

From Infrastructure Assets to Infrastructure Intelligence: Reimagining Civil Engineering for an Era of Artificial Intelligence, Climate Uncertainty, and Autonomous Infrastructure Systems

Thabo Jena¹

School of Civil Engineering, Anhui University of Technology, Ma'anshan, Anhui 243002, China

DOI: <https://doi.org/10.51583/IJLTEMAS.2026.150600237>

Received: 15 July 2026; Accepted: 20 July 2026; Published: 01 August 2026

ABSTRACT

Civil engineering enters the second quarter of the twenty-first century carrying a conceptual inheritance it has rarely interrogated directly. This paper advances a single, falsifiable thesis: the defining failure of twentieth-century civil engineering was not a failure of materials, computation, or design method, but a category error in the treatment of infrastructure as a class of physical assets to be optimized rather than as a class of adaptive, decision-making systems operating under irreducible uncertainty.

We call this inheritance the Asset Paradigm and argue that it persists because it is embedded in codes, procurement, liability law, curricula, and professional identity, not because it remains effective. Building on, and formally sharpening, the emerging literature on infrastructure intelligence and resilience, we contribute: (i) necessary and sufficient conditions that distinguish Infrastructure Intelligence from smart infrastructure, digital twins, and cyber-physical systems, together with a quantitative formalization; (ii) a five-stage Infrastructure Evolution Model, classifying real systems by the locus of run-time decision authority and justified against maturity-model precedents from other engineering domains; (iii) four empirically testable propositions, each with explicit variables, candidate datasets, and validation methods; (iv) an illustrative bridge structural-health-monitoring application; (v) a phased, institutionally feasible transition pathway; and (vi) a ten-item Grand Research Agenda for 2025–2050. We position the framework against complex adaptive systems, socio-technical, resilience, adaptive-governance, and deep-uncertainty theory, and identify precisely where each breaks down under conditions of artificial intelligence, climate non-stationarity, and machine-speed autonomy. We conclude with a set of concrete claims and institutional commitments.

Keywords: infrastructure intelligence; asset paradigm; complex adaptive systems; climate non-stationarity; adaptive governance; autonomous infrastructure; deep uncertainty; decision-making under deep uncertainty; antifragility; digital twins

INTRODUCTION

A Discipline at the Limit of Its Own Assumptions

Civil engineering is the oldest of the engineering disciplines and, by a wide margin, the most consequential for human survival. It built the aqueducts, the grids, the flood defenses, and the transport arteries on which every other human activity is parasitic.

It is also, we argue, the discipline least prepared epistemologically and institutionally for the conditions now arriving. The reason is not a deficit of talent or computation. It is that the discipline's foundational mental model has gone unexamined for so long that practitioners no longer recognize it as a model at all. They mistake it for reality.

That model is the treatment of infrastructure as a portfolio of assets: discrete physical objects, designed once to a specification, commissioned, depreciated, maintained against a predictable decay curve, and eventually renewed. The Asset Paradigm is so deeply naturalized that the phrase “asset management” is taught, regulated, and procured as though it were a neutral description of what infrastructure is, rather than one historically contingent way of construing it [1, 27]. This paper’s central claim is that this construal is a category error whose costs were tolerable in the twentieth century and are becoming catastrophic in the twenty-first.

The argument is organized around a thesis stated plainly so that it can be attacked: the fundamental failure of twentieth-century civil engineering was not a failure of materials, computation, or design methods, but a category error treating infrastructure as a class of physical assets to be optimized rather than as a class of adaptive, decision-making systems operating under irreducible uncertainty. If this thesis is correct, then the proliferation of sensors, digital twins, and machine learning models across the built environment is not the cause of the coming transformation but a symptom of the Asset Paradigm’s collapse. And the appropriate response is not a technological upgrade but a reconstruction of the engineer’s relationship to the built environment along three axes: from command to negotiation, from prediction to adaptation, and from control to stewardship.

The label “infrastructure intelligence” is not introduced here for the first time, and we do not claim that prior research has overlooked artificial intelligence in civil systems the opposite is true, and Section 6 surveys a fast-moving literature on exactly this. What that literature has generally supplied, however, is technology- or component-level: specific digital-twin architectures, specific sensing or learning methods, specific AI applications within a single asset class [22, 14, 20, 34].

“Infrastructure intelligence” has, in parallel, emerged as an organizing label for this activity, most visibly through a dedicated Elsevier journal of that name launched in 2022. What has not, to our knowledge, been supplied is a systems-level paradigm: a domain-general account, expressed as necessary and sufficient conditions, of what makes any of these technologies constitute intelligence in our sense, independent of which sensor, model, or control method is used. That is the gap Section 4 addresses.

We present what follows as a conceptual and theoretical framework: a proposed set of definitions, a proposed evolution model, and falsifiable propositions offered for future empirical test, not as findings already established. Section 13 states this boundary explicitly and outlines how the framework could be moved from proposal toward validated theory.

The remainder of the paper proceeds as follows. Section 2 situates the argument against established bodies of theory and identifies where each becomes insufficient. Section 3 deconstructs the Asset Paradigm. Section 4 defines Infrastructure Intelligence, including a variable-based formalization. Section 5 develops the Infrastructure Evolution Model and relates it to established maturity-model precedents. Section 6 reframes enabling technologies, including recent foundation-model and agentic-AI developments, as expressions of a deeper shift. Section 7 presents the conceptual figures. Section 8 works through an illustrative bridge structural-health-monitoring application. Section 9 formulates falsifiable research questions and testable propositions with candidate variables, datasets, and validation methods. Section 10 examines domain-specific implications, transition barriers, and a phased transition pathway. Section 11 addresses ethics, governance, and cybersecurity. Section 12 critiques the profession itself and evaluates objections. Section 13 states the framework’s limitations and validation pathway, before the Grand Research Agenda and a concluding set of claims.

Contributions

This paper makes seven specific contributions:

1. A category-error diagnosis of the Asset Paradigm and an account of the institutional mechanisms that keep it in place after its founding assumptions have failed (Sections 3 and 12).

2. Necessary and sufficient conditions for Infrastructure Intelligence, together with a quantitative formalization intended as a basis for future measurement rather than a validated instrument (Section 4).
3. A five-stage Infrastructure Evolution Model, justified against established maturity-model precedents from other engineering domains, with a proposed empirical validation pathway (Section 5).
4. Four falsifiable propositions, each with explicit variables, candidate datasets, and applicable validation methods (Section 9).
5. An illustrative bridge structural-health-monitoring application that traces the framework from current practice through to expected engineering benefit (Section 8).
6. A phased, institutionally feasible transition pathway spanning policy, standards, procurement, and certification (Section 10).
7. A ten-item Grand Research Agenda for 2025–2050 (Section 14).

Theoretical Foundations and Their Limits

A framework paper earns its standing not by inventing concepts from nothing but by showing precisely where the existing conceptual apparatus fails. We therefore engage nine bodies of theory, stating for each what remains valid and what breaks down under the joint conditions of artificial intelligence, climate non-stationarity, deep uncertainty, and machine-speed autonomy.

Complex Adaptive Systems (CAS) theory [2, 11] supplies the indispensable insight that infrastructure networks exhibit emergence, non-linearity, and adaptation that cannot be reduced to component behavior. What CAS theory does not supply is a normative engineering account: it describes how complex systems behave but is silent on how a professional discipline bearing legal liability should intervene in a system whose components now include learning algorithms that adapt faster than their human supervisors can comprehend. CAS theory assumes an observer outside the system; infrastructure intelligence collapses the observer into the system.

