



EPIC-PP-2026-007 · RESEARCH ARTICLE · AI GOVERNANCE

Algorithmic Shura: Consultation, Consent, and the Advice–Decision Distinction in Public-Sector AI

EPIC – Emirati Perspectives and Insight Centre

Seeing what others miss...

20 October 2026

CC BY 4.0 · Creative Commons Attribution 4.0 International · creativecommons.org/licenses/by/4.0

Algorithmic Shura: Consultation, Consent, and the Advice–Decision Distinction in Public-Sector AI

EPIC – EMIRATI PERSPECTIVES AND INSIGHT CENTRE

EPIC-PP-2026-007 · AI Governance · Published 20 October 2026

© 2026 EPIC – Emirati Perspectives and Insight Centre. Released under Creative Commons Attribution 4.0 International (CC BY 4.0): <https://creativecommons.org/licenses/by/4.0/>

EPIC is an independent research initiative in formation. It is not yet a registered or licensed entity. EPIC does not accept donations, payments or contracts.

Abstract

When artificial intelligence (AI) systems rank, flag, and recommend inside public administration, the philosophically interesting question is not only whether their outputs are accurate but where advice ends and authoritative decision begins. Existing accounts, including the threat of “algocracy”, the responsibility gap, and meaningful human control, diagnose how machine outputs can erode the conditions of legitimate public authority. This article contributes a conceptual resource from outside the usual Euro-American canon: the Islamic principle of consultation (*shura*). Read alongside consent-based theories of legitimacy, *shura* yields what we call *algorithmic shura*: a set of three jointly necessary conditions for government AI that informs significant decisions. First, meaningful consultation among accountable humans must precede deployment. Second, machine output must be institutionally treated as counsel rather than command, with an advice–decision wall that makes refusal no costlier than acceptance. Third, authorisation must be renewable rather than fossilised at launch. We reconstruct the framework, distinguish it from claims that machines can be consultative partners, and test it against AI governance instruments in the United Arab Emirates and neighbouring Gulf states, which already articulate fragments of the view. We argue that algorithmic *shura* sharpens familiar human-oversight principles by locating the core problem in the political relation between counsel and authority, and that it offers a model for cross-traditional work in the philosophy of technology.

Keywords: algorithmic governance; *shura*; consent; legitimacy; meaningful human control; comparative political philosophy

1. Introduction

When a government system ranks applicants, prioritises cases, flags risk, drafts a recommendation, or routes a citizen to one pathway rather than another, something politically consequential has already occurred, even if a human later clicks “approve”. The political problem of AI in government is therefore not only accuracy. It concerns the distribution of advice, the location of decision, and the conditions under which the public authority that deploys a system may still be said to govern with open eyes.

Philosophers of technology have given this worry several names. Danaher (2016) describes the “threat of algocracy”: the prospect that decision procedures become so dependent on opaque algorithms that those subject to them can no longer participate in, or even comprehend, the reasons that govern them. Matthias (2004) identified a “responsibility gap” opened by learning systems whose behaviour their operators cannot fully predict or control. Santoni de Sio and van den Hoven (2018) responded with an account of “meaningful human control” requiring that systems track the relevant reasons of human agents and that outcomes be traceable to someone who understands both the system and their own moral role. These accounts share a diagnosis: machine output can quietly displace the human practices of reason-giving on which legitimate authority depends.

This article adds a conceptual resource that these debates have largely not drawn upon: the Islamic principle of consultation, *shura*. Our claim is not that *shura* is superior to consent theory or to meaningful human control, nor that classical *shura* was a doctrine of digital governance. It is that *shura* names, with unusual precision, a relation that the AI governance debate needs: the relation between *counsel* and *command*. A tradition that has reflected for centuries on how rulers should take advice without surrendering responsibility, and how advisers may be heeded without becoming rulers, has something to offer a world in which the most prolific advisers are statistical models.

