

FROM MONITORING DASHBOARDS TO GOVERNANCE TWINS: AN ACCOUNTABILITY-CENTRED ARCHITECTURE FOR HUMAN-IN-THE-LOOP AI IN COMPETITIVE CITY ECONOMIC REGIONS.

Abstract:

Background: City economic regions are delivered through portfolios of mutually dependent infrastructure, land, service and regulatory actions. Conventional dashboards improve visibility but often fail to preserve the evidence, authority, assumptions and distributional consequences behind consequential decisions.

Methods: This conceptual study applies design-science research, structured synthesis and scenario-based stress testing to develop an accountability-centred governance artifact. The proposed architecture was evaluated *ex ante* against lifecycle completeness, decision traceability, cross-agency usability, AI oversight, auditability and capacity for institutional learning.

Results: The study introduces the Accountability-Centred Governance Twin for City Economic Regions (ACGT-CER). It mirrors seven governance states—strategic intent, readiness, dependencies, delivery, outcomes, legitimacy and learning—through a decision-traceability spine. Three closed assurance loops connect operational control, authorization and adaptive learning. Eight Decision Assurance Points regulate irreversible commitments. Four operational measures are specified: Decision Traceability Ratio, Dependency Exposure Index, Human Oversight Sufficiency Score and Adaptive Closure Rate. Artificial intelligence is limited to sensing, anomaly detection, forecasting, consistency checking and prioritization; sanction, contract award, safety certification, land acquisition, payment and public commitment remain non-delegable human functions.

Conclusion: The governance twin extends the idea of a digital twin from physical representation to institutional state. It offers a practical way to connect data, judgment, authority, action, outcomes and learning without transferring public responsibility to algorithms. Field validation is required, but the architecture provides a testable control layer for competitive, inclusive and resilient regional development.

Key words:

City economic regions; governance twin; project assurance; institutional accountability; human-in-the-loop artificial intelligence; regional digital governance; decision traceability; infrastructure delivery

1. Introduction

Competitive city economic regions (CERs) are functional territories in which labour markets, supply chains, mobility systems, housing demand, ecological services and investment decisions extend across municipal boundaries. Their economic logic is networked, whereas their administrative logic is usually fragmented. A single regional programme may therefore contain transport corridors, industrial nodes, affordable housing, water and wastewater systems, energy connections, environmental restoration, digital platforms, skills initiatives and regulatory reforms, each controlled by a different institution and timetable. The central implementation problem is not the absence of projects; it is the absence of a shared mechanism for governing their interdependence.

Public authorities increasingly respond to this complexity through common dashboards, geographic information systems, building information models, sensors and artificial intelligence. These tools can improve situational awareness, but they do not automatically create accountability. A dashboard may show that physical progress is 72 per cent while concealing that the underlying design was changed through an informal instruction, a utility interface remains unowned, a benefit assumption has expired, or an algorithmic forecast is being used outside its validated context. Visibility is therefore a necessary but insufficient condition for governability.

This distinction becomes critical when digital-twin language is applied to cities. A conventional urban digital twin seeks a dynamic digital representation of physical or socio-technical conditions and may combine geometry, real-time data, semantic information and simulation (Batty, 2018; Ketzler et al., 2020). Yet the regional implementation failures that most directly damage public value frequently arise in the institutional layer: who was authorized to decide, which evidence was accepted, which alternatives were rejected, whether affected groups were considered, how uncertainty was treated, who owns an interface, and whether the decision can be reversed. A technically rich representation of the city can coexist with an institutionally weak decision system.

The present study develops a different artifact: an Accountability-Centred Governance Twin for City Economic Regions (ACGT-CER). The artifact is not intended to reproduce the city. It is designed to reproduce the current accountability state of the regional programme. Its purpose is to maintain a live, auditable relationship among evidence, interpretation, authority, action, outcomes and learning. The governance twin is therefore a control layer above project-management and urban-data systems, not a replacement for them.

The study addresses four research questions. First, what institutional states must be mirrored for a CER programme to remain governable across jurisdictions and lifecycle stages? Second, how can decision assurance be embedded without turning every gate into bureaucratic delay? Third, how can artificial intelligence assist regional foresight while preserving non-delegable public accountability? Fourth, which operational measures can reveal weaknesses that cost, schedule and output dashboards do not expose?

The contribution is fivefold. The article defines the governance twin as a distinct socio-technical artifact; specifies seven mirrored governance states and a decision-traceability spine; develops three closed assurance loops; introduces eight Decision Assurance Points and four operational metrics; and analytically stress-tests the architecture against recurrent CER failure scenarios. These components move the discussion from dashboard design to decision-system design.

2. Conceptual Foundations and Research Gap

2.1 CERs as interdependent regional systems

Functional urban areas are identified through economic and mobility relationships rather than administrative contiguity alone (Dijkstra et al., 2019). This perspective is relevant to CERs because firms and households experience a region through travel time, logistics reliability, land markets, service access and environmental conditions, not through the formal boundary of a single local body. Agglomeration can increase productivity through labour matching, specialized suppliers, knowledge diffusion and scale; however, congestion, unaffordable land, fragmented transport and ecological stress can reverse those gains.

Institutional fragmentation is therefore not a peripheral administrative matter. Comparative evidence indicates that fragmented metropolitan structures can impose a productivity penalty and that coordinating institutions may reduce part of that penalty (Ahrend et al., 2014). The practical implication is that a CER must be governed as a portfolio of linked commitments. An apparently successful project can remain regionally ineffective when a complementary investment, service reform or land decision is delayed.

2.2 The limits of conventional implementation dashboards

Most implementation dashboards are organized around what can be counted regularly: expenditure, physical quantities, milestone dates, contract status, grievances, incidents and selected outcomes. Their visual design often assumes that a red or amber indicator will prompt corrective action. In practice, an indicator does not close a control loop. Closure requires a defined owner, an authorized response, a deadline, evidence of completion and a test that the action was effective.

Three structural weaknesses follow. First, dashboards tend to represent objects and events but not the quality of the decision process that produced them. Second, they aggregate across projects without preserving causal dependencies and conditions. Third, they separate delivery reporting from institutional learning, allowing repeated failures to be treated as isolated exceptions. The result is a paradox: more data may be available while the institution's capacity to explain and correct its own decisions remains unchanged.

