

POLICY ANALYSIS

Algorithmic Government in a Decentralised State: Capacity, Accountability and the Procurement Gap

Farhan Qudratulloh Ginanjar^{1*}

¹ Diponegoro University, Semarang, Indonesia

* Corresponding author: farhanqudratullah@students.undip.ac.id

Received 14 January 2025 | Revised 22 January 2025 | Accepted 28 January 2025

ABSTRACT

Governments increasingly delegate consequential administrative decisions to algorithmic systems, and the accountability mechanisms available to citizens have not kept pace. This study examines where in the life of such a system accountability can actually attach, and what follows for a decentralised state in which adoption occurs at subnational level. The method combines a structured review of research published between 2017 and 2024 with comparative analysis of two extensively documented failures and a lifecycle-based design analysis. Three findings emerge. First, accountability is effective mainly *ex ante*: analysis of the Dutch risk-indication case concludes that accountability configurations depend predominantly on the prior sensitivity and responsiveness of political forums, and that where these prove insufficient, mitigation during or after deployment is very difficult. Second, the decisions that matter most are made during procurement rather than during policymaking, since vendors embed optimisation targets, data selection and thresholds that agencies treat as technical rather than political. Third, the legal remedies commonly assumed to exist are weaker than supposed: no general right to explanation of automated decisions exists in European data protection law, and administrative law has no settled account of how judicial review applies to automated public decisions. The Australian and Dutch failures show what follows, including a settlement returning over one billion Australian dollars to approximately four hundred thousand people. For decentralised states the implication is structural: the forum that must act *ex ante* is dispersed across hundreds of subnational bodies with uneven capacity, which is where the safeguards proposed here are directed.

KEYWORDS: algorithmic decision-making; public accountability; government procurement; decentralisation; administrative discretion

ABSTRAK

Pemerintah semakin banyak mendelegasikan keputusan administratif yang berkonsekuensi kepada sistem algoritmik, sementara mekanisme akuntabilitas yang tersedia bagi warga tidak mengikutinya. Penelitian ini menelaah di titik mana dalam daur hidup sistem semacam itu akuntabilitas benar-benar dapat melekat, dan apa akibatnya bagi negara terdesentralisasi yang adopsinya berlangsung di tingkat daerah. Metodenya memadukan telaah terstruktur riset terbitan 2017 hingga 2024 dengan analisis komparatif atas dua kegagalan yang terdokumentasi luas serta analisis desain berbasis daur hidup. Tiga temuan muncul. Pertama, akuntabilitas efektif terutama secara *ex ante*: analisis atas kasus sistem indikasi risiko di Belanda menyimpulkan bahwa konfigurasi akuntabilitas bergantung terutama pada kepekaan dan daya tanggap forum politik sebelumnya, dan bahwa ketika hal itu tidak memadai, mitigasi saat atau sesudah penerapan menjadi sangat sulit. Kedua, keputusan yang paling menentukan diambil pada tahap pengadaan, bukan pada tahap perumusan kebijakan, karena vendor menanamkan sasaran optimasi, pemilihan data, dan ambang batas yang diperlakukan instansi sebagai soal teknis alih-alih politis. Ketiga, pemulihan hukum yang lazim diandaikan ada ternyata lebih lemah dari dugaan: tidak ada hak umum atas penjelasan keputusan otomatis dalam hukum perlindungan data Eropa, dan hukum administrasi belum memiliki pemahaman mapan tentang penerapan uji peradilan atas keputusan publik yang terotomasi. Kegagalan di Australia dan Belanda menunjukkan akibatnya, termasuk penyelesaian yang mengembalikan lebih dari satu miliar dolar Australia kepada sekitar

empat ratus ribu orang. Bagi negara terdesentralisasi, implikasinya bersifat struktural: forum yang harus bertindak ex ante tersebar pada ratusan lembaga daerah dengan kapasitas tidak merata, dan ke sanalah pengaman yang diusulkan di sini diarahkan.