Socio-technical systems (STS) theory [9, 33] correctly insists that technical artefacts and social institutions co-evolve. Its limit is the implicit assumption that the technical subsystem is a passive tool wielded by social actors. When the technical subsystem becomes an autonomous decision-maker, the symmetry STS theory presumes between human and artefact dissolves: the artefact is no longer mediated by human agency but is itself an agent. STS theory has no vocabulary for a technical actor with goals.

Resilience theory [7, 12] re-framed engineering from resisting disturbance to absorbing and recovering from it. But resilience theory presupposes a known disturbance regime, a definable envelope of shocks against which recovery is measured. Climate non-stationarity destroys precisely this presupposition: the disturbance distribution is itself drifting and unknowable [24]. Resilience to what becomes undefined.

Antifragility [31] pushes further, describing systems that gain from disorder. It remains among the least operationalized of our foundations for civil systems specifically: to our knowledge, no mainstream civil engineering design code yet contains an antifragile provision, and the discipline still has no established method for procuring or certifying a structure that should improve under stress. This is not, however, a purely theoretical vacuum: a nascent technical literature has begun formalizing antifragile design principles for closed-loop traffic-control, robotic, and cyber-physical systems [3], and it is a natural next step named explicitly in our Grand Research Agenda, Challenge 2, to test whether these formalisms transfer to load-bearing civil structures.

Cyber-physical systems (CPS) and systems engineering theory [13, 17] provide the architecture for tightly

coupled computation and physical process. Their limit is philosophical: both assume design-time completeness that the system's requirements can be specified before it is built. Infrastructure intelligence is defined by the abandonment of this assumption.

Decision-making under deep uncertainty (DMDU) theory [10, 16, 18] is the strand of planning scholarship most directly relevant to this paper's propositions. Robust decision-making and dynamic adaptive policy pathways support infrastructure investment when the future probability distribution of loads and demands cannot be specified in advance, with an established track record in water and coastal planning. DMDU supplies the operational vocabulary Proposition 1 (Section 9) needs adaptation tipping points, monitored signposts, contingency pathways and is the closest existing analogue to the adaptation-rate criterion formalized there. Its limit is that DMDU pathways are specified and re-evaluated at human deliberative speed: a monitored signpost is reviewed at a planning board, not by the infrastructure itself. It therefore shares adaptive governance theory's human-speed limit below, and adds a design-time limit of its own, since its pathways are enumerated in advance as a finite decision tree. Infrastructure Intelligence extends DMDU's central insight plan for change, not only for a fixed future to systems whose own decision authority operates inside the interval between deliberative reviews.

Adaptive governance theory [8, 28] and sustainability transitions theory [9, 21] describe how institutions learn and how socio-technical regimes shift. Both assume governance operates at human deliberative speed. Neither contemplates a regime in which the governed system acts faster than any deliberative body can convene.

Infrastructure systems theory [1, 19] supplies interdependency and cascade analysis but retains the asset ontology this paper rejects.

The pattern is consistent: each theory assumes one of an external observer, a passive technical substrate, a stationary disturbance regime, design-time completeness, or human-speed deliberation. Infrastructure intelligence violates all five simultaneously. The Infrastructure Intelligence Framework developed below is positioned exactly in this gap.

Deconstructing the Asset Paradigm

Naming the paradigm

We define the Asset Paradigm as the tacit ontology under which infrastructure is construed as a set of discrete, ownable, physical objects whose value is realized through optimized provision of a fixed function across a predictable service life. Its four load-bearing assumptions are: (1) functional fixity what the asset is for is known and stable; (2) stationarity the environment of loads and demands fluctuates within a stable, knowable distribution; (3) deterministic degradation deterioration follows curves estimable from material science; and (4) separability each asset can be analyzed, financed, and managed as an isolable unit.

How it emerged and why it became dominant

The paradigm is a child of the mid-twentieth century. Post-war reconstruction and the great infrastructure expansion of 1945–1975 rewarded exactly the capabilities the Asset Paradigm formalizes: the rapid, standardized, high-volume delivery of discrete works to deterministic specifications [27]. The intellectual scaffolding arrived in parallel: the codification of limit-state design, the institutionalization of the factor of safety, and decisively the importation of stationarity as an unexamined axiom of hydrological and structural practice [24]. Engineering optimization culture, inheriting the mathematics of operations research, supplied the final ingredient: the reframing of every infrastructure question as a constrained optimization over a known objective. The paradigm became dominant because, in its era, it was true enough.

Why it persists despite obsolescence

The paradigm's persistence is the central puzzle. It survives not on its merits but through institutional embedding. It is written into design codes that specify deterministic return-period loads as though the return periods were stable. It is enforced by procurement systems that contract for delivery of a specified object, not for the adaptive provision of a service. It is protected by liability frameworks that assign accountability to discrete design decisions at a fixed point in time. It is reproduced by university curricula that teach optimization against fixed requirements and rarely teach decision-making under non-stationarity. And it is internalized in professional identity: the self-image of the engineer as the one who computes the right answer and commits to it.

Why it is now dangerous

What changed is that all four load-bearing assumptions have failed at once. Functional fixity fails as demand patterns shift faster than asset lives. Stationarity fails under anthropogenic climate change, the foundational result that stationarity is dead [24]. Deterministic degradation fails as compound climate hazards interact non-linearly with ageing materials. Separability fails as digitalization couples previously isolable assets into interdependent networks whose failures cascade [19]. A paradigm that was an efficient simplification under its founding conditions becomes, when those conditions vanish, a systematic generator of error. This is the category error of the thesis, now made precise: optimizing the wrong ontology.

Defining Infrastructure Intelligence

The need for a rigorous definition

“Smart” has been so thoroughly diluted by marketing that it can no longer do conceptual work. The related term “infrastructure intelligence” fares only slightly better: it now names an active research area visible, for instance, in the dedicated journal noted in Section 1 and in a fast-growing digital-twin literature applying AI to civil structures [14, 22, 32], but is generally used descriptively rather than with the necessary and sufficient conditions a design code or certification regime would require. We therefore define Infrastructure Intelligence by its function of decision-making under uncertainty rather than by its components, and supply the conditions that existing usage leaves implicit.

Formal definition

Definition. Infrastructure Intelligence is the capacity of an infrastructure system to (i) sense its own state and environment, (ii) maintain an updatable internal model of itself and its operating context, (iii) evaluate alternative future actions against goals that include its own continued function, and (iv) act upon, or materially shape, decisions that alter its physical or operational configuration all at timescales and under uncertainty conditions where ex-ante human specification of the correct action is impossible.

Necessary and sufficient conditions

A system possesses Infrastructure Intelligence if and only if it satisfies all four of the following necessary conditions, whose conjunction is sufficient:

Reflexive self-modelling: it maintains a model of its own state that it updates from data, not merely a model held by external operators.

Goal-directed evaluation: it evaluates actions against explicit objectives that include self-preservation of function.

Run-time decision authority: it exercises, or directly conditions, decisions during operation that were not fully specified at design time.

Uncertainty-native operation: its decision procedures represent and act under irreducible uncertainty rather than assuming a known environment.