Table 1. Conventional dashboard versus accountability-centred governance twin

Dimension	Conventional dashboard	Governance twin	Decision value
STATE	Reports project status	Mirrors project and institutional decision state	Shows whether information is decision-ready
TIME	Compares actual and planned dates	Tracks deadlines, dependencies, evidence expiry and reversibility	Exposes where delay becomes irreversible
AUTHORITY	Lists responsible agency	Records mandate, decision owner, risk acceptance and conditions	Prevents collective responsibility from becoming no responsibility
EVIDENCE	Displays indicators	Preserves source, quality, assumptions, alternatives and uncertainty	Makes challenge and audit possible
AI	Presents scores or alerts	Links each model output to intended use, confidence, limits and human disposition	Reduces automation bias
OUTCOME	Adds selected outcome indicators	Maintains results chain, distribution and attribution limits	Prevents output completion from being treated as public value
LEARNING	Archives reports	Converts recurring causes into changed standards, contracts and models	Creates institutional memory

Source: Author's conceptual synthesis. The two artifacts may coexist; the governance twin adds the accountability state that a dashboard normally omits.

2.3 Digital twins, public purpose and governance

Digital twins are commonly described as connected and evolving digital representations of physical systems. In urban contexts, their meaning varies from 3-D city models to platforms integrating real-time data, simulation, semantic models and participatory interfaces (Ketzler et al., 2020). The Gemini Principles place public good, value creation, security, openness, quality, curation and evolution at the centre of digital-twin development (Bolton et al., 2018). These principles establish that a twin is not neutral infrastructure; it embodies purpose and governance choices.

The literature nevertheless concentrates primarily on the fidelity of the digital representation, technical interoperability and analytical capability. Less attention is given to whether the representation preserves the legal and ethical state of decisions. A forecast may be technically accurate but institutionally unusable because no authorized officer can explain its drivers, affected groups were not considered, or the model's purpose does not cover the proposed action. The governance twin addresses this missing layer by making institutional state a first-class object of representation.

2.4 Human oversight and public accountability

Responsible-AI frameworks converge on governance, risk mapping, measurement, management, transparency and human oversight (National Institute of Standards and Technology, 2023; NITI Aayog, 2021; UNESCO, 2021). Yet the phrase human-in-the-loop is frequently under-specified. A person positioned after an algorithm does not constitute meaningful oversight when the person lacks domain competence, access to underlying evidence, authority to reject the recommendation, time to deliberate or an obligation to record reasons.

Public accountability requires an identifiable actor, a duty to explain and justify conduct, and the possibility of scrutiny or consequence (Bovens, 2007). Collaborative governance additionally requires dialogue, trust and shared understanding across autonomous institutions (Ansell & Gash, 2008). The design challenge is therefore dual: collaboration must be sufficiently horizontal to resolve dependencies, while authorization must remain sufficiently vertical to preserve answerability.

2.5 Research gap

Existing bodies of knowledge provide strong but separate components: functional-region analysis, metropolitan coordination, public-investment management, stage-gated assurance, urban digital twins, responsible AI and policy evaluation. What remains underdeveloped is an integrated artifact that represents the institutional state of a regional programme in real time and connects that state to authorized action. The research gap lies between data visibility and decision traceability.

The ACGT-CER is designed to fill that gap. It treats a high-consequence decision as a traceable chain rather than a single approval. The chain begins with evidence and ends only when the action is verified, the outcome is evaluated and the learning is incorporated.

3. Materials and Methods

3.1 Design-science research strategy

The study adopts a conceptual design-science methodology because the intended contribution is an artifact for addressing an organizational problem (Hevner et al., 2004). The process followed six activities: problem diagnosis; definition of design objectives; construction of the artifact; analytical demonstration; ex ante evaluation; and communication, consistent with the design-science research process described by Peffers et al. (2007).

Problem diagnosis drew on the CER implementation context and a purposive synthesis of literature on regional governance, public investment, project assurance, data governance, digital twins, responsible AI and evaluation. The objective was not to produce a comprehensive bibliometric review, but to derive a coherent set of design requirements for a cross-jurisdictional public programme.

3.2 Design requirements and artifact construction

Eight design requirements were established: (1) lifecycle coverage from strategic problem to benefits review; (2) preservation of evidence provenance and expiry; (3) explicit decision rights and risk acceptance; (4) representation of dependencies among projects, policies and agencies; (5) time-bounded human review of AI outputs; (6) distributional and sustainability visibility; (7) verified closure of corrective actions; and (8) conversion of repeated failure into changed institutional routines.

The artifact was constructed through morphological mapping. The regional programme was decomposed into governance objects—decisions, evidence, dependencies, risks, actions, outcomes, model outputs and lessons—and each object was mapped against lifecycle stage, owner, permissible state, escalation rule and audit evidence. This produced the seven-state governance model, the decision-traceability spine, three assurance loops, eight Decision Assurance Points and four operational measures.

138 **3.3 Analytical demonstration through scenario stress testing**

139 The artifact was demonstrated against five stylized but recurrent failure scenarios: premature investment sanction; a
140 cascading utility-interface delay; contradictory inter-agency communication; an AI forecast with distributional bias; and a
141 completed project that fails to produce intended benefits. Each scenario was traced through four questions: what the
142 conventional dashboard would show, what additional state the governance twin would represent, which human authority would
143 act, and what institutional update would follow.

144 This is an ex ante analytical evaluation, not empirical validation. The assessment criteria were completeness, internal
145 coherence, accountability clarity, implementability, AI risk control, auditability and learning capacity. No claim is made that the
146 proposed thresholds or indices have predictive validity.

147 **3.4 Unit of analysis, assumptions and boundaries**

148 The primary unit of analysis is the CER programme: a portfolio of interdependent projects and policy actions intended to
149 produce regional economic, social and environmental outcomes. Individual projects remain units of contract delivery, but their
150 status is interpreted through the regional dependency network. The architecture is technology-neutral and can be hosted by a
151 state department, regional authority, special-purpose vehicle, metropolitan institution or inter-governmental council.

152 The artifact assumes that each participating institution has a lawful mandate, a minimum digital record system and an
153 identified officer who can own decisions or evidence. It does not prescribe a statutory form, procurement model, database
154 vendor or AI technique. It also does not replace engineering assurance, financial appraisal, environmental assessment or public
155 consultation; it connects their outputs into a traceable decision system.

156 **4. Results: The Accountability-Centred Governance Twin**


157 **4.1 Definition and operating proposition**

158 The ACGT-CER is defined as a continuously updated institutional representation of the evidence, authority, dependency,
159 action, outcome and learning states associated with consequential decisions in a city economic region. Its fidelity is judged not
160 by geometric similarity to the physical region but by the completeness and reliability of decision traceability.

161 A governance twin is considered current only when each high-consequence decision has a valid evidence package, an
162 identified authority, a documented rationale, explicit conditions, an action owner, a review date and a closure status. Missing
163 elements are represented as governance defects rather than hidden as blank fields. The model therefore treats incompleteness as
164 information.

