ChatGPT

Weaving Understanding Under Uncertainty: A Modular Framework for Comprehending Complex, Multi-Species Systems Across Physical, Digital, and Quantum Domains

Noel Michael Conrad and OpenAI
White Paper (MLA-Style)
27 Feb. 2026

Abstract

Complex systems often exceed what can be fully understood “in the moment” because they exhibit high element interactivity (many simultaneously relevant parts), nonlinear behavior, and emergent properties that depend on context and history. [1] This paper develops a coherent framework for designing *comprehension itself* as an operational capability: how humans and AI systems can (a) recognize when they do not understand; (b) construct metaphors, analogs, and direct mappings that bridge from known domains to unknown ones; (c) increase the likelihood of “aha”/insight moments without confusing fluency for truth; and (d) build reusable modular mechanisms that allow sustainable scaling of multi-species systems—physical, digital, and (by analogy) quantum—while maintaining auditable clarity and versioned evolution. [2]

The core contribution is the **Cinch-and-Weave Framework**, a learning-and-design architecture that treats understanding as a fabric: stable “cinches” (high-confidence anchors, interfaces, invariants) combined with “weaves” (interleaved expansions, iterative testing, and periodic audits) that progressively strengthen coherence. The framework integrates evidence from cognitive science of analogy and metaphor (structure-mapping and constraint satisfaction), neuroscience of insight and metacognitive monitoring, and systems engineering principles of modularity, information hiding, architecture description, and life-cycle governance. [3]

A detailed case study illustrates how to design **modular building envelopes and structural systems for biodiversity**, leveraging multi-species design research (“ecolopes”), green roof ecology, bird-collision mitigation science, and statutory biodiversity measurement approaches to create scalable, auditable biodiversity-supporting architecture. [4]

Foundations: why “understanding” fails in complex moments

The practical problem is not simply missing facts; it is that **complexity can exceed real-time cognitive capacity**. Cognitive load theory distinguishes intrinsic load (driven by element interactivity) from extraneous load (driven by avoidable presentation and design choices). When element interactivity is high,

understanding becomes difficult precisely because many components must be processed together. [5] In these conditions, people often default to partial models, simplified narratives, or surface cues of coherence.

A second barrier is **bounded rationality**: even highly capable agents have limited time, attention, and information, implying that "optimal" understanding is frequently unattainable. Bounded rationality frameworks emphasize designing procedures and representations that support effective action under constraints rather than assuming full optimization. [6]

A third barrier is metacognitive: humans can be sincerely unaware that they do not understand. The "illusion of explanatory depth" demonstrates that people systematically overestimate how well they can explain mechanisms until asked to produce a detailed explanation; confidence falls when confronted with the need for causal detail. [7] This matters for both individual learning and organizational design reviews: confidence and fluency are not reliable proxies for comprehension.

Finally, complex adaptive systems—ecological, sociotechnical, and increasingly computational—are often **path-dependent and nonlinear**, so local intuitions can fail. Ecosystems specifically have been described as prototypical complex adaptive systems in which higher-level patterns emerge from localized interactions, with multiple possible outcomes and historical dependency. [8] Resilience thinking further emphasizes that systems can persist through variability by retaining function, but can also reorganize abruptly when thresholds are crossed. [9]

Together, these findings imply a design mandate: when real-time understanding is impossible, **agents need disciplined scaffolds** for (a) rapid partial comprehension, (b) safe uncertainty handling, and (c) structured expansion of understanding over time.

Building bridges: metaphors, analogs, and reliable "aha" moments

Analogy as a disciplined mapping process

A productive way to comprehend unfamiliar and complex domains is to map them onto a known domain through **relational structure**, not superficial similarity. Structure-mapping theory formalizes this: analogical reasoning primarily aligns relations among entities, prioritizing systematic relational correspondences. [10] This supports a robust method for building "direct connections" to higher-level scenarios: choose a base domain with well-understood relational structure, then map that structure onto the target domain to generate candidate inferences—followed by disciplined validation.

Complementarily, multi-constraint theories propose that analogical mapping is guided by interacting constraints—structural consistency, semantic similarity, and pragmatic relevance to goals—making analogy a constrained optimization rather than free association. [11] This is operationally important: it explains why metaphors and analogs can mislead if semantic similarity is high but structural correspondence is low (a common failure mode in strategy and product analogies).

Metaphor as schema: powerful, partial, and risky

Conceptual metaphor frameworks argue that metaphor is not merely decorative language but a mechanism for structuring thought by reusing patterns from embodied and social experience to

understand abstract domains. [12] Experimental work also supports that metaphors can function as schemas that influence how information is structured and interpreted, shaping reasoning beyond immediate language use. [13]

Because metaphors act as schemas, they can accelerate comprehension by providing ready-made structure—but they can also **import hidden assumptions**. This motivates a standard practice for complex-systems work: treat metaphors as **hypotheses** to test, not truths to adopt.

