

April 2026

The Association for Computing Machinery (ACM) is the world's longest-established professional society of individuals involved in all aspects of Computing. It annually bestows the ACM A.M. Turing Award, often popularly referred to as the "Nobel Prize of Computing." ACM's Europe Technology Policy Committee ("Europe TPC") is charged with, and committed to, providing policymakers and the public with sound technical information to support sound public policymaking. Europe TPC has previously responded to European Union stakeholder consultations in the context of the AI Act, the Data Act, the Digital Services Act, the Digital Citizen Principles, and the Cyber Resilience Act, amongst others. ACM and Europe TPC are non-profit, non-political, and non-lobbying organizations.

We are pleased to provide comments on the proposed Dutch Technical Agreement on profiling algorithms. This consultation, sponsored by the Royal Netherlands Standardization Institute, provides timely and impactful recommendations at both national and international levels. The draft document and background are available [here](#).

Section 3.2

Comment:

Section 3.2 does not distinguish between different levels of data abstraction in profiling algorithms. The literature identifies four main categories: provided, observed, derived, and inferred data (Abrams, 2014). This distinction is critical because derived and particularly inferred data introduce additional error sources beyond input data quality. Inferred data relies on statistical correlations that may be inaccurate, discriminatory, or contextually inappropriate for individual cases. Without acknowledging these abstraction levels, the standard inadequately addresses scenarios where faulty inferences cannot be effectively challenged due to the lack of transparency in derivation processes.

Reference: Abrams, M. (2014). The origins of personal data and its implications for governance. *Available at SSRN 2510927*.

Alteration:

Add after the current Section 3.2 content:

"Profiling algorithms may operate on different data abstraction levels: provided (directly supplied), observed (recorded from behavior), derived (computed deterministically), and inferred (predicted statistically) [Abrams, 2014]. Derived and inferred data introduce additional error sources. Organizations should: (a) document which data categories are used in profiling, (b) separately assess accuracy of derived and inferred attributes, and (c) enable individuals to challenge inferred attributes meaningfully."

Section 5.1

Comment

The framing of "objective justification" for differential treatment in sections 5.1.2-5.1.4 is contested. The minority position from Amnesty International, Bits of Freedom, and Control Alt Delete in Annex A demonstrates fundamental disagreement on whether objective justification adequately protects against discriminatory outcomes in algorithmic systems. This minority position raises concerns that the current framing may legitimize discrimination under technical rationality while failing to account for structural inequalities and power imbalances. The standard should acknowledge this ongoing debate and provide guidance that does not presume consensus on what constitutes acceptable differential treatment.

Alteration

Add note to Section 5.1.2:

"NOTE: The concept of 'objective justification' for differential treatment is subject to ongoing debate, as reflected in the minority position in Annex A. Organizations should be aware that technical or statistical justifications may not be sufficient to address concerns about structural discrimination or societal fairness."

Section 5.2.4

Comment:

Section 5.2.4 defines effectiveness as the evaluation of the decision process, including the profiling algorithm, but does not address multi-stakeholder contexts where different parties have conflicting values and objectives. In platform ecosystems, buyers, sellers, and platform operators may have fundamentally different criteria for what constitutes "effective" profiling (Burke et al., 2016). The current framing assumes a single measure of effectiveness, which may obscure whose interests are prioritised. Additionally, transparency regarding algorithmic outcomes and their role in decision processes is essential for meaningful effectiveness evaluation (Mittelstadt, 2016; Sonboli et al., 2021), yet it is not explicitly required in this section.

References:

- Burke, R. D., Abdollahpouri, H., Mobasher, B., & Gupta, T. (2016). Towards multi-stakeholder utility evaluation of recommender systems. *UMAP (Extended Proceedings)*, 750.

- Mittelstadt, B. (2016). Auditing for transparency in content personalization systems. *International Journal of Communication*.
- Sonboli, N., Smith, J. J., Cabral Berenfus, F., Burke, R., & Fiesler, C. (2021). Fairness and transparency in recommendation: The users' perspective. *Proceedings of the 29th ACM Conference on User Modeling, Adaptation and Personalization*, 274-279.

Alteration:

Revise Section 5.2.4 to add:

"Effectiveness evaluation should acknowledge that different stakeholders (e.g., service users, service providers, platform operators) may have different and potentially conflicting values regarding what constitutes effective profiling [Burke et al., 2016]. Organizations should: (a) identify relevant stakeholder groups, (b) document whose interests are prioritized in effectiveness measures, and (c) ensure transparency regarding algorithmic outcomes and their role in decision processes [Mittelstadt, 2016; Sonboli et al., 2021]."