Accountability-Centred Governance Twin for a City Economic Region

A live institutional mirror of evidence, authority, dependencies, outcomes and learning



The twin does not reproduce the city; it reproduces the accountability state of consequential regional decisions.

165
166 Fig. 1. Accountability-centred governance twin for a city economic region. Source: Author's formulation.

4.2 Seven mirrored governance states

The governance twin mirrors seven states. Strategic intent records the regional problem, intended public value and boundary of intervention. Readiness records whether the evidence and enabling conditions are sufficient for the next commitment. Dependencies record predecessor relationships, interfaces and reciprocal obligations. Delivery records baseline, performance, variance and corrective action. Outcomes record benefits, distribution and attribution limits. Legitimacy records authority, reasons, consultation, uncertainty and public communication. Learning records overrides, recurring causes and changes to institutional rules.

The seven states are not sequential database folders. They are concurrent views of the same programme. A change in one state can invalidate another: a revised land boundary may change readiness, cost, distributional impact, procurement documents and the validity of an AI forecast. The twin must therefore propagate material changes through linked objects and require re-authorization where consequence or reversibility changes.

Table 2. Seven mirrored governance states in the ACGT-CER

State	Governing question	Minimum live evidence	Failure prevented
S1 : INTENT	What regional problem and public value justify intervention?	Problem statement, outcomes, spatial/economic logic, exclusions	Project pipeline driven by announcement rather than need
S2 : READINESS	Is the next commitment supportable now?	Land, design, clearances, finance, capacity, safeguards, data quality	Premature sanction and avoidable post-award change
S3: DEPENDENCY	What must happen first, by whom and by when?	Predecessors, interfaces, owner, float, reciprocal commitments	Local progress masking regional non-functionality
S4 : DELIVERY	Is implementation within an authorized and achievable baseline?	Cost, schedule, scope, quality, safety, change and corrective action	Persistent variance and informal rebaselining
S5: OUTCOME	Are outputs producing intended and fairly distributed benefits?	Baseline, result chain, accessibility, jobs, service, inclusion, ecology	Completion being mistaken for success
S6 : LEGITIMACY	Was the decision lawful, reasoned, explainable and reviewable?	Mandate, alternatives, uncertainty, affected groups, reasons, review date	Displaced accountability and contradictory commitments
S7 : LEARNING	What must change because of experience?	Override logs, root causes, recurrence, revised standards/models/contracts	Repeated failure without institutional adaptation

Source: Author's formulation. A state may be "unknown," but it may not be silently assumed to be satisfactory.

4.3 The decision-traceability spine

The seven states are connected through a decision-traceability spine. For each high-consequence decision, the spine preserves six linked records: evidence, interpretation, authority, action, outcome and learning. Evidence identifies source, quality, custodian, timestamp and missing data. Interpretation records assumptions, alternatives, uncertainty and affected groups. Authority identifies legal mandate, risk acceptance and conditions. Action assigns owner, deadline and verification method. Outcome records benefits and burdens. Learning records whether the decision or model should be changed.

The spine addresses a recurrent audit weakness: a final approval may be available even when the reasoning chain is not. Traceability makes it possible to distinguish a poor outcome produced by an unreasonable decision from a poor outcome produced by an unforeseeable event, and to distinguish a valid model warning that was ignored from an invalid warning that was correctly overridden.

4.4 Three closed assurance loops

The ACGT-CER operates through three nested loops. The delivery control loop senses actual conditions, compares them with an authorized baseline, initiates correction and verifies closure. The accountability assurance loop tests evidence, challenges interpretation, authorizes commitment and records reasons. The adaptive learning loop evaluates outcomes, diagnoses recurring causes, updates standards or models and tests whether the change is effective.

The loops have different cadences. Delivery control may operate daily or weekly; accountability assurance occurs at gates and material changes; adaptive learning operates at quarterly, annual and post-implementation intervals. A programme is not governable when any loop remains open. For example, rapid correction without authorization can create informal scope change, while authorization without feedback can preserve a defective policy.

Three Assurance Loops: Control, Accountability and Adaptation

A regional programme remains governable only when all three loops close with evidence




Fig. 2. Three assurance loops connecting delivery control, accountability and institutional adaptation. Source: Author's formulation.

4.5 Decision Assurance Points

Decision Assurance Points (DAPs) regulate commitments according to consequence and reversibility. They differ from routine progress reviews. A DAP asks whether the evidence is sufficient for a specified commitment and whether the authorized institution is prepared to accept the residual risk. The decision is limited to proceed, proceed with conditions, recycle for further evidence, defer or reject.

Eight DAPs cover the CER lifecycle: problem entry, option selection, readiness certification, investment sanction, market authorization, delivery continuation, operational opening and benefits disposition. Evidence thresholds increase as commitments become less reversible. A budget allocation, previous announcement or percentage of expenditure does not itself satisfy a DAP.

Decision Assurance Points across the CER Lifecycle

Each point authorizes a specific commitment; no gate is passed merely because expenditure or announcement pressure exists




Fig. 3. Decision Assurance Points across the CER lifecycle. Source: Author's formulation.

Table 3. Decision Assurance Points, evidence thresholds and hard-stop conditions

DAP	Commitment being authorized	Core evidence threshold	Decision output	Illustrative hard stop
DAP0	Enter regional pipeline	Problem, affected population, strategic fit and initial system boundary	Enter, redefine or reject	No identifiable regional problem or mandate
DAP1	Select option for preparation	Demand, alternatives, distributional and environmental screen	Select option and preparation plan	Preferred solution chosen before credible alternatives
DAP2	Certify readiness	Land, design, utilities, clearances, finance, capacity, safeguards and data	Certify, condition, recycle or defer	Unresolved statutory approval or critical interface
DAP3	Sanction investment baseline	Business case, affordability, benefits, risk allocation and assurance	Sanction with conditions or reject	Benefits not linked to measurable baseline
DAP4	Enter market / procure	Tender completeness, interface ownership, data and AI clauses, market test	Authorize procurement	Scope or decision rights materially ambiguous
DAP5	Continue or materially change delivery	Integrated schedule, variance causes, safety, quality, claims and forecast	Continue, correct, rebaseline or stop	Unsafe condition or unowned critical-path dependency
DAP6	Open for operation	Testing, service integration, staffing, maintenance, asset data and user readiness	Open, open with conditions or defer	Safety or operating responsibility not certified
DAP7	Close, adapt or scale	Outcome evidence, distribution, sustainability, attribution and lessons	Close, remediate, adapt or scale	Outputs complete but intended benefits unverified

Source: Author's formulation. "Hard stop" means that aggregate readiness scoring cannot compensate for a critical control failure.