Insight ("eureka") cannot be continuous, but it can be cultivated

The neuroscience of insight characterizes "aha" moments as sudden reinterpretations that yield a nonobvious solution or meaning (including comprehension of metaphors). Reviews describe insight as involving patterns such as coarse semantic coding and shifts toward internally focused attention preceding solution, with measurable individual differences in insight tendency. [14]

These findings are consistent with a practical constraint: insight depends on a background process of restructuring and association. Attempting to force "eureka moments without pause" conflicts with known limits of attention and working memory. [15] However, insight frequency can be increased by designing workflows that alternate **analytic constraint-building** with **incubation and recombination**—not as mysticism, but as a deliberate alternation of modes around clear goals and representations.

A disciplined method for generating metaphors and analogs

The following "Analogy-to-Action" protocol operationalizes the research above into a reusable mechanism for both human designers and AI agents.

```
ANALOGY-TO-ACTION PROTOCOL

1) Target Frame (T): define what must be understood or decided.
2) Candidate Base Set (B): list 3–5 known domains with strong relational
structure.
3) Mapping Draft (M): map relations (not objects) from B → T; write 5–10
relational correspondences.
4) Inference Projection (I): list predictions or implications if M is correct.
5) Constraint Check (C): test M for (a) structural consistency, (b) semantic
plausibility, (c) goal relevance.
6) Falsification Hooks (F): define what observation would break the analogy.
7) Minimal Test (E): run a cheap test (simulation, small experiment,
counterexample search).
8) Adoption Level (A): label the analogy as: "explanatory," "design heuristic,"
or "decision rule."
```

This protocol is grounded in structure-mapping's emphasis on relational alignment and Holyoak-style constraint satisfaction, while embedding an explicit falsification step to reduce metaphor-driven overreach. [16]

Knowing when you don’t know: metacognition for humans and uncertainty systems for AI

Human-side: monitoring understanding without self-deception

Metacognition is the capacity to reflect on, evaluate, and control cognition, including judging uncertainty and confidence. Review work emphasizes that the *accuracy* of metacognition—how well confidence tracks performance—is central to adaptive behavior, and that miscalibration can lead to failure even when aptitude is sufficient. [17]

The illusion of explanatory depth demonstrates a particularly relevant failure mode: people are more overconfident about explanatory (mechanistic) knowledge than about other kinds of knowledge, and confidence drops when detailed explanation is demanded. [7] For complex system design, this implies a rule: **force mechanism-producing outputs** early and often (teach-back, causal diagrams, runnable specifications), because they expose hidden gaps.

Evidence-based “gap detection” practices for humans therefore emphasize:

- **Generate-then-check**: produce a causal explanation or procedure, then test it against reality (a prototype, counterexample, or peer critique) to reveal overconfidence. [7]
- **Load-aware chunking**: when element interactivity is high, reduce extraneous load by reorganizing information into schemas and modules, because working memory limits obstruct comprehension. [18]
- **Distributed reinforcement**: use spacing and structured revisiting to stabilize knowledge and reduce false familiarity. A large meta-analysis of distributed practice found extensive evidence across hundreds of assessments and experiments, supporting the practical value of spacing for retention. [19]

AI-side: knowledge limits as a design requirement, not a feature request

For AI systems, “do I understand?” must be formalized as **uncertainty, out-of-distribution detection, and abstention behavior**.

Calibration research shows that modern neural networks can be poorly calibrated: predicted probabilities may not reflect true correctness likelihood, motivating post-hoc calibration (e.g., temperature scaling) and systematic evaluation of confidence as an operational safety handle. [20] Out-of-distribution detection work provides baseline methods for detecting misclassified or OOD examples by using softmax probabilities, highlighting the feasibility—and limitations—of “this input looks unfamiliar” detection. [21]

Conformal prediction provides a complementary approach: it produces prediction sets with validity guarantees under exchangeability assumptions, explicitly separating (1) guarantees of being right at a chosen error rate from (2) efficiency goals. This is directly relevant to safe system behavior when the system must know when not to commit. [22]

Moreover, selective classification (“reject option”) formalizes the risk–coverage tradeoff: a system can abstain on uncertain cases to reduce risk. [23]

For explainability and interface obligations, a standards-oriented perspective proposes principles for explainable AI systems and explicitly includes “knowledge limits” as a principle: explainable systems should know and communicate when they are operating outside their reliable regime. [24]

Finally, large language model hallucination research (including an ACM-published survey) underscores that fluent generation can diverge from factual grounding and faithfulness, making explicit uncertainty handling and grounding mechanisms essential for reliable outputs in open-ended tasks. [25]

