

AI-Driven Smart Contracts for Automated Compliance and Governance in Blockchain-Based Public-Sector Platforms.

Abstract:

Public sector digital frameworks are increasingly combining distributed ledgers and programmable intelligent contracts to make administrative transactions verifiable, traceable and automatically executable. But traditional intelligent contracts are rigid: they can enforce pre-defined logic but cannot interpret ambiguous regulations, identify changing risk or explain why a transaction should be escalated. This disadvantage can be mitigated to some degree by using artificial intelligence for assisting with the interpretation of regulatory text, anomaly detection, predictive risk scoring, document classification, and adaptive workflow control. This review offers an in-depth analysis of the potential for AI-enabled smart contract technology to automate compliance and governance in blockchain-based public-sector platforms while maintaining legality, accountability, privacy, security and administrative discretion. The integrative review design was developed by synthesising peer-reviewed work published between 2020 and 2025, mainly through discovery pathways in Scopus, Web of Science, IEEE Xplore, ScienceDirect, SpringerLink, and Google Scholar. The review describes the search logic, eligibility criteria, quality assessment and thematic synthesis process instead of invented screening numbers. Five themes are identified: architectural convergence of AI and intelligent contracts; machine-readable compliance and regulatory semantics; governance, accountability and public value; security, privacy and oracle assurance; and organisational adoption and interoperability. The evidence suggests that the strongest model is not autonomous algorithmic government but bounded automation, where AI proposes classifications or risk signals and intelligent contracts execute formally authorised rules. Accountable officials may override, appeal and change control. A reference architecture and governance-assurance life cycle links technical controls to legal and managerial obligations. The review finds gaps in empirical public-sector evaluations, explainability standards, rule-update governance, cross-agency interoperability and assurance of AI-produced contract actions. Future research must test measurable effects such as compliance accuracy, false detection rates, audit latency, appeal resolution, administrative burden, and citizen trust under realistic regulatory conditions.

Key words: artificial intelligence, smart contract, blockchain, public sector, compliance automation, digital government, algorithmic governance, regtech

1.Introduction

Digital government has expanded from providing services online to creating platforms that bring together data, identity, payments, licensing, procurement, benefits, and regulatory reporting across many institutions. Blockchain research suggests that such shared ledgers can improve some aspects of traceability, tamper evidence, and inter-organizational coordination, and intelligent contracts add programmability to such shared records (Ante, 2021; Hewa et al., 2021). Public sector applications are particularly attractive in situations where several agencies need to share a transaction history without the need to continually reconcile multiple databases. However, the reviews of the government blockchain projects indicate that the technological transparency does not automatically lead to the organisational accountability, managerial efficiency, and public trust. These outcomes depend on the architecture, governance, law, organisational readiness, and quality of the underlying data (Verma and Sheel, 2022; Tan et al., 2022). Smart agreements increase the opportunity and the risks. They can bake eligibility requirements, procurement milestones, reporting requirements, payment triggers, or licensing constraints into software that runs when conditions are met. In public administration this could mean less time on routine processing and a reliable audit trail. But the same automation can also entrench errors, entrench contested interpretations or scale an obsolete rule. Technical surveys repeatedly find vulnerabilities, verification issues, privacy limitations, and irreversible outcomes when contract logic is defective (Sayeed et al., 2020; Almakhour et al., 2020; Hu et al., 2021). The central governance challenge is not whether code can enforce a rule, but how a public institution can demonstrate that the encoded rule is legally valid, situationally appropriate, technically secure, reviewable and changeable through legitimate authority. Artificial intelligence changes this problem because it can provide capabilities that deterministic contracts do not have. Machine-learning models can flag suspicious transactions, evaluate compliance risk, sort documents, extract obligations from regulatory text, or flag outlier cases for human review. Intelligent contracts research views AI and blockchain as complementary technologies: the blockchain provides persistent and auditable state, while the AI provides inference and adaptive decision support (Ouyang et al., 2022; Hussain and Al-Turjman, 2021). But the combination of