KATA KUNCI: pengambilan keputusan algoritmik; akuntabilitas publik; pengadaan pemerintah; desentralisasi; diskresi administratif

INTRODUCTION

A government decision that once required an official now often requires a system. Eligibility for a benefit, selection for an audit, a risk score that determines whether a household is investigated — these are increasingly produced by software, and the official's role has shifted from deciding to confirming. The administrative literature has examined what this does to discretion. It has examined rather less carefully the question of when, in the life of such a system, anyone can still be held to account for what it does.

The question is not rhetorical, because the documented failures suggest an uncomfortable answer. Two cases have been examined in unusual depth. In Australia, an automated debt-raising scheme operated from 2016 to 2019, was found unlawful by multiple investigations, and ended in a class action settlement returning more than one billion Australian dollars to approximately four hundred thousand people (Whiteford, 2021). In the Netherlands, a court in 2020 invalidated the legislation underpinning a welfare-fraud risk-indication system on human rights grounds, holding the system insufficiently transparent and verifiable.

In both cases the systems ran for years. Accountability arrived, but it arrived after the harm, through litigation and inquiry rather than through the ordinary mechanisms of administrative oversight. Wieringa (2023), analysing the Dutch case specifically for what it teaches about accountability, reaches a conclusion that organises the present study: accountability configurations hinge predominantly on the ex ante sensitivity and responsiveness of political forums, and where these prove insufficient, mitigation during or after the fact is very difficult for other actors.

If that is right, then the design question is not how to make algorithms explainable after they have decided, but which forum, at which moment, has both the standing and the information to stop a system before it operates. This study argues that the moment is procurement and that the forum is usually absent — a problem that a decentralised state multiplies, because the adopting body is subnational, numerous, and unevenly resourced.

Indonesia supplies the relevant setting. Wadipalapa et al. (2024), studying artificial intelligence policy through cases in Jakarta, Central Java and East Java together with elite interviews, find significant misalignment between central and subnational priorities and implementation capacity, central regulatory control constraining local flexibility, and regional leaders adopting national initiatives partly for political return. Adoption, in other words, is occurring at the tier where the accountability forums are weakest.

Three questions organise the analysis. First, what does the evidence establish about where accountability for algorithmic public decisions can effectively attach? Second, what do the available legal remedies actually provide, as distinct from what they are assumed to provide? Third, what safeguards follow for a decentralised state, and at which stage of the system lifecycle must each be placed?

LITERATURE REVIEW AND CONCEPTUAL FRAMEWORK

Artificial discretion and bureaucratic form

Young et al. (2019) introduce the concept of artificial discretion, defining discretion as the execution of tasks requiring non-trivial decisions and assessing algorithmic systems as a governance tool against effectiveness, efficiency, equity, manageability and political feasibility. They identify gains in scalability, cost and quality alongside serious risks to equity, manageability and political feasibility — a balance sheet rather than an endorsement or a warning.

Bullock et al. (2020) ask how artificial intelligence both changes and is changed by bureaucratic organisations, and find through contrasting cases in health insurance auditing and policing that the effect on discretion is non-linear and depends on bureaucratic form, with a tendency to push organisations toward centralised screen-level bureaucracy even where they had historically resisted it. Meijer et al. (2021), comparing predictive policing in Berlin and Amsterdam, find near-identical technology producing two divergent patterns — an algorithmic cage that constrains professionals and an algorithmic colleague that preserves their judgement — with the difference driven by pre-existing administrative traditions rather than by technical specification.

This last result matters for the present argument. If administrative context rather than software design determines what a system does to an organisation, then a system procured identically by two subnational governments with different administrative capacity will not behave identically, and evaluation conducted centrally will not predict local outcomes. Busch and Henriksen (2018), reviewing 44 studies of information technology and street-level discretion, reach a compatible conclusion: digital systems reinforce ethical and democratic values while eroding professional and relational ones, with professionalisation and organisational context shaping the result.

