

ARCHITECTURAL HYPOTHESIS

The Cognitive Machinery Limit Hypothesis

*A proposed principle concerning the developmental boundaries of human
and synthetic minds*

***“The limit of a mind is the completeness of its reasoning machinery,
not the infinity of its knowledge.”***

Proposed by Charles Dooner

Expanded definitive formulation · Version 3.0 · September 2026

Abstract

The Cognitive Machinery Limit Hypothesis (CMLH) proposes that, for finite cognitive systems operating under specified and comparable opportunity conditions, long-run adaptive reach is constrained by the range, integration and revisability of the cognitive operations through which experience is converted into new understanding. Initial knowledge determines part of what a mind can do immediately. Cognitive machinery constrains which further knowledge states it can construct, test, correct and transfer. A system may therefore be ignorant yet developmentally capable, or extensively informed yet developmentally bounded.

The hypothesis contains a conceptual claim and an empirical claim. Conceptually, additional stored content cannot by itself substitute for an absent operation required to transform that content into a task-relevant representation, inference or revision. Empirically, after access to evidence, embodiment, time and computational resources are controlled, the functional coverage, causal integration and revisability of knowledge-generating operations should predict adaptation, transfer and recovery under novelty beyond what is predicted by initial knowledge quantity alone. The proposal neither promises omniscience nor treats machinery as the only limit. It identifies machinery as the principal internal constraint on the set of knowledge states reachable from experience.

Central proposition A lack of knowledge is recoverable when relevant evidence and the required cognitive operations are available. A lack of the required operation is a deeper internal boundary because additional information alone cannot perform the missing transformation.

1. Status of the Proposal

This document presents a hypothesis rather than an established scientific theory. Its central distinction overlaps with earlier work on cognitive architecture, learning, bounded rationality, constructivism and general intelligence, but the specific claim advanced here is that the boundary of a mind's developmental reach lies principally in its knowledge-generating machinery rather than in the finite contents of its knowledge at a given time.

The hypothesis becomes a theory only if its terms are operationalised, its predictions are tested across competing systems, and evidence consistently favours it over rival explanations. The present document supplies the conceptual structure required for that process.

1.1 Epistemic discipline

No empirical hypothesis can be made immune to criticism or disproof without becoming scientifically empty. The defensibility of the CMLH therefore rests on the opposite standard: narrow claims, explicit controls, operational definitions, risky predictions and declared failure conditions. Its strongest form is not an absolute claim that machinery explains everything. It is a controlled claim about what additional knowledge can and cannot compensate for, and about which internal properties should predict continued adaptation.

1.2 Scope

The CMLH applies to finite systems that maintain internal states and can, at least in principle, acquire or use information over time. It concerns developmental cognition: changes in what a system can represent, infer, revise or transfer. It does not attempt to define consciousness, subjective experience, moral status, wisdom, creativity in every sense, or intelligence as a whole. Those phenomena may depend on additional properties.

1.3 Explicit non-claims

- The hypothesis does not claim that knowledge is unimportant or that an ignorant system will necessarily outperform a knowledgeable one.
- It does not claim that reasoning can recover unavailable empirical facts or overcome inadequate sensing.
- It does not claim that one architecture is complete for every possible environment.
- It does not claim that cognitive mechanisms are independent of learned knowledge; learning may alter their effective operation.
- It does not claim that every limitation can be removed by architectural expansion.
- It does not claim that a mind with broader machinery will realise its potential without motivation, opportunity and resources.
- It does not claim that NEL currently confirms the hypothesis.

2. Formal Statement

Primary scientific formulation For finite cognitive systems with comparable access to evidence, embodiment and computational resources, long-run adaptive reach is constrained by the range, causal integration and revisability of the cognitive operations available to acquire, represent, relate, test, correct and transfer knowledge. Increasing stored knowledge alone cannot compensate for the absence of an operation required to transform that knowledge into the relevant form of understanding.

For a mind M operating over interval t , let:

- K_0 denote its initial usable knowledge state;
- O denote the set of cognitive operations it can execute;
- I denote the causal integration structure governing which operations can use the outputs of others;
- V denote revisability: the conditions under which representations, beliefs, strategies or operations can change;
- E_t denote the evidence and interactions made available during t ;
- R_t denote finite resources, including memory, computation, energy and time;
- B denote embodiment and interface constraints;
- $\text{Reach}(M \mid K_0, O, I, V, E_t, R_t, B)$ denote the set of cognitive states and competencies M can attain under those conditions.

The CMLH makes two separable claims. First, Reach is operation-constrained: if attaining state s requires transformation τ , and neither τ nor a functionally equivalent composition is executable within O , I and V , then s is not reachable merely by enlarging K_0 . Second, in controlled comparisons under novelty, measures of O , I and V should explain variance in long-run adaptation beyond measures of K_0 . The first is the non-compensation claim. The second is the empirical architectural-reach claim.