The Gulf states provide an apt testing ground. They have adopted AI in public administration early and ambitiously, and their governance instruments already contain fragments of the view defended here: human final authority, accountability that does not rest with the machine, and an explicit demand for consistency with local cultural and religious norms (Smart Dubai, 2019; Ministry of Communications and Information Technology [MCIT], 2019). The case is philosophically useful because it shows a political community translating global principles into its own normative vocabulary, and it is with that translation, rather than with regional policy for its own sake, that this article is concerned.

The argument proceeds as follows. Section 2 reconstructs algorithmic shura from three strands: shura as consultation, consent as renewable authorisation, and the advice–decision distinction. Section 3 examines UAE and Gulf instruments through that lens and identifies characteristic failure modes. Section 4 considers objections and the limits of the framework. Section 5 draws design implications, and Section 6 concludes.

1.1 Approach

This is a work of conceptual analysis supported by documentary evidence. Policy instruments are cited as evidence of how governing institutions articulate their commitments, not as philosophical authorities.

2. Framework: Shura, Consent, and the Advice–Decision Wall

2.1 Shura as consultation, not machine rule

In the Qur’an, consultation appears both as a virtue of the believing community and as an instruction to the Prophet in worldly affairs. Surah al-Shura (42:38) praises those “whose affairs are [a matter of] consultation among themselves”, and Surah Al Imran (3:159) instructs: “consult them in the matter”.^[1] Classical and modern scholars have treated shura as a principle that resists unilateralism without fixing a single constitutional form (Al-Raysuni, 2011; Crone, 2001).

Patricia Crone’s historical work is useful precisely because it refuses anachronism. She notes that shura ordinarily means consultation between a person in authority and subordinates, and that in early Islamic political practice it also acquired a specialised sense as a procedure for selecting a ruler (Crone, 2001). Ahmad Al-Raysuni, writing from a contemporary reformist standpoint, treats shura as a Qur’anic principle of mutual consultation that should be institutionalised as a way of life, while acknowledging that Muslims disagree over whether it is obligatory in every domain and how far it resembles democracy (Al-Raysuni, 2011). Comparative political theorists have likewise warned against collapsing shura into a single Western template (Shavit, 2010; Thielicke-Witt, 2025). The careful move is to recover what consultation demands, namely deliberation before decision, openness to expertise, and resistance to solitary rule, without pretending that one institutional form exhausts the concept.

Three implications follow for AI governance. First, shura concerns counsel before command; it is not a licence for the counsellor to become the commander. Second, consultation is not noise: it presupposes that those consulted can affect the outcome, or at least that ignoring them requires justification. Third, shura is compatible with specialised expertise. Consulting the knowledgeable is still consultation. What it rejects is the fiction that an unexamined recommendation is already a decision.

This last point bears on the philosophy of technology directly. Debates about automated decision-making often oscillate between treating AI as a mere tool, whose output carries no normative weight, and treating it as a quasi-agent whose output must be either trusted or resisted. Shura suggests a third model: AI output is *counsel*, which has genuine epistemic weight and may properly shape a decision, but which acquires authority only through the act of a competent human who takes responsibility for heeding it. The counsellor’s role is real but derivative; the decider’s role cannot be delegated away by the fact that counsel was good.

2.2 Consent and renewable authorisation

Consent theory, in its Lockean form, locates political legitimacy in the idea that government is entrusted with authority to protect rights and promote the public good, and that people are not naturally subject to absolute dominion without their consent (Tuckness, 2020). Other traditions emphasise collective self-legislation or public justification rather than individual consent (Peter, 2023). We do not adjudicate between actual, tacit, and hypothetical consent. We borrow a narrower lesson: where public power is exercised over persons, legitimacy is strained when those persons have no intelligible stake in how the power was authorised, reviewed, or contested.

Algorithmic systems complicate consent in two ways. They can make the exercise of power opaque, so that affected persons cannot tell whether a human judgment or a model score did the decisive work. And they can outlast the conditions under which they were authorised: the population changes, the policy goal changes, the vendor updates the model, or the error profile drifts. Consent, in this thin administrative sense, is not a one-time click. It is a renewable authorisation by a competent public authority, informed by evidence and open to challenge where significant decisions are at stake.^[2]

2.3 Advice versus decision

The third strand is the advice–decision wall. Machine learning systems generate scores, rankings, classifications, drafts, and alerts. None of these is by itself a public decision. A decision, in the sense that matters for political responsibility, is an authoritative determination by a person or body with competence to bind. The wall between advice and decision is therefore not a metaphor about attitudes; it is a requirement about who may finalise an outcome and how finalisation is recorded.