Accountability and the problem of many hands

Busuioc (2021) examines algorithmic accountability specifically through a public administration lens, identifying where existing accountability mechanisms break down when applied to algorithmic decisions. Wieringa (2020), reviewing 242 articles narrowed to 93 core works, applies the standard accountability framework — actor, forum, relationship, standards, consequences — to algorithmic systems, and the exercise exposes how many of these elements are indeterminate when a decision is produced by a system assembled from many contributors.

Cooper et al. (2022) name the underlying difficulty directly, revisiting the classic account of accountability eroding through ubiquitous delegation of consequential functions to computational systems and updating the problem of many hands for contemporary machine learning. When a decision emerges from a vendor's model, an agency's data, a contractor's integration and an official's confirmation, each participant can accurately say that the outcome was not theirs alone.

Criado et al. (2020), studying an irregularity-detection early warning system in a Spanish regional administration, add an operational distinction between transparency as accessibility and transparency as explainability, and find that positive outcomes for civil servants correlate with their active involvement in development and with auditing processes being recognised in law. Participation and legal recognition of audit, not disclosure alone, are what made the difference.

What the documented failures establish

Whiteford (2021) analyses the Australian automated debt scheme and argues it was not an ordinary policy fiasco, since the design was intentional rather than the result of implementation error, and it failed even to deliver its projected savings. Van Krieken (2024), drawing on inquiry evidence, reframes the same episode as organised ignorance functioning as ordinary bureaucratic operation rather than as institutional breakdown — an interpretation with an uncomfortable implication, namely that the accountability gap was produced rather than encountered.

On the Dutch case, van Bekkum and Zuiderveen Borgesius (2021) analyse the 2020 judgment invalidating the risk-indication legislation and note, importantly for balance, that its immediate practical effect is limited because most existing fraud-detection systems operate within narrower constraints than the invalidated one. Rachovitsa and Johann (2022) argue that digital welfare state initiatives are likely to fall short of basic requirements of legality and protection against arbitrariness. Wieringa (2023) draws the accountability lesson quoted above.

Calo and Citron (2021) generalise from a body of such cases, arguing that agencies derive legitimacy from expertise and that automation erodes that basis, and documenting benefits-denial and rights-restriction decisions producing outcomes they describe as bizarre and unintelligible. Their target is legitimacy rather than procedural due process, which distinguishes their claim from the remedies literature.

Legal responses and their actual reach

It is commonly assumed that European data protection law confers a right to an explanation of automated decisions. Wachter et al. (2017) argue that it does not: the instrument creates a right to be informed, delivering meaningful but properly limited information, while the provision on automated decision-making is ambiguous and restricted in scope. Any argument built on an assumed right to explanation must engage this.

Kaminski and Malgieri (2021) respond by shifting ground, arguing that the explanation framing is too narrow and that mandatory impact assessments should produce multi-layered explanations serving both individual rights and systemic governance. Cobbe (2019) addresses the administrative law side and states plainly that there is no clear understanding of how English administrative law applies to automated public-sector decision-making, mapping the potential grounds of judicial review rather than reporting settled doctrine.

The combined position is that the two remedies most often invoked — a data protection right to explanation and judicial review of the decision — are respectively narrower than assumed and unsettled. This is a reason to locate safeguards earlier in the lifecycle rather than to rely on remedies after a decision has been made.

Procurement as policymaking

Mulligan and Bamberger (2019) supply the argument this study builds on most directly. Agencies, they contend, wrongly treat the adoption of machine learning as procurement rather than as policymaking, which allows vendors to embed consequential policy choices — what the system optimises for, which data it is trained on, where thresholds are set — without agency deliberation or public input. They propose subjecting adoption decisions to the standards applied to arbitrary and capricious agency action.

The claim is important because it locates the policy decision at a stage where administrative law, public consultation and accountability forums are all typically absent. A procurement is a contract award. It is evaluated on price, specification compliance and vendor capability, and neither the evaluation committee nor the appeal mechanism attached to it is constituted to ask what population the model will misclassify.