probabilistic AI and deterministic smart contract execution does create a new accountability boundary. Model output can be uncertain, data drift sensitive, difficult to explain, or biased. If that output directly results in an irreversible administrative action, traditional procedural safeguards may be weakened rather than enhanced. In the public sector, these concerns are more stringent than in many commercial deployments. Government decisions can impact legal status, access to services, financial entitlements, regulatory sanctions and rights of appeal. This means that public organisations must balance automation with principles such as legality, proportionality, transparency, privacy, reviewability and equal treatment. Empirical and conceptual work on blockchain in government also notes that public value is a function of governance choices as much as it is of decentralised infrastructure (Cagigas et al., 2022; Semenzin et al., 2022). Yet AI governance research adds requirements such as organisational responsibility, training, standards, documented processes, and ethical oversight (Almeida & dos Santos Júnior, 2025). This review discusses the intersection of these literatures. It is about AI-enabled intelligent contracts as a socio-technical governance mechanism, not as a computational novelty. The focus is on public sector platforms built on blockchain where automated compliance has to be auditable and contestable. The review distinguishes three often conflated functions: artificial intelligence is able to score or interpret information, intelligent contracts can execute explicit authorised rules, and institutions have to govern both the models and the executable rules. This distinction provides a basis for deciding when automation can improve public administration and when human authorisation, review, or override is still needed.

2. Purpose of Study

This paper seeks to critically synthesise recent scholarship on the integration of artificial intelligence, smart contracting, and blockchain into public-sector platforms and to develop an evidence-informed framework for automating compliance without displacing public accountability. It examines how AI-supported interpretation and risk analysis can be integrated with deterministic contract execution, what technical and institutional controls are necessary, and how public organisations can preserve legal review, privacy, security, and human accountability across the automation lifecycle.

3. Research Goals

The review has 5 objectives. It begins by looking at the technical roles that artificial intelligence can play in relation to blockchain smart contract technology, such as the interpretation of regulatory text, anomaly detection, risk scoring, and workflow management. Second, it assesses the potential of intelligent contracts to translate compliance requirements into executable controls in the public sector while maintaining traceability to authoritative legal sources. Third, it compares the governance and accountability challenges of explainability, oversight by humans, appeals, rule changes, and institutional responsibility. Fourth, it incorporates security, privacy, oracle, interoperability, and verification requirements that decide if automated compliance is trustworthy. Fifth, it uncovers gaps of empirical and methodological research, and extracts pragmatic design principles and evaluation metrics for future deployments in the public-sector.

4. Research Questions:

The synthesis is guided by five research questions. RQ1: Which AI capabilities best support compliance in public administration by means of intelligent contracts? RQ2: How to transform regulatory obligations into machine-executable rules, without losing legal meaning and administrative discretion. RQ3 addresses the question of what governance arrangements can render AI-driven contract decisions accountable, explainable, reviewable and appealable. RQ4 What security, privacy, oracle and interoperability controls are needed for reliable deployment across government platforms? RQ5 asks what evidence gaps, performance measures, and institutional conditions should guide future research and implementation.

5. Methodology for Review

We selected an integrative structured review because the topic spans computer science, information systems, public administration, regulation and AI governance. The approach was

intended to integrate heterogeneous evidence, not to calculate a pooled effect size. The attached reference paper served only as a structure for transparent search reporting, thematic organisation and explicit quality appraisal; its wording, findings and numerical screening claims have not been reused. The searches were based on Scopus, Web of Science, IEEEExplore, ScienceDirect, SpringerLink and Google Scholar as a secondary discovery channel. The primary search combinations were “AI smart contract compliance”, “intelligent contracts blockchain”, “smart contract regulatory compliance”, “blockchain public sector governance”, “government blockchain adoption”, “smart contract verification”, “blockchain oracle”, “blockchain privacy compliance” and “AI governance public sector”. Citation chaining was used to identify directly related peer-reviewed studies that may have been missed by keyword variation. The eligibility criteria were focused on peer-reviewed articles in English and rigorously reviewed conference papers mainly published between 2020 and 2025. Previous literature was only reviewed in cases where it was necessary to define a technical or governance concept. The contemporary window contained the final reference set. Studies were included if they covered at least one of five domains: smart-contract architecture or verification; AI-blockchain integration; automated or machine-readable compliance; public-sector blockchain governance or adoption; or assurance issues such as privacy, security, identity, or oracles. Studies purely on cryptocurrency price prediction, speculative token markets, or unrelated private sector automation were excluded. The final evidentiary base also excluded non-peer reviewed promotional reports, blogs, vendor claims and unverifiable preprints. The screening process was conceptually designed in three steps: relevance of titles and abstracts, relevance of full texts to the research questions and, final suitability of the evidence. We do not report numerical flow counts, as presenting such numbers without a reproducible exported search log would give a false sense of precision. The quality assessment was informed by source credibility, methodological transparency, clarity of research design, relevance to public governance or smart-contract assurance, and whether claims were supported by empirical data, formal analysis, systematic review methods, or clearly bounded conceptual reasoning. Technical studies were weighted more heavily for reproducibility and explicit threat models, and public-management studies were weighted for transparency of sampling, analytical procedure, and institutional context. The synthesis was done using iterative thematic coding. Initial codes were AI function, contract function, regulatory semantics, governance authority, oversight by humans, security, privacy,