2.1 Why the claim is not circular

The hypothesis does not define a required operation as “whatever produces success.” Operations must be specified independently of the outcome—for example, belief revision under contradictory evidence, relational abstraction across changed surface labels, or active selection of discriminating evidence. A task class must then be designed whose solution requires the independently specified transformation. Only after those commitments are fixed may performance be measured.

2.2 Why the claim is not a tautology

It is trivially true that a system cannot execute an operation it cannot execute. The substantive hypothesis goes further. It predicts that identifiable differences in functional coverage, integration and revisability create characteristic, repeatable boundaries in learning; that additional stored content does not erase those boundaries; that restoring the missing transformation changes the class of reachable outcomes; and that these architectural measures add predictive power under controlled novelty. Those are contingent empirical claims and may be false.

3. Definitions

Term	Operational meaning
Mind	An organised system capable of maintaining internal states and using them to guide perception, learning, inference, valuation or action. The term is substrate-neutral and may refer to biological or synthetic systems.
Knowledge	Stored or currently available representations, relationships, models, skills or justified expectations that the mind can use. Knowledge is always finite in an implemented mind.
Cognitive machinery	The implemented operations and causal pathways that transform input, stored states and prior outputs into usable and revisable representations, inferences, predictions, decisions or new operations. This is broader than deduction.
Reasoning machinery	The original philosophical term. In the scientific formulation it is replaced by knowledge-generating and knowledge-revising cognitive operations to prevent an artificially narrow interpretation.
Functional coverage	The range of independently specified cognitive transformations an architecture can perform, either directly or through compositions of existing operations.
Compositional closure	The transformations executable directly or by reliable compositions of available operations, within a declared resource budget and error tolerance.
Minimal sufficient coverage	A set-inclusion-minimal collection of required transformations that meets a preregistered performance criterion for a declared task class. More than one minimal set may exist.
Causal integration	The degree to which outputs of one operation are available at the right time and in a usable form to alter other operations and live behaviour. Mere co-location of modules is not integration.

Term	Operational meaning
Revisability	The architecture's capacity to alter representations, confidence, relationships, strategies or learned operations under specified evidence conditions while preserving appropriate stability and provenance.
Completeness	A philosophical shorthand retained in the original quotation. It does not mean absolute logical completeness. Scientifically, it means functional sufficiency relative to a declared task or environment class, and must be replaced by measured coverage, integration and revisability.
Developmental reach	The range of cognitive states, concepts, explanations, skills and models that a mind could attain through its own operations when given relevant evidence, interaction and finite resources.
Reachable knowledge	Knowledge that the mind can construct under declared evidence, embodiment, time and resource conditions—not merely information that exists somewhere in the world or could be inserted by an external system.
Knowledge quantity	A task-relative measure of usable stored content fixed before testing. Raw parameter count, file size or number of statements is not automatically a valid measure because content differs in relevance, compression and accessibility.
Required operation	A transformation specified independently of system success and shown through task analysis, intervention or ablation to be necessary for a defined task class, allowing for functionally equivalent implementations.
Functional equivalence	Different physical or computational mechanisms that produce the same task-relevant transformation under the tested conditions.
Machinery growth	A persistent expansion or reorganisation of functional closure, causal integration or revisability that makes a previously unavailable transformation executable.
Meta-machinery	The learning and construction processes by which a mind can acquire, compose, integrate or revise its own cognitive operations.

Term	Operational meaning
Machinery gap	A missing, isolated or defective mechanism that prevents available evidence from being transformed into a required kind of understanding.
Knowledge gap	A recognised absence of information or confidence that the mind could potentially reduce using machinery it already possesses.

3.1 Operationalising functional coverage

Let T_D be a preregistered set of transformations required by domain or task class D . Let $CI_R(O, I, V)$ be the transformations a system can execute directly or through reliable composition of its operations, integration structure and revision policies within resource budget R . Unweighted task-relative coverage is:

Task-relative functional coverage $FC_D = |T_D \text{ intersect } CI_R(O, I, V)| / |T_D|$

When transformations differ in importance, weights w_tau must be fixed independently of system outcomes:

Weighted functional coverage $WFC_D = [\text{sum over tau in } T_D \text{ of } w_tau \text{ multiplied by indicator}(\text{tau in } CI_R)] / [\text{sum over tau in } T_D \text{ of } w_tau]$

Minimal coverage is the proportion of individually required transformations available. Task-relative coverage is indexed to a declared D , R and reliability threshold. Compositional closure counts a transformation only when a direct operation or reproducible composition performs it within those limits. A minimally sufficient set S is any set-inclusion-minimal subset of T_D that meets a preregistered criterion P_D ; it is not necessarily unique. Where transformations are not commensurable, coverage must be reported as a profile rather than collapsed into one score.

Anti-circularity rule T_D , weights, success thresholds, acceptable compositions and resource ceilings must be declared before outcomes are observed. Output success alone cannot be used to invent the transformation list after testing.