Table 1: Infrastructure Intelligence against adjacent concepts. $\checkmark$ = required and present; , = not required.

Concept	Self-model	Goal eval.	Run-time authority	Unc.-native
Monitoring system	–	–	–	–
Smart infrastructure	$\checkmark$	–	–	–
Digital twin	$\checkmark$	–	–	–
ITS (fixed model)	$\checkmark$	$\checkmark$	–	–
Cyber-physical sys.	$\checkmark$	$\checkmark$	$\checkmark$	–
Infra. Intelligence	$\checkmark$	$\checkmark$	$\checkmark$	$\checkmark$

Exclusion criteria and boundary conditions

The definition deliberately excludes systems widely mis-labelled as intelligent. A digital twin that mirrors an asset’s state but issues no decisions fails condition 3. A structural-health-monitoring array that streams data to human engineers fails conditions 2 and 3. Intelligent transportation systems executing fixed signal-timing optimizations fail condition 4 (they assume a known demand model). A cyber-physical system with hard-coded control laws fails condition 1 if it does not update its self-model. Smart infrastructure, as marketed, is typically connected and instrumented but decision-mute. Table 1 formalizes these distinctions. We take the boundary condition to be decisive: Infrastructure Intelligence begins precisely where design-time specification ends and run-time judgement begins.

Toward a quantitative formalization

The four necessary conditions of Section 4.3 are stated categorically so that they can be applied by inspection. For future empirical and design work, however, a graded, measurable form is more useful than a yes/no test. We propose the following formalization as a basis for future operationalization and instrumentation, not as a validated measurement instrument, which would itself require the expert-elicitation and comparative-case validation described in Section 5.

For an infrastructure system i observed at time t , define five variables, each normalized to $[0, 1]$ where a continuous scale is meaningful:

- $S_i(t)$ *sensing capability*: the fraction of state variables relevant to the system’s decision function that are directly observed, rather than assumed, at time t .
- $M_i(t)$ *self-model quality*: one minus the normalized out-of-sample prediction error of the system’s internal model of its own state (e.g., $1 - \text{NRMSE}$ between modelled and

subsequently observed structural response).

- $G_i \in \{0, 1\}$, *goal-directedness*: an indicator equal to 1 if and only if action selection explicitly optimizes, or is explicitly constrained by, an objective function that includes the system's own continued function, and 0 otherwise.

- $R_i(t) \in [0, 1]$, *run-time decision authority*: the fraction of operationally consequential decisions, within a defined decision-relevant latency window, that are executed without an intervening human sign-off and that were not fully specified at design time.
- $U_i \in \{0, 1\}$, *uncertainty-native operation*: an indicator equal to 1 if and only if the decision procedure represents and propagates uncertainty explicitly (e.g., via a posterior distribution, ensemble, or robust-optimization formulation) rather than acting on a point estimate.

The categorical definition of Section 4.3 is then the conjunction

$$II_i(t) = \mathbb{I}[M_i(t) \text{ is data up-dated}] \wedge \mathbb{I}[G_i = 1] \wedge \mathbb{I}[R_i(t) > 0] \wedge \mathbb{I}[U_i = 1], \quad (1)$$

where $II_i(t) = 1$ if and only if system i possesses Infrastructure Intelligence at time t in the strict, necessary-and-sufficient sense of Section 4.3; sensing $S_i(t)$ is a precondition for a data-updated $M_i(t)$ rather than an independent necessary condition, consistent with Table 1.

For the graded maturity classification of Section 5, where the question is not whether a system qualifies but how far along the evolution axis it sits, we propose a companion continuous index,

$$III_i(t) = w_S S_i(t) + w_M M_i(t) + w_G G_i + w_R R_i(t) + w_U U_i, \sum_k w_k = 1, w_k \geq 0, \quad (2)$$

where the weights w_k are not asserted here but are left as free parameters to be fixed empirically, for example via the Delphi and comparative-case procedures of Section 5. Equation 2 is deliberately weaker than Equation 1: it supports ranking and stage-boundary estimation, but because it is a weighted sum it can assign a high score to a system that is high on some variables and zero on others, which Equation 1's conjunction correctly excludes. The two are intended to be used together: Equation 1 to classify, Equation 2 to rank within a class.

Proposition 1 in Section 9 can now be stated in these terms. Let $\alpha_i = dM_i/dt$ denote the system's model-updating rate and let

β denote the rate of drift of the environment's governing distribution (e.g., the rate of change of design-relevant load or hazard

statistics). Proposition 1 is then the claim that Infrastructure Intelligence improves resilience relative to a robust, low-precision design if and only if $\alpha_i > \beta$, which is a directly estimable crossover condition rather than a qualitative claim.

The Infrastructure Evolution Model

We now develop the paper's principal analytical instrument: a five-stage model of infrastructure evolution. We stress at the outset that this is not a technology-maturity ladder. The stages are differentiated by the locus of decision-making authority and the governance structure, not by the sophistication of the hardware. A bridge stuffed with sensors that changes nothing about who decides remains Infrastructure 1.0 with telemetry.

The five stages

Infrastructure 1.0 Physical Assets. Decision locus: human, at design time. Knowledge model: deterministic, codified ex ante. Uncertainty model: safety factors absorb a stationary envelope. Human role: designer-commander. Governance: hierarchical, code-compliant. Failure mode: brittle exceedance of design loads. Resilience: robustness via over-specification. Vulnerability: unmodelled non-stationary hazards.

Infrastructure 2.0, Connected Infrastructure. Decision locus: human, informed by telemetry. Knowledge model: empirical, data-augmented. Uncertainty: quantified post hoc. Human role: monitor. Governance: hierarchical with feedback. Failure mode: data overload, alarm fatigue. Vulnerability: the illusion that measurement equals understanding.

Infrastructure 3.0, Intelligent Infrastructure. Decision locus: shared; algorithms recommend, humans decide. Knowledge model: predictive, learned. Uncertainty: modelled probabilistically. Human role: supervisor. Governance: socio-technical. Failure mode: automation bias, model misspecification. Vulnerability: brittle confidence in learned models outside their training distribution.

Infrastructure 4.0, Adaptive Infrastructure. Decision locus: system adapts within human-set envelopes. Knowledge model: reflexive, self-updating. Uncertainty: treated as irreducible and managed. Human role: steward, setting goals and constraints. Governance: adaptive, polycentric. Failure mode: goal misalignment, adaptation to the wrong objective. Resilience: graceful degradation, antifragile potential. Vulnerability: value misspecification.

Infrastructure 5.0, Autonomous Infrastructure Systems. Decision locus: the system, within constitutional limits. Knowledge model: generative, self-revising. Uncertainty: native. Human role: constitutional designer and final arbiter. Governance: machine-speed with human constitutional oversight. Failure mode: emergent collective behavior across autonomous systems; un-auditable decisions. Vulnerability: the accountability vacuum.

Classification, hybrids, and conditions

The model is designed so that an independent researcher can classify a real system by asking four observable questions corresponding to the four necessary conditions of Section 4. Real systems are overwhelmingly hybrid: a motorway may run

5.0 ramp-metering over a 1.0 deck. The model treats stage as a property of decision functions, not whole assets, so hybridity is expected, not anomalous. Necessary condition for stage $n + 1$: the decision authority demonstrably resides one level further from design-time human specification than at stage n . Sufficient condition: that authority is exercised in operation under uncertainty. Boundary condition: the human role transforms qualitatively (commander $\rightarrow$ monitor $\rightarrow$ supervisor $\rightarrow$ steward $\rightarrow$ constitutional arbiter). Failure condition: a stage is unstable when its governance structure lags its decision locus.