The empirical literature on automation bias explains why the wall is fragile. Humans over-rely on automated cues even when they know the system can err (Parasuraman & Manzey, 2010). A wall that exists in policy but not in practice is precisely the condition Danaher (2016) warns of: formal human authority hollowed out by functional dependence. It is also the condition that meaningful human control is meant to exclude, since a decider who merely ratifies a recommendation they cannot interrogate does not satisfy the tracing condition (Santoni de Sio & van den Hoven, 2018).

International instruments converge on versions of the wall. The OECD AI Principles emphasise safeguards for human agency and oversight (Organisation for Economic Co-operation and Development [OECD], 2024); the NIST AI Risk Management Framework treats governance as a function that continues throughout the system lifecycle (National Institute of Standards and Technology [NIST], 2023); and the EU AI Act requires human oversight of high-risk systems together with post-market monitoring (European Union, 2024).^[3]

Algorithmic shura binds the three strands. Consultation without an advice–decision wall becomes theatre: experts and citizens speak, then the model quietly decides. A wall without consultation becomes technocracy: humans rubber-stamp a recommendation they never interrogated. Consent without renewal becomes fossilised authorisation: a pilot approved in one year becomes invisible infrastructure the next. Algorithmic shura therefore requires (1) meaningful consultation before deployment of significant systems, (2) a documented wall between machine advice and human decision, and (3) renewable authorisation by a competent authority in light of evidence.

3. Gulf AI Governance Through an Algorithmic-Shura Lens

3.1 Strategy and institutional architecture

In October 2017 the UAE appointed a Minister of State for Artificial Intelligence, widely reported as the first such portfolio (Arabian Business, 2017; The National, 2017). The UAE Cabinet later adopted the National Strategy for Artificial Intelligence 2031, which sets the ambition of world leadership in AI by 2031 and lists eight objectives, including, as Objective 8, strong governance and effective regulation (UAE Cabinet, 2019; United Arab Emirates, n.d.-b). The strategy assigns the UAE Council for Artificial Intelligence and Blockchain a remit that includes reviewing national approaches to data management, ethics, and cybersecurity (United Arab Emirates Council for Artificial Intelligence and Blockchain, n.d.).

From an algorithmic-shura perspective, this architecture treats governance as a strategic objective rather than an afterthought and creates a standing consultative-administrative body with a review remit. The Council is not a classical shura assembly. But a body tasked with ethics review and cross-government coordination is a site where consultation among specialists and officials can be organised before systems harden into infrastructure. UAE public law also already distinguishes consultation from enactment: the Federal National Council discusses federal bills and general affairs within constitutional limits (United Arab Emirates, n.d.-a; UAE Government, n.d.).^[4] That distinction is the one algorithmic shura seeks to preserve when machines enter the advisory chain.

3.2 Ethical instruments: advice, challenge, and human finality

Dubai's AI Ethics Principles and Guidelines are among the clearest regional statements of the advice–decision wall. Under the humanity principle, society should be consulted in a representative fashion to inform AI development, and humanity should retain the power to make the final decision, with AI in an assisting role. Under accountability, out-

comes are not to be attributed to the system itself; subjects should be able to challenge significant automated decisions and, where appropriate, opt out; and systems informing significant decisions should not make value judgments on people's behalf without prior consent (Smart Dubai, 2019). These provisions map closely onto the three conditions: consultation, advice rather than decision, and a consent threshold for value-laden automation.

The UAE Charter for the Development and Use of Artificial Intelligence (2024) reiterates human oversight, governance, and accountability (UAE Government, 2024; WAM, 2024). National ethics guidance has been described as non-binding and accompanied by self-assessment tools, complementary to binding instruments such as the Personal Data Protection Law (Federal Decree-Law No. 45 of 2021; Regulations.AI, 2022). We treat these as authoritative statements of aspiration, not as statutes that by themselves create justiciable duties.

