Original Article

law& pass International Journal of Law, Public Administration and Social Studies https://lawpass.org

https://doi.org/10.47353/lawpass.v3i4.131

From Soft-Law AI Ethics to Risk-Based Public Sector AI Procurement: Recalibrating Malaysia's Post-PDPA Governance Model

Rozaini Binti Rosli
Lincoln University College, Malaysia
*Corresponding author: rozaini@lincoln.edu.my

Abstract

Public-sector artificial intelligence (AI) is increasingly acquired through procurement rather than developed entirely in-house. This shifts a central governance question from whether public institutions endorse ethical principles to whether those principles become enforceable duties across the system life cycle. Malaysia has introduced national AI governance and ethics guidance and public-sector AI adoption guidance, while the Personal Data Protection (Amendment) Act 2024 recalibrates parts of the data-protection framework. Yet privacy compliance and high-level ethical principles alone cannot fully govern the acquisition and operation of AI used in public administration. This article asks how Malaysia can translate soft-law principles and existing legal duties into procurement controls proportionate to risk. Using doctrinal legal analysis and functional comparison with the European Union's AI Act, the OECD AI Principles, UNESCO's Recommendation on the Ethics of AI, Canada's Directive on Automated Decision-Making, and the NIST AI Risk Management Framework, it develops a normative Risk-to-Contract Accountability Framework (RCAF). The framework links risk classification to six controls: pre-award disclosure; impact assessment; contractual auditability; contestability and human review; incident reporting and material-change control; and continuing monitoring, remedies, and exit. The contribution is not a claim that AI procurement has never been studied, but a Malaysia-specific legal design connecting governance principles to the contract as a practical instrument of public accountability. The article argues that procurement should allocate responsibilities before deployment, preserve the authority's ability to scrutinise supplier claims, and prevent contractual arrangements from insulating consequential decisions from review.

Keywords: artificial intelligence; public procurement; data protection; administrative law; algorithmic accountability; risk-based regulation.

Introduction

Artificial intelligence is often discussed as a problem of innovation policy, data protection, or ethics. In public administration, however, much of the practical governance problem arises when an agency chooses, purchases, configures, and operates a system supplied by another organisation. Public bodies may procure automated eligibility screening, fraud detection,

Received: August 3, 2026 | **Revised:** Sept 24, 2026 | **Accepted:** Oct 6, 2026 | **Publication:** Oct 31, 2026

Copyright (c) 2026 Rozaini Binti Rosli
Creative Commons Attribution 4.0 International License.

document classification, predictive maintenance, citizen-service chatbots, decision-support tools, or systems that rank applications and allocate scarce resources. Even where a model cannot make the final decision autonomously, its output can shape the information available to officials and influence the result. Procurement therefore determines not only which technology is acquired but also which evidence the government receives, which rights it retains, and which risks it can realistically control.

This relationship creates a distinctive accountability problem. The public authority remains responsible for exercising public power lawfully, but technical knowledge and system documentation may be concentrated in the supplier. Model updates may occur after award; training data and performance evidence may be proprietary; and contracts may narrowly define access to logs, explanations, testing environments, and audit rights. If safeguards are considered only after purchase, agencies may discover that the contract does not permit the scrutiny needed to investigate errors or explain adverse outcomes. Procurement is therefore part of AI governance, not merely an administrative step that follows it.

Research on algorithmic accountability shows that accountability requires more than a general commitment to fairness or transparency. It requires identifiable actors, access to relevant evidence, procedures for challenge, and consequences when systems fail (Kroll et al., 2017; Raji et al., 2020). Public-sector machine-learning scholarship also highlights tensions between technical capacity, bureaucratic discretion, and public values (Veale & Brass, 2019; Wirtz et al., 2019). Procurement makes these concerns concrete: principles must become requirements bidders can understand, evaluators can assess, contract managers can enforce, and affected people can invoke.

Malaysia's AI governance environment has developed through general legislation, national policy, and non-binding guidance. The National Guidelines on AI Governance and Ethics (AIGE), introduced in 2024, articulate principles intended to guide responsible AI development and use. Public-sector adoption guidance issued by Jabatan Digital Negara (JDN) adds an operational orientation for government agencies, including risk differentiation and principles such as accountability, transparency, fairness, privacy and security, inclusivity, reliability, and sustainability. These instruments are important because they articulate expectations before a comprehensive AI-specific statutory regime is in place. Their practical effect, however, depends on how agencies incorporate them into decisions, workflows, and enforceable instruments.