identity, oracle dependence, interoperability, adoption, and public value. These codes were then grouped into five themes and compared across disciplinary perspectives. The contradictions were not averaged out, but were preserved. For example, the immutability of blockchain is both an audit advantage and a governance constraint when rules, data or legal interpretations need correction. Like-wise, the adaptability of AI is a value in risk detection but problematic when automatic administrative results require stable and explainable legal bases. The resulting synthesis is therefore analytical, not celebratory. The methodological limitations were clearly recognised. The coverage of databases and terminology varies. Many government projects are still in pilot stage. References and DOIs were checked against indexing records and verified independently for citation accuracy.

6. Critical Analysis & Thematic Literature Review

Smart-contract research provides the technical basis for automated compliance, but also shows why public deployment cannot rest on code alone. In the literature, intelligent contracts have been described as event-driven programmes stored and executed on blockchain platforms to support transparent and repeatable process automation (Ante, 2021; Khan et al., 2021). Their design space covers contract languages, off-chain components, privacy extensions, execution frameworks and verification tools (Hu et al., 2021). The relevant benefit for public platforms is not decentralisation for its own sake but the ability to tie administrative actions to a shared, tamper-evident state and to execute authorised rules consistently across organisational boundaries. AI is a second layer of computation. Ouyang et al. (2022) distinguish intelligent-contract approaches, where AI models are encoded into or coordinated by intelligent contracts. The idea of a fully on-chain AI is generally not feasible for governments to implement as complex models are expensive, data-intensive, and hard to update. A more defensible architecture is one in which most model training and inference takes place off-chain, while model identifiers, input hashes, risk outputs, approvals, and execution events are recorded on a permissioned ledger. This provides traceability for AI decisions, without claiming that probabilistic inference is the same as deterministic legal logic. Hussain and Al-Turjman (2021) also considered blockchain and AI as complementary technologies. But public governance demands a better division of responsibility: AI should be used to help classify or forecast, but

180 legal enforcement needs to rely on authorised rules and explicit thresholds.
181 One of the big use cases is reading regulations. Natural language processing can identify duties,
182 entities, deadlines, exceptions, and evidence requirements from statutes, regulations, policies,
183 and guidance. But legal text is ambiguous, cross-referenced, hierarchical and discretionary. This
184 is a risk of false formalisation: translating such text directly into code. Smart agreements are
185 better seen as executable representations of agreed operational rules, connected to authoritative
186 legal sources and version-controlled interpretations. Bartoletti (2020) advocates for stronger
187 specification of smart-contract behaviour. Almakhour et al. (2020) demonstrate the necessity of
188 verification prior to immutable deployment. Thus, for public compliance, formal verification
189 should examine not only software properties, but also the traceability between the policy intent,
190 executable conditions, allowable data inputs, and exception paths. This requirement translates
191 into a layered compliance model. The first level is responsible specialists who interpret
192 authoritative sources of law and policy. At the second level requirements are then converted into
193 structured rule objects with provenance, effective dates, jurisdiction, exceptions and responsible
194 authority. The third layer encodes deterministic rules in intelligent contracts. AI proposes
195 classifications, finds anomalies, ranks cases and signals potential rule conflicts across these
196 layers. New rule versions are approved by human reviewers' Human reviewers resolve material
197 uncertainty cases. This limited automation avoids the category error of a statistical prediction
198 silently becoming a legal determination.

199
200 This institutional view is strongly supported in the public-sector literature. Triana Casallas et al.
201 (2020) find transparency and automated transactions relevant for intelligent contracts in the
202 public sector. Alexopoulos et al. (2021) show that government blockchain applications vary
203 widely in terms of purpose and architecture. According to Lykidis et al. (2021), e-government
204 blockchain can enhance integrity and trust but raises concerns about privacy, scalability and legal
205 issues. Kassen (2022) argues that blockchain can automate processes in the public-information,
206 but also points to regulatory risk and information leakage. Together, these studies suggest that
207 intelligent contracts add value through controlled process redesign rather than wholesale
208 replacement of administrative judgement. Governance architecture is as important as the process
209 design (Ellul et al., 2020). Tan et al. (2022) differentiate between blockchain governance
210 decisions at the infrastructure, application, and broader institutional levels. Liu et al. (2022) also