Relation to existing maturity models, and why five stages

A five-level, decision-authority-defined staging is not without precedent. The closest analogue is SAE International's taxonomy of driving-automation levels [29], which defines six levels (0–5) by the allocation of the dynamic driving task between human and system, explicitly rejecting sensor count or software sophistication as the classification basis the same design decision defended in Section 5.2, adopted by regulators because a decision-authority criterion is externally observable and legally actionable where a technology-sophistication criterion is not. We adopt five stages rather than six because civil infrastructure has no equivalent of SAE's Level 0: Infrastructure 1.0; already presumes basic engineering knowledge encoded at design time.

A second precedent is the Capability Maturity Model Integration (CMMI) framework [5], which also uses five levels differentiated by the locus and formality of process control rather than by tooling. CMMI's levels are not directly transferable they describe organizational process maturity, not decision authority over a physical system but the precedent supports two choices made here: that five levels is a defensible resolution for an engineering maturity model (finer-grained models have generally not achieved adoption, since assessors rarely discriminate more than five to seven ordinal levels reliably), and that stage boundaries should be qualitative changes in control structure rather than thresholds on a continuous score, which is why Equation 1 rather than Equation 2 anchors the stage definitions.

The Infrastructure Evolution Model has not yet been empirically validated; Sections 5.1–5.2 present an

analytical instrument, not a measured result. Three validation pathways would allow the stage boundaries to be tested and, if necessary, revised:

Delphi study. Structured, iterated expert elicitation [6] with panels of civil engineers, asset managers, and AI-systems engineers, converging on inter-rater agreement about stage boundaries for anonymized real system descriptions.

Comparative document analysis. Independent coding of published SHM and digital-twin deployment reports against the four necessary conditions of Section 4, with inter-rater reliability (e.g., Cohen's or Fleiss' κ) reported.

Comparative literature analysis. Mapping the maturity constructs used in adjacent frameworks (smart-city, BIM, and industrial control/SCADA security maturity models) onto the four conditions, to test whether the proposed boundaries are a special case of a general pattern or an artefact of this paper's framing.

A minimal measurement protocol for pathway (1) would specify, in advance: a panel of no fewer than 15–20 practitioners balanced across the three named specialisms; a convergence criterion (e.g., stopping when round-on-round change in the median stage-boundary rating falls below a pre-set threshold, typically after two to three Delphi rounds); and a pre-registered inter-rater reliability target (Fleiss' $\kappa \geq 0.6$ is a common threshold for “substantial” agreement) below which the boundary in question would be reported as unresolved rather than forced to consensus. Pathway (2) requires the same reliability threshold applied to independent coders before any classification is used in analysis, with disagreements adjudicated by a third coder rather than by the original two. None of these has been carried out here; we flag them as the immediate next step (Grand Research Agenda, Challenge 1) and return to them in Section 13.

Technologies as Symptoms, Not Causes

The enabling technologies artificial intelligence, digital twins, the Internet of Things, advanced sensing, predictive analytics, cyber-physical integration are routinely narrated as drivers of transformation. We argue the reverse. They are symptoms of a single underlying shift whose deeper cause is the failure of the Asset Paradigm's assumptions: the migration from design-time completeness to run-time adaptability.

The twentieth-century engineering contract was a promise of design-time completeness: specify all requirements, verify against them, commission, and the system is done. Every assumption of the Asset Paradigm depends on this. Once stationarity fails, design-time completeness becomes impossible in principle no specification written today can enumerate the load and demand regimes of 2075. The system must therefore retain the capacity to decide at run time what could not be specified at design time. AI is not the cause of this; it is the discipline reaching for the only available tool once completeness became impossible. Digital twins are the instruments of run-time self-modelling, and recent civil-engineering implementations increasingly close this loop by feeding monitored structural state back into decision-relevant models rather than into static visualizations alone [32, 22, 14]. IoT and sensing are the afferent nerves of run-time state estimation. Predictive analytics is the substitution of learned, updatable models for fixed code provisions. Read this way, the technology wave is not an opportunity the discipline may choose to adopt; it is the visible surface of a paradigm collapse already underway.

The current technology wave, read against the framework

Three developments since 2023 sharpen this reading and deserve explicit treatment, since a reviewer familiar with the AI literature would otherwise reasonably ask why a 2026 paper on infrastructure and AI omits them.

Foundation models and large language models [4] have moved from a general AI concept to specific civil-engineering use within two years: automated structural analysis and code generation from natural-language specification, LLM agents that verify structural-analysis workflows [20], multi-agent LLM systems for foundation design [34], and post-event damage assessment from LLM-based reasoning [15]. In our terms, these are overwhelmingly Infrastructure 2.0–3.0: they raise M , and in agentic configurations R over bounded

sub-tasks, but typically retain a human in the loop for physical action and rarely report U as a first-class property. An LLM agent that recommends a repair action fluently but without calibrated uncertainty is Infrastructure 3.0 with an articulate interface, not 4.0.

Reinforcement learning, multi-agent systems, and agentic AI have matured furthest in traffic-signal and network control, where recent reviews report substantial, empirically measured gains over fixed-time baselines [23], the clearest existing instance of Infrastructure 3.0–4.0 in daily operation at scale, with an explicit reward function ($G = 1$) and coordinating-agent architectures now marketed elsewhere as “agentic AI.” We treat that label, in this paper’s vocabulary, not as a new category

but as a naming convention for systems combining goal-directed evaluation (G) with widened run-time authority (R); the framework’s contribution is to insist the label be accompanied by an audit of U and of the Proposition 4 escalation conditions, which much current commercial agentic-AI literature does not yet report.

Edge AI and distributed sensing move inference from centralized servers to on-structure or roadside compute, motivated by latency and connectivity constraints on safety-relevant decisions; fleet-based SHM using instrumented connected vehicles is an early instance [30]. This matters to the framework because it changes where condition 3 is physically exercised, and therefore where liability and auditability must be engineered in (Section 11).

None of this changes the paper’s thesis; if anything, the speed of foundation-model and agentic-AI diffusion into civil engineering over 2024–2026 [11] strengthens the urgency of Sections 11 and 12: technical capacity for 3.0–4.0 operation is arriving faster than the certification and liability mechanisms that would make its use accountable.

Conceptual Framework: Figures and Claims

The Infrastructure Evolution Model

Figure 1 claims the five stages are separated by shifts in decision authority, not by year or technology




Figure 1: The Infrastructure Evolution Model. Analytical schema (not derived from measured data). Claim: stages are defined by displacement of decision authority from the design-time human, not by technological vintage.

The Predictive Precision, Systemic Resilience Trade-off

Figure 2 advances the paper's most contestable theoretical claim: under climate non-stationarity, maximizing predictive precision can reduce systemic resilience, because precision is purchased by overfitting to a historical distribution that no longer holds. The curves shown are a conceptual schematic illustrating the hypothesized relationship of Proposition 3, not an empirical result.




Figure 2: Predictive precision vs. systemic resilience (conceptual schematic, illustrative). Claim: non-stationarity shifts the resilience-maximizing operating point toward lower precision and higher adaptive reserve.