The Personal Data Protection Act 2010 (PDPA) and Personal Data Protection (Amendment) Act 2024 remain central where an AI system processes personal data within the statute's scope. The amendment is a significant recalibration, but it should be described precisely: commencement and application depend on relevant commencement instruments, and the PDPA is not a general AI statute. Data-protection compliance is necessary in many deployments, but it cannot alone answer whether an AI tool is suitable for a public function, whether its outputs are reliable enough for a particular use, how an agency should justify reliance on a score, or what contractual access is needed to investigate a disputed outcome.

This article therefore treats Malaysia's post-amendment data-protection setting as part of a wider governance model, rather than assuming that the PDPA amendment itself creates a

complete AI regulatory framework. A procurement-centred approach can connect data protection to legality, proportionality, public reason-giving, procedural fairness, and institutional responsibility. It can also address a weakness in soft-law governance: principles can be widely endorsed while remaining unevenly operationalised in the contracts through which public institutions acquire systems.

The international literature on algorithmic regulation, automated administration, and AI accountability is extensive. It questions whether opaque systems can be reconciled with due process, how bias can emerge from seemingly neutral abstractions, and how institutions should evaluate claims made by developers and vendors (Mittelstadt et al., 2016; Selbst et al., 2019; Yeung, 2018). Work on public-sector AI identifies organisational capacity, legitimacy, and public value as conditions for responsible adoption (Culot et al., 2026; Wirtz et al., 2019). More recent procurement scholarship directly addresses responsible AI acquisition and the incorporation of sustainability and accountability into public purchasing (Lau et al., 2025; Sciarrone & Calabrese, 2026). It would therefore be inaccurate to claim that AI procurement has never been studied.

The more defensible gap is narrower and legally actionable. General AI ethics frameworks describe values; technical risk frameworks organise risk-management activities; and procurement research identifies operational and legitimacy challenges. A remaining design question is how to connect a risk classification to coherent contractual duties suitable for public-sector acquisition in Malaysia, while recognising the limits of current data-protection law and the non-binding character of much AI guidance. This article proposes a normative framework for that connection. It does not claim empirical proof that the proposed controls reduce harm; it develops a legally reasoned architecture that future empirical work can test.

The article asks: (1) what can Malaysia's current data-protection and AI-governance instruments reasonably achieve, and where do procurement-specific gaps remain? (2) How can risk-based principles become procurement and contract-management duties without treating every AI system as equally dangerous? (3) What lessons can be adapted from international frameworks without assuming their legal structures can simply be transplanted?

The contribution is threefold. Conceptually, procurement is treated as a governance gateway between public-law responsibility and privately supplied technical systems. Doctrinally, the article distinguishes the legal function of the PDPA from the policy and administrative function of AI guidance. Practically, it proposes the Risk-to-Contract Accountability Framework (RCAF), linking system risk to pre-award evidence, impact assessment, auditability, contestability, incident and change reporting, and post-award remedies. A low-risk internal drafting assistant should not automatically attract the same controls as a system materially affecting benefits, enforcement, education, health, or other important public services.

The analysis is normative and legal, not an empirical assessment of Malaysian procurement practice. It does not assert that all Malaysian public-sector contracts omit AI safeguards, nor that every AI use falls within the PDPA. The comparison focuses on the EU AI Act, OECD AI Principles, UNESCO's Recommendation on the Ethics of AI, Canada's Directive on Automated Decision-Making, and the NIST AI Risk Management Framework. Section 2 explains the

method; Section 3 evaluates the governance gap and develops the RCAF; Section 4 concludes with implementation priorities and limitations.

Method

This article uses doctrinal legal research combined with functional comparative analysis. The doctrinal component examines the legal and policy roles of Malaysia's PDPA 2010, the Personal Data Protection (Amendment) Act 2024, national AI governance and ethics guidance, and public-sector AI adoption guidance. It distinguishes binding legal obligations from policy commitments and administrative guidance, and does not treat a policy instrument as legislation merely because it is issued by a government body.

The comparative component evaluates five reference models: the EU AI Act as a risk-based statutory regime; the OECD AI Principles and UNESCO Recommendation as international normative instruments; Canada's Directive on Automated Decision-Making as a public-administration control; and the NIST AI Risk Management Framework as a voluntary operational risk tool. Each is assessed against six questions: how risk is classified; what evidence is required before deployment; how responsibility is allocated; what documentation and audit access are preserved; how affected people can contest outcomes; and how systems are monitored or changed after deployment.