5. Measurement Architecture

5.1 Decision Traceability Ratio

The Decision Traceability Ratio (DTR) measures the proportion of high-consequence decisions that contain a complete reasoning and accountability chain. A decision is traceable only when minimum evidence, authority, rationale and review-date requirements are satisfied. Weighted calculation prevents a large number of low-consequence administrative decisions from masking poor traceability in irreversible commitments.

$$\text{DTR: } DTR = [\Sigma(wd \times qd) / \Sigma wd] \times 100$$

For decision d, wd is a consequence weight and qd equals 1 only when the required evidence, authority, rationale, conditions and review date are complete; otherwise qd equals 0. Suggested bands are provisional: below 60 = weak, 60–79 = controlled, 80–94 = accountable, and 95–100 = high assurance.

223 **5.2 Dependency Exposure Index**

224 The Dependency Exposure Index (DEI) represents the weighted exposure of the regional programme to unresolved
225 predecessor and interface conditions. Unlike a simple count of open dependencies, it considers consequence, criticality,
226 remaining float and ownership. The index can be calculated at project, corridor or CER level and should be accompanied by a
227 network view to prevent an aggregate score from hiding a single catastrophic interface.

228 **DEI:** $DEI = [\sum(wi \times ci \times ui) / \sum wi] \times 100$

229 For dependency i, wi represents consequence, ci represents criticality and float consumption on a 0–1 scale, and ui
230 represents unresolved status. A falling DEI is desirable only when closures are verified; deletion or reclassification of
231 dependencies is not closure.

232 **5.3 Human Oversight Sufficiency Score**

233 The Human Oversight Sufficiency Score (HOSS) tests whether a nominal human review is substantively capable of
234 governing a model-supported decision. Six dimensions are scored from zero to five: domain competence, legal authority,
235 deliberation time, access to explanation, override capability and independent verification. A model output may not be used as
236 the principal basis for a high-consequence decision when any dimension scores zero.

237 **HOSS:** $HOSS = [(C + A + T + X + O + V) / 30] \times 100$

238 C = competence, A = authority, T = time, X = explainability and evidence access, O = override capability, and V =
239 verification and audit. The score is a safeguard diagnostic, not a certification of model accuracy.

240 **5.4 Adaptive Closure Rate**

241 The Adaptive Closure Rate (ACR) distinguishes completion of an action from verified institutional change. An action
242 counts as adaptively closed only when the immediate correction is complete, effectiveness is tested and any required change to
243 standards, contracts, data definitions or models is implemented. This prevents recurring failures from being administratively
244 closed while their cause remains intact.

245 **ACR:** $ACR = (Verified\ adaptive\ closures\ within\ target / Adaptive\ actions\ due) \times 100$

246 Each closure record must identify evidence, reviewer and recurrence test. A high ACR obtained by repeatedly extending
247 target dates is invalid and must be reported separately as re-planning.

248 **Table 4. Operational metrics for the governance twin**

Metric	What it reveals	Primary evidence	Interpretation	Use limitation
DTR	Completeness of consequential decision chains	Decision packs, authority records, reasons and review dates	Higher is better; weight irreversible commitments	Does not prove that the decision was substantively correct
DEI	Exposure to unresolved cross-project and cross-agency conditions	Dependency network, float, interface actions and owner status	Lower is better after verified closure	Aggregate score must not hide a single critical dependency
HOSS	Whether human oversight is meaningful rather than nominal	Competence, mandate, time, explanation, override and validation	No zero dimension for high-consequence use	Does not measure model performance
ACR	Whether corrective action changes the institution, not only the symptom	Closure evidence, effectiveness review, revised rule/model/contract	Higher is better; recurrence must fall	Can be gamed by weak closure definitions
OUTCOME LINK	Whether outputs are causally connected to public value	Results chain, baseline, counterfactual reasoning and distribution	Strength of inference must be stated	Regional outcomes are affected by external conditions

249 *Source: Author's formulation. Thresholds should be calibrated through expert review and field evidence before use for comparative ranking.*

250 **5.5 Decision-oriented dashboard design**

251 The user interface should be designed around decisions that are due, not around the volume of available data. The primary
252 CER view should show upcoming DAPs, evidence defects, critical dependencies, overdue human dispositions of AI alerts,
253 unverified closures and benefits at risk. Drill-down should preserve a path from regional outcome to project, contract,
254 milestone, decision, evidence source and responsible officer.

255 Every red or amber condition requires an owner, action, deadline and closure test. Repeated amber cannot become a stable
256 reporting category. It must be resolved, formally re-baselined through authorized change control or escalated. Public views
257 should disclose verified scope, progress, material changes, expected benefits, grievance channels and limitations while
258 protecting sensitive personal, security and commercial information.

259 **6. Human-in-the-Loop AI Governance**

260 **6.1 Permissible analytical functions**

261 Within the ACGT-CER, AI is an attention-allocation and analytical aid. Permissible functions include anomaly detection
262 across physical, financial and sensor records; schedule and cost-risk forecasting; natural-language extraction of commitments
263 and unresolved actions from approved records; consistency checking across plans, contracts and public statements; and scenario
264 comparison. These functions may prioritize review but may not determine statutory or contractual action.

265 Each model must have a stated intended use, validated data domain, performance measures, known limitations, model
266 owner, change history and retirement rule. Outputs must show confidence, key drivers, data gaps and expiry. Where the
267 consequence of error is high and the model cannot provide an adequate explanation, its output may support inquiry but cannot
268 be the primary basis for decision.


269 **6.2 Non-delegable decisions and the alert evidence card**

270 Investment sanction, contract award, safety certification, land acquisition, acceptance of liability, release of payment,
271 operational opening and public commitment are non-delegable. The authorized institution may use AI-supported evidence but
272 must independently evaluate the legal, technical, fiscal, social and ethical consequences.

273 Every significant model output is issued as an alert evidence card. The card identifies the project or interface, signal,
274 severity, confidence, source, data freshness, explanatory drivers, potential impact, model limits, required reviewer and expiry
275 time. The alert remains open until a human disposition is recorded. Acceptance of the alert is separated from acceptance of any
276 proposed response: an officer may agree that a risk exists while choosing a different intervention.