3.2 Machinery growth and the knowledge-to-machinery transition

Knowledge growth changes usable content K while leaving the relevant operational closure approximately unchanged. Machinery growth occurs when learning persistently changes what

transformations can be executed, connected or revised. Its strongest operational signature is CI_R at time $t+1$ strictly containing a transformation absent from CI_R at time t , under matched resource and reliability conditions. Machinery can also grow by improving integration or revision even when the operation inventory is unchanged.

- Structural growth: a new operation, module or causal pathway is constructed.
- Compositional growth: existing operations become a new stable and reusable procedure.
- Integration growth: a formerly isolated capacity gains live causal access to memory, uncertainty, inference or action.
- Control growth: the system learns when and where to invoke an available transformation.
- Representational growth: a new internal format makes previously unavailable transformations executable.
- Revisability growth: the system acquires a new evidence-sensitive policy for changing representations, strategies or operations.

Machinery growth may begin as knowledge but is not identical to knowledge. A relation may be discovered and stored declaratively, compiled into a reusable procedure, integrated into live control and eventually alter future functional closure. It counts as machinery growth only when intervention shows that the acquired structure causally enables later cognition, transfers beyond the original examples and survives removal of mere answer retrieval. Thus machinery is not a fixed box opposed to learning; learning can change machinery when the architecture contains meta-machinery capable of making that conversion.

Growth boundary A mind's developmental reach is constrained both by its current machinery and by the meta-machinery through which it can construct, integrate and revise new machinery. Knowledge may seed that process, but information does not convert itself into an operation.

3.3 Functional equivalence classes

To remain substrate-neutral and avoid counting implementation duplicates as broader coverage, the unit of analysis is an equivalence class of transformations. For declared domain D , resources R , evidence conditions E and tolerance ϵ , operations a and b are functionally equivalent when substituting one for the other preserves the preregistered task-relevant transformation and its downstream causal role within ϵ . The relation is context-indexed, not universal.

- Direct equivalence: distinct mechanisms perform the same declared transformation under matched conditions.
- Equivalence under composition: one operation and a reliable composition of other operations perform the same transformation within resource, latency and error bounds.

- Equivalence under abstraction: mechanisms preserve the same relational invariants across changed labels, representations or surface forms.
- Causal-role qualification: matching final answers is insufficient when error profile, resource cost or downstream effects materially differ for the claim being tested.

Within fixed conditions and tolerance, equivalence classes should satisfy reflexivity, symmetry and transitivity. If approximate equivalence fails transitivity, researchers must use a similarity relation or narrower classes rather than claiming strict equivalence. Functional coverage is therefore measured over declared equivalence classes, while implementation details remain available for causal analysis.

4. The Core Distinction

4.1 Ignorance is not incapacity

A mind may truthfully represent “I do not know” while retaining the ability to seek evidence, construct alternatives and revise its model. Such ignorance is a temporary state, not necessarily a cognitive boundary. The capacity to recognise the gap is itself evidence of machinery operating beyond the missing fact.

4.2 Information is not understanding

A system may store or reproduce a vast quantity of information without possessing mechanisms that allow it to relate that information to experience, expose contradictions, generate novel explanations or transfer structure into unfamiliar situations. Its apparent competence can therefore exceed its developmental independence.

4.3 Current performance is not developmental potential

A knowledgeable but rigid mind may outperform an initially ignorant but adaptive mind in the short term. The hypothesis predicts a possible reversal over sufficiently rich experience: the adaptive mind continues to reorganise and extend its understanding, while the rigid mind approaches the boundary of its stored contents.

5. Proposed Cognitive Machinery

The hypothesis does not require one universal implementation. It predicts, however, that open-ended developmental reach depends on an integrated cycle containing at least the following functional capacities:

1. Perceive or otherwise receive differentiated evidence from an environment.

2. Preserve episodes while constructing higher-level internal representations.
3. Discover temporal, causal, spatial, categorical and relational structure.
4. Represent uncertainty, conflict, novelty and missing knowledge.
5. Generate alternative hypotheses or candidate explanations.
6. Predict consequences that can distinguish among those alternatives.
7. Seek or select informative evidence rather than passively accumulate input.
8. Evaluate outcomes and assign confidence without converting one success into an inflexible rule.
9. Revise beliefs, relationships and representations when evidence conflicts with them.
10. Generalise structure across different objects, situations or environments.
11. Form abstractions that reorganise how earlier and later evidence is perceived.
12. Monitor the mind's own failures and, where possible, change its strategies or cognitive organisation.

Integration requirement A list of mechanisms is insufficient. Developmental reach depends on whether the output of one process becomes causally available to the others. Unintegrated capabilities do not form complete cognitive machinery.