emphasise stakeholders, accountability, legal compliance, and ethical responsibility in blockchain governance. At a minimum, we need to differentiate between four authorities in the case of AI-driven intelligent contracts: the authority that owns the policy; the team that translates the policy into executable rules; the technical operator that maintains the platform; and the oversight function that audits performance and handles challenges or appeals. While this may be efficient during a pilot, the combination of these authorities in one technical team creates weak separation of duties for production government systems. Public value research also cautions against equating immutable records with trustworthy government. In a study of Estonia, Semenzin et al. (2022) show that “blockchain” can accommodate competing institutional narratives and may reinforce centralised practices. Cagigas et al. (2022) state that the acceptance of blockchain by public officials is dependent on the way it is configured and perceived and not only on technical capability. Falwadiya and Dhingra (2022) identify trust, transparency, costs, security and enabling conditions as adoption factors. The findings are especially relevant for AI-driven compliance, because the legitimacy of the public sector depends on whether officials and citizens understand who is left responsible when automated systems act. Another important theme is security and correctness. Sayeed et al. (2020) describe common classes of smart-contract attacks and limitations of detection tools. Kemmoe et al. (2020) review design and access-control issues in smart-contract systems. Formal and semi-formal assurance is important because a defective contract can commit an error forever. Software assurance in government should include secure coding, independent review, formal specification for high impact logic, test coverage, role-based permissions, key management, emergency pause functions, and controlled upgrade mechanisms. The rule of immutability should be confined to the evidentiary history, not to the inability to correct a defective administrative rule. Oracles are the critical weak spot whenever a contract depends on off-chain facts. Caldarelli and Ellul (2021) demonstrate how the design of oracles brings back trust assumptions into blockchain systems, and Sheldon (2021) defines oracle providers as assurance-relevant service organisations. Public sector oracles might include registries, identity systems, tax records, inspection databases, IoT sensors or AI inference services. The fact that an external datum is delivered cryptographically does not tell a smart contract if the datum is correct. So government platforms need provenance, source authorisations, freshness checks, multi-source reconciliation for high-risk facts, and audit procedures on systems feeding external data.

Privacy is another constraint. Permissioned architectures may restrict participation but do not address the tension between transparent audit trails and personal-data obligations. Zero-knowledge mechanisms can provide selective proof of compliance without disclosing full transaction histories, as shown by Buterin et al. (2024). The public-sector implication is wider than financial privacy: systems should lower personal data on-chain, have sensitive content off-chain where correction and retention controls are available under law, and put hashes, proofs, permissions and event references on-chain. Similarly, identity-management research points to the benefits of verifiable credentials, but warns that complexity of implementation and governance remain significant (Sung and Park, 2021). Studies of government adoptions show that interoperability and institutional fit cannot be left as later implementation details. Zhang et al. (2023) survey blockchain for government information sharing and discuss challenges in cross-chain, throughout, and supervision. Verma and Sheel (2022) observe a field still in flux from experimentation to operational integration. Khan et al. (2022) explored the Dubai case to show how regulatory checks, document authenticity, accountability and service efficiency can facilitate government blockchain adoption. But it also suggests the need for empirical validation beyond a single deployment. Rodriguez Müller et al. (2025) report similar findings for European public sector blockchain cases that aim to generate trust, efficiency, accountability, transparency, security and cooperation, while calling for more evidence on realised rather than expected public values.

First and foremost, AI governance must be integrated, not appended as a separate ethics checklist. Almeida and dos Santos Júnior (2025) reveal how AI governance is applied in public organisations through processes, standards, training and organisational practices. Model governance for AI-enabled intelligent contracts should encompass training data provenance, intended use, validation, bias testing, performance thresholds, drift monitoring, explainability, incident handling, and retirement. Contract governance should include rule origin, legal authorisation, verification, deployment, versioning, and emergency intervention. The two governance regimes intersect at the decision interface, which is the precise point at which an AI output becomes an input to an executable rule. The interface should be well documented and auditable.




Figure 1. Reference architecture for bounded AI-driven smart-contract compliance in public-sector platforms.

Technology feature	Potential public value	Primary risk	Required safeguard
Smart-contract automation	Consistent execution and faster processing	Errors scale automatically	Formal specification, testing, human exception path
Blockchain immutability	Auditable history and shared evidence	Incorrect data or rules are difficult to correct	Off-chain sensitive data, governed upgrades, versioned history
AI adaptiveness	Risk detection and handling of unstructured inputs	Opacity, bias, drift, uncertainty	Validation, confidence thresholds, explainability, monitoring
Cross-agency sharing	Reduced reconciliation and duplicated verification	Unclear ownership and interoperability failure	Data standards, node governance, responsibility matrix
Oracle connectivity	Use of real-world facts in automated decisions	External data may be inaccurate or compromised	Authorized sources, provenance, freshness checks, assurance

Table 1. Core design tensions and safeguards identified in the literature.