Sources comprise legislation and official guidance, peer-reviewed scholarship, and academic books. The framework is developed through synthesis rather than interviews, surveys, or quantitative analysis. Accordingly, the article makes no empirical claims about the frequency of procurement failures, agency compliance, or measured effectiveness of the proposed controls. Legal instruments and commencement status should be checked against the latest official publications before submission.

Soft law can guide conduct but cannot substitute for allocated duties

AI ethics instruments perform valuable agenda-setting and interpretive functions. They identify concerns that might otherwise be omitted from technical specifications: privacy, fairness, transparency, reliability, inclusivity, human oversight, and accountability. Flexibility can be an advantage in a changing technical field. Yet a principle such as transparency does not specify what a bidder must disclose, which records an agency must retain, who may inspect them, or what happens when a supplier refuses access.

The issue is not that soft law is useless; it is that soft law needs an institutional mechanism to influence conduct. In procurement, that mechanism can combine tender requirements, evaluation criteria, contractual warranties, service-level obligations, audit rights, reporting duties, and termination or remediation provisions. Without translation, an agency may endorse a policy but lack the contractual rights needed to implement it against a vendor. The RCAF is designed to bridge this implementation gap.

The EU AI Act demonstrates a formal risk-based model, imposing obligations that vary by system category and actor. It is not directly interchangeable with Malaysian law, and its

provisions have phased application dates. OECD and UNESCO instruments offer broader principles, while NIST supplies a voluntary structure for governing, mapping, measuring, and managing risk. Canada's directive is useful because it connects automated decision-making with government administrative processes. These models suggest that principles, legal duties, and operational controls are complementary.

The PDPA is important but does not answer every AI-governance question

The PDPA is a key legal anchor where personal data is processed in commercial transactions within its scope. Procurement documents should identify the parties' data-protection roles, permitted purposes, security measures, retention arrangements, breach-notification responsibilities, sub-processing conditions, and cooperation duties. Agencies should consider whether data used to train, tune, test, or operate a system is lawfully obtained and whether its proposed use is compatible with applicable legal requirements.

However, a system can raise public-law concerns even when it processes little or no personal data. An AI tool prioritising infrastructure repairs may influence public-resource allocation; a model generating legal or policy summaries may introduce errors into official advice; and a system classifying tenders may create procedural-fairness risks. Conversely, a system processing personal data may be relatively low impact if it supports a reversible task with robust human review. Risk cannot be inferred from personal-data processing alone.

The PDPA should be integrated into, rather than equated with, a wider assessment. Agencies should separately ask whether the use is lawful and within their authority; what public function it influences; who may be adversely affected; what errors would mean; whether outputs can be explained, reviewed, and corrected; and whether the contract permits investigation. This connects data protection with legality, procedural fairness, reliability, and accountability.

Risk classification should trigger differentiated controls

A risk-based model should avoid under-regulation and over-regulation. Treating every AI acquisition as ordinary software procurement may overlook effects on rights or access to public services. Imposing burdensome assessment and audit procedures on every minor productivity tool may waste scarce capacity. Proportionality requires considering purpose, degree of influence on decisions, vulnerability of affected groups, reversibility of error, scale, data sensitivity, and the ability of a human official to detect and correct problems.

Malaysia's public-sector AI adoption guidance provides a basis for risk differentiation. The RCAF does not replace official classifications or claim statutory status for its own labels. It proposes that agencies use applicable official categories as a starting point and map them to a procurement control schedule. Consequential uses should attract stronger evidence, impact assessment, human oversight, audit access, and monitoring. Lower-impact and readily reversible uses may justify lighter controls.

Classification should be revisited when purpose, model, data, user population, or operational context changes. A tool initially acquired for internal summarisation may later be

repurposed to score staff or screen citizens. That change can alter the risk profile even if the supplier describes it as a routine update. Contracts should define material changes and require notification, reassessment, and approval before materially different use.

The Risk-to-Contract Accountability Framework (RCAF)

The RCAF links risk classification to six control groups across the procurement lifecycle. Its central proposition is that risk classification should not remain a label in an assessment form: it should determine what evidence is requested, what contract rights are retained, and what remedies are available. It is a normative proposal subject to applicable Malaysian procurement rules, agency mandates, and system-specific context.