6. Causal Model

The proposed developmental cycle is recursive rather than linear:

Cognitive development loop Experience → representation → relationship discovery → uncertainty or knowledge gap → hypothesis → prediction → evidence seeking → outcome evaluation → belief revision → concept formation → generalisation → abstraction → altered future perception and reasoning

New knowledge changes subsequent cognition. A later concept may allow a mind to revisit unchanged sensory evidence and discover relationships it could not previously represent. Knowledge and machinery therefore interact recursively: machinery generates knowledge, while acquired knowledge changes the effective operation of that machinery. The hypothesis does not separate them absolutely; it identifies architecture as the deeper constraint on what this recursion can eventually produce.

7. Principal Claims

- Finite Knowledge Claim: No implemented mind requires infinite stored knowledge to exhibit broad developmental intelligence.

- Architectural Reach Claim: The kinds of understanding a mind can attain depend upon the transformations its cognitive machinery can perform.
- Mechanism Deficit Claim: Repeatedly supplying information will not reliably compensate for a missing mechanism required to interpret, test or reorganise that information.
- Revisability Claim: A mind whose beliefs cannot be corrected possesses a narrower developmental reach than one capable of principled revision, even if the former initially knows more.
- Integration Claim: Cognitive breadth depends upon causal integration among mechanisms, not merely their isolated presence.
- Developmental Reversal Claim: Under sufficiently varied experience, an initially less knowledgeable but more adaptively complete mind may surpass a more knowledgeable but structurally rigid mind.
- Reinterpretation Claim: Growth in representation and abstraction can enable a mind to derive new structure from previously encountered evidence.
- Substrate-Neutral Claim: The principle concerns functional cognitive organisation and therefore applies in principle to biological, synthetic and hybrid minds.
- Machinery-Growth Claim: Learning can enlarge developmental reach when it converts knowledge into persistent, causally integrated operations.
- Meta-Machinery Claim: A system's capacity to grow new machinery is itself constrained by the operations through which it constructs, validates and integrates that machinery.

7.1 Hierarchy of hypotheses

The CMLH is best tested as a family of increasingly demanding propositions. Failure of a stronger proposition need not invalidate the narrower logical core.

Label	Proposition	Evidential burden
H1: Non-compensation	Stored knowledge alone cannot perform an independently specified missing transformation.	Demonstrate a stable failure boundary despite added task-relevant content.
H2: Restoration	Adding or integrating the required transformation expands the relevant reachable task class.	Intervention must selectively remove the predicted boundary.
H3: Integration	Causally integrated operations produce greater adaptive reach than the same operations operating in isolation.	Hold component availability constant while varying live causal access.

Label	Proposition	Evidential burden
H4: Revisability	Systems capable of evidence-sensitive revision recover better from changed contingencies than otherwise comparable rigid systems.	Pre-register change points and recovery measures.
H5: Incremental prediction	Coverage, integration and revisability predict long-run adaptive reach beyond initial knowledge, evidence and resources.	Use controlled comparisons and model comparison on held-out novelty.
H6: Machinery growth	A learned change that expands functional closure or integration creates transfer beyond ordinary content acquisition.	Show pre/post expansion, cross-surface transfer and selective loss after intervention on the acquired operation.
H7: Growth boundary	The kinds of machinery a system can acquire are constrained by its meta-machinery.	Vary construction and integration capacities while matching knowledge and training opportunity.
H8: Equivalence invariance	Functionally equivalent implementations preserve the predicted reach boundary under matched conditions.	Use preregistered substitution tests across direct, composed and abstract implementations.

7.2 The irreducible core

Irreducible claim When a defined class of cognitive outcomes requires transformation τ , increasing stored content without adding τ or a functionally equivalent transformation will not make that outcome class reachable. If the system acquires τ through learning, the learning process is itself part of the machinery and must be included in the architectural account.

8. Predictions

If the hypothesis is substantially correct, the following outcomes should be observable:

- Across unfamiliar environments, integrated cognitive mechanisms will predict transfer and continued learning better than the volume of preloaded domain facts.

- A system with a specific machinery gap will repeatedly fail at tasks requiring that mechanism even after receiving additional relevant examples.
- Adding or integrating the missing mechanism will produce a qualitative expansion in what the system can learn, not merely a local performance improvement.
- Systems capable of uncertainty detection and active evidence seeking will correct false models more reliably than systems receiving comparable data passively.
- Systems with belief revision will recover from environmental change better than systems whose prior successes become fixed rules.
- A developing system will sometimes reinterpret stored episodes after forming new concepts, producing conclusions not available when those episodes were first experienced.
- Broader machinery will increase learning transfer between structurally related but superficially different environments.
- The long-run performance gap between adaptive and rigid minds will widen as environmental novelty increases.
- Systems that convert learned representations into integrated procedures will show stepwise expansion of task classes, not only smoother improvement within trained instances.
- Two implementations in the same preregistered functional equivalence class will exhibit comparable reach boundaries, unless relevant resource cost or causal role differs.
- A system with broader meta-machinery will acquire new transformations from the same curriculum that leaves a knowledge-matched system operationally unchanged.