The Paradigm Collapse Model

Figure 3 claims that the Asset Paradigm fails as a phase transition once accumulated assumption-violations exceed the buffering capacity of safety factors; again, the curve is a conceptual schematic rather than a fitted empirical trajectory.




Figure 3: The Paradigm Collapse Model (conceptual schematic, illustrative). Claim: paradigm failure is a phase transition, not a gradual decline.

Governance, transition, and maturity

Three further conceptual models are consolidated in Table 2: the Governance Evolution Framework (hierarchical → socio-technical → polycentric → constitutional), the Infrastructure Intelligence Maturity Model (mapping the five stages to assessable organizational capabilities), and the Asset-to-Intelligence Transition Framework (Figure 1 read as a governance trajectory).

Table 2: Infrastructure Intelligence Maturity Model.

Stage	Governance	Ethical frame	Dominant failure
1.0	Hierarchical	Duty of care	Brittle exceedance
2.0	Hier. + feedback	Data stewardship	Alarm fatigue
3.0	Socio-technical	Accountable rec.	Automation bias
4.0	Adaptive/polycentric	Value alignment	Goal misspecification
5.0	Constitutional	Machine accountability	Accountability vacuum

Illustrative

Application: Bridge Structural Health Monitoring

The framework above is deliberately domain-general. This section works a single domain through it end to end, to show that the abstractions of Sections 4–7 constrain something concrete rather than accommodating anything. We use a representative medium-span highway bridge (a common composite steel-girder or prestressed-concrete typology subject to periodic heavy-vehicle loading and, increasingly, scour and thermal loading outside its original design envelope) as a worked, hypothetical scenario grounded in established structural-health-monitoring (SHM) and digital-twin practice [32, 22, 14, 30], rather than as a report of a specific, named, implemented project. Figure 4 summarizes the flow described below.

feedback: monitored outcome updates M




Figure 4: Illustrative application to bridge structural health monitoring: current Asset Paradigm practice, Infrastructure Intelligence instrumentation, the bounded decision process, and expected engineering benefit. Hypothetical, literature-grounded scenario, not a reported implemented case.

Current Asset Paradigm state (Infrastructure 1.0–2.0)

Under conventional practice, the bridge is managed as a discrete asset: a fixed inspection interval (typically biennial visual inspection), a rating factor computed once and held constant between assessments, and a maintenance schedule set by a deterministic deterioration curve rather than measured condition. Where sensors exist, they typically populate a dashboard for human review (Infrastructure 2.0: $S > 0$, but M is held by external operators, so $G = 0$ and $R = 0$).

Infrastructure Intelligence instrumentation (target: Infrastructure 3.0–4.0)

An SHM sensor network (strain, acceleration, tilt, and displacement sensors, plus, following [30], data opportunistically gathered from instrumented vehicles crossing the structure) raises S . A digital-twin model physics-based, data-driven, or hybrid, consistent with recent implementations [32, 22], is continuously updated against this stream, raising M . The model is required to carry calibrated uncertainty (a posterior over remaining capacity, not a single rating factor), setting $U = 1$; the operator defines an explicit objective balancing safety margin, availability, and lifecycle cost, setting $G = 1$.

The bounded decision process

Run-time authority R is not granted unconditionally: it is bounded by the Proposition 4 test. *Below-threshold, reversible actions* flagging a span for targeted inspection, or issuing a temporary, automatically-lifted weight advisory when a monitored anomaly exceeds a pre-agreed threshold, execute at run time without sign-off, consistent with 3.0 moving toward 4.0 for this decision class. *Value-laden, irreversible, or unauditable actions* closing the bridge, ordering a major intervention, revising the public load

rating remain with a licensed engineer, since they satisfy all three Proposition 4 conditions and sit in the non-delegable core. This is the hybridity anticipated in Section 5.2: the same bridge is 4.0 for bounded monitoring and 1.0–2.0 for major interventions, by design.

Expected engineering benefits, and their evidentiary status

Directionally: fewer blanket inspections and more targeted ones; fewer unplanned closures, since degradation is flagged before a safety-critical threshold rather than at the next scheduled inspection; and longer effective service life at a given margin, since the margin is monitored rather than assumed. We report these directionally, not as a fabricated point estimate for this hypothetical structure: comparable RL-based control deployments in adjacent domains report large, empirically measured gains over static baselines (traffic-signal RL achieving reductions of several tens of per cent in travel time and queue length over fixed-time control [23]), which is suggestive, not a substitute for a bridge-specific study. Establishing the equivalent magnitude for SHM-driven bridge management is exactly the kind of study the Grand Research Agenda (Challenge 1) calls for, and a natural next step beyond this paper.

Falsifiable Questions and Testable Propositions

A framework that cannot be wrong is worthless. We therefore commit to propositions that future empirical work can refute, each stated with variables from Section 4.5 where applicable, a candidate test and data source, an applicable validation method, and a stated limitation.

Proposition 1 (adaptation-rate crossover). Infrastructure Intelligence improves resilience if and only if $\alpha_i > \beta$, where $\alpha_i = dM_i/dt$ is the system's model-updating rate and β is the rate of drift of the environment's governing distribution; where $\alpha_i < \beta$, intelligence reduces resilience relative to a robust, low-precision design. This formalizes, as a single testable criterion, the adaptation-tipping-point logic used qualitatively in dynamic adaptive policy pathways [10, 16], extended from pre-specified human decision points to continuously updating machine models. *Test and data:* estimate α_i from the update frequency and out-of-sample error trajectory of deployed SHM/digital-twin models, and β from rolling-window trends in governing hazard statistics (e.g., extreme discharge, wind, or traffic-load records in national bridge inventories or utility SCADA archives); test whether resilience outcomes are better predicted by $\text{sign}(\alpha_i - \beta)$ than by a static baseline. *Validation method:* changepoint detection (CUSUM or Bayesian online changepoint detection) to estimate α_i and β independently, followed by survival analysis (Cox proportional hazards, with $\alpha_i - \beta$ as a time-varying covariate) on time-to-incident. *Limitation:* β is not directly observable and must itself be estimated from a finite, non-stationary record, introducing a circularity the changepoint approach only partly resolves; falsifiable in principle, data-hungry in practice.

Proposition 2 (timescale-dependent authority reversal). Autonomous infrastructure (5.0) outperforms human-centred governance (3.0) where the decision timescale τ_{decision} is shorter than the human deliberative cycle τ_{human} (e.g., grid-frequency response, flood-gate actuation) and underperforms it on outcomes laden with contested values (e.g., equity of allocation under scarcity), independent of τ_{decision} . *Test and data:* compare outcome quality (frequency-deviation recovery time; distributional equity metrics) between matched autonomous and human-supervised deployments using grid-operator dispatch logs and utility curtailment records, stratified by $\tau_{\text{decision}}/\tau_{\text{human}}$. *Validation method:* matched difference-in-differences or regression discontinuity around $\tau_{\text{decision}} \approx \tau_{\text{human}}$, since a genuine reversal implies a sign change in treatment effect either side of it. *Limitation:* autonomous and human-supervised deployments are rarely randomly assigned, so confounding by which assets get automated threatens identification; a stepped-wedge or natural-experiment design would be needed.