Human-in-the-Loop Alert Evidence Card and Decision Path

AI ranks attention; authorized institutions retain interpretation, choice and consequence



277 Fig. 4. Human-in-the-loop alert evidence card and decision path. Source: Author's formulation.
278

279 **6.3 Oversight roles and model lifecycle**

280 Oversight roles are separated to prevent concentration of control. The data steward is accountable for data definition, quality
281 and lawful access; the model owner for intended use and maintenance; the analyst for interpretation; the domain officer for
282 contextual judgment; the authorized public official for the decision; and an independent validator for performance, bias and drift
283 review. A vendor may supply technology but cannot become the bearer of public responsibility.

284 Model monitoring includes calibration, false-positive and false-negative rates, subgroup performance where applicable,
285 drift, override patterns and incidents. Repeated justified overrides indicate a possible model defect or context shift; repeated
286 disregard of accurate alerts indicates an institutional-behaviour problem. Both patterns enter the learning state of the governance

287 twin. Procurement terms should preserve public access to model documentation, assumptions, logs, interfaces and exit
288 arrangements.

289 **7. Scenario-Based Stress Test**

290 **7.1 Analytical scenarios**

291 The five scenarios were selected because they expose different weaknesses: decision haste, dependency blindness,
292 communication conflict, algorithmic bias and outcome illusion. The purpose is not to estimate probability but to demonstrate
293 whether the artifact maintains traceability and closes the appropriate loop.

294 In each scenario, the governance twin adds a state that is normally absent from routine reporting. It does not eliminate
295 uncertainty; it makes uncertainty, authority and unresolved conditions visible enough to govern.

296 **Table 5. Scenario stress test of the ACGT-CER**

Scenario	Failure mechanism	Governance-twin response	Required decision	human	Institutional learning
PREMATURE SANCTION	Political or financial pressure precedes land, utility and safeguard readiness	DAP2 shows critical no-go condition; DTR fails because evidence chain is incomplete	Defer, condition or redefine package despite announcement pressure		Readiness criteria and announcement protocol updated
UTILITY CASCADE	One unowned interface consumes float across several projects	DEI identifies network criticality and integration-owner defect	Assign integration owner, recovery sequence and authorized rebaseline		Interface clauses and dependency templates revised
CONTRADICTIONARY SPEECH	Agencies issue inconsistent dates or scope commitments	Legitimacy state compares approved baseline with public and internal records	Withdraw, correct or formally authorize changed commitment		Single-source and change-confirmation rule strengthened
BIASED FORECAST	AI predicts demand or access using historically unequal data	Alert card shows subgroup performance and model limits; HOSS blocks unreviewed use	Commission alternative analysis, modify model or reject recommendation		Data, feature and validation standards revised
BENEFIT ILLUSION	Asset is completed but access, affordability or employment outcomes do not improve	Outcome state separates output from benefit and exposes attribution limits	Remediate service integration, adapt policy or stop scaling		Benefits model and future DAP3 evidence updated

297 *Source: Author’s analytical demonstration. Empirical testing is required to determine effectiveness and administrative burden.*

298 **7.2 Stress-test findings**

299 The stress test produced three findings. First, conventional delivery metrics are insufficient when the dominant failure
300 occurs before contract execution or across institutional boundaries. Readiness and dependency state must be represented with
301 equal prominence. Second, AI governance cannot be separated from project governance. A technically valid model can still
302 create an illegitimate decision when authority, affected groups or intended use are misaligned. Third, learning requires a formal
303 object and owner. Without a learning state, recurring causes disappear into narrative completion reports.

304 The architecture also reveals a design trade-off. Greater traceability can increase administrative burden if every low-
305 consequence action is treated as a high-consequence decision. Proportionality is therefore essential. Decision classes should be
306 defined by consequence, reversibility, legal significance, distributional effect and use of automated analysis. Routine actions
307 may use simplified records; irreversible or high-impact commitments require the full spine.

308 **8. Implementation Pathway**

309 **8.1 Minimum viable governance twin**

310 A pilot should begin with a small portfolio of projects whose benefits depend on one another, rather than with an attempt to
311 digitize the entire region. The minimum viable governance twin requires a governance charter, unique identifiers, a decision
312 register, a dependency network, DAP templates, defined roles, a data dictionary and a limited dashboard. AI should initially be
313 restricted to non-binding anomaly, schedule-risk and consistency alerts.

314 The first implementation objective is not predictive sophistication. It is reliable institutional record. If source data, authority
315 and closure evidence are weak, adding more advanced models will amplify uncertainty rather than reduce it.

316 **Table 6. Indicative 180-day pilot roadmap**

Phase	Days	Core deliverables	Acceptance test	Decision to proceed
FOUNDATI ON	0–30	Governance charter, portfolio boundary, roles, decision classes and ethics/data rules	Every consequential decision has one authorized owner	Approve pilot scope
MIRROR	31–60	Unique identifiers, seven-state schema,	Cross-system records reconcile for	Authorize data

Phase	Days	Core deliverables	Acceptance test	Decision proceed to
		decision register and dependency map	selected projects	integration
ASSURE	61–100	DAP2–DAP5 operation, DTR/DEI baseline, corrective-action closure tests	At least one gate and one dependency escalation completed end-to-end	Continue or redesign controls
AUGMENT	101–140	Non-binding AI alerts, evidence cards, HOSS checks and independent validation	Every alert receives timely human disposition and audit trail	Expand, limit or suspend model use
LEARN	141–180	Outcome baseline, ACR review, officer training and pilot evaluation	Verified policy/template/model changes from observed failures	Scale, adapt or stop

Source: Author's formulation. The roadmap is indicative and should be adapted to legal, institutional and data readiness.

8.2 Capacity, procurement and public reporting

Officials require capability in systems thinking, project assurance, data interpretation, AI limitations, facilitation and reasoned decision writing. The objective is not to convert every officer into a data scientist, but to ensure that officers can challenge evidence, recognize model limits and document proportionate decisions. Local bodies with weaker technical capacity require embedded support so that the regional system is not dominated by the institution with the most advanced information technology.

Digital and consultancy procurement must specify data ownership, interoperability, cybersecurity, documentation, model performance, bias testing, audit access, change control, continuity and exit. Contracts should distinguish a technology service level from an accountability obligation. The public authority remains answerable even when the technology is externally supplied.

Public reporting should state what is approved, what remains conditional, what has changed, which benefits are expected, which groups bear material burdens, and when the next review will occur. Where AI materially influences prioritization or review, that use should be disclosed in an intelligible form and affected persons should have access to human reconsideration.

9. Discussion

9.1 Theoretical contribution

The governance twin extends digital-twin theory from representation of physical or socio-technical conditions to representation of institutional accountability state. It proposes that a public-sector twin is incomplete when it can simulate the region but cannot explain the authority, evidence and conditions behind action. This shifts fidelity from resemblance alone to traceability.