9. Falsification Conditions

The hypothesis would be weakened or falsified in its strong form if controlled studies consistently showed any of the following:

- Initial knowledge quantity predicts long-run generalisation and transfer as well as or better than differences in cognitive machinery.
- Systems lacking a proposed essential mechanism nevertheless develop the relevant capacity through no alternative mechanism and without architectural change.
- Adding integrated hypothesis formation, evidence seeking, revision or abstraction produces no measurable expansion of reachable knowledge beyond what additional stored data alone produces.
- A highly knowledgeable but non-revising system remains more developmentally capable across open-ended novelty than a less knowledgeable system with integrated learning and self-correction.
- No operational distinction can be established between a knowledge gap and a machinery gap.
- The claimed machinery factors cannot predict which classes of knowledge or competence will remain unreachable.

- Purported machinery growth disappears when answer retrieval is blocked, does not transfer across surface variation, or cannot be selectively disrupted by intervening on the acquired procedure.
- Preregistered functionally equivalent implementations systematically yield incompatible reach predictions after relevant resource and causal-role differences are controlled.
- Meta-machinery measures fail to predict which new operations systems can construct under matched curricula and resources.

A weaker version may survive even if no single ordering of mechanisms proves universal. The essential claim would remain that developmental reach is constrained by available transformations, although different architectures might realise equivalent transformations through different means.

10. Experimental Programme

10.1 Knowledge-machinery crossover experiment

Construct two otherwise comparable minds. Give Mind A extensive task-relevant knowledge but restrict revision, abstraction or active inquiry. Give Mind B sparse initial knowledge but integrated learning, uncertainty, hypothesis testing and revision. Expose both to familiar and progressively novel environments. Measure initial performance, learning rate, transfer, error recovery, conceptual novelty and asymptotic competence. The predicted result is an early advantage for A followed by a developmental crossover favouring B.

10.2 Mechanism ablation experiment

Remove or disconnect one mechanism at a time from an integrated cognitive architecture. Preserve data, training exposure and resources. Test whether each ablation creates a characteristic class of unreachable understanding. Restore the mechanism and determine whether the boundary changes qualitatively.

10.3 Old-evidence reinterpretation experiment

Preserve original episodes separately from higher-level interpretations. After the mind forms a new concept, allow it to reconsider earlier episodes without adding new sensory evidence. Test whether it discovers relationships that were previously unrepresentable and whether those discoveries improve prediction in novel cases.

10.4 Cross-world structural transfer

Allow a mind to learn relationships in one environment and then move the same mind into a superficially different environment containing homologous structure. Success requires first-attempt or

accelerated inference based upon discovered relationships rather than repeated trial and error, label matching or hardcoded mappings.

10.5 Functional-coverage mapping

Before testing, define T_D through task analysis and validate each required transformation using independent intervention or ablation. Measure direct and composed execution under fixed R , then calculate FC_D or report a coverage profile. Test whether the measured gaps predict task-class failures better than knowledge quantity, compute and exposure. Repeat with independently designed held-out task families to prevent a benchmark-specific definition of coverage.

10.6 Machinery-growth discrimination

Test the same system before and after learning on novel instances and surface-shifted tasks. Block retrieval of trained answers, measure whether a formerly unavailable transformation is now executable, intervene on the acquired procedure or pathway, and test whether the new competence selectively disappears. Persistence, transfer, causal disruption and expansion of CI_R jointly distinguish machinery growth from memorisation or transient strategy use.

10.7 Functional-equivalence substitution

Implement a declared transformation directly, through a composition and through a representation-independent abstraction. Substitute the implementations while matching evidence and relevant resources. Compare success, reliability, latency, error profile and downstream use. Equivalent implementations should preserve the predicted reach boundary; systematic differences identify an omitted condition and require refinement of the equivalence class rather than post hoc redefinition.

11. Evaluation Measures

Measure	What it tests
Learning slope	Rate at which predictive or behavioural competence improves from experience.
Transfer efficiency	Reduction in evidence or attempts required in a structurally related new environment.
Revision recovery	Speed and correctness of recovery after prior relationships cease to hold.

Measure	What it tests
Knowledge-gap accuracy	Ability to distinguish uncertainty from genuine confidence and identify what evidence would resolve it.
Conceptual expansion	Formation of useful representations or categories not explicitly supplied.
Counterfactual reach	Ability to evaluate alternatives beyond directly experienced sequences.
Reinterpretation yield	New valid inferences extracted from stored episodes after cognitive development.
Architectural dependence	Degree to which removing a mechanism selectively eliminates a class of learning outcomes.
Task-relative coverage	Proportion or profile of preregistered transformation classes executable within declared limits.
Closure expansion	Previously unavailable transformations made reliably executable after learning.
Equivalence stability	Preservation of reach predictions when implementations in a declared equivalence class are substituted.
Meta-learning reach	Range of new operations autonomously constructed and integrated under matched curricula and resources.