Proposition 3 (precision, fragility trade-off). Increasing predictive precision raises systemic fragility wherever it is obtained by fitting a non-stationary process as if stationary (the formal content of Figure 2). *Test and data:* out-of-distribution stress-testing of precision-optimised controllers or capacity models train on a historical window, evaluate across a documented regime shift (e.g., a pre-/post-extreme-event split in SHM, hydrological, or traffic-demand data) and test whether higher in-sample precision predicts larger out-of-sample degradation. *Validation method:* adversarial/distribution-shift benchmarking, with fragility operationalised as the slope of the precision-vs-out-of-sample-error curve. *Limitation:* requires an identified

shift of known timing to test cleanly; absent one, the test degrades to a weaker correlational comparison between model complexity and forecast error

Proposition 4 (the non-delegable core). Human oversight remains strictly necessary wherever a decision is (i) value-laden, (ii) irreversible, and (iii) un-auditable at machine speed; the simultaneous presence of all three defines the non-delegable core of engineering judgement. *Test and data:* structured coding of infrastructure failure post-mortems and near-miss databases (e.g., ASCE and equivalent national failure case archives) for the three properties, cross-tabulated against whether human oversight was present at the point of failure. *Validation method:* content analysis with a pre-registered coding scheme and inter-rater reliability (Cohen's κ) between at least two independent coders, since “value-laden” and “un-auditable” require judgement calls that must themselves be shown reliable. *Limitation:* failure databases are selected on failure, not the full decision population, so the test can show association but not, alone, establish that absence of the three properties is sufficient for safe delegation; a matched comparison against successful autonomous decisions would be needed.

Anchor Research Questions

RQ1. Can a quantitative adaptation-rate versus non-stationarity-rate criterion (P1) be estimated ex ante for real infrastructure classes, or is it knowable only after failure?

RQ2. Is antifragility constructible in load-bearing civil systems, or is it confined to informational subsystems while physical structure remains, at best, resilient?

RQ3. Do networks of autonomous infrastructure systems exhibit emergent collective failure modes absent from any individual system, and can these be bounded a priori?

RQ4. What is the maximum decision latency at which democratic oversight of machine-speed infrastructure remains meaningful rather than ceremonial?

Domain Implications and Transition Barriers

The framework's value is tested by its bite across domains. In transportation, the 3.0→4.0 barrier is regulatory: signal and ramp control can adapt technically but liability law forbids ceding the decision. In rail and airports, safety-case regimes built for design-time completeness cannot certify run-time learning systems a structural, not technical, barrier.

For bridges, dams, and offshore structures, the 4.0→5.0 barrier is epistemic and ethical: we do not know how to certify a load-bearing structure that revises its own operating envelope, and the failure consequences forbid learning by trial. Water and flood defence are where non-stationarity bites hardest and adaptive (4.0) intelligence is most justified yet financing instruments demand fixed specifications.

Energy infrastructure is furthest into 5.0 (autonomous grid balancing) and the leading indicator of the accountability problems other domains will face. Smart cities and national critical infrastructure expose the interdependency barrier: stage progression in one system imposes uncompensated risk on coupled systems. Across domains the barriers rank, in descending tractability: technical < financial < regulatory < institutional < cultural.

A feasible transition pathway

Section 15 states four institutionally targeted actions as commitments with fixed deadlines. Table 3 operationalizes those same commitments as a phased pathway across the practical dimensions a transition actually requires: policy, standards, procurement, certification, and industry adoption, so that the roadmap can be assessed for feasibility rather than read as aspiration alone.

The phasing is deliberately conservative: each phase only assumes the decision authority already demonstrated in the phase before it, mirroring the stage-hybridity logic of Section 5.2.

Table 3: A phased transition pathway from Asset Paradigm practice toward Infrastructure Intelligence.

Dimension	Near-term (2026–2030)	Medium-term (2031–2035)	Long-term (2036–2050)
Policy	Position papers; model-code drafting for probabilistic provisions	Adoption of adaptive design chapters (Section 15, action 1)	Harmonization across jurisdictions
Standards / certification	Pilot decision-integrity security guidance; certification pathway for bounded, low-risk run-time learning (e.g., signal timing)	Certification extended to Infrastructure 4.0 bounded decisions	Constitutional governance standards for 5.0 systems
Procurement	Performance-based pilots; DMDU-style adaptive-pathway contracts (Section 15, action 3)	Adaptive-service procurement as a standard option	Default procurement mode for eligible asset classes
Practice/education	Elective stewardship modules; accreditation requirement drafted (Section 15, action 2)	Stewardship competency mandatory in accredited programs	First fully trained cohort in post (Section 15, action 4)
Industry adoption	Instrumented pilots per domain; domain-specific playbooks	Cross-domain interoperability standards	Majority of new-build assets Infrastructure 3.0+

The pathway is intentionally front-loaded with low-risk, reversible actions (pilots, position papers, elective modules) and defers governance of irreversible, value-laden decisions to later phases, consistent with Proposition 4. It is also intentionally domain-agnostic in its early phases and domain-specific from Phase 2 onward, because Section 10 shows the barriers themselves are domain-specific: a single global timeline would understate how far ahead energy and traffic control already are, and how far behind bridges and dams necessarily remain.

Ethics, Governance, Cybersecurity, Digital Inequality

The defining governance problem of Infrastructure 4.0–5.0 is the accountability vacuum: when infrastructure makes consequential decisions that no human fully authored or can audit at the speed they occur, the liability frameworks of the Asset Paradigm have no subject to attach to. Democratic oversight cannot be preserved by inserting a human into a machine-speed loop a human who must approve in milliseconds approves nothing. It must instead migrate to the constitutional layer: humans set, audit, and revise the goals, constraints, and inviolable limits within which autonomous systems act, and retain the authority to suspend.

We distinguish two security problems. Cybersecurity of connected infrastructure (2.0–3.0) protects data and actuation integrity, the classical confidentiality-integrity-availability triad. Cybersecurity of autonomous decision-making infrastructure (4.0–5.0) must additionally protect decision integrity: the system can be compromised not by falsifying its data but by corrupting its objectives or poisoning the models through which it learns. To our knowledge, no current civil-engineering-specific security standard yet contemplates this; general-purpose AI risk frameworks are only beginning to engage the problem for critical infrastructure. A national standards body's April 2026 concept note toward a dedicated critical-infrastructure profile of its AI risk management framework [25] indicates the field is still at the scoping stage, not standard-setting. Finally, digital inequality becomes infrastructural inequality: when adaptive flood defense protects wealthy coasts while poorer regions retain brittle 1.0 assets, the technology stratifies

Critique of the Profession and of the Framework

Mechanisms that preserve the Asset Paradigm

The profession reproduces the paradigm through identifiable mechanisms. Design codes encode stationary return periods as deterministic targets. Procurement contracts for delivered objects, structurally excluding adaptive service provision. Liability frameworks attach accountability to fixed design decisions, penalizing run-time adaptation. Educational models teach optimization against fixed requirements. Accreditation certifies mastery of the old ontology. Funding systems reward discrete capital projects over adaptive capacity. Academic incentives reward incremental optimization within the paradigm over replacing it. These mechanisms are mutually reinforcing: a lock-in in the precise sense of transitions theory [21].

Surviving the strongest objections

We confront the objections most likely to be raised.

Technological determinism: our thesis inverts the deterministic claim that technology is symptom, not cause; the objection misfires. *Overreliance on AI/loss of accountability*: a structural risk, not a contradiction, and precisely what constitutional governance is designed to bound; P4 specifies the non-delegable core.