Adoption in decentralised and middle-income settings

Sun and Medaglia (2019), studying an artificial intelligence implementation in Chinese public healthcare across government officials, hospital managers and technology firms, find stakeholders holding diverse and sometimes contradictory framings of the challenges, and note that very little empirical research underpins widespread claims about public sector artificial intelligence. Zuiderwijk et al. (2021), reviewing 26 studies, characterise the field as exploratory, conceptual, qualitative and practice-driven, and call for empirical, explanatory, public-sector-focused work.

Kuziemski and Misuraca (2020) examine three cases of automated decision-making in democratic settings — immigration process control in Canada, employment services optimisation in Poland, and a national artificial intelligence programme in Finland — and argue that such systems may amplify existing power asymmetries, calling for a common framework to evaluate their impact.

For Indonesia, Wadipalapa et al. (2024) provide the most directly relevant evidence, finding through cases in Jakarta, Central Java and East Java significant misalignment between central and subnational artificial intelligence priorities and implementation capacity, central regulatory control constraining local flexibility, and regional leaders adopting national initiatives partly for electoral return. The combination — adoption driven by political signalling, capacity uneven, central control regulatory rather than operational — describes conditions under which the procurement stage is least likely to receive scrutiny.

Citizen response

Three experimental studies establish what citizens actually object to. Grimmelikhuijsen (2023), testing visa denial and welfare fraud investigation scenarios, finds explainability having a more pronounced effect on perceived trustworthiness than accessibility, and concludes that transparency as accessibility is insufficient to foster citizen trust. Schiff et al. (2022), surveying 1,460 participants across child welfare and criminal justice

scenarios, find substantial negative reactions where fairness and transparency are not realised, reactions that cut across political ideology and persist among respondents facing no personal consequence. Ingrams et al. (2022), in a tax auditing setting, find decisions attributed to artificial intelligence rated less trustworthy than human ones.

Within the Indonesian literature, Zulkifli et al. (2019) argue conceptually that the use of information and communication technology raises strategic quality and thereby local government performance. Their paper is propositional rather than empirical, and it is cited here for what its framing assumes rather than for what it demonstrates: that the relevant decision is the adoption of technology, and that performance follows from adoption.

That assumption is precisely what the procurement argument developed in this study questions. If the choices determining what a system does to people are made when a vendor is selected and a specification written, then adoption is not a single decision but a bundle of policy choices presented as a technical one, and a framing that treats adoption as the unit of analysis will not reach them. The Indonesian conceptual literature and the international accountability literature therefore locate the decision at different points, and the difference is consequential for where scrutiny should attach.

Research gap

Two gaps are addressed. The accountability literature has concentrated on explanation and audit of systems already deployed, while the evidence indicates that effective accountability is *ex ante*; the implications of that timing for institutional design have not been worked out. And no located study examines what the *ex ante* requirement means where adoption is decentralised across many subnational bodies of uneven capacity. Verified empirical work on algorithmic procurement in developing and middle-income countries is very thin, and this study does not claim to fill that gap; it states where the argument is transferring evidence from other settings.

METHOD

The study is a structured review with comparative case analysis and a lifecycle-based design analysis. It generates no primary data, examines no specific Indonesian system, and makes no claim about the extent of algorithmic adoption in Indonesian subnational government.

Source identification and verification

Peer-reviewed studies published between 2017 and 2024 on algorithmic decision-making in government, algorithmic accountability, automated administrative decisions, public sector artificial intelligence adoption and citizen trust were identified and retained only where bibliographic details could be verified against a publisher page, indexing record or the journal's own document. Conference proceedings are identified as such. Works in outlets without demonstrable peer review were excluded, as were works whose publication year was contested across records in a way that placed them outside the review window.

Case selection and treatment

Two failures were selected because they are documented in peer-reviewed analysis rather than in press coverage alone: the Australian automated debt scheme and the Dutch risk-indication system. Only facts appearing in the peer-reviewed sources or in an official court communication are reported. Figures circulating in secondary sources without an accessible primary record are omitted, which is why several widely quoted numbers about both cases do not appear here.