A unified “Ignorance Handling Contract” for humans and AI

The shared requirement is not perfect accuracy; it is **safe behavior under uncertainty**. The following contract specifies how a system should behave when it detects (or suspects) insufficient understanding.

```
IGNORANCE HANDLING CONTRACT (IHC)

Signal:
- “UNCERTAIN” when confidence or calibration metrics fall below threshold.
- “UNKNOWN” when input appears OOD, requirements are underspecified, or
constraints conflict.

Action:
- Ask for the minimum clarifying information needed to proceed safely; or
- Abstain and escalate; or
- Proceed only with a clearly labeled “best-effort” mode plus verification
hooks.

Evidence:
- Attach the basis for uncertainty: missing parameters, violated assumptions, or
low-support evidence.

Verification:
- Define what would confirm or falsify the output, and how to check it cheaply.
```

This contract operationalizes metacognitive accuracy goals for humans and uncertainty/abstention goals for AI, aligning with explainability principles that include communicating knowledge limits. [26]

The Cinch-and-Weave Framework: replacing linear learning with scalable coherence

Why linear modules are insufficient

Linear learning modules and case studies are valuable, but they often assume that understanding grows in a straight line. Complex systems instead require **iterative integration across overlapping contexts**, because high element interactivity creates comprehension that depends on relationships, not isolated facts. [1]

Research on effective learning supports the value of non-linear sequencing:

- **Spacing** improves retention across broad literatures, with meta-analytic evidence showing that timing variables (spacing intervals and retention interval) jointly determine optimal retention. [19]
- **Interleaving** can improve discrimination and long-term performance compared with blocked practice, even when blocked practice looks better during training (a critical warning for perceived progress). [27]

These findings naturally motivate a "knitting" metaphor: revisit key threads at increasing depth and in varied contexts, tightening understanding at "cinch points," then expanding from them.

What a "cinch" is: anchor concepts, invariants, and interfaces

A **cinch** is a stable anchor that can support expansion without collapsing into ambiguity. Cinches come in three practical types:

1. **Conceptual cinches**: definitions, causal mechanisms, or invariants that remain stable across contexts.
2. **Interface cinches**: explicit boundaries, contracts, and hidden design decisions (information hiding) that stabilize collaboration and reuse. [28]
3. **Measurement cinches**: standardized metrics and tests that define success and failure, enabling auditing and scaling.

Information hiding research argues that modules should be characterized by design decisions they hide from others, with interfaces chosen to reveal as little as possible about inner workings—stabilizing the system against change and supporting independent development. [28] This principle generalizes beyond software: in physical and ecological modules, interfaces and protected assumptions equally control how reuse scales.

The Cinch-and-Weave cycle

```
CINCH-AND-WEAVE CYCLE

(1) CINCH: define anchors (invariants, interfaces, measures)
      |
      v
(2) WEAVE: expand via interleaving across contexts + analogy mapping
      |
      v
(3) TIGHTEN: run verification hooks (tests, prototypes, peer review)
      |
      v
(4) AUDIT: check coherence to origin intent + detect gaps/contradictions
      |
      v
(5) VERSION: record changes, assumptions, and what is now known/unknown
      |
```

```
      v
Repeat at deeper scale (nested cycles)
```

This cycle is consistent with learning research on interleaving and spacing (weave), with cognitive load constraints (cinches reduce element interactivity by chunking), and with systems thinking that emphasizes iterative adaptation in complex systems. [29]

Operationalizing “audit pauses” for clarity and cohesion to origin

Audits are not bureaucratic pauses; they are **complexity management events**. Two families of evidence support this:

- In complex adaptive systems, nonlinear dynamics and multiple potential outcomes imply that drift from intent can accumulate invisibly until thresholds are crossed. [30]
- In human cognition, overconfidence in explanation can persist until forced to produce mechanistic detail. [7]

Thus, audit pauses should be triggered by **signals**, not schedules alone:

- Increasing contradiction counts between modules or stakeholder viewpoints
- Rising uncertainty rates (“UNKNOWN/UNCERTAIN” signals) in outputs
- Interface churn: frequent changes at boundaries
- Safety or sustainability metrics degrading relative to baseline

An architecture description standard emphasizes that architecture descriptions should capture stakeholder concerns and required content structured by viewpoints, which supports a formal mechanism for auditing “cohesion to origin” by checking whether concerns remain satisfied as the system evolves. [31]

Designing scalable modular systems: from 3D-printed concrete to AI interfaces to quantum modules

Early-stage components that enable constant scalable standardization

“Scaling constantly” without coherence collapse requires that standardization be designed from the start. Across domains, five core components recur:

Interfaces with explicit assumptions
In product architecture research, product architecture can be precisely defined to include (1) arrangement of functional elements, (2) mapping from functional elements to physical components, and (3) specification of interfaces among interacting components. [32] This definition generalizes to software modules and ecological modules: scaling occurs when mappings and interfaces are explicit and stable.