7. Discussion

The literature supports a bounded-automation model of public-sector compliance. Deterministic intelligent contracts are at their most powerful when the rules are stable, objectively testable and based on reliable data: for example, confirming whether a license is current, whether a filing came in before a deadline, or whether a payment milestone has been formally approved. AI works best when the inputs are unstructured, patterns are uncertain, or where prioritisation is required. Examples include classifying documents, detecting anomalous procurement behaviour, estimating the risk that a filing requires inspection, or matching regulatory language to a rule library. The problems arise when the boundary is concealed and an AI prediction is treated as legal fact.

Thus, the proposed architecture decouples sensing, inference, authorisation, execution and review. Data comes from controlled agency systems, trusted external services, or citizen submissions. AI services produce structured outputs, such as model metadata and confidence values. A policy engine then decides whether the output can satisfy a contract condition, require corroboration, or needs human review. The smart contract does only allowed operations, and creates an evidentiary trail. Oversight services continuously monitor outcomes, exceptions, overrides, appeals, model drift and contract versions. In Figure 1, this architecture is shown as a sequence of accountable transformations, not as a single autonomous “AI smart contract.” The separation also makes clear who is responsible. The model developer is responsible for the quality of a model, but not for the legal effect of a score. The policyholder determines what evidence is acceptable and what thresholds are legally or administratively enforceable. The smart-contract developer implements that approved logic and demonstrates that it is technically correct. The platform operator provides availability, keys, interfaces and logging. An independent assurance function to determine if the system remains compliant and if citizens or regulated entities can challenge adverse outcomes. Such role separation operationalizes the governance concerns identified by Tan et al. (2022) and Liu et al. (2023), but without losing the efficiency advantages described in smart-contract research. The review also highlights designing explainability for multiple audiences. Developers need technical explanations of the behaviour of models and the state transitions of contracts. Case officers need short explanations of flags, confidence and the rules that apply. Auditors need repeatable evidence that connects inputs, model version, rule version, approvals and execution.

Citizens need to understand the reasons, have easy access to appeal channels and clarity on whether the outcome was automated. There is no single, generic explanation that can satisfy all four needs. So public platforms should generate layered explanations from the same decision record.

Governance will also have to be able to change. Regulations change, policies are re-interpreted, models drift, vulnerabilities are found, administrative priorities shift. This fact is not compatible with a rigid immutable contract. The best practice is fixed history + controlled mutability of active logic. That is, each rule version is approved, time-bounded, cryptographically identified, and associated with the authority that approved it. Upgrades are multi-party approved, tested and controlled with rollback or pause. it's a way of treating the ledger as a durable record of governance and not an excuse to make governance impossible to revise. Table 1 summarises the key design tensions. In the literature, the apparent strengths of blockchain, smart contract technology, and AI are only public value when paired with corresponding safeguards. Automation reduces manual handling. But, it can scale error. Transparency strengthens auditability. But, it can expose sensitive information. Adaptiveness improves detection. But, it can weaken predictability. Immutability preserves evidence. But, it complicates correction. Thus, instead of maximum automation, the right goal is controlled, measurable and reviewable automation.

8. Future research and research gaps

The most important gap is empirical. Much of the literature describes architectures, adoption factors or expected benefits. Comparatively few studies evaluate mature AI-smart-contract government services under real regulatory conditions. Future research should employ longitudinal case studies, controlled pilots, and comparative designs to measure whether automation actually leads to reduced processing time, administrative cost, error, fraud exposure, or audit effort, without increasing unjustified denials or unresolved appeals. Studies should report baseline processes so that the effects of blockchain and artificial intelligence can be separated from digitisation in general. A second gap is the legal semantics of executable rules. There's a need for research on traceable pipelines from legislation and policy to structured requirements, verified contract code and