Table 1. Risk-to-Contract Accountability Framework (RCAF). Source: author's normative synthesis.

Control group	Core procurement / contract requirement	Accountability purpose
1. Pre-award disclosure	Require intended purpose, system boundaries, limitations, data provenance, validation evidence, model/version information, and foreseeable failure modes proportionate to risk.	Enables informed evaluation and limits reliance on unsupported vendor claims.
2. Impact assessment	Document necessity, alternatives, affected groups, foreseeable harms, legal basis, error consequences, mitigation, and reasons for selection.	Connects acquisition to legality, proportionality, and public-value justification.
3. Auditability and records	Secure access to relevant logs, technical and user documentation, test results, change history, and independent audit rights, with appropriate safeguards.	Preserves the authority's ability to verify performance and investigate disputes.
4. Contestability and human review	Specify notice, explanation, escalation, human review, correction, and appeal routes where outputs materially affect people.	Prevents technical outputs from becoming practically unchallengeable.
5. Incident and change control	Require prompt reporting of material incidents, performance degradation,	Prevents silent drift and prompts reassessment when risk changes.

Control group	Core procurement / contract requirement	Accountability purpose
	security events, model changes, and changes in purpose or data.	
6. Monitoring, remedies, and exit	Set performance thresholds, periodic reviews, remediation deadlines, suspension rights, data return/deletion, transition assistance, and termination for serious failures.	Makes accountability continuous and gives the agency enforceable options.

Pre-award disclosure and evidence

A procurement authority cannot evaluate risk if it receives only a demonstration and a general assurance that a product is ethical or secure. Tender documents should require bidders to describe intended use, technical limits, system dependencies, data inputs, expected outputs, known performance constraints, and conditions under which the system should not be used. For higher-impact applications, evidence should include testing methods, performance results in relevant contexts, known failure modes, human-factors assumptions, and how outputs can be challenged or corrected.

Disclosure must be proportionate. Agencies need not demand public release of source code in every case, and intellectual-property or cybersecurity interests may justify controlled access. But commercial confidentiality should not prevent the government from verifying a system used to exercise public functions. Contracts can distinguish public-facing transparency from confidential disclosure to authorised auditors, while requiring record preservation and cooperation with investigations.

Evaluation criteria should reward evidence, not promises. Claims of high accuracy should be assessed against the actual task, operating conditions, error costs, and validation quality. For consequential decisions, false-positive and false-negative consequences matter, not merely aggregate accuracy. A high average score can conceal shortcomings for a smaller group or specific context, consistent with warnings that fairness claims fail when abstract metrics are detached from sociotechnical conditions (Selbst et al., 2019).

Impact assessment as a reason-giving instrument

An AI impact assessment should help an agency decide whether a system is necessary, whether a non-AI alternative would be safer, and whether expected public benefit justifies residual risk. It should identify the public function, legal authority, affected people, foreseeable failure modes, and mitigation options. It should record uncertainty rather than imply that all risk can be eliminated.

For high-impact uses, the assessment should explain the role of human officials. A human-in-the-loop label is insufficient if staff must approve outputs under time pressure, lack the competence to question them, or routinely defer to the system. Meaningful oversight requires training, access to relevant information, authority to depart from recommendations, and organisational support. Procurement can require training materials, override functionality, escalation protocols, and periodic review of oversight in practice.

The assessment should be updated when deployment conditions change. A model may behave differently with a new agency, population, dataset, or purpose. Contracts should allocate responsibility for reassessment and define which changes require prior approval. A supplier should not be able to materially alter the risk profile while leaving the authority contractually blind.

Auditability, contestability, and procedural fairness

Accountability requires the ability to reconstruct what happened. Contracts should address logs, version records, data lineage where applicable, configuration history, user actions, output timestamps, and the evidence available to the decision-maker. Recordkeeping must be designed in advance; it is often impossible to reconstruct a decision if relevant information was never retained. Access should be secure and role-based, but the authority must retain enough information to discharge its legal duties.

Contestability is distinct from transparency. A person affected by an AI-assisted decision may need to know that a system materially influenced the outcome, understand the basis of the decision to an appropriate degree, submit correcting information, and obtain review by a competent human decision-maker. The appropriate explanation depends on context and legitimate constraints, including privacy, security, and intellectual property. Yet merely repeating a model output is not an effective remedy: review must be capable of changing the result.