Protected decisions (information hiding) and change insulation
Modules should hide volatile decisions; interfaces should minimize leaked details to avoid unnecessary restriction and enable alternative implementations. [28]

Architecture descriptions and viewpoint management
Formal architecture description standards define required contents and emphasize stakeholder concerns, enabling systematic handling of "individual perspectives" as first-class design objects. [31]

Verification hooks and explanation channels
Explainable systems should provide explanations, ensure explanations are meaningful, ensure explanation accuracy, and communicate knowledge limits—positioning "explainability" as an interface requirement rather than a post-hoc feature. [24]

Life-cycle governance and versioning discipline
Systems life-cycle standards explicitly cover full life cycles from conception through retirement and emphasize iterative application of processes across elements and stages. [33] Semantic versioning provides a widely adopted scheme that makes compatibility information legible through version number changes, supporting scalable collaboration if consistently applied. [34]

Physical domain: extrusion-based 3D concrete printing as modular manufacturing

Extrusion-based 3D concrete printing (3DCP) is frequently framed as enabling formwork-free construction and new geometric freedom, while also presenting unresolved challenges in materials, process reliability, reinforcement, and standardization. Reviews describe current progress in workability, mechanical properties, plan design, and applications, emphasizing that properties depend on material selection and printing parameters. [35]

Sustainable construction-focused reviews characterize 3D printed concrete in categories (functional, sustainable, special concretes) and highlight that printing can introduce anisotropy and defects, making optimal structural design and material tailoring essential. [36]

A key modular-risk issue is **interlayer bond strength**: because printed structures are layered, interface quality becomes a structural interface. Review work on interlayer bond strength synthesizes the state of the art and calls attention to factors influencing bonding and performance. [37]

Sustainability claims must be grounded in life-cycle assessment (LCA). Systematic reviews of LCAs in 3D concrete printing examine ecological impacts across studies, reinforcing that outcomes depend on system boundaries and assumptions. [38]

Implication for modular standardization: 3DCP scalability depends less on "printing faster" and more on standardizing **interfaces**: nozzle and extrusion system specifications, material rheology requirements, interlayer time-gap constraints, reinforcement strategies, and inspection/verification procedures. [39]

Digital domain: explainable interfaces and coding architecture as modular cognition

In AI, explainability is not merely for trust; it is a mechanism for *auditing coherence*. Research on explanation in AI argues that explainable AI should draw from how humans generate and evaluate explanations, rather than relying only on designers' intuition of what counts as "good." [40] A complementary position paper emphasizes that "interpretability" lacks consensus and should be evaluated rigorously relative to context and purpose. [41]

A standards-based framework proposes explainability principles, including “knowledge limits,” which directly supports safe interface design: systems should know and communicate boundaries where they are unreliable. [24]

For code design and implementation, modular decomposition research emphasizes that modularization improves comprehensibility and flexibility when modules hide design decisions; it also warns that revealing unnecessary interface details restricts evolution and constitutes a design error. [28] Formal architecture description standards further support traceable designs by requiring explicit architecture descriptions structured around stakeholder concerns and viewpoints. [31]

Quantum domain: modularity as a scaling strategy under coherence constraints

Quantum systems introduce constraints that intensify the need for modular architectures: coherence is fragile, connectivity is expensive, and error correction imposes overhead. A widely cited vision of quantum networking describes quantum internet architectures as networks of nodes and channels requiring new capabilities for generating and characterizing quantum coherence and entanglement, including quantum interconnects connecting different physical systems. [42]

Modular quantum computing proposals analyze architectures with separate registers (modules) connected by photonic interconnects, aiming to scale to large numbers of qubits while confining noise and enabling fault tolerance. [43] A 2025 survey of distributed quantum computing further frames the field as integrating communication protocols and entanglement-based methods to address limitations of monolithic architectures, reinforcing the modularity principle as a cross-domain scaling strategy. [44]

Implication for this paper’s thesis: modularity is not domain-specific; it is a general response to bounded rationality and constrained reliability. Whether the constraint is human cognition, concrete rheology, software complexity, ecological dynamics, or quantum decoherence, scaling requires explicit interfaces, robust verification, and disciplined evolution.

Case study: designing modular architecture for biodiversity in multi-species systems

Case study goal and design premise

The case study designs a *reusable modular building system* that embeds biodiversity as a first-class requirement, not an afterthought, while remaining compatible with scalable manufacturing (including extrusion-based concrete printing) and auditable ecological outcomes.