Saudi Arabia's AI Ethics Principles include a humanity principle calling for human-centric design that allows human choice and determination, together with lifecycle accountability for system owners (Saudi Data and Artificial Intelligence Authority [SDAIA], 2023). Qatar's National AI Strategy warns that AI methods can acquire black-box characteristics and inherit bias, and recommends an ethics and governance framework consistent both with Qatari social, cultural, and religious norms and with international guidelines (MCIT, 2019). That dual insistence is philosophically significant. It is not a rejection of global standards but an invitation to translate them, and algorithmic shura is one such translation.

3.3 Where practice can fall short

If these instruments already sound like algorithmic shura, why introduce a new concept? Because principles on paper can coexist with practices that quietly invert them. Three failure modes stand out.

The first is *consultative deficit*. A ministry may procure a scoring tool because a vendor demonstrates benchmark accuracy, without structured consultation of frontline officials, affected communities, or domain experts about the values encoded in the objective. Without recorded dissent and a justification for adoption despite it, consultation remains ornamental.

The second is *advice collapse*. Interfaces that present a single recommended action with a confidence score invite automation bias (Parasuraman & Manzey, 2010). If a recommendation is harder to refuse than to accept, the wall exists only on paper. Algorithmic shura therefore requires symmetry of effort: refusing machine advice should not be costlier, procedurally or reputationally, than accepting it.

The third is *frozen consent*. A system approved for a pilot can become permanent without renewed authorisation. Lifecycle monitoring addresses part of this problem (NIST, 2023; European Union, 2024). Algorithmic shura adds a political layer: monitoring informs re-authorisation; it does not replace it. A dashboard is not consent.

4. Objections and Limits

4.1 Is the machine a consultee?

One might read “algorithmic shura” as a claim that AI systems can be parties to consultation. It is not. Machines do not deliberate in the sense shura presupposes; they do not bear responsibility for counsel given, and they cannot be held to account for it. The consultees are human: officials, experts, affected persons, and oversight bodies. The machine is a source of advice and an object of consultation, not a member of the consultative circle. This keeps the framework on the side of those who resist attributing agency to AI in order to fill the responsibility gap (Matthias, 2004).

4.2 Does shura carry unwanted political commitments?

Shura has been invoked both to support and to resist democratic vocabularies (Shavit, 2010), and its concepts are plural (Thielicke-Witt, 2025). Algorithmic shura brackets that controversy. It uses shura for a narrower purpose: to insist that public AI remain under consultative human authority. One need not settle the relation between shura and democracy to accept that machine advice should not quietly become machine rule. Nor is the framework exclusionary: it does not claim that only Islamic concepts can ground AI ethics, and it does not license departures from fairness principles that the same Gulf instruments adopt.

4.3 Is this merely human oversight renamed?

The strongest objection is that algorithmic shura restates familiar human-oversight principles in religious vocabulary. We offer two replies. First, the shura frame shifts the locus of the problem from the individual operator's cognitive control to the *political relation* between counsel and authority. Meaningful human control asks whether a human can track and trace system behaviour; shura additionally asks whether those with standing were consulted before the system was adopted and whether authority to continue is periodically re-earned. Second, conceptual translation has independent value. A principle that a political community recognises as its own is more likely to be operationalised than one perceived as imported, a point Gulf strategies themselves make (MCIT, 2019). Cross-traditional work in the philosophy of technology is not decoration; it can reveal structural features of a problem that a single tradition leaves implicit.

4.4 Proportionality

Algorithmic shura does not demand that every routine automation be deliberated. Dubai's guidelines already distinguish significant decisions from lesser uses (Smart Dubai, 2019). The stricter the consequences for rights, livelihood, liberty, or equal treatment, the thicker the consultative and renewal requirements should be. Nor is algorithmic shura a substitute for law; it is a normative grammar for how advice enters public decision-making alongside statutory protections.