Lifecycle analysis

The design component divides the life of an algorithmic public decision system into five stages, and asks for each stage which actor holds the decision, which forum could scrutinise it, and what information that forum would need. Safeguards are then assigned to the stage at which the decision they address is actually made. The output is a design proposal that has not been implemented or tested.

RESULTS AND DISCUSSION

What the failures have in common

Table 1. Two documented failures compared

Source: compiled by the authors from the cited peer-reviewed analyses and the court's official communication. Figures circulating in secondary sources without accessible primary records are excluded.

The final row is the finding. In neither case did the ordinary machinery of administrative accountability — internal audit, parliamentary oversight, complaint handling, individual appeal — stop the system. Accountability arrived from outside, through courts and commissions, after populations measured in hundreds of thousands had been affected. Wieringa's (2023) conclusion that mitigation during or after deployment is very difficult for actors other than the responsible political forum is not an abstract claim; it is a description of what happened twice.

This has an implication that runs against the direction most reform proposals take. Explainability requirements, appeal rights and audit obligations all operate on systems that are already deciding. They are valuable, but the evidence indicates they are not where the decisive intervention lies.

Where the decisions are actually made

Figure 1. Stages of an algorithmic public decision system and where accountability attaches (authors' construction)

Figure 1 places the two observations against each other. The choices that determine what the system will do to people — what it optimises for, which data it learns from, where its thresholds sit — are made at the procurement and design stages, as Mulligan and Bamberger (2019) argue. The forums that can call anyone to account are concentrated at the review stage, where the same authors' analysis and both documented failures indicate intervention is hardest.

Table 2. Decision, forum and information by lifecycle stage

Source: authors' analysis. The middle rows are where the consequential choices are made and where forum availability is weakest.

The two middle rows constitute the procurement gap. A procurement appeal exists, but it is constituted to hear a disappointed bidder complain about process, not to hear a citizens' organisation argue that the specified optimisation target will misclassify a particular population. At the model design stage, in most arrangements, no forum exists at all, because the work is contractual and commercially confidential.

Cooper et al.'s (2022) treatment of the many-hands problem explains why later remedies cannot recover this. By the review stage the decision has been distributed across a vendor's design, an agency's data and an official's confirmation, and each contributor's account of its own role is accurate. There is no point at which a single actor can be shown to have made the decision that caused the harm, which is precisely the condition under which accountability forums fail.

Why decentralisation multiplies the problem

If effective accountability must operate *ex ante*, and if the decisive stage is procurement, then the capacity that matters is the capacity of the procuring body. In a centralised system this is one institution, which can be resourced, trained and scrutinised. In a decentralised system it is many.

The Indonesian evidence indicates the conditions under which this occurs. Wadipalapa et al. (2024) find central and subnational priorities misaligned, implementation capacity uneven, central control exercised through regulation rather than through operational involvement, and adoption motivated partly by political signalling. A regional government adopting a system for its demonstrative value, under a central framework that regulates but does not staff the decision, is procuring policy without the capacity to recognise that it is doing so.

Meijer et al.'s (2021) finding compounds this. Because identical technology produced an algorithmic cage in one administration and an algorithmic colleague in another, depending on administrative tradition, a system

evaluated favourably in one subnational government supplies no assurance about another. Centralised approval of a vendor or product — the natural administrative response to dispersed capacity — would therefore be a partial safeguard at best, since it certifies the technology and not the organisation that will operate it.

Safeguards placed by stage

The analysis implies that safeguards should be assigned to the stage at which the decision they address is made, rather than concentrated where forums happen to exist.

Table 3. Safeguards mapped to the stage at which the decision occurs

Source: authors' derivation. Each safeguard is placed at the stage where the decision it addresses is made rather than where scrutiny is currently available.