Multi-species design research argues that biodiversity should become a key driving force in architectural design and proposes a “multi-species design paradigm” that explicitly includes non-human needs, introducing the concept of an “ecolope”: a multi-criteria designed building envelope that modulates environmental conditions and creates/connects habitats. [45] Complementary building-scale multispecies design literature maps concepts and gaps, emphasizing the need to integrate ecological knowledge into architectural design processes. [46]

Modular system architecture: habitat as an interface specification

The system is specified as a family of interoperable modules (physical components) with defined ecological functions (functional elements) and explicit interfaces—matching product architecture definitions that emphasize functional arrangement, mapping to components, and interface specification. [32]

```
BIODIVERSITY ARCHITECTURE MODULE FAMILY

A) Roofscape Modules
   - Extensive green roof habitat trays
   - Water retention + drought refugia pockets
   - Invertebrate substrate diversity zones

B) Facade Ecolope Modules
   - Nest/roost cavities (birds/bats) with thermally stable chambers
   - Micro-planter ledges for pollinator-supporting flora
   - Shaded moisture microhabitats for invertebrates

C) Ground Interface Modules
   - “Permeability corridors” supporting small fauna movement
   - Native planting beds + microtopography
   - Lighting and glare control elements to reduce ecological traps

D) Monitoring & Maintenance Modules
   - Inspection ports and replaceable inserts
   - Sensor-ready bays (temperature, humidity, occupancy indicators)
   - Standardized maintenance access points
```

This modularization is governed by information-hiding rules: each module hides internal construction details and exposes only interfaces needed for integration and maintenance, improving evolvability and independent development. [28]

Evidence-driven ecological design choices

Green roofs as habitat and ecosystem service infrastructure
Green roofs are described as urban ecosystems that can provide stormwater management, building temperature regulation, reduced heat island effects, and increased wildlife habitat; ecological function depends on interactions among substrate, biota, and vegetation. [47] Empirical case study evidence also shows substantial increases in species richness and diversity on green versus conventional roofs (including multiple taxa), supporting the plausibility of roof modules as biodiversity infrastructure when designed intentionally. [48]

Bird-safe design as a biodiversity requirement, not a decorative add-on
Bird-window collisions can be mitigated through window treatments and bird-friendly glazing and by avoiding ecological traps created by landscape design and reflective surfaces; mitigation efficacy has been demonstrated in field studies. [49] A 2025 critical review emphasizes that uniform window patterns following

commonly cited spacing rules can be effective deterrents, reflecting convergence toward design heuristics that can be standardized as interface constraints. [50] Government guidance documents also provide straightforward recommendations for reducing bird collisions with buildings, supporting translation into design checklists. [51]

Multi-species envelope ("ecolope") and computational workflow integration
The ecolope concept explicitly frames the building envelope as a dynamic living space for animals, plants, and microbiota, and argues that ecological modeling and computational design workflows should become intrinsic parts of design platforms that interface with users and key performance indicator optimization. [45] This supports treating ecological performance as a modular testable specification rather than an aspirational narrative.

Implementation using extrusion-based 3D concrete printing

The case study assumes cavity-rich, geometry-flexible modules can be produced using extrusion-based concrete printing. Reviews indicate that 3D printed concrete can be deposited layer-by-layer without formwork, with properties influenced by materials and printing parameters, and that the technology has potential but requires careful attention to structural performance and design optimization. [52]

However, layered fabrication makes **interlayer interfaces** a structural and durability risk; interlayer bond strength is a known focus of research, implying that module interfaces and printing process constraints must be explicit in the manufacturing specification. [37] Sustainability evaluation should be supported with LCA-based approaches because impacts vary by assumptions and system boundaries. [38]

Auditing biodiversity outcomes and preventing "paper biodiversity"

Scaling biodiversity architecture requires standardized measurement and enforcement mechanisms. A statutory biodiversity metric tool and guidance (developed for biodiversity net gain evaluation) exemplifies how biodiversity value can be calculated and standardized across projects. [53] Even if regulatory contexts differ globally, the methodological principle holds: biodiversity design must be coupled to auditable metrics and verification steps, otherwise promised features can degrade into unverified claims.

Case study control loop: biodiversity as a life-cycle process

```
BIODIVERSITY LIFE-CYCLE CONTROL LOOP

Design Intent:
- Define target taxa guilds + habitat functions + constraints (human safety,
maintenance).

Build:
- Manufacture modules with documented interfaces + installation procedures.

Verify:
- Post-install inspection + collision-risk review + habitat feature validation.
```

```
Monitor:
- Seasonal biodiversity surveys + sensor indicators + maintenance log.

Adapt:
- Replace inserts, adjust plantings, update glazing treatments, revise lighting
policies.

Version:
- Record changes as module family versions; persist assumptions and metrics.

Retire:
- Decommission modules responsibly; preserve habitat continuity where feasible.
```

This loop aligns with life-cycle standards that cover conception through retirement and stress iterative application of processes across system elements. [33]

Governance: project endings, versioning, and sustainable scaling

When to define “end of project” in evolving modular systems

For complex adaptive systems, “end” is often a governance decision, not a natural conclusion. Resilience and adaptive cycle perspectives emphasize that systems move through phases of growth, conservation, reorganization, and renewal, implying that “end” may mean transition to a maintenance and adaptation regime rather than cessation. [9]