11.1 Minimum decisive experiment

A strong first test should use a factorial design rather than a showcase. Construct multiple systems varying initial task-relevant knowledge (low/high) and one independently defined cognitive transformation (absent/present), while matching sensors, actions, compute budgets, feedback and exposure. Use at least three task families: one solvable through retrieval, one requiring the target transformation, and one requiring transfer after surface features change. Pre-register success criteria and evaluate held-out environments.

- H1 is supported if additional knowledge improves retrieval tasks but does not remove the predicted boundary on transformation-dependent tasks.
- H2 is supported if adding the target operation selectively removes that boundary without supplying the answers.

- H3 is supported if the operation affects live cognition only when causally integrated with memory, uncertainty and action selection.
- H5 is supported if architecture variables improve held-out prediction after knowledge, exposure and compute are included in the statistical model.
- H6 is supported only if the acquired capacity transfers beyond trained answers, alters measured closure and is causally disrupted by intervention on the learned operation.
- H7 is supported if preregistered differences in meta-machinery predict which new transformations arise from the same evidence and resources.
- H8 is supported if substitution among declared equivalents preserves the predicted boundary within tolerance.
- The strong form is weakened if knowledge enrichment eliminates the supposed machinery boundary under valid controls, or if the architectural intervention provides no selective benefit.

11.2 Required reporting

- Exact operational definition of every knowledge and machinery variable.
- Proof that the environment does not leak the answer through labels, ordering or privileged observations.
- Compute, memory, evidence and action budgets for every condition.
- Ablation and restoration results, not performance of the complete system alone.
- Individual trials, failure cases and confidence intervals, not only average success.
- Tests on unseen structures or changed contingencies, not merely new instances of memorised patterns.
- Alternative explanations considered before interpreting the result as architectural evidence.

12. Boundary Conditions and Limitations

The hypothesis is not a claim of omnipotence. Even exceptionally broad cognitive machinery remains constrained by conditions outside and inside the mind:

- Evidence: empirical facts cannot be reliably inferred when no discriminating evidence is available.
- Observability: sensors and interfaces determine which differences in the world can enter cognition.
- Embodiment: some concepts may require forms of interaction or feedback unavailable to a particular body.
- Computation and memory: finite resources restrict search, resolution, retention and the complexity of representable models.
- Time and energy: reachable knowledge in principle may remain unreachable within the lifespan or operating budget of the mind.
- Formal limits: some questions may be undecidable, underdetermined or computationally intractable.

- Values and motivation: a mind may possess the machinery to learn something yet never direct attention or inquiry toward it.
- Social and environmental access: tools, language, culture and other minds can radically alter the evidence and cognitive scaffolding available.

Accordingly, cognitive machinery is proposed as the primary internal limit on developmental reach, not the sole cause of realised intelligence.

13. Relationship to Existing Thought

The hypothesis shares territory with several traditions without being identical to them. Bounded rationality emphasises finite information, computation and time. Constructivist theories emphasise the active formation and restructuring of knowledge. Cognitive architectures describe interacting mechanisms underlying perception, memory, reasoning and action. Accounts of general intelligence evaluate competence across diverse environments. The present hypothesis joins these concerns around a specific distinction: accumulated content determines what a mind can use now, while integrated cognitive machinery determines what further understanding can become reachable.

The proposal is deliberately neutral on whether cognition must be symbolic, neural, dynamical, embodied or hybrid. Its predictions concern functional transformations and integration. Competing implementations may therefore instantiate the same machinery at different physical or representational levels.

14. Implications for Human Minds

For human cognition, the hypothesis implies that education aimed only at information accumulation addresses one layer of intelligence. Developmental reach also depends upon cultivating mechanisms and habits of inquiry: distinguishing observation from interpretation, recognising uncertainty, comparing explanations, testing predictions, revising beliefs, transferring relationships and reflecting upon one's own reasoning. It does not follow that human cognitive machinery is unlimited or freely redesignable. The claim is that a person's present ignorance should not automatically be mistaken for an incapacity to understand.

15. Implications for Synthetic Minds

For synthetic cognition, the hypothesis shifts the design target from building a system that already contains answers toward building a mind that can responsibly construct, test and revise answers through continuing experience. A database may enlarge current knowledge; it does not by itself enlarge

the architecture's developmental reach. Likewise, externally generated outputs do not demonstrate that the receiving system possesses the machinery that produced them.

A synthetic mind approaching the hypothesis would preserve raw or episodic evidence separately from interpretations; ground concepts through interaction; represent uncertainty and knowledge gaps; seek evidence for internally generated questions; form relationships and abstractions; revise models without erasing provenance; transfer structure across environments; and allow new concepts to change future perception and reasoning.

