance as age estimation continues to scale across digital and in-person environments.



Commissioned by the **Australian Department of Infrastructure, Transport, Regional Development, Communications, Sport and the Arts**, the Trial assessed 48 vendors and over 60 distinct technologies across various sectors, including social media, gaming, adult content and online retail. Through lab-based testing, interviews, analysis, school-based trials and mystery shopper evaluations, the Trial investigated how well different solutions could confirm, estimate or imply a user's age in ways that are secure, privacy-preserving and inclusive.

**Can age assurance be done?** The answer – based on thousands of data points, stakeholder interviews and international standards – is **yes, it can.** While no single solution fits all contexts, the Trial found that a wide variety of technologies already meet meaningful thresholds for accuracy, security and privacy when carefully selected and implemented. The report offers a comprehensive evidence base to support regulators, industry leaders and the broader public in shaping a safer, age-appropriate digital environment for all Australians.

@AgeCheckCert



AVID Certification Services Ltd t/a Age  
Check Certification Scheme, registered in  
England 14865982 • Unit 321 Broadstone  
Mill, Broadstone Road, Stockport, SK5 7DL,  
United Kingdom • ABN 76 211 462 157