citizen-facing explanations. Such work should explore how exceptions, discretionary standards, conflicting rules and temporal changes are represented. The version control needs to include the legal authority, interpretation memorandum, effective date and responsible approver for each executable rule, not just the software commits. Third, explainability research should shift from generic AI explanations to administrative due-process requirements. Experiments should test which formats of explanation enable case officers, auditors and affected persons to identify errors and exercise review rights. Metrics should be completeness of explanations, time to reconstruct a decision, success of an appeal due to a system error, and consistency between the machine explanation and the governing rule. Fourth, assurance research should target the AI-to-contract interface. The current smart-contract verification shows properties of the code, and the AI validation estimates properties of the model; neither one alone shows that the combined system behaves according to the law. New assurance methods should verify end-to-end properties such as that low-confidence predictions are always escalated, that protected attributes are omitted when needed, that oracle freshness is enforced, and that adverse actions are blocked when required evidence is missing. Fifth, further study is needed on cross-agency interoperability. Public platforms will often cross identity, licensing, procurement, tax and sector regulators. Future work should compare shared data models, verifiable credentials, privacy-preserving proofs, and governance arrangements for interagency nodes. Failure scenarios where agencies have different ideas about data quality, rule interpretation or responsibility should also be investigated. Finally, public-value studies should investigate distributional effects, not just average efficiency. Even if it is operationally efficient, a system that is efficient in most cases but systematically burdens some groups is institutionally unacceptable.


9. Implications for practice, management and policy

The main lesson for public sector leaders is to first tackle a governance problem, not a blockchain purchase. Candidate processes must be multi-party, have significant reconciliation, have strong auditability and have rules that can be made sufficiently explicit. Traditional databases work best when you can have one trusted authority handle the process efficiently and there is no shared verification problem. Permissioned networks better satisfy government requirements around identity, access control, confidentiality, and accountable node governance

when blockchain makes sense.

Before implementation, managers should establish a joint policy-technology control board. The board should approve rule semantics, AI use cases, data sources, risk thresholds, model and contract changes, emergency procedures and appeal mechanisms. High impact decisions should be performed with human-in-the-loop or human-on-the-loop controls, depending on the legal requirements and the consequences of an error. Procurement specifications should specify model documentation, smart-contract source availability, independent security testing, interoperability standards, audit exports, key-management procedures, and exit arrangements to prevent vendor lock-in.

Policy makers should distinguish three objects of regulation: the administrative rule, the AI inference process and the blockchain execution mechanism. Each needs a different approach to assurance. The rule is bound by legal authority and procedural fairness; the model by statistical validation and responsible-AI controls; and the contract and ledger by software assurance and cryptographic controls. Treating the combined platform as a single opaque technology under regulation will make accountability harder. Figure 2 thus depicts a governance-assurance cycle where policy authorisation, translation, testing, deployment, monitoring, challenge and revision are continuous and documented. Performance management also needs to be multidimensional. Agencies should measure transaction time and cost as well as false positives, false negatives, frequency of overrides, unexplained outcomes, oracle failures, security incidents, latency in rule-updating, volume of appeals, time to resolve appeals, and disparities across relevant service populations. Table 2 shows these metrics against governance controls. Publishing appropriately aggregated performance indicators can enhance transparency without disclosing personal or security sensitive data.



399 **Figure 2. Governance-assurance lifecycle for AI-driven smart-contract public-sector**
400 **systems.**

401 **Table 2. Suggested operational metrics and corresponding governance controls.**

Metric	Governance purpose	Example control
Compliance false-positive / false-negative rate	Detect harmful model or rule errors	Threshold review and periodic validation
Human override frequency	Reveal mismatch between automation and case reality	Reason-coded overrides and policy review

Rule-update latency	Measure responsiveness to legal change	Effective-date register and controlled deployment
Oracle failure / stale-data rate	Protect factual reliability	Source redundancy, freshness limits, incident escalation
Appeal volume and resolution time	Assess contestability and procedural fairness	Accessible appeal workflow and audit reconstruction
Security incidents and emergency pauses	Assess operational resilience	Key controls, monitoring, pause authority, post-incident review

402

403

404 10. Limitations

405

406

There are four limitations to this review. It also relies on several disciplines, each with its own terminology and evaluation standards, so that direct comparison is at times interpretive, rather than statistical. Second, the evidence from the public sector remains patchy: many blockchain projects are still pilots, and the compliance of AI-driven intelligent contracts is less mature than either blockchain government research or AI governance research in isolation. Third, our review focuses on peer-reviewed literature in English and may overlook relevant public-administration studies published in other languages. Fourth, the rapid development of AI models, privacy technology, and blockchain platforms means that particular technical implementations can quickly become obsolete. These limitations are partially compensated by an emphasis on governance principles, assurance requirements and measurable decision interfaces rather than any particular platform or programming language.

11. Conclusions

AI-powered intelligent contracts can enhance public-sector compliance, but they should be treated as regulated elements of administrative systems rather than independent replacements for government decisions. The literature proposes a plausible division of labour: artificial intelligence can interpret unstructured information, identify anomalies and prioritise risk; intelligent contracts can execute explicit, authorised and verifiable rules; blockchain can maintain shared evidence and transaction history; and public institutions must retain responsibility for authorisation, oversight, correction and appeal. This is more defensible than directly embedding opaque model outputs in irreversible administrative actions.