Liability allocation between engineers and AI systems: existing liability law attaches responsibility to a named professional at a fixed decision point (Section 3), which is exactly what run-time, machine-executed decisions disrupt. We do not resolve this here, but P4's three-part test (value-laden, irreversible, un-auditable) is offered as the criterion a liability framework would need to operationalize: it is precisely on decisions failing that test that human professional liability should, we argue, remain attached.

Cyber vulnerability: a genuine structural limitation, which is why we elevate decision-integrity security to first-order status.

Computational cost: continuous sensing, model updating, and uncertainty propagation (Section 4.5) are not free, and a full lifecycle-cost accounting of Infrastructure 3.0–4.0 operation against Asset Paradigm maintenance is an open empirical question. Where that accounting is un-favourable, the framework predicts (via Proposition 1) that the rational choice is to remain at a lower stage, not to intelligence-wash a brittle asset.

Public acceptance: citizens and elected officials may resist ceding visible infrastructure decisions to automated systems regardless of demonstrated performance a legitimacy constraint distinct from the technical and governance barriers of Section 10. We treat this as a contested-values case inside Proposition 4's non-delegable core, not as a barrier to be engineered around.

Regulatory inertia and institutional resistance: real and severe, but contingent barriers rather than refutations.

Economic infeasibility: a temporary challenge the costs of brittle assets under non-stationarity already exceed adaptive alternatives on whole-life accounting.

Ethical contradiction: the strongest objection, and the reason the framework confines autonomous authority to value-neutral, reversible, auditable decisions (P4) and reserves the rest for human stewardship.

The framework is strengthened, not defeated, by each.

Limitations of the Proposed Framework

Consistent with its stated aim, this paper is a conceptual and theoretical contribution: it proposes a framework and a set of falsifiable propositions, and it does not itself report empirical validation of either.

We state the resulting limitations here as boundaries on what has been established, and treat each as a concrete next step rather than an open-ended gap. We considered, and deliberately did not attempt, a small proof-of-concept pilot in place of the illustrative case study of Section 8.

A genuine pilot even a minimal one requires instrumented access to a real asset, monitored outcome data over a period long enough to observe the adaptation-rate dynamics of Proposition 1, and institutional partners willing to share operational records; assembling this was judged likely to narrow the framework prematurely to whatever single system such access permitted, before the cross-domain and stage-boundary questions of Sections 5 and 5.3 had been worked out conceptually. We regard this as a sequencing choice, not a decision to avoid empirical accountability, and Section 14 commits the framework to exactly that pilot work as the immediate next step. The necessary and sufficient conditions (Section 4) and quantitative formalization (Section 4.5) are proposed constructs: the weights w_k in Equation 2 are left free, and neither form has been fitted or tested against measured system behavior.

The boundaries between the five evolution stages (Section 5) rest on conceptual argument and analogy to maturity models in adjacent domains; whether the true boundaries are sharp or graded is an empirical question the Delphi and comparative-case procedures of Section 5.3 have not yet answered. The four propositions (Section 9) are stated with candidate variables and test designs, but none has been executed; the bridge case study (Section 8) illustrates the framework's logic and is not offered as evidence for the propositions themselves.

Finally, the framework is domain-general by design, which is a source of its explanatory reach but also means every domain-specific claim in Section 10 requires separate confirmation before it can be relied upon in that domain.

From proposal to validated theory. We envisage the framework's validation following the same progression by which comparable engineering maturity constructs (CMMI, SAE J3016) moved from proposal to standard:

Conceptual Framework → *Operational Definitions* → *Pilot Engineering Studies* → *Multi-Domain Validation* → *Standardization* → *Professional Adoption*

Concretely: (1) the definitions of Section 4 must be operationalized into measurement protocols for S , M , R , and U on real instrumented assets; (2) pilot studies, beginning with a single domain such as bridge SHM, should test Proposition 1's crossover criterion against monitored outcomes; (3) results should be replicated across at least two further domains (e.g., water networks, traffic control) to test the domain-general claims of Section 10; (4) convergent pilot results would support the stage-boundary standardization exercises of Section 5.3; and (5) standardization would be the precondition for the accreditation and certification steps already proposed in Table 3.

We regard the framework as complete for its stated purpose a conceptual foundation others can build on and, most importantly, attempt to refute and incomplete, by design, as an empirical theory.

The Grand Research Agenda, 2025–2050

We identify ten Grand Challenges, each multidisciplinary, each capable of anchoring multiple doctoral programs and shaping policy.

The Adaptation-Rate Theorem. Establish whether the P1 crossover criterion (α_i vs. β , Section 4.5) can be derived analytically for classes of infrastructure, and validate the Infrastructure Evolution Model's stage boundaries via the Delphi and comparative-analysis pathways of Section 5.3.

Constructible Antifragility. Determine whether load-bearing civil systems can be made to gain from disorder, and how to certify them, building on emerging technical-antifragility formalisms [3].

Emergent Network Autonomy. Characterize and bound emergent failure across populations of autonomous infrastructure.

Machine-Speed Democracy. Design constitutional governance that preserves meaningful human authority over systems acting faster than deliberation.

Decision-Integrity Security. Build a security science for infrastructure whose objectives, not just data, are attack surfaces.

Certification of Run-Time Learning. Replace design-time safety cases with frameworks that certify bounded learning in operation.

Equitable Intelligence. Prevent infrastructure intelligence from stratifying survival across regions and incomes.

The Stewardship Competency. Define, teach, and accredit the skills of infrastructure stewardship as distinct from design optimization.

Procurement for Adaptation. Invent contractual instruments that procure adaptive service provision instead of fixed deliverables, extending DMDU-style pathway contracts [10] beyond water infrastructure.

Epistemology of Run-Time Judgement. Formalize a theory of valid judgement under deep uncertainty to replace the deterministic epistemology of the Asset Paradigm.

CONCLUSION

We close not with a summary but with a set of claims and commitments.

The discipline that built the modern world by learning to command matter is now being asked by a changing climate, by machines that learn, and by systems too interdependent to optimize in isolation to cede its command and learn instead to negotiate.

This is not a downgrade. It is among the most intellectually demanding transformations the profession has faced since it separated from military engineering. The engineer of 2050 will not be measured by the load a beam carries but by the wisdom with which she sets the constitutional constraints within which infrastructure thinks for itself, by the clarity with which she distinguishes the delegable from the non-delegable, and by her refusal to outsource responsibility for outcomes that machines cannot comprehend.

What is at stake is authorship. The infrastructure that emerges over the next quarter-century will either be co-constructed by a profession that has earned the right to steward it, or it will be constructed by default by technology firms, by optimization algorithms, and by the accumulated stresses of a non-stationary planet that does not negotiate. The profession chooses, now, whether it will write the next chapter of its own story or discover that it has been written for it.

We therefore call for four institutionally targeted, time-bound actions within the next decade: model codes replacing deterministic, stationary provisions with probabilistic, adaptive frameworks, adopted by the lead professional bodies by 2030; an infrastructure-stewardship certification requirement added to accredited undergraduate programs by 2032; adaptive-service procurement instruments piloted by national infrastructure agencies and development banks by 2028; and the Grand Research Agenda of Section 14 jointly funded by the research councils of at least ten nations, yielding a first trained cohort by 2035.

Table 3 sets these four actions inside a fuller, phased pathway across policy, standards, procurement, and practice.