5. Design Implications

5.1 An advice–decision wall in the interface

For systems that inform significant decisions, guidance should require a written allocation of roles: what the system may recommend, who may decide, what reasons must be recorded, and how overrides are logged. Four artefacts can implement this: a role matrix stating whether each output is informational, recommendatory, or exceptionally automatic, and on what legal basis; an interface rule that recommendations appear with alternatives and an equally accessible refusal path; a reason log; and periodic audits checking whether reason logs have become empty rituals. Where a system cannot explain a recommendation in terms intelligible to the decider, it should not present that recommendation as the default.

5.2 Pre-deployment consultation dossiers

Before deploying a significant system, the responsible entity should record who was consulted, what harms and distributional effects were raised, what alternatives were considered, and how affected persons may challenge outcomes. Three political prompts should supplement existing ethics self-assessment: which public purpose authorises the system; who is most exposed if it errs; and what evidence would justify pausing or withdrawing it. The dossier must go to a named person who can refuse deployment. A dossier that nobody can reject is not consultation; it is paperwork.

5.3 Renewable authorisation and responsibility metrics

Significant systems should carry a review date at which continued use requires affirmative re-approval, informed by performance, complaints, drift, and legal change. Performance reports should include not only accuracy but override rates, challenge rates, and whether reviews ever change anything. Accuracy can improve while political responsibility deteriorates: a more accurate tool that no official feels able to refuse is a governance regression.

6. Conclusion

Algorithmic shura holds that machines may advise, humans must decide, and the authority to keep a system in public use must be consulted into being and periodically renewed. The framework is deliberately modest. It does not equate shura with electoral democracy, it does not treat algorithms as moral agents, and it does not claim that existing Gulf institutions already practise it in full. It claims that consultation, the advice–decision wall, and renewable authorisation are jointly necessary if government AI is to remain an instrument of public authority rather than a substitute for it.

Three tests keep the concept honest. The consultation test asks whether anyone who disagreed before deployment left a trace that decision-makers had to answer. The wall test asks whether a competent official can refuse a recommendation without disproportionate cost. The renewal test asks whether a system that fails its metrics or outlives its purpose is in fact paused or redesigned. A government that fails all three may still have an AI programme; it does not have algorithmic shura.

More broadly, the argument suggests that the philosophy of technology gains from traditions that have thought hard about counsel and authority. Historians may insist on sharper boundaries around classical shura, and lawyers may demand tighter links to positive rules. Those critiques would improve the framework. What would impoverish the debate is to proceed as if machine advice raised no distinctive question of public authority.