Two entries deserve comment. Publishing the optimisation target and thresholds with the tender is the single change that most directly converts a procurement into a policy act, because it exposes to comment the choices that Mulligan and Bamberger (2019) identify as being made invisibly. And the recorded override rate is proposed because it is the one measure that reveals whether the confirming official is exercising judgement or ratifying output; an override rate near zero indicates that the human in the loop is decorative, which is a fact an oversight body can act on without any technical access to the model.

The trust evidence supports placing explanation at the individual level. Grimmelikhuijsen (2023) finds explainability affecting perceived trustworthiness more than accessibility, and Schiff et al. (2022) find negative citizen reactions to unrealised fairness and transparency crossing ideological lines and persisting among people not personally affected. Publishing system documentation therefore satisfies the accessibility condition without satisfying the one that citizens respond to. Wachter et al. (2017) establish that the legal right to obtain an individual explanation is weaker than commonly assumed, which makes this an administrative commitment to be made deliberately rather than a legal obligation to be complied with.

Limitations

Four limitations qualify this analysis. First, the two failure cases are drawn from high-income administrations with strong courts and free press, and both were exposed by mechanisms that may be less available elsewhere; the argument that ex ante accountability is decisive may therefore understate the problem in other settings rather than describe it. Second, verified empirical work on algorithmic procurement in developing and middle-income countries is very thin, and the extension of the procurement argument to such settings is an inference from administrative similarity rather than a finding. Third, the safeguards are a design proposal that has not been implemented or evaluated anywhere. Fourth, no Indonesian algorithmic decision system was examined, and the Indonesian material describes adoption conditions rather than any specific system's operation.

CONCLUSION

The accountability problem in algorithmic government is usually posed as a problem of opacity: the system cannot explain itself, so no one can be held responsible. The evidence assembled here suggests the problem is better posed as one of timing. Analysis of the Dutch case concludes that accountability depends predominantly on the prior sensitivity and responsiveness of political forums, and that once these have failed, mitigation is very difficult for anyone else. Two extensively documented failures bear this out: both systems operated for years, both were stopped from outside the administration, and one ended in a settlement returning over a billion Australian dollars to approximately four hundred thousand people.

If accountability is effective mainly before deployment, the decisive stage is procurement, because that is where the choices are made — what the system optimises for, which data it learns from, where its thresholds sit — and those choices are policy dressed as specification. Yet procurement is the stage at which no accountability forum with standing to raise such questions exists. That is the gap this study names. The remedies usually relied on to close it are weaker than assumed: no general right to explanation of automated decisions exists in European data

protection law, and administrative law has no settled account of how review applies to automated public decisions.

For a decentralised state the consequence is structural rather than technical. The body that must exercise ex ante scrutiny is subnational, numerous and unevenly resourced, and Indonesian evidence indicates adoption proceeding under precisely those conditions, with priorities misaligned between tiers and adoption motivated partly by political signalling. Because identical technology has been shown to produce opposite effects in administrations with different traditions, central certification of products cannot substitute for capacity in the adopting organisation. The safeguards proposed here are therefore placed by stage rather than by convenience: publish the optimisation target and thresholds with the tender, give affected groups standing at procurement, require impact assessment before training, involve the officials who will operate the system, record the override rate, and give reasons at the level of the individual decision. None requires access to a model's internals, and all of them act before the harm rather than after it.

DECLARATIONS

Author contributions (CRediT)

Conceptualisation: Author 1. Methodology: Author 1. Formal analysis and investigation: Author 1. Writing — original draft preparation: Author 1. Writing — review and editing: Author 1. Supervision: Author 1. All named authors have read and approved the final manuscript and accept responsibility for its content.

Conflict of interest

The authors declare that they have no financial or non-financial conflict of interest in relation to the subject matter of this article. No author holds a position in, or receives support from, any organisation examined or evaluated in this study.

Data availability

This study generated no primary data. It draws exclusively on published peer-reviewed literature, publicly available legal instruments and publicly available official statistics, all of which are identified in the text and listed in the references. No confidential administrative record, individual taxpayer record, permit file or personal datum was accessed, and no individual person or business is identified. Every quantity reported was traced to a single identified public source and can be verified by any reader from that source.