16. Application to NEL

The hypothesis emerged from the NeuralEdge Limitless (NEL) project and provides an architectural compass for it. Under this principle, NEL's objective is not to be filled with the world's knowledge. Its objective is to possess increasingly complete and integrated mechanisms through which encounters with the world can become experience, structure, uncertainty, inquiry, prediction, revision, concepts and generalisation.

NEL would support the hypothesis only through causal evidence. Individual mechanisms must not merely exist as isolated modules; they must participate in live cognition and demonstrably alter what NEL can learn or understand. Cross-world transfer, mechanism ablation, old-evidence reinterpretation and the distinction between knowledge gaps and machinery gaps are therefore central tests.

Architectural consequence When NEL fails, the first question should not always be “What knowledge must be added?” It should also be “What cognitive transformation is absent, disconnected or incapable of using the evidence already available?”

17. Competing Explanations

- Scale hypothesis: sufficiently large stores of examples and associations alone account for apparent generality.
- Knowledge-first hypothesis: breadth of prior knowledge is the dominant determinant of future competence because new reasoning is largely recombination of stored content.
- Environment-first hypothesis: differences in learning are primarily caused by the richness of experience and external scaffolding rather than internal machinery.
- Resource-first hypothesis: computation, memory and time explain developmental differences without reference to qualitative cognitive mechanisms.

- Embodiment-first hypothesis: sensorimotor organisation, rather than abstract reasoning machinery, establishes the fundamental limits of concepts and intelligence.

These explanations are not mutually exclusive. Experiments must vary knowledge, architecture, environment, embodiment and resources independently wherever possible. The Cognitive Machinery Limit Hypothesis earns support only if machinery explains distinctive variance and predicts characteristic failure boundaries beyond those alternatives.

17.1 Confounds that must be controlled

- Knowledge relevance: equal amounts of information are not equivalent if one set contains the solution or a privileged representation.
- Accessibility: stored content must be equally retrievable; otherwise a memory-access difference may masquerade as a reasoning difference.
- Compute leakage: a supposedly knowledge-only condition must not receive additional search, inference or external processing.
- Training exposure: examples, feedback and environmental opportunities must be matched or explicitly modelled.
- Interface quality: sensing and action bandwidth must be held constant when architecture is the independent variable.
- Goal and motivation: systems must have comparable opportunity to engage the task rather than one being prevented from acting.
- Implementation competence: failure caused by a programming defect is not evidence for a theoretical machinery boundary.
- Equivalent mechanisms: a system must not be declared deficient merely because it implements the required transformation differently.

17.2 Adversarial objections and replies

Strong objection	Defence or concession
"The hypothesis is true by definition."	H1 contains a conceptual necessity, but H2-H5 are empirical: they predict identifiable failure signatures, intervention effects and incremental predictive power. The complete hypothesis can fail even if its logical core survives.
"You have defined machinery as everything cognition does."	The scientific formulation restricts machinery to independently specified transformations and causal pathways. Evidence, stored content, resources, embodiment and motivation remain separate variables.

Strong objection	Defence or concession
“Knowledge itself changes machinery.”	Conceded. The architecture includes learning operations that can proceduralise knowledge or construct new operations. The question becomes whether the system contains a pathway capable of that conversion; external insertion must not be misclassified as autonomous construction.
“Machinery growth is merely more knowledge under another name.”	No. Knowledge growth changes usable content; machinery growth requires a persistent, causally demonstrated change in executable closure, integration or revision. A learned rule counts as machinery only after it becomes a reusable operation that transfers and alters later cognition.
“Functional coverage is circular because successful systems define the required functions.”	Coverage requirements, weights, tolerances and resources must be fixed independently before performance is observed. Held-out tasks and intervention tests then determine whether the map predicts failure boundaries.
“Functional equivalence erases meaningful implementation differences.”	Equivalence is indexed to a declared domain, resources, evidence and causal role. Differences in latency, errors or downstream effects invalidate equivalence whenever those differences matter to the tested claim.
“Enough knowledge could contain every answer.”	A lookup system may cover a closed test distribution, but coverage is not developmental reach. Under unbounded novelty, contradiction or representational change, stored answers do not by themselves supply operations for determining which answer applies or how beliefs should change.
“The environment, not architecture, sets the limit.”	External opportunity is indispensable and explicitly conditioned. The CMLH concerns the principal internal constraint after evidence, embodiment and resources are specified or controlled.
“A system can simulate a missing mechanism.”	If it can reliably simulate the required transformation using available operations, then the transformation lies within its functional closure and is not missing. The hypothesis is implementation-neutral.