The artifact also links collaborative governance and accountability. Cross-agency coordination is represented through dependencies and shared evidence, while final authorization remains identifiable. This avoids two extremes: siloed decision-making that cannot resolve regional interfaces, and diffuse collaboration in which no institution is answerable for consequence.

9.2 Practical and policy contribution

For practitioners, the ACGT-CER provides a control architecture that can sit above existing project-management, finance, GIS and asset systems. It does not require a single centralized database. A federated implementation can preserve source systems while using common identifiers, metadata, decision records and interfaces. The priority is consistency of meaning and accountability, not uniformity of software.

For policy, results-based financing should reward verified readiness, closure of critical dependencies, decision traceability and benefit realization rather than expenditure alone. Common minimum standards for DAPs, AI oversight, data provenance and public reporting could support benchmarking across CERs without prescribing identical project choices.

9.3 Relationship to public trust

Public trust is affected not only by whether a project succeeds but by whether institutions communicate uncertainty and correction honestly. The legitimacy state requires approved and public records to remain aligned. Premature claims, undocumented changes and algorithmic opacity are treated as governance defects because they constrain future choices and weaken scrutiny.

Trust is not equivalent to the absence of bad news. A programme that discloses a material delay, assigns responsibility, explains the decision and verifies correction may be more trustworthy than one that maintains a green dashboard through reclassification.

10. Limitations and Validation Agenda

The proposed architecture is conceptual and has not been field-tested in a live CER. The seven-state model may require adaptation to statutory arrangements, and the four metrics may be vulnerable to gaming, inconsistent scoring and unequal

evidence capacity. Additional controls can also create administrative load or encourage formal compliance without genuine challenge.

Validation should proceed in stages. A multidisciplinary Delphi study can test content validity and refine decision classes, DAP evidence and hard-stop conditions. Pilot studies can measure administrative burden, inter-rater reliability, decision timeliness, alert response and user behaviour. DTR, DEI, HOSS and ACR require sensitivity analysis and association with independent project outcomes. AI modules require out-of-sample testing, calibration, subgroup performance analysis, drift monitoring and incident review.

Comparative longitudinal research could examine whether governance-twin-enabled portfolios exhibit fewer post-award changes, faster closure of cross-agency decisions, lower recurrence of known failure modes and more credible benefit reporting than portfolios using conventional dashboards alone. Qualitative research should also examine whether traceability improves deliberation or merely changes documentation practices.

11. Conclusion

Competitive city economic regions require connected decision systems as much as connected infrastructure systems. Conventional dashboards can show progress, but they rarely preserve the full institutional chain from evidence to authority, action, outcome and learning. The Accountability-Centred Governance Twin addresses that omission by mirroring seven governance states, linking them through a decision-traceability spine and closing delivery, accountability and adaptation loops.

The architecture places artificial intelligence inside a bounded public-governance process. AI may detect patterns, forecast risk and prioritize attention; it may not assume statutory responsibility or convert a probabilistic signal into an automatic public commitment. Meaningful human oversight requires competence, authority, time, explanation, override and verification.

The main proposition is practical: a regional programme becomes governable when every consequential decision has valid evidence, one authorized owner, explicit conditions, a deadline, verified closure and a learning consequence. Field validation will determine how the architecture performs in practice, but the need it addresses is immediate. The transition from dashboard to governance twin is ultimately a transition from seeing projects to governing decisions.

Declarations

Ethics approval and consent to participate. Not applicable. The study is conceptual and uses no human participants, personal data or animals.

Consent for publication. Not applicable.

Data availability. No empirical dataset was generated. The conceptual artifact, definitions, formulas and analytical scenarios are contained in the manuscript.

Funding. The author received no external funding for this study.

Competing interests. The author declares no competing interests.

Author contribution. The sole author undertook conceptualization, methodology, analysis, writing, visualization and final approval.

Prior dissemination and originality. The broad theme of CER implementation and human-in-the-loop accountability has been explored in the author's earlier work. This manuscript presents a distinct research artifact—the ACGT-CER—with new constructs, metrics, stress-test analysis, implementation protocol, figures and tables. Full bibliographic disclosure of the related publication should accompany submission to the editor in accordance with journal policy.

Acknowledgement

The author acknowledges the practical lessons derived from public-sector infrastructure planning, project delivery, safety governance and digital-system implementation that informed the design orientation of this study.

References

- Ahrend, R., Farchy, E., Kaplanis, I., & Lembcke, A. C. (2014). What makes cities more productive? Evidence on the role of urban governance from five OECD countries. OECD Regional Development Working Papers, 2014/05. <https://doi.org/10.1787/5jz432cf2d8p-en>
- Ansell, C., & Gash, A. (2008). Collaborative governance in theory and practice. *Journal of Public Administration Research and Theory*, 18(4), 543–571. <https://doi.org/10.1093/jopart/mum032>
- Argyris, C., & Schön, D. A. (1978). *Organizational learning: A theory of action perspective*. Addison-Wesley.
- Batty, M. (2018). Digital twins. *Environment and Planning B: Urban Analytics and City Science*, 45(5), 817–820. <https://doi.org/10.1177/2399808318796416>
- [6:15 pm, 29/07/2026] Dr Alok Kumar Bhargava: 5 author booked
1 is remaining
- [9:06 pm, 29/07/2026] Dr Alok Kumar Bhargava: Bhargava, A. K. (2025). *TrayiVāṇī: Eternal Verses on Peace, Silence and Discernment*. ISBN: 978-81-994680-3-0.

411 Bhargava, A. K. (2026, February). Rebuilding the Inner Infrastructure of Humanity. Hindenburg Times (Special Edition).

412 Bhargava, A. K. (2026, March). The Architecture of Thought: How Three Sanskrit Verses Are Redefining Speech & Inner Clarity. London Herald.

413 Bhargava, A. K. (2026, March). The Quiet Force of Global Leadership: Redefining Leadership Through Silence, Thought & Ethical Power. USA Herald.

414

415 Bhargava, A. K. (2026, March). TrayiVāṇī: A Contemporary Voice of Ethical Clarity in a Noisy World. Washington Post Magazine.

416 Bhargava, A. K. (2026, April). Leading Through Silence: The Thought-Centric Philosophy of Dr. Alok Kumar Bhargava. The CEO Magazine, 11 (Special Edition).

417

418 Bhargava, A. K. (2026, May). The Quiet Power of Thought: Blueprint for Conscious Leadership, Leading Through Silence: The Thought-Centric Philosophy of Dr. Alok Kumar Bhargava. The C-Connects – Global C-Suite Community Platform, 33.