Systems life-cycle standards explicitly include retirement/disposal processes as part of the full system life cycle, supporting a disciplined approach to defining project completion and transition criteria. [54]

A pragmatic definition of project end for modular systems includes meeting three conditions:

1. **Interface stabilization**: core interfaces are unlikely to change except through controlled version increments. [55]
2. **Verification sufficiency**: key claims are supported by tests, audits, or measurement cinches, reducing reliance on narrative. [56]
3. **Operational handoff readiness**: monitoring, maintenance, and escalation pathways are defined (human operators and AI agents can identify “unknown/unsafe” states and respond appropriately). [57]

Versioning as epistemic discipline

Versioning is often treated as a software administrative detail; in complex-system design it is an epistemic artifact: it records what changed, what assumptions were violated, and what new evidence supports the current state.

Semantic versioning formalizes how version number changes communicate compatibility and meaning about modifications, enabling distributed coordination when used consistently. [34] For modular multi-

domain systems, a “semantic version” should attach to *module families* and to *knowledge models* (the analogies and explanations used), not only to code repositories.

A recommended practice is **Dual Versioning**:

- **Implementation version**: concrete artifacts (code, CAD, mix designs, manufacturing routines).
- **Understanding version**: the current best explanatory model, including metaphors/analogs in use, their falsification hooks, and known limits.

This practice directly counters illusion-of-understanding risks by forcing explicit tracking of model adequacy over time. [58]

Scaling with sustainability: avoiding brittle global standardization

Global standardization supports collaboration, but brittle standardization can fail when local ecological and sociotechnical contexts differ. Complex adaptive systems research highlights path dependence and multiple possible outcomes, implying that modules should support local adaptation within standardized interfaces. [8]

Multi-species building design research similarly emphasizes integrating ecological knowledge into design processes and treating non-human stakeholders as design constraints, implying that “one global biodiversity module” is less realistic than a standardized *module interface schema* with locally adapted content. [59]

Thus, the recommended standardization approach is:

- Standardize **interfaces, metrics, and audit procedures**.
- Localize **species targets, planting palettes, microhabitats, and maintenance regimes**.

This is the same strategy used in scalable modular computing and in modular product architecture: stable boundary rules enable innovation inside modules. [60]

Works Cited

(All entries are selected for peer-reviewed or standards-body provenance where feasible; citations above link to source records and full texts where accessible.)

Adam L. Alter [61], et al. “Missing the Trees for the Forest: A Construal Level Account of the Illusion of Explanatory Depth.” *Psychological Science*, 2010. PDF. [62]

Robert A. Bjork [63] and Elizabeth L. Bjork [64]. “Creating Desirable Difficulties to Enhance Learning.” In *Psychology and the Real World: Essays Illustrating Fundamental Contributions to Society*, 2011. PDF. [65]

B. B. Brown [66], et al. “Bird-window collisions: Mitigation efficacy and risk factors across two years of monitoring.” *Ecology and Evolution*, 2021. [49]

U.S. Fish and Wildlife Service [67] . *Reducing Bird Collisions with Buildings and Building Glass Best Practices.* PDF. [51]

Nicholas J. Cepeda [68] , et al. “Distributed Practice in Verbal Recall Tasks: A Review and Quantitative Synthesis.” *Psychological Bulletin*, vol. 132, no. 3, 2006, pp. 354–380. PubMed record. [19]

John Sweller [69] . “Cognitive load theory, learning difficulty, and instructional design.” *Learning and Instruction*, vol. 4, no. 4, 1994, pp. 295–312. [5]

Dedre Gentner [70] . “Structure-mapping: A theoretical framework for analogy.” *Cognitive Science*, 1983. [10]

Keith J. Holyoak [71] . “Analogical mapping by constraint satisfaction.” *Cognitive Science*, 1989. [11]

John Kounios [72] and Mark Beeman [73] . “The Cognitive Neuroscience of Insight.” PubMed record, 2014. [14]

Leonid Rozenblit [74] and Frank Keil [75] . “The misunderstood limits of folk science: an illusion of explanatory depth.” *Cognitive Science*, 2002. Full text (PMC). [7]

Stephen A. Levin [76] . “Ecosystems and the Biosphere as Complex Adaptive Systems.” *Ecosystems*, 1998. PDF. [77]

C. S. Holling [78] . “Resilience and Stability of Ecological Systems.” *Annual Review of Ecology and Systematics*, 1973. PDF. [79]

E. Oberndorfer [80] , et al. “Green Roofs as Urban Ecosystems: Ecological Structures, Functions, and Services.” *BioScience*, vol. 57, no. 10, 2007. [81]

Wolfgang W. Weisser [82] , et al. “Creating ecologically sound buildings by integrating ecology, architecture and computational design.” *People and Nature*, 2022/2023. [45]