Strong objection	Defence or concession
“Human cognition is too entangled for clean decomposition.”	This may restrict direct human testing but does not invalidate intervention studies in synthetic systems or convergent evidence from neuropsychology and developmental change.
“Completeness is impossible.”	Absolute completeness is not asserted. The quotation uses philosophical compression; experiments use functional coverage and sufficiency relative to declared task classes.
“Greater machinery may simply mean greater compute.”	Resource-matched ablation, integration and architecture comparisons are required. A result that disappears under compute matching does not support the machinery claim.
“The theory cannot measure ultimate potential.”	Conceded. The scientific dependent variable is long-run adaptive reach within stated conditions, never metaphysical or unlimited potential.

17.3 Protection against unfalsifiable rescue

The hypothesis must not be protected after failure by claiming that an unobserved mechanism was secretly absent, that the test was never long enough, or that every successful alternative was “really” the same machinery. Before testing, researchers must declare the target transformation, acceptable functionally equivalent implementations, opportunity conditions, resource ceiling, measurement interval and predicted failure signature. Unregistered reinterpretation after results cannot count as confirmation.

17.4 Evidence grading

- Demonstration: one system exhibits the predicted pattern. This establishes possibility, not generality.
- Controlled intervention: ablation or integration selectively changes the predicted task class.
- Replication: the effect recurs across implementations, seeds and environment variants.
- Discriminant evidence: machinery measures outperform knowledge-only, scale-only and resource-only explanations on held-out tests.
- Generality evidence: the same architectural variables predict adaptation across multiple superficially different task families.
- Theoretical maturity: accumulated results support a quantitative model specifying when knowledge, machinery and opportunity dominate.

18. Refined Thesis

The Cognitive Machinery Limit Hypothesis For a finite mind under declared evidence, embodiment and resource conditions, present knowledge contributes to immediate competence, while task-relative functional coverage, causal integration, revisability and meta-machinery constrain long-run adaptive reach. If a class of understanding requires a transformation outside the mind's resource-bounded compositional closure, and its learning processes cannot construct or integrate a functionally equivalent operation, enlarging stored knowledge alone will not make that class reachable. Machinery growth occurs only when learning persistently changes that closure, integration or revisability and causally enables transfer beyond stored answers. Under controlled novelty, these architectural properties should predict adaptation, transfer, recovery and self-expansion beyond initial knowledge quantity.

19. Concise Formulations

Original formulation

“The limit of a mind is the completeness of its reasoning machinery, not the infinity of its knowledge.”

Scientific formulation

For finite cognitive systems with comparable opportunity and resources, long-run adaptive reach is constrained by the range, integration and revisability of available cognitive operations; stored knowledge alone cannot compensate for an absent operation required to transform that knowledge into the relevant form of understanding.

Expanded operational formulation

A finite mind's developmental reach is bounded by its resource-constrained closure over executable transformation classes and by the meta-machinery through which that closure can grow. Knowledge can become machinery, but only when it is converted into a persistent, integrated and causally effective operation.

Plain-language formulation

Not knowing is not always the deepest limit. The deeper limit is lacking a way to turn evidence into the understanding that is missing.

20. Conclusion

The Cognitive Machinery Limit Hypothesis reframes intelligence as developmental capacity rather than informational possession. It distinguishes a mind that contains answers from a mind that can originate justified understanding. Knowledge remains essential, but it is treated as the changing product and working material of cognition, not as the final measure of cognitive possibility.

The hypothesis is strongest when interpreted neither as a promise of unlimited intelligence nor as a dismissal of evidence and resources. It is a claim about the location of the deeper internal boundary. Finite ignorance can be crossed. A missing cognitive operation cannot be crossed until the mind acquires, constructs or integrates machinery capable of performing it.

The completeness issue is therefore resolved operationally rather than asserted absolutely. No mind is claimed to possess complete reasoning in an unrestricted sense. Every empirical claim is indexed to a task class, equivalence standard, resource budget, evidence condition and reliability threshold. The theory succeeds only insofar as those preregistered architectural measures predict and causally explain reach boundaries better than knowledge, scale, environment or resources alone.

Selected Intellectual Context

- Legg, S., & Hutter, M. (2007). Universal Intelligence: A Definition of Machine Intelligence. <https://arxiv.org/abs/0712.3329>
- Nilsson, N. J. (2007). The Physical Symbol System Hypothesis: Status and Prospects. <https://ai.stanford.edu/~nilsson/OnlinePubs-Nils/PublishedPapers/pssh.pdf>
- Rescorla, M. Computational Theory of Mind. Stanford Encyclopedia of Philosophy. <https://plato.stanford.edu/entries/computational-mind/>
- Wheeler, G. Bounded Rationality. Stanford Encyclopedia of Philosophy. <https://plato.stanford.edu/entries/bounded-rationality/>
- Piagetian assimilation and accommodation overview. NCBI Bookshelf. <https://www.ncbi.nlm.nih.gov/books/NBK448206/>

Authorship note: The central proposition and original formulation were proposed by Charles Dooner in the course of developing NEL. The references above identify related intellectual territory; they are not presented as sources of the proposition itself.