419

420 Bhargava, A. K. (2026, May). Dr. Alok Kumar Bhargava – Globally Recognized Award-Winning Voice in Ethical Leadership. Next India, 25.

421 Bhargava, A. K., et al. (2026). The Inner Engine of Leadership: Foundations of Conscious Decision-Making and Speech Governance (Pt-I). ISBN: 978-81-994680-4-7.

422

423 Bhargava, A. K., et al. (2026). The Inner Engine of Leadership: Operational Excellence, AI, Crisis Leadership and Sectoral Applications (Vol-II). ISBN: 978-81-994680-5-4.

424

425 Bhargava, A. K., et al. (2026). The Inner Engine of Leadership: Total Leadership Index, 360-Degree Audit, Training and Institutional Implementation (Pt-III). ISBN: 978-81-687626-0-2.

426

427 Bhargava, A. K. (2026). Rewriting the Economics of Leadership: From Reaction to Reflective Intelligence. International Economic Times.

428 TrayiVāṇī Foundation Communications Cell. (2026). Media Coverage Report: TrayiVāṇī – Eternal Verses on Peace, Silence & Discernment.

429 Profile Document. (2026). Dr. Alok Kumar Bhargava (IRSE): Chief Engineer, Northern Railway—Infrastructure Leader, Author, Philosopher, and Public Intellectual.

430

431 Bhargava, A. K. (2026). Reflective Intelligence as Cognitive Capital: A TrayiVāṇī-Based Framework for Ethical Leadership, Decision Stillness and Communication Governance in High-Speed Economies. EPH – International Journal of Business and Management, 12(2s), 59–72.

432

433 Bhargava, A. K. (2026). Founder Speech Governance as an Entrepreneurial Capability for Sustainable Growth in Asian SMEs. The Journal of Asia Entrepreneurship and Sustainability (JAES), 22(3s), 115–125.

434

435 Bhargava, A. K. (2026). The Pre-Decisional Accountability Gap in Accounting Governance: Developing the Inner Engine as a Theoretical Architecture for Audit Judgment, Ethical Disclosure and Communication Control. The Journal of Theoretical Accounting Research (JTAR), 22(5), 1–13.

436

437

438 Bhargava, A. K. (2026). Conscious Operational Excellence: Integrating the Inner Engine with Lean, Six Sigma, Agile and Total Quality Management. EPH – International Journal of Business and Management, 12(3s), 1–10.

439

440 Bhargava, A. K. (2026). Conscious Governance in AI-Enabled Institutions: A Human-in-the-Loop Cultural Framework for Decision Intelligence, Ethical Speech and Institutional Trust. Scientific Culture, 12(4), 11407–11418.

441

442 Bhargava, A. K. (2026). Measuring Inclusive Education Leadership: Development of a Total Leadership Index for Special Education Governance, Family Trust and Multidisciplinary Coordination. International Journal of Special Education, 41(7), 1000–1014.

443

444 Bhargava, A. K. (2026). From Crisis Trigger to Institutional Clarity: A First-Hour Leadership Protocol for Verified Communication and Trust Recovery. International Research Journal of Scientific Studies (IRJSS), 12(2s), 59–72.

445

446 Bhargava, A. K. (2026). Readiness-Gate Governance in Public Infrastructure Projects: Preventing Premature Milestones, Infrastructure Gridlocks and Trust Loss. The Journal of Asia Entrepreneurship and Sustainability (JAES), 22(3s).

447

448 Bhargava, A. K. (2026). Technology Incident Leadership: Governed Communication, Root-Cause Reflection and Trust Recovery After Digital System Failure. International Journal of Research and Innovation in Social Science (IJRISS), 10(6), 1–13.

449

450 Bhargava, A. K. (2026). Executive Language and Compliance Culture: How Leadership Speech Shapes Ethical Behaviour in Banking and Regulated Organizations. Jana Nexus: Journal of Business and Management Studies, 1(12), 39–53.

451

452 Bhargava, A. K. (2026). Policy Timing and Governed Public Communication: A Conscious Leadership Framework for Administrative Accuracy and Citizen Trust. International Journal of Current Research, 18(6), 37656–37670.

453

454 Bhargava, A. K. (2026). Trust-Preserving Safety Leadership: Near-Miss Reporting, Incident Communication and Learning in Rail, Healthcare and Aviation Systems. International Journal of Research and Innovation in Social Science (IJRISS), 10(6), 12108–12124.

455

456 Bhargava, A. K. (2026). Leadership Governance Dashboards: Designing Scorecards, Heat Maps and Data-Ethics Protocols for Executive Behavioural Risk Monitoring. Jana Nexus, 2(7), 1–11.

457

458 Bhargava, A. K. (2026). Audit-Linked Executive Education: Training, Simulation and Certification Pathways for Conscious Leadership Development. The Journal of Theoretical Accounting Research (JTAR), 22(3s), 147–159.

459

460 Bhargava, A. K. (2026). Institutionalizing Conscious Leadership: Adoption Roadmap, Performance Improvement Pathways and Validation Agenda for the Total Leadership Index. International Journal of Aquatic Research and Environmental Studies (IJARES), 6(2s), 801–811.

461

462 Bhargava, A. K., et al. (2026). Artificial Intelligence and Machine Learning in Special Education: Enhancing Personalized Learning and Assistive Technologies for Diverse Learners. International Journal of Special Education, 41(8s), 246–257.

463

464 Bhargava, A. K., et al. (2026). Regulatory Frameworks and Their Impact on Microfinance Institutions: A Political–Legal Analysis in Emerging Economies. International Journal of Special Education, 41(8s), 234–245.

465

466 Bhargava, A. K., et al. (2026). Artificial Intelligence in Human Resource Management: Transforming Employee Engagement and Organizational Performance. International Journal of Special Education, 41(13s), 99–109.

467

468 Bhargava, A. K., et al. (2026). Business Leadership and Organisational Sustainability: Evidence from Emerging Asian Economies. The Journal of Asia Entrepreneurship and Sustainability (JAES), 22(6), 30–39.

469

470 Bhargava, A. K., et al. (2026). Strategic Management Practices and Their Impact on Sustainable Enterprise Development in Emerging Economies. International Journal of Aquatic Research and Environmental Studies (IJARES), 6(S2), 791–800.

471

472 Bhargava, A. K., et al. (2026). Strategic Human Resource Management in Startups: A Catalyst for Entrepreneurial Growth in Emerging Asian Economies. The Journal of Asia Entrepreneurship and Sustainability (JAES), 22(6), 61–70.