Funding

This research received no specific grant from any funding agency in the public, commercial or not-for-profit sectors.

Declaration on the use of artificial intelligence

In accordance with COPE, ICMJE and WAME positions, and with the RCP AI Use Declaration requirement, the authors disclose that generative artificial intelligence tools were used in preparing this manuscript, specifically to assist with literature identification, synthesis of the reviewed sources, and language editing of the draft. Artificial intelligence was not used to generate data, results or citations. Every reference in this article was verified by the authors against the publisher's own record, a digital object identifier record or an indexing database record before inclusion, and sources whose bibliographic details or indexing status could not be confirmed were excluded. No artificial intelligence tool is listed as an author. The authors take full responsibility for the accuracy, integrity and originality of the entire content of this article.

REFERENCES

Bullock, J. B., Young, M. M., & Wang, Y.-F. (2020). Artificial intelligence, bureaucratic form, and discretion in public service. *Information Polity*, 25(4), 491-506. <https://doi.org/10.3233/IP-200223>

- Busch, P. A., & Henriksen, H. Z. (2018). Digital discretion: A systematic literature review of ICT and street-level discretion. *Information Polity*, 23(1), 3-28. <https://doi.org/10.3233/IP-170050>
- Busuioc, M. (2021). Accountable artificial intelligence: Holding algorithms to account. *Public Administration Review*, 81(5), 825-836. <https://doi.org/10.1111/puar.13293>
- Calo, R., & Citron, D. K. (2021). The automated administrative state: A crisis of legitimacy. *Emory Law Journal*, 70(4), 797-844.
- Cobbe, J. (2019). Administrative law and the machines of government: Judicial review of automated public-sector decision-making. *Legal Studies*, 39(4), 636-655. <https://doi.org/10.1017/lst.2019.9>
- Cooper, A. F., Moss, E., Laufer, B., & Nissenbaum, H. (2022). Accountability in an algorithmic society: Relationality, responsibility, and robustness in machine learning. In *Proceedings of the 2022 ACM Conference on Fairness, Accountability, and Transparency* (pp. 864-876). <https://doi.org/10.1145/3531146.3533150>
- Criado, J. I., Valero, J. M., & Villodre, J. (2020). Algorithmic transparency and bureaucratic discretion: The case of SALER early warning system. *Information Polity*, 25(4), 449-470. <https://doi.org/10.3233/IP-200260>
- Grimmelikhuisen, S. (2023). Explaining why the computer says no: Algorithmic transparency affects the perceived trustworthiness of automated decision-making. *Public Administration Review*, 83(2), 241-262. <https://doi.org/10.1111/puar.13483>
- Ingrams, A., Kaufmann, W., & Jacobs, D. (2022). In AI we trust? Citizen perceptions of AI in government decision making. *Policy & Internet*, 14(2), 390-409. <https://doi.org/10.1002/poi3.276>
- Kaminski, M. E., & Malgieri, G. (2021). Algorithmic impact assessments under the GDPR: Producing multi-layered explanations. *International Data Privacy Law*, 11(2), 125-144. <https://doi.org/10.1093/idpl/ipaa020>
- Kuziemski, M., & Misuraca, G. (2020). AI governance in the public sector: Three tales from the frontiers of automated decision-making in democratic settings. *Telecommunications Policy*, 44(6), 101976. <https://doi.org/10.1016/j.telpol.2020.101976>
- Meijer, A., Lorenz, L., & Wessels, M. (2021). Algorithmization of bureaucratic organizations: Using a practice lens to study how context shapes predictive policing systems. *Public Administration Review*, 81(5), 837-846. <https://doi.org/10.1111/puar.13391>
- Mulligan, D. K., & Bamberger, K. A. (2019). Procurement as policy: Administrative process for machine learning. *Berkeley Technology Law Journal*, 34(3), 773-850.
- Rachovitsa, A., & Johann, N. (2022). The human rights implications of the use of AI in the digital welfare state: Lessons learned from the Dutch SyRI case. *Human Rights Law Review*, 22(2), ngac010. <https://doi.org/10.1093/hrlr/ngac010>
- Schiff, D. S., Schiff, K. J., & Pierson, P. (2022). Assessing public value failure in government adoption of artificial intelligence. *Public Administration*, 100(3), 653-673. <https://doi.org/10.1111/padm.12742>
- Sun, T. Q., & Medaglia, R. (2019). Mapping the challenges of artificial intelligence in the public sector: Evidence from public healthcare. *Government Information Quarterly*, 36(2), 368-383. <https://doi.org/10.1016/j.giq.2018.09.008>
- van Bekkum, M., & Zuiderveen Borgesius, F. (2021). Digital welfare fraud detection and the Dutch SyRI judgment. *European Journal of Social Security*, 23(4), 323-340. <https://doi.org/10.1177/13882627211031257>
- van Krieken, R. (2024). The organization of ignorance: The Australian 'Robodebt' affair, bureaucracy, law and politics. *Critical Sociology*, 50(7-8), 1379-1398. <https://doi.org/10.1177/08969205241245257>
- Wachter, S., Mittelstadt, B., & Floridi, L. (2017). Why a right to explanation of automated decision-making does not exist in the General Data Protection Regulation. *International Data Privacy Law*, 7(2), 76-99. <https://doi.org/10.1093/idpl/ipx005>