The twentieth-century engineer commanded concrete, steel, and the deterministic calculus of safety. The twenty-first-century engineer must learn a harder art: to build systems that think, to set the constraints within

which they decide, and to accept that the final service of her intelligence is to make herself, in some decisions, unnecessary. That is not the end of civil engineering. It is the beginning of a larger discipline than the one we inherited.

REFERENCES

1. M. Amin, "Toward self-healing energy infrastructure systems," *IEEE Computer Applications in Power*, vol. 14, no. 1, pp. 20–28, 2001.
2. W. B. Arthur, "Complexity and the economy," *Science*, vol. 284, no. 5411, pp. 107–109, 1999.
3. C. Axenie, M. Akbarzadeh, M. A. Makridis, M. Saveriano, and A. Stancu, *Applied Antifragility in Technical Systems: From Principles to Applications*. Cham: Springer Nature, 2025.
4. R. Bommasani et al., "On the Opportunities and Risks of Foundation Models," *arXiv preprint arXiv:2108.07258*, 2021.
5. CMMI Product Team, *CMMI for Development, Version 1.3* (CMU/SEI-2010-TR-033). Pittsburgh, PA: Software Engineering Institute, Carnegie Mellon University, 2010.
6. N. Dalkey and O. Helmer, "An experimental application of the Delphi method to the use of experts," *Management Science*, vol. 9, no. 3, pp. 458–467, 1963.
7. C. Folke, "Resilience: the emergence of a perspective for social, ecological systems analyses," *Global Environmental Change*, vol. 16, no. 3, pp. 253–267, 2006.
8. C. Folke, T. Hahn, P. Olsson, and J. Norberg, "Adaptive governance of social-ecological systems," *Annual Review of Environment and Resources*, vol. 30, pp. 441–473, 2005.
9. F. W. Geels, "From sectoral systems of innovation to socio-technical systems," *Research Policy*, vol. 33, no. 6–7, pp. 897–920, 2004.
10. M. Haasnoot, J. H. Kwakkel, W. E. Walker, and J. ter Maat, "Dynamic adaptive policy pathways: a method for crafting robust decisions for a deeply uncertain world," *Global Environmental Change*, vol. 23, no. 2, pp. 485–498, 2013.
11. J. H. Holland, "Complex adaptive systems," *Daedalus*, vol. 121, no. 1, pp. 17–30, 1992.
12. C. S. Holling, "Resilience and stability of ecological systems," *Annual Review of Ecology and Systematics*, vol. 4, pp. 1–23, 1973.
13. INCOSE, *Systems Engineering Handbook: A Guide for System Life Cycle Processes and Activities*, 4th ed. Hoboken, NJ: Wiley, 2015.
14. F. Jiang, L. Ma, T. Broyd, and K. Chen, "Digital twin and its implementations in the civil engineering sector," *Automation in Construction*, vol. 130, art. 103838, 2021.
15. Y. Jiang, J. Wang, X. Shen, and K. Dai, "Large language model for post-earthquake structural damage assessment of buildings," *Computer-Aided Civil and Infrastructure Engineering*, vol. 40, no. 31, pp. 6324–6342, 2025.
16. J. H. Kwakkel, M. Haasnoot, and W. E. Walker, "Developing dynamic adaptive policy pathways: a computer-assisted approach for developing adaptive strategies for a deeply uncertain world," *Climatic Change*, vol. 132, no. 3, pp. 373–386, 2015.
17. E. A. Lee, "Cyber physical systems: design challenges," in *Proc. 11th IEEE Int. Symp. Object and Component-Oriented Real-Time Distributed Computing*, 2008, pp. 363–369.
18. R. J. Lempert, S. W. Popper, and S. C. Banks, *Shaping the Next One Hundred Years: New Methods for Quantitative, Long-Term Policy Analysis*. Santa Monica, CA: RAND Corporation, 2003.
19. R. G. Little, "Controlling cascading failure: understanding the vulnerabilities of interconnected infrastructures," *Journal of Urban Technology*, vol. 9, no. 1, pp. 109–123, 2002.
20. J. Liu, Z. Geng, R. Cao, L. Cheng, P. Bocchini, and M. Cheng, "A large language model-empowered agent for reliable and robust structural analysis," *Structure and Infrastructure Engineering*, 2026.
21. J. Markard, R. Raven, and B. Truffer, "Sustainability transitions: an emerging field of research and its prospects," *Research Policy*, vol. 41, no. 6, pp. 955–967, 2012.
22. J. Michael, J. Blankenbach, J. Derksen, B. Finklenburg, R. Fuentes, T. Gries, S. Hendiani, S. Herlé, S. Hesseler, M. Kimm, et al., "Integrating models of civil structures in digital twins: state-of-the-art and challenges," *Journal of Infrastructure Intelligence and Resilience*, vol. 3, no. 3, art. 100100, 2024.
23. P. Michailidis, I. Michailidis, C. R. Lazaridis, and E. Kosmatopoulos, "Traffic signal control via

- reinforcement learning: a review on applications and innovations,” *Infrastructures*, vol. 10, no. 5, art. 114, 2025.
24. P. C. D. Milly, J. Betancourt, M. Falkenmark, R. M. Hirsch, Z. W. Kundzewicz, D. P. Lettenmaier, and R. J. Stouffer, “Stationarity is dead: whither water management?” *Science*, vol. 319, no. 5863, pp. 573–574, 2008.
25. National Institute of Standards and Technology (NIST), *AI Risk Management Framework (AI RMF 1.0)*. Gaithersburg, MD: U.S. Department of Commerce, 2023.
26. T. Nyokum and Y. Tamut, “Artificial intelligence in civil engineering: emerging applications and opportunities,” *Frontiers in Built Environment*, vol. 11, art. 1622873, 2025.
27. A. O’Connor, *Rethinking Infrastructure as Assets: Policy, Regulation and Design*. London: Routledge, 2019.
28. E. Ostrom, “Beyond markets and states: polycentric governance of complex economic systems,” *American Economic Review*, vol. 100, no. 3, pp. 641–672, 2010.
29. SAE International, *J3016_202104: Taxonomy and Definitions for Terms Related to Driving Automation Systems for On-Road Motor Vehicles*. Warrendale, PA: SAE International, 2021.
30. H. Shokravi, M. Vafaei, B. Samali, and N. Bakhary, “In-fleet structural health monitoring of roadway bridges using connected and autonomous vehicles’ data,” *Computer-Aided Civil and Infrastructure Engineering*, vol. 39, no. 14, pp. 2122–2139, 2024.
31. N. N. Taleb, *Antifragile: Things That Gain from Disorder*. New York: Random House, 2012.
32. M. Torzoni, M. Tezzele, S. Mariani, A. Manzoni, and K. E. Willcox, “A digital twin framework for civil engineering structures,” *Computer Methods in Applied Mechanics and Engineering*, vol. 418, art. 116584, 2024.
33. E. L. Trist and K. W. Bamforth, “Some social and psychological consequences of the longwall method of coal-getting,” *Human Relations*, vol. 4, no. 1, pp. 3–38, 1951.
34. S. Youwai, D. Phim, V. G. Murcia, and R. C. Onas, “Large language model-based multi-agent systems for automated foundation design: router-driven task classification and expert selection framework,” *AI in Civil Engineering*, vol. 5, no. 1, art. 5, 2026.