473

474 Bhargava, A. K., et al. (2026). Architectural Image Modeling as a Predictive Tool for Civil Engineering Stress and Load Distribution Analysis. International Journal of Aquatic Research and Environmental Studies (IJARES), 6(S5), 1511–1518.

475

476 Bhargava, A. K., et al. (2026). The Impact of Educational Technology on Cognitive and Emotional Learning Outcomes. *International Journal of*
477 *Special Education*, 41(11s), 384–394.

478 Bhargava, A. K., et al. (2026). Conceptualizing Marketing Intangibles in Accounting: Challenges in Measurement and Reporting. *The Journal of*
479 *Theoretical Accounting Research (JTAR)*, 22(3s), 305–313.

480 Bhargava, A. K., et al. (2026). Entrepreneurial Management and Sustainable Development: A Theoretical Perspective. *The Journal of Asia*
481 *Entrepreneurship and Sustainability (JAES)*, 22(4s), 340–348.

482 Bhargava, A. K., et al. (2026). Interdisciplinary Perspectives on Literature, Education, and Financial Literacy in Contemporary Society. *International*
483 *Journal of Special Education*, 41(9s), 893–906.

484 Bhargava, A. K., et al. (2026). Mathematical Modelling and Artificial Intelligence Techniques for Predictive Maintenance in Mechanical Systems.
485 *Genetics and Molecular Research*, 25(7s).

486 Bhargava, A. K., et al. (2026). Policy Challenges in Implementing Equal Educational Opportunities for Special Needs Students. *International Journal*
487 *of Special Education*, 41(9s), 842–857.

488 Bhargava, A. K., et al. (2026). An Intelligent Decision Support Framework Integrating Artificial Intelligence, Computer Information Systems, and
489 Management Analytics for Industry. *International Journal of Computer Information Systems and Industrial Management Applications*, 18(7s), 932–
490 943.

491 Bhargava, A. K., et al. (2026). Artificial Intelligence-Driven Workforce Optimization: An Operational Research Framework for Strategic Human
492 Resource Management and Organizational Decision-Making. *Journal of Intelligent Decision Making and Information Science*, 3(3s). (In Press)

493 Bhargava, A. K. (2026). Sustainable Human Resource Strategies for Heritage Organizations: Exploring Ethical Leadership and Workforce
494 Engagement. *Heritage and Sustainable Development*, 8(1).

495 Bhargava, A. K., et al. (2026). Legal Compliance, Strategic Management, and Ethical Marketing in Educational Institutions: A Framework for
496 Sustainable Stakeholder Engagement. *International Journal of Special Education*, 41(13s). (In Press).

497 Bhargava, A. K., et al. (2026). Machine Learning and Big Data Frameworks for Sustainable Economic and Environmental Decision-Making.
498 *International Journal of Special Education*, 41(15s). (In Press)

499 Bhargava, A. K., et al. (2026). Integration of Artificial Intelligence, Internet of Things, and Robotics for Intelligent Automation: Advances in Signal
500 Processing and Image Analysis for Smart Systems. *International Journal of Special Education*, 41(15s). (In Press).

501 Bhargava, A. K., et al. (2026). Artificial Intelligence-Driven Personalized Learning Models for Special Education. *International Journal of Special*
502 *Education*, 41(15s). (In Press).

503 Bhargava, A. K., et al. (2026). An Artificial Intelligence and Machine Learning Framework for Predictive Decision Analytics in Business and
504 Industrial Operations. *Journal of Intelligent Decision Making and Information Science*, 3(3s). (In Press).

505

506

507 Bolton, A., Enzer, M., & Schooling, J. (2018). The Gemini Principles: Guiding values for the national digital twin and information management
508 framework. Centre for Digital Built Britain. <https://doi.org/10.17863/CAM.32260>

509 Bovens, M. (2007). Analysing and assessing accountability: A conceptual framework. *European Law Journal*, 13(4), 447–468.
510 <https://doi.org/10.1111/j.1468-0386.2007.00378.x>

511 Dijkstra, L., Poelman, H., & Veneri, P. (2019). The EU–OECD definition of a functional urban area. *OECD Regional Development Working Papers*,
512 2019/11. <https://doi.org/10.1787/d58cb34d-en>

513 Government of India, Ministry of Finance. (2026). Budget speech 2026–2027. Government of India.

514 Hevner, A. R., March, S. T., Park, J., & Ram, S. (2004). Design science in information systems research. *MIS Quarterly*, 28(1), 75–105.
515 <https://doi.org/10.2307/25148625>

516 Infrastructure and Projects Authority. (2021). Project Routemap: Governance. UK Government.

517 International Monetary Fund. (2022). PIMA handbook: Public Investment Management Assessment (1st ed.).
518 <https://doi.org/10.5089/9781513571829.071>

519 Ketzler, B., Naserentin, V., Latino, F., Zangelidis, C., Thuvander, L., & Logg, A. (2020). Digital twins for cities: A state of the art review. *Built*
520 *Environment*, 46(4), 547–573. <https://doi.org/10.2148/benv.46.4.547>

521 Kitchin, R. (2014). The real-time city? Big data and smart urbanism. *GeoJournal*, 79, 1–14. <https://doi.org/10.1007/s10708-013-9516-8>

522 National Institute of Standards and Technology. (2023). Artificial Intelligence Risk Management Framework (AI RMF 1.0) (NIST AI 100-1). U.S.
523 Department of Commerce. <https://doi.org/10.6028/NIST.AI.100-1>

524 NITI Aayog. (2021). Responsible AI: Approach document for India, Part 2—Operationalizing principles for responsible AI. Government of India.

525 OECD. (2019). Better criteria for better evaluation: Revised evaluation criteria definitions and principles for use. OECD Development Assistance
526 Committee.

527 OECD. (2023). Smart city data governance: Challenges and way forward. *OECD Urban Studies*. <https://doi.org/10.1787/e57ce301-en>

528 Peffers, K., Tuunanen, T., Rothenberger, M. A., & Chatterjee, S. (2007). A design science research methodology for information systems research.
529 *Journal of Management Information Systems*, 24(3), 45–77. <https://doi.org/10.2753/MIS0742-1222240302>

530 School of Planning and Architecture, New Delhi. (2026). International Conference on Competitive City Economic Regions for Future Urbanization:
531 Call for papers.

532 UNESCO. (2021). Recommendation on the ethics of artificial intelligence. UNESCO.

533 World Bank. (2020). International practices of metropolitan governance: A compendium of collaborative arrangements in metropolitan areas. World
534 Bank.

535 World Bank. (2023). Infrastructure Governance Assessment Framework (InfraGov). World Bank.