- Wadipalapa, R. P., Katharina, R., Nainggolan, P. P., Aminah, S., Apriani, T., Ma'rifah, D., & Anisah, A. L. (2024). An ambitious artificial intelligence policy in a decentralised governance system: Evidence from Indonesia. *Journal of Current Southeast Asian Affairs*, 43(1), 65-93. <https://doi.org/10.1177/18681034231226393>
- Whiteford, P. (2021). Debt by design: The anatomy of a social policy fiasco - Or was it something worse? *Australian Journal of Public Administration*, 80(2), 340-360. <https://doi.org/10.1111/1467-8500.12479>
- Wieringa, M. (2020). What to account for when accounting for algorithms: A systematic literature review on algorithmic accountability. In *Proceedings of the 2020 Conference on Fairness, Accountability, and Transparency* (pp. 1-18). <https://doi.org/10.1145/3351095.3372833>
- Wieringa, M. (2023). 'Hey SyRI, tell me about algorithmic accountability': Lessons from a landmark case. *Data & Policy*, 5, e2. <https://doi.org/10.1017/dap.2022.39>
- Young, M. M., Bullock, J. B., & Lecy, J. D. (2019). Artificial discretion as a tool of governance: A framework for understanding the impact of artificial intelligence on public administration. *Perspectives on Public Management and Governance*, 2(4), 301-313. <https://doi.org/10.1093/ppmgov/gvz014>
- Zuiderwijk, A., Chen, Y.-C., & Salem, F. (2021). Implications of the use of artificial intelligence in public governance: A systematic literature review and a research agenda. *Government Information Quarterly*, 38(3), 101577. <https://doi.org/10.1016/j.giq.2021.101577>
- Zulkifli, Nurhayatie, T., & Junaidi, M. (2019). Penelitian konseptual: Peningkatan kualitas strategis melalui penggunaan teknologi informasi dan komunikasi dalam meningkatkan kinerja pemerintah daerah. *Jurnal Litbang Sukowati: Media Penelitian dan Pengembangan*, 3(1), 8. <https://doi.org/10.32630/sukowati.v3i1.113>

© 2025 The Author(s). Published by Research Center Progresif, Semarang City, Central Java, Indonesia.

This is an open access article distributed under the terms of the Creative Commons Attribution-ShareAlike 4.0 International Licence (CC BY-SA 4.0), which permits use, sharing and adaptation provided the original work is properly cited.

How to cite: Farhan Qudratulloh Ginanjar (2025). Algorithmic Government in a Decentralised State: Capacity, Accountability and the Procurement Gap. JLPPF, 1(1), 34-43.