Contracts should avoid making a supplier's refusal to disclose relevant information the practical end of an individual's challenge. Agencies may need confidential access, independent expert review, or a process for resolving disputes over proprietary information. The objective is not unlimited disclosure but sufficient transparency for lawful administration and effective review (Kroll et al., 2017; Pasquale, 2015; Wachter et al., 2021).

Monitoring, incidents, and model change

Risk may change after procurement because data distributions shift, software is updated, users develop new practices, or deployment expands. Contracts should establish monitoring responsibilities, indicators, review intervals, and escalation thresholds. Indicators should reflect the system's purpose: error rates, complaint patterns, override rates, service delays, performance across relevant groups, security incidents, or other measures appropriate to the use case. One generic service-level metric is unlikely to capture every public harm.

Incident reporting should cover more than cybersecurity breaches. It may include material accuracy failures, bias indicators, unexplained degradation, serious complaints, unauthorised use,

and failures of human review. Contracts should set reporting deadlines, preserve evidence, allocate investigation responsibilities, and require corrective action. For serious incidents, agencies should be able to suspend use pending investigation. Contractual reporting must be coordinated with, not assumed to replace, applicable statutory reporting duties.

Material-change control is especially important for AI-as-a-service. Contracts should define changes requiring notice or approval, including model versions, training or reference data, integrations, purpose, user population, and decision thresholds. Agencies should be able to require re-testing or reassessment. If adequate evidence is unavailable, the authority should be able to restrict the new version, continue with a validated version where feasible, or exit.

Comparative lessons and limits of transplantation

The EU AI Act illustrates how risk categories can trigger differentiated legal obligations, but Malaysia should not copy categories without considering local procurement rules, institutional capacity, and the statutory environment. Its obligations also have phased application dates, so it should not be described as uniformly applicable from one date. Its useful lesson is architectural: risk can determine the intensity of governance, and responsibility should be allocated across actors rather than placed entirely on the end user.

The OECD AI Principles offer a flexible vocabulary around inclusive growth, human-centred values, transparency, robustness, and accountability. UNESCO's Recommendation expands the lens to human rights, dignity, diversity, environmental concerns, and governance. These instruments can shape tender principles and evaluation criteria, but do not themselves provide the detailed contractual machinery needed in a particular acquisition. The RCAF translates broad principles into evidence, rights, procedures, and remedies.

Canada's Directive on Automated Decision-Making is relevant because it connects automated systems with public-sector decisions and differentiated impact assessment. NIST's AI Risk Management Framework helps organise governance, context mapping, measurement, and risk management over time. Neither should be represented as binding Malaysian law. Their role is comparative: they offer design patterns that can be adapted within Malaysia's own legal and administrative arrangements.

Recent procurement scholarship also cautions against viewing responsible acquisition solely as a technical task. Public value, organisational capability, legitimacy, data rights, audit trails, and the ability to challenge supplier claims are interdependent concerns (Culot et al., 2026; Sciarrone & Calabrese, 2026). The RCAF emphasises the contract as the connecting instrument. Risk classification matters only if it changes what an agency asks for, what it can verify, and what it can do when a system fails.

Implementation barriers and anti-box-ticking safeguards

The framework faces practical constraints. Procurement teams may lack AI expertise; agencies may be under pressure to deploy quickly; suppliers may resist disclosure; and fragmented responsibility may leave no single office accountable for ongoing performance.

Smaller agencies may be unable to conduct sophisticated testing. These limitations argue for shared templates, central guidance, training, and access to independent technical expertise—not abandonment of risk-based controls.

A central procurement or digital-government authority could maintain model tender clauses, impact-assessment templates, risk-classification guidance, and standard definitions for material changes and incidents. Agencies should still own context-specific decisions because the same product can create different risks in different uses. Shared tools can reduce transaction costs while permitting stronger controls where the stakes justify them.

Formal compliance can become a substitute for substantive accountability. A completed checklist, an audit clause that cannot realistically be exercised, or a human-review provision without time and authority may create the appearance of governance while leaving risks untouched. Agencies should test whether a reviewer can reconstruct a sample decision, whether a complaint can lead to correction, whether material changes are reported, and whether the authority can suspend the system when necessary.

Bounded novelty and future research

The novelty claim is deliberately bounded. This article does not claim to be the first study of AI procurement, nor that prior Malaysian work has never proposed risk-based safeguards. It advances a specific legal synthesis: a risk-to-contract mapping for Malaysia's combination of general data-protection legislation, national AI guidance, and public-sector procurement. Its contribution is the integration of six control groups into a lifecycle architecture linking pre-award evidence to contestability, change control, remedies, and exit.