References

- Al-Raysuni, A. (2011). *Al-Shura: The Qur'anic principle of consultation*. International Institute of Islamic Thought. <https://iiit.org/en/book/al-shura-the-quranic-principle-of-consultation/>
- Arabian Business. (2017, October). *UAE appoints first Minister for Artificial Intelligence*. <https://www.arabianbusiness.com/politics-economics/381648-uae-appoints-first-minister-for-artificial-intelligence>
- Crone, P. (2001). Shura as an elective institution. *Quaderni di Studi Arabi*, 19, 3–39. https://www.ias.edu/sites/default/files/hs/Crone_Articles/Crone_Shura_as_an_Elective_Institution.pdf
- Danaher, J. (2016). The threat of algocracy: Reality, resistance and accommodation. *Philosophy & Technology*, 29(3), 245–268. <https://doi.org/10.1007/s13347-015-0211-1>
- European Union. (2024). *Regulation (EU) 2024/1689 of the European Parliament and of the Council laying down harmonised rules on artificial intelligence (Artificial Intelligence Act)*. Official Journal of the European Union. <https://artificialintelligenceact.eu/article/112/>
- Federal Decree-Law No. 45 of 2021 concerning the Protection of Personal Data. (2021). United Arab Emirates. <https://uaelegislation.gov.ae/en/legislations/1972/>
- Matthias, A. (2004). The responsibility gap: Ascribing responsibility for the actions of learning automata. *Ethics and Information Technology*, 6(3), 175–183. <https://doi.org/10.1007/s10676-004-3422-1>
- Ministry of Communications and Information Technology, State of Qatar. (2019). *National artificial intelligence strategy for Qatar*. https://www.mcit.gov.qa/-/media/mcit/documents/strategies/national_artificial_intelligence_strategy_for_qatar_2019_en.pdf
- 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>
- Organisation for Economic Co-operation and Development. (2024). *AI principles*. <https://www.oecd.org/en/topics/ai-principles.html>
- Parasuraman, R., & Manzey, D. H. (2010). Complacency and bias in human use of automation: An attentional integration. *Human Factors*, 52(3), 381–410. <https://doi.org/10.1177/0018720810376055>
- Peter, F. (2023). Political legitimacy. In E. N. Zalta & U. Nodelman (Eds.), *The Stanford encyclopedia of philosophy* (Winter 2023 ed.). <https://plato.stanford.edu/entries/legitimacy/>
- Regulations.AI. (2022). *AI ethics principles and guidelines (UAE)*. <https://regulations.ai/regulations/RAI-AE-NA-AEPGUX-2022>
- Santoni de Sio, F., & van den Hoven, J. (2018). Meaningful human control over autonomous systems: A philosophical account. *Frontiers in Robotics and AI*, 5, Article 15. <https://doi.org/10.3389/frobt.2018.00015>
- Saudi Data and Artificial Intelligence Authority. (2023). *AI ethics principles* (Version 1.0). <https://dgp.sdaia.gov.sa/wps/wcm/connect/3fcea0d-bbb6-4ead-9cab-c141c2eca42d/ai-principles-EN.pdf?MOD=AJPERES>
- Shavit, U. (2010). Is shura a Muslim form of democracy? Roots and systemization of a polemic. *Middle Eastern Studies*, 46(3), 349–374. <https://doi.org/10.1080/00263200902917085>
- Smart Dubai. (2019). *AI ethics principles & guidelines*. <https://www.digitaldubai.ae/docs/default-source/ai-principles-resources/ai-ethics.pdf>
- The National. (2017). *Sheikh Mohammed bin Rashid announces UAE Cabinet reshuffle*. <https://www.thenationalnews.com/uae/sheikh-mohammed-bin-rashid-announces-uae-cabinet-reshuffle-1.668631>
- Thielicke-Witt, V. (2025). Comparing concepts of shura: Insights for a further deparochialization of democracy. *Comparative Political Theory*, 5(1), 1–24. <https://doi.org/10.1163/26669773-bja10075>

Tuckness, A. (2020). Locke's political philosophy. In E. N. Zalta (Ed.), *The Stanford encyclopedia of philosophy*. <https://plato.stanford.edu/entries/locke-political/>

UAE Cabinet. (2019, April 21). *UAE Cabinet adopts National Artificial Intelligence Strategy 2031*. <https://uaecabinet.ae/en/news/uae-cabinet-adopts-national-artificial-intelligence-strategy-2031>

UAE Government. (n.d.). *The Federal National Council*. <https://u.ae/en/about-the-uae/the-uae-government/the-federal-national-council->

UAE Government. (2024). *The UAE Charter for the Development and Use of Artificial Intelligence*. <https://u.ae/en/about-the-uae/strategies-initiatives-and-awards/policies/Ai/The-UAE-Charter-for-the-Development-and-Use-of-Artificial-Intelligence>

United Arab Emirates. (n.d.-a). *Constitution of the United Arab Emirates* (Unofficial English text, WIPO Lex). <https://www.wipo.int/wipolex/en/text/577564>

United Arab Emirates. (n.d.-b). *UAE national strategy for artificial intelligence 2031*. https://wp.oecd.ai/app/uploads/2021/12/United_Arab_Emirates_National_Strategy_for_Artificial_Intelligence_2017-2031.pdf

United Arab Emirates Council for Artificial Intelligence and Blockchain. (n.d.). *AI Council*. https://ai.gov.ae/ai_council/

WAM. (2024, July 30). *Artificial Intelligence Office launches “UAE Charter for Development and Use of AI”*. <https://www.wam.ae/en/article/b4f7n7s-artificial-intelligence-office-launches-%E2%80%98uae>