The review also suggests that technical trust cannot replace institutional trust. Secure code, immutable records, cryptographic proofs, all good things. But they do not determine whether a rule is lawful, whether a model is fair, whether an oracle is accurate, or whether an affected person received a meaningful explanation. The value in the public sector is created where technical mechanisms are linked to accountable roles, traceable legal sources, controlled change, privacy safeguards, independent assurance and accessible review procedures. Therefore, the next phase of research should progress beyond feasibility demonstrations. Mature evaluation must test end-to-end administrative outcomes and failures, including model uncertainty, rule changes, interagency disagreement, privacy constraints, security incidents, and appeals. The core design principle is bounded, contestable automation: automate what can be specified and assured, escalate what cannot, keep evidence of every transformation, and ensure that legitimate public authority is able to review and correct the system. In such a context, AI-enabled intelligent contracts could contribute to the acceleration, standardisation and transparency of public services, without jeopardising the legal and democratic guaranties that distinguish public administration from purely technical transaction processing.

1. References

2. Ante, L. (2021). Smart contracts on the blockchain: A bibliometric analysis and review. *Telematics and Informatics*, 57, 101519. <https://doi.org/10.1016/j.tele.2020.101519>
3. Hewa, T., Ylianttila, M., & Liyanage, M. (2021). Survey on blockchain based smart contracts: Applications, opportunities and challenges. *Journal of Network and Computer Applications*, 177, 102857. <https://doi.org/10.1016/j.jnca.2020.102857>
4. Hu, B., Zhang, Z., Liu, J., Liu, Y., Yin, J., Lu, R., & Lin, X. (2021). A comprehensive survey on smart contract construction and execution: Paradigms, tools, and systems. *Patterns*, 2(2), 100179. <https://doi.org/10.1016/j.patter.2020.100179>
5. Khan, S. N., Loukil, F., Ghedira-Guegan, C., Benkhelifa, E., & Bani-Hani, A. (2021). Blockchain smart contracts: Applications, challenges, and future trends. *Peer-to-Peer Networking and Applications*, 14(5), 2901–2925. <https://doi.org/10.1007/s12083-021-01127-0>
6. Kemmoe, V. Y., Stone, W., Kim, J., Kim, D., & Son, J. (2020). Recent advances in smart contracts: A technical overview and state of the art. *IEEE Access*, 8, 117782–117801. <https://doi.org/10.1109/ACCESS.2020.3005020>
7. Sayeed, S., Marco-Gisbert, H., & Caira, T. (2020). Smart contract: Attacks and protections. *IEEE Access*, 8, 24416–24427. <https://doi.org/10.1109/ACCESS.2020.2970495>
8. Almakhour, M., Sliman, L., Samhat, A. E., & Mellouk, A. (2020). Verification of smart contracts: A survey. *Pervasive and Mobile Computing*, 67, 101227. <https://doi.org/10.1016/j.pmcj.2020.101227>
9. Bartoletti, M. (2020). Smart contracts contracts. *Frontiers in Blockchain*, 3, 27. <https://doi.org/10.3389/fbloc.2020.00027>
10. Caldarelli, G., & Ellul, J. (2021). The blockchain oracle problem in decentralized finance: A multivocal approach. *Applied Sciences*, 11(16), 7572. <https://doi.org/10.3390/app11167572>
11. Sheldon, M. D. (2021). Auditing the blockchain oracle problem. *Journal of Information Systems*, 35(1), 121–133. <https://doi.org/10.2308/ISYS-19-049>