The framework yields propositions for future empirical research. Agencies using risk-linked clauses may obtain more relevant supplier evidence than agencies relying on general ethics statements. The effectiveness of audit and contestability clauses may depend on staff capability, recordkeeping, and the practical independence of review. Standard templates and shared technical expertise may reduce costs, particularly for lower-capacity agencies. Finally, meaningful accountability may depend less on the existence of a clause alone than on the combination of contractual rights, organisational capacity, and willingness to use remedies. These propositions require interviews, procurement-document analysis, case studies, or evaluation of deployed systems; the present article has not conducted those studies.

Conclusion

Malaysia's emerging AI governance model combines data-protection law, national ethical guidance, and public-sector adoption guidance. These instruments provide an important foundation, but do not automatically ensure that AI risks are identified, allocated, and controlled through every government acquisition. The PDPA remains central where personal data is processed within its scope, yet responsible public-sector AI procurement must also address legality, reliability, public value, procedural fairness, auditability, and the ability to challenge consequential outcomes.

This article has proposed the Risk-to-Contract Accountability Framework as a normative bridge between risk classification and enforceable procurement practice. It comprises six connected controls: pre-award disclosure; impact assessment; auditability and recordkeeping; contestability and meaningful human review; incident reporting and material-change control; and continuing monitoring, remedies, and exit. The central proposition is that risk classification should alter the evidence demanded from bidders, the rights retained by the public authority, and the actions available when a system fails or changes.

Implementation should begin with standard tender clauses and assessment templates, differentiated requirements for consequential uses, staff capability-building, and periodic checks that safeguards work in practice. The framework's effectiveness remains to be tested empirically, and the proposed controls must be aligned with applicable Malaysian procurement requirements and the current status of laws and guidance. Properly adapted, procurement can become more than a purchasing procedure: it can be a practical site where public accountability is secured before AI systems shape public decisions.

References

Ada Lovelace Institute. (2021). Regulating AI in the UK: A framework for AI regulation. https://www.adalovelaceinstitute.org/report/regulating-ai-in-the-uk/

Benjamin, R. (2019). Race after technology: Abolitionist tools for the new Jim code. Polity.

Busuioc, M. (2021). Accountable artificial intelligence: Holding algorithms to account. Public Administration Review, 81(5), 825–836. https://doi.org/10.1111/puar.13293

Canada, Treasury Board of Canada Secretariat. (2019, as amended). Directive on Automated Decision-Making. https://www.tbs-sct.canada.ca/pol/doc-eng.aspx?id=32592

Calo, R. (2015). Robotics and the lessons of cyberlaw. California Law Review, 103(3), 513–563.

Citron, D. K., & Pasquale, F. (2014). The scored society: Due process for automated predictions. Washington Law Review, 89(1), 1–33.

Culot, G., Podrecca, M., Nassimbeni, G., & Sartor, M. (2026). Artificial intelligence in public procurement: Aligning operational capacity, legitimacy, and public value. Journal of Business Logistics, 47(4), e70088. https://doi.org/10.1111/jbl.70088

European Parliament & Council of the European Union. (2024). Regulation (EU) 2024/1689 laying down harmonised rules on artificial intelligence (Artificial Intelligence Act). https://eur-lex.europa.eu/eli/reg/2024/1689/oj

Jabatan Digital Negara. (n.d.). Guidelines for the adoption of AI in the public sector. Government of Malaysia. https://www.jdn.gov.my/en/guidelines-for-the-adoption-of-ai-in-the-public-sector/

Kroll, J. A., Huey, J., Barocas, S., Felten, E. W., Reidenberg, J. R., Robinson, D. G., & Yu, H. (2017). Accountable algorithms. University of Pennsylvania Law Review, 165(3), 633–705.

Lau, D. K., Samy, G. N., Abdul Rahim, F., Selvanathan, M., Maarop, N., Krishnan, M. R., & Perumal, S. (2025). Responsible procurement of AI applications: A risk-based framework for Malaysian government agencies. https://doi.org/10.24191/mij.v6i2.9172.

Lodge, M., & Wegrich, K. (Eds.). (2014). The problem-solving capacity of the modern state: Governance challenges and administrative capacities. Oxford University Press.