Section 7.3

Comment:

The statement "Een datagebaseerd algoritme bestaat uit een gekozen model getraind op data" does not explicitly acknowledge that data-based models, particularly AI/ML models, operate partially as black boxes and may produce incorrect outcomes. This is a critical omission, as it affects risk assessment and explainability requirements.

Additionally, Section 7.3's treatment of the relationship between model complexity and explainability is oversimplified. The standard appears to suggest a straightforward trade-off, but this relationship is technically subtle: some complex models may be more interpretable than simpler ones in specific contexts, post-hoc explanations may not faithfully represent model behavior, and explainability requirements differ across use cases and stakeholders. The current framing risks promoting misleading guidance that equates simple models with adequate explainability or assumes all complex models are equally opaque.

Alteration:

Revise the opening sentence of Section 7.3:

"A data-based algorithm involves a chosen model trained on data. Such models, particularly AI/ML models, operate partially as black boxes and may produce incorrect outcomes even when properly trained. Organizations should assess model opacity and error risk when determining explainability requirements."

Add clarifying note:

"NOTE: The relationship between model complexity and explainability is context-dependent. Model choice should balance performance requirements, explainability needs for different stakeholders, and the specific risks associated with the application domain. Post-hoc explanations of complex models do not necessarily provide faithful representations of model behavior."

Sections 7.4.2 and 8.3.2

Comment:

Sections 7.4.2 and 8.3.2 discuss separating training and test datasets but do not mention cross-validation techniques, which are often superior to simple train-test splits, particularly for limited datasets. K-fold cross-validation provides more robust performance estimates by using all available data for both training and testing through multiple iterations, reducing variance in performance metrics and better detecting overfitting. For smaller datasets or when maximizing data utility is critical, cross-validation is standard practice in machine learning. The current text may lead organizations to adopt suboptimal validation strategies.

Reference: Cross-validation (statistics). [https://en.wikipedia.org/wiki/Cross-validation_\(statistics\)](https://en.wikipedia.org/wiki/Cross-validation_(statistics))

Alteration:

Section 7.4.2 - Add:

Organizations should consider using k-fold cross-validation techniques when dividing datasets, particularly when data is limited. Cross-validation provides more robust estimates of model performance than simple train-test splits by iteratively using different data subsets for training and validation."

Section 8.3.2 - Add:

"For robust model evaluation, organizations should consider k-fold cross-validation as an alternative or complement to simple train-test splits. Cross-validation uses all available data for both training and validation through multiple iterations, providing more reliable performance estimates. See: [https://en.wikipedia.org/wiki/Cross-validation_\(statistics\)](https://en.wikipedia.org/wiki/Cross-validation_(statistics))"

Section: 8.6

Comment:

Section 8.6 addresses independent assessment and audit of algorithmic systems, but does not acknowledge that audit methodology for algorithmic systems remains underdeveloped globally. There is currently no consensus on standardized audit procedures, appropriate expertise requirements for auditors, or effective mechanisms for verifying algorithmic behavior in production environments. Key challenges include access to proprietary systems, reproducibility of findings, auditing dynamic models that change over time, and assessing emergent behaviors in complex systems. The current text may create false confidence that established audit practices exist when, in fact, organizations and auditors are operating in a methodologically immature field. This risks superficial compliance exercises rather than meaningful independent assessment.

Alteration:

Add to Section 8.6:

"NOTE: Audit methodology for algorithmic systems is still developing globally. Organizations and auditors should be aware that standardized audit procedures, auditor qualification requirements, and verification mechanisms are not yet well-established. Independent assessment should be understood as an evolving practice requiring ongoing methodological development rather than application of fully mature audit frameworks."

Section Annex A

Comment

Annex A presents the minority position from Amnesty International, Bits of Freedom, and Control Alt Delete, but does not provide explicit engagement with the fundamental tension it reveals in applying EU non-discrimination law to algorithmic decision-making. The disagreement on "objective justification" for differential treatment is not merely a difference of opinion among stakeholders, but reflects unresolved questions about whether frameworks designed for human decision-making adequately address algorithmic systems that operate at scale, with opacity, and with the potential to systematize discrimination. The minority position challenges whether technical or statistical rationales constitute legitimate objective justification when they may perpetuate or amplify structural inequalities. Without explicit engagement with this tension, the standard risks appearing to resolve by majority vote what is actually an ongoing legal and ethical debate requiring substantive analysis.

Alteration:

Add an introductory paragraph to Annex A:



"The minority position presented here reflects a fundamental and unresolved tension in applying EU non-discrimination law to algorithmic decision-making. The disagreement centers on whether 'objective justification' frameworks developed for human decision-making adequately address algorithmic systems that operate at scale, with limited transparency, and with potential to systematize discrimination. This standard does not resolve this tension, which requires ongoing legal, technical, and ethical analysis."