12. Ellul, J., Galea, J., Ganado, M., McCarthy, S., & Pace, G. J. (2020). Regulating blockchain, DLT and smart contracts: A technology regulator's perspective. *ERA Forum*, 21(2), 209–220. <https://doi.org/10.1007/s12027-020-00617-7>
13. Triana Casallas, J. A., Cueva Lovelle, J. M., & Rodríguez Molano, J. I. (2020). Smart contracts with blockchain in the public sector. *International Journal of Interactive Multimedia and Artificial Intelligence*, 6(3), 63–72. <https://doi.org/10.9781/ijimai.2020.07.005>
14. Ouyang, L., Zhang, W., & Wang, F.-Y. (2022). Intelligent contracts: Making smart contracts smart for blockchain intelligence. *Computers & Electrical Engineering*, 104, 108421. <https://doi.org/10.1016/j.compeleceng.2022.108421>
15. Hussain, A. A., & Al-Turjman, F. (2021). Artificial intelligence and blockchain: A review. *Transactions on Emerging Telecommunications Technologies*, 32(9), e4268. <https://doi.org/10.1002/ett.4268>
16. Buterin, V., Illum, J., Nadler, M., Schär, F., & Soleimani, A. (2024). Blockchain privacy and regulatory compliance: Towards a practical equilibrium. *Blockchain: Research and Applications*, 5(1), 100176. <https://doi.org/10.1016/j.bcra.2023.100176>
17. Zhang, L., Mao, J., An, Y., Zhang, T., Ma, J., Feng, C., & Zhou, X. (2023). A systematic review of blockchain technology for government information sharing. *Computers, Materials & Continua*, 74(1), 1161–1181. <https://doi.org/10.32604/cmc.2023.032452>
18. Verma, S., & Sheel, A. (2022). Blockchain for government organizations: Past, present and future. *Journal of Global Operations and Strategic Sourcing*, 15(3), 406–430. <https://doi.org/10.1108/JGOSS-08-2021-0063>
19. Tan, E., Mahula, S., & Cromptvoets, J. (2022). Blockchain governance in the public sector: A conceptual framework for public management. *Government Information Quarterly*, 39(1), 101625. <https://doi.org/10.1016/j.giq.2021.101625>
20. Semenzin, S., Rozas, D., & Hassan, S. (2022). Blockchain-based application at a governmental level: Disruption or illusion? The case of Estonia. *Policy and Society*, 41(3), 386–401. <https://doi.org/10.1093/polsoc/puac014>
21. Cagigas, D., Clifton, J., Díaz-Fuentes, D., Fernández-Gutiérrez, M., Echevarría-Cuenca, J., & Gilsanz-Gómez, C. (2022). Explaining public officials' opinions on blockchain

- adoption: A vignette experiment. *Policy and Society*, 41(3), 343–357.
<https://doi.org/10.1093/polsoc/puab022>
22. Lykidis, I., Drosatos, G., & Rantos, K. (2021). The use of blockchain technology in e-government services. *Computers*, 10(12), 168.
<https://doi.org/10.3390/computers10120168>
23. Alexopoulos, C., Charalabidis, Y., Loutsaris, M. A., & Lachana, Z. (2021). How blockchain technology changes government: A systematic analysis of applications. *International Journal of Public Administration in the Digital Age*, 8(1), 1–20.
<https://doi.org/10.4018/IJPADA.20210101.0a10>
24. Kassen, M. (2022). Blockchain and e-government innovation: Automation of public information processes. *Information Systems*, 103, 101862.
<https://doi.org/10.1016/j.is.2021.101862>
25. Falwadiya, H., & Dhingra, S. (2022). Blockchain technology adoption in government organizations: A systematic literature review. *Journal of Global Operations and Strategic Sourcing*, 15(3), 473–501. <https://doi.org/10.1108/JGOSS-09-2021-0079>
26. Khan, S., Shael, M., Majdalawieh, M., Nizamuddin, N., & Nicho, M. (2022). Blockchain for governments: The case of the Dubai government. *Sustainability*, 14(11), 6576.
<https://doi.org/10.3390/su14116576>
27. Rodriguez Müller, A. P., Martin Bosch, J., & Tangi, L. (2025). An overview of the expected public values arising from blockchain adoption in the European public sector. *International Journal of Public Sector Management*, 38(1), 53–76.
<https://doi.org/10.1108/IJPSM-12-2023-0363>
28. Sung, C. S., & Park, J. Y. (2021). Understanding of blockchain-based identity management system adoption in the public sector. *Journal of Enterprise Information Management*, 34(5), 1481–1505. <https://doi.org/10.1108/JEIM-12-2020-0532>
29. Liu, Y., Lu, Q., Zhu, L., Paik, H.-Y., & Staples, M. (2023). A systematic literature review on blockchain governance. *Journal of Systems and Software*, 197, 111576.
<https://doi.org/10.1016/j.jss.2022.111576>
30. Liu, Y., Lu, Q., Yu, G., Paik, H.-Y., & Zhu, L. (2022). Defining blockchain governance principles: A comprehensive framework. *Information Systems*, 109, 102090.
<https://doi.org/10.1016/j.is.2022.102090>

- 535 31. Almeida, P. G. R. de, & dos Santos Júnior, C. D. (2025). Artificial intelligence
536 governance: Understanding how public organizations implement it. Government
537 Information Quarterly, 42(1), 102003. <https://doi.org/10.1016/j.giq.2024.102003>
538

UNDER PEER REVIEW IN IJAR