Malaysia. (2010). Personal Data Protection Act 2010 (Act 709). Laws of Malaysia.

Malaysia. (2024). Personal Data Protection (Amendment) Act 2024 (Act A1727). https://www.pdp.gov.my/ppdpv1/en/akta/personal-data-protection-amendment-act-2024/

Maaz, Z. N., et al. (2026). Explainable artificial intelligence (XAI) model for transparent and trustworthy tender evaluation of construction projects. International Journal of Technology, 17(3), 953–972. https://doi.org/10.14716/ijtech.v17i3.8288

Mittelstadt, B. D., Allo, P., Taddeo, M., Wachter, S., & Floridi, L. (2016). The ethics of algorithms: Mapping the debate. Big Data & Society, 3(2). https://doi.org/10.1177/2053951716679679

National Institute of Standards and Technology. (2023). Artificial intelligence risk management framework (AI RMF 1.0) (NIST AI 100-1). https://doi.org/10.6028/NIST.AI.100-1

National Institute of Standards and Technology. (2024). Artificial intelligence risk management framework: Generative artificial intelligence profile (NIST AI 600-1). https://doi.org/10.6028/NIST.AI.600-1

Nik Mohammad Wafiy Azmi, M. F. H. M. Z., Tarudin, M. D. A., & Shahidan, N. F. (2025). AI and public governance in Malaysia: A literature review on ethical, legal, and operational challenges. Journal of Mathematics and Computing Science, 11(1). https://doi.org/10.24191/jmcs.v11i1.8090

OECD. (2019, updated 2024). Recommendation of the Council on Artificial Intelligence (OECD/LEGAL/0449). https://legalinstruments.oecd.org/en/instruments/OECD-LEGAL-0449

OECD. (2024). AI, data governance and privacy: Synergies and areas of international co-operation. OECD Publishing.

Pasquale, F. (2015). The black box society: The secret algorithms that control money and information. Harvard University Press.

Raji, I. D., Smart, A., White, R. N., Mitchell, M., Gebru, T., Hutchinson, B., Smith-Loud, J., Theron, D., & Barnes, P. (2020). Closing the AI accountability gap: Defining an end-to-end framework for internal algorithmic auditing. Proceedings of the 2020 Conference on Fairness, Accountability, and Transparency, 33–44. https://doi.org/10.1145/3351095.3372873

Sciarrone, F., & Calabrese, A. (2026). Governing AI in sustainable public procurement: Towards an ESG-oriented conceptual architecture. VINE Journal of Information and Knowledge Management Systems. https://doi.org/10.1108/VJIKMS-07-2025-0309

Selbst, A. D., boyd, d., Friedler, S. A., Venkatasubramanian, S., & Vertesi, J. (2019). Fairness and abstraction in sociotechnical systems. Proceedings of the 2019 Conference on Fairness, Accountability, and Transparency, 59–68. https://doi.org/10.1145/3287560.3287598

Tabassi, E. (2023). Artificial intelligence risk management framework (AI RMF 1.0). National Institute of Standards and Technology. https://doi.org/10.6028/NIST.AI.100-1

UNESCO. (2021). Recommendation on the ethics of artificial intelligence. https://www.unesco.org/en/legal-affairs/recommendation-ethics-artificial-intelligence

Veale, M., & Brass, I. (2019). Administration by algorithm? Public management meets public sector machine learning. In K. Yeung & M. Lodge (Eds.), Algorithmic regulation. Oxford University Press. https://doi.org/10.1093/oso/9780198838494.003.0006

Wachter, S., Mittelstadt, B., & Russell, C. (2021). Why fairness cannot be automated: Bridging the gap between EU non-discrimination law and AI. Computer Law & Security Review, 41, 105567. https://doi.org/10.1016/j.clsr.2021.105567

Wirtz, B. W., Weyerer, J. C., & Geyer, C. (2019). Artificial intelligence and the public sector—Applications and challenges. International Journal of Public Administration, 42(7), 596–615. https://doi.org/10.1080/01900692.2018.1498103

World Economic Forum. (2022). AI procurement in a box: A guide to AI procurement for the public sector. World Economic Forum.

Yeung, K. (2018). Algorithmic regulation: A critical interrogation. Regulation & Governance, 12(4), 505–523. https://doi.org/10.1111/rego.12158

Zuboff, S. (2019). The age of surveillance capitalism: The fight for a human future at the new frontier of power. PublicAffairs.