Yasha J. Grobman [83] , et al. “Architectural Multispecies Building Design: Concepts, Challenges, and Design Process.” *Sustainability*, 2023. [46]

EIF Wooster [84] , et al. “Urban green roofs promote metropolitan biodiversity: A comparative case study.” *Urban Forestry & Urban Greening*, 2022. [48]

Jingchuan Zhang [85] , et al. “A review of the current progress and application of 3D printed concrete.” *Composites Part A*, 2019. [35]

X. Wang [86] , et al. “Concrete 3D printing technology for sustainable construction: A review on raw material, concrete type and performance.” *Case Studies in Construction Materials*, 2024. [36]

A. Motalebi [87] , et al. “A systematic review of life cycle assessments of 3D concrete printing.” *Cleaner Engineering and Technology*, 2024. [88]

D. L. Parnas [89]. “On the criteria to be used in decomposing systems into modules.” 1972. PDF. [28]

Karl Ulrich [90]. “The role of product architecture in the manufacturing firm.” *Research Policy*, vol. 24, 1995, pp. 419–440. Screenshots. [91]

Tim Miller [92]. “Explanation in artificial intelligence: Insights from the social sciences.” *Artificial Intelligence*, 2019. [40]

Finale Doshi-Velez [93] and Been Kim [94]. “Towards A Rigorous Science of Interpretable Machine Learning.” *arXiv*, 2017. [41]

National Institute of Standards and Technology [95]. Four Principles of Explainable Artificial Intelligence [96]. 2021. PDF. [24]

Glenn Shafer [97], et al. “A Tutorial on Conformal Prediction.” *Journal of Machine Learning Research*, 2008. PDF. [22]

Yair Geifman [98] and Ran El-Yaniv [99]. “Selective Classification for Deep Neural Networks.” NeurIPS, 2017. PDF. [100]

Dan Hendrycks [101] and Kevin Gimpel [102]. “A Baseline for Detecting Misclassified and Out-of-Distribution Examples in Neural Networks.” ICLR, 2017. PDF. [21]

H. J. Kimble [103]. “The Quantum Internet.” *Nature*, 2008. PDF. [104]

Christopher Monroe [105], et al. “Large Scale Modular Quantum Computer Architecture with Atomic Memory and Photonic Interconnects.” *arXiv*, 2012. [43]

International Organization for Standardization [106]. *ISO/IEC/IEEE 42010: Architecture description.* ISO record. [107]

———. *ISO/IEC/IEEE 15288: Systems and software engineering—System life cycle processes.* ISO record. [108]

———. *ISO/IEC/IEEE 12207: Software life cycle processes.* ISO record. [109]

semver.org [110]. *Semantic Versioning 2.0.0.* [34]

UK Government [111]. *Statutory biodiversity metric tools and guides.* [112]

---

[1] [5] [18] [97] Cognitive load theory, learning difficulty, and instructional design - ScienceDirect
https://www.sciencedirect.com/science/article/pii/0959475294900035

[2] [14] [15] The cognitive neuroscience of insight - PubMed
https://pubmed.ncbi.nlm.nih.gov/24405359/

3 10 16 95 Structure-mapping: A theoretical framework for analogy - ScienceDirect
https://www.sciencedirect.com/science/article/pii/S0364021383800093

4 45 59 93 (PDF) Creating ecologically sound buildings by integrating ecology, architecture and computational design
https://www.researchgate.net/publication/366095485_Creating_ecologically_sound_buildings_by_integrating_ecology_architecture_and_computational_design

6 Bounded Rationality - Stanford Encyclopedia of Philosophy
https://plato.stanford.edu/entries/bounded-rationality/?utm_source=chatgpt.com

7 58 64 68 86 The misunderstood limits of folk science: an illusion of explanatory depth - PMC
https://pmc.ncbi.nlm.nih.gov/articles/PMC3062901/

8 30 https://link.springer.com/article/10.1007/s100219900037
https://link.springer.com/article/10.1007/s100219900037

9 79 106 https://pure.iiasa.ac.at/id/eprint/26/1/RP-73-003.pdf
https://pure.iiasa.ac.at/id/eprint/26/1/RP-73-003.pdf

11 https://www.sciencedirect.com/science/article/pii/0364021389900165
https://www.sciencedirect.com/science/article/pii/0364021389900165

12 https://press.uchicago.edu/ucp/books/book/chicago/M/bo3637992.html
https://press.uchicago.edu/ucp/books/book/chicago/M/bo3637992.html

13 75 https://www.tandfonline.com/doi/abs/10.1207/s15327868ms1001_4
https://www.tandfonline.com/doi/abs/10.1207/s15327868ms1001_4

17 26 72 annualreviews.org
https://www.annualreviews.org/doi/pdf/10.1146/annurev-psych-022423-032425

19 101 Distributed practice in verbal recall tasks: A review and quantitative synthesis - PubMed
https://pubmed.ncbi.nlm.nih.gov/16719566/

20 98 On Calibration of Modern Neural Networks
https://arxiv.org/abs/1706.04599?utm_source=chatgpt.com

21 https://arxiv.org/pdf/1610.02136
https://arxiv.org/pdf/1610.02136

22 78 111 jmlr.csail.mit.edu
https://jmlr.csail.mit.edu/papers/volume9/shafer08a/shafer08a.pdf

23 74 90 100 https://papers.neurips.cc/paper/7073-selective-classification-for-deep-neural-networks.pdf
https://papers.neurips.cc/paper/7073-selective-classification-for-deep-neural-networks.pdf

24 56 57 Four Principles of Explainable Artificial Intelligence
https://nvlpubs.nist.gov/nistpubs/ir/2021/nist.ir.8312.pdf

25 https://dl.acm.org/doi/10.1145/3703155
https://dl.acm.org/doi/10.1145/3703155

27 29 65 73 83 CH05.qxp:FABBS_DESIGN_NE
https://bjorklab.psych.ucla.edu/wp-content/uploads/sites/13/2016/04/EBjork_RBjork_2011.pdf

28 55 60 66 89 On the criteria to be used in decomposing systems into modules
https://prl.khoury.northeastern.edu/img/p-tr-1971.pdf

31 102 103 107 https://www.iso.org/standard/50508.html
https://www.iso.org/standard/50508.html

32 91 105 ktulrich.com
https://ktulrich.com/wp-content/uploads/2022/09/ulrich-architecture-rp.pdf

33 54 108 https://www.iso.org/standard/81702.html
https://www.iso.org/standard/81702.html

34 63 https://semver.org/
https://semver.org/

35 52 82 99 A review of the current progress and application of 3D printed concrete - ScienceDirect
https://www.sciencedirect.com/science/article/abs/pii/S1359835X19302829

36 61 Concrete 3D printing technology for sustainable construction: A review on raw material, concrete type and performance - ScienceDirect
https://www.sciencedirect.com/science/article/pii/S2666165924000590

37 71 76 https://www.mdpi.com/2071-1050/13/13/7137
https://www.mdpi.com/2071-1050/13/13/7137

38 69 88 https://www.sciencedirect.com/science/article/pii/S2666412723000132
https://www.sciencedirect.com/science/article/pii/S2666412723000132

39 https://pmc.ncbi.nlm.nih.gov/articles/PMC10095855/
https://pmc.ncbi.nlm.nih.gov/articles/PMC10095855/

40 67 70 94 96 https://www.sciencedirect.com/science/article/pii/S0004370218305988
https://www.sciencedirect.com/science/article/pii/S0004370218305988

41 110 https://arxiv.org/abs/1702.08608
https://arxiv.org/abs/1702.08608

42 104 https://ixnovi.people.wm.edu/papers/kimble08.pdf
https://ixnovi.people.wm.edu/papers/kimble08.pdf

43 87 [1208.0391] Large Scale Modular Quantum Computer Architecture with Atomic Memory and Photonic Interconnects
https://arxiv.org/abs/1208.0391

44 https://www.sciencedirect.com/science/article/pii/S1574013725000231
https://www.sciencedirect.com/science/article/pii/S1574013725000231

46 Architectural Multispecies Building Design: Concepts, Challenges, and Design Process
https://www.mdpi.com/2071-1050/15/21/15480

47 81 85 academic.oup.com
https://academic.oup.com/bioscience/article/57/10/823/232363

48 80 Urban green roofs promote metropolitan biodiversity: A comparative case study - ScienceDirect
https://www.sciencedirect.com/science/article/abs/pii/S0360132321008544

49 92 https://pmc.ncbi.nlm.nih.gov/articles/PMC8325430/
https://pmc.ncbi.nlm.nih.gov/articles/PMC8325430/

50 https://www.tandfonline.com/doi/full/10.1080/15594491.2025.2542043
https://www.tandfonline.com/doi/full/10.1080/15594491.2025.2542043

51 https://www.fws.gov/sites/default/files/documents/reducing-bird-collisions-with-buildings.pdf
https://www.fws.gov/sites/default/files/documents/reducing-bird-collisions-with-buildings.pdf

53 84 112 https://www.gov.uk/government/publications/statutory-biodiversity-metric-tools-and-guides
https://www.gov.uk/government/publications/statutory-biodiversity-metric-tools-and-guides

62 https://pages.stern.nyu.edu/~aalter/jpspioed.pdf
https://pages.stern.nyu.edu/~aalter/jpspioed.pdf

77 https://www1.cs.columbia.edu/~traub/sloan/LevinEcos.pdf
https://www1.cs.columbia.edu/~traub/sloan/LevinEcos.pdf

109 https://www.iso.org/standard/63712.html
https://www.iso.org/standard/63712.html