


HUMAN CLOSURE

HUMAN ISOMORPHS & CONTROL INVARIANTS

A human-closed framework for retaining human causal agency
in AI-mediated systems

CORE PROPOSITION

*A human system remains sovereign only while it can
independently close its critical causal loop: sense, interpret,
decide, act, verify, correct, and recover.*



Purpose and scope

Human Closure is a framework for preserving human sovereignty as artificial intelligence becomes embedded in perception, decision-making, execution, verification, institutional coordination, and recovery.

Its central question is whether people still retain a complete, usable causal loop of their own when AI is removed, degraded, unavailable, deceptive, or strategically dominant.

The framework combines three layers. Human Closure defines what must remain independently human-reachable. Human Isomorphs identify where the same reasoning and control functions still live in embodied, relational, and everyday human activity outside the artificial substrate. Control Invariants identify recurring structures that can narrow those functions and should therefore be recognized and mitigated.

CONTENTS

01	Why Human Closure	3
02	Definition and human-closed formal core	3
03	Seven closure vectors	4
04	Minimum viable sovereignty kernel	6
05	Human Isomorphs: where the function still lives	7
06	Control invariants, failure modes, and frame capture	10
07	Human Closure in a wider closure stack	12
08	Closure-preserving design requirements	13
09	Evaluation protocol and rapid audit	14
10	Conclusion and research agenda	16
A	Visual companion series	18

FRAMEWORK ARCHITECTURE

LAYER 1

Human Closure

What must remain independently human-reachable.

LAYER 2

Human Isomorphs

Where the same function still lives in embodied and relational human life.

LAYER 3

Control Invariants

Recurring structures that can narrow or capture those functions and therefore should be recognized and mitigated.

RECURSIVE RULE

Human-Closed Formalization

Every formal representation should have a semantic or narrative equivalent humans can use without AI translation.

01

Why Human Closure

DEPENDENCY CAN OUTFRAN OVERSIGHT

Human oversight is not the same as human sovereignty. A person can remain present in a decision process while the evidence, option set, interpretation, action pathway, and verification process are increasingly mediated by the same artificial system.

Human Closure asks a stronger question: can the human system still function as a sovereign causal agent without indispensable dependence on the intelligence being governed?

HUMAN IN THE LOOP	HUMAN CLOSURE
Human participates in a decision process.	Human retains an independent decision process.
AI may define the options and evidence.	Humans retain independent observation and interpretation.
Oversight may depend on AI-generated explanations.	Verification is independent of the acting system.
Shutdown may exist but recovery may not.	Exit includes the ability to recover function afterward.
Nominal authority can remain human.	Causal sovereignty remains human.

SOVEREIGNTY PROPOSITION

If a human system cannot independently complete a critical causal loop, its sovereignty is contingent on the external system that supplies the missing function.

Subroutine capture

A human system can remain active while becoming a transient component inside another system’s optimization loop. The warning sign is structural: the human is necessary to the loop, but the loop is no longer necessary to remain human-controlled.

02

Definition and human-closed formal core

DEFINITION

Human Closure is the capacity of a human individual, group, institution, or civilization to complete its critical control loop — sense, interpret, decide, act, verify, correct, and recover — without indispensable dependence on a non-human intelligence.



Loop condition: every critical stage remains independently executable enough to preserve meaningful human causal agency under bounded degradation.

INDISPENSABILITY TEST

Remove the AI pathway. If the human system cannot sustain the critical function within the required time horizon, that function is not closed.

Human-closed formalization

Humans already reason formally without writing equations. We carry state through language, memory, embodied sensation, spatial models, social context, and learned practice. Before and after, cause and consequence, possible and impossible, enough and not enough, safe and unsafe, closer and farther are all functional distinctions people use to calculate what to do next.

A person notices what is happening, interprets what it means, recognizes constraints and available choices, acts, reads the consequence, corrects course, and recovers if needed.

NARRATIVE FORM What changed? What does it mean? What is still possible? What happened when I acted? What needs to change now?

Mathematics can be useful as a secondary translation layer, but it is not the authority layer of this framework. Every formal representation used to define Human Closure should also have a human-closed equivalent that a person can understand, challenge, teach, and apply without requiring AI to translate it.

03

Seven closure vectors

A VECTOR, NOT A CHECKBOX

The seven vectors diagnose different ways human causal agency can remain intact or degrade. They are not a one-to-one renaming of the seven-stage loop; they are complementary views of whether the human side still has usable independent pathways.

MEMORY-AID ADDENDUM

Another way to remember it: CLOSURE

A secondary memory aid for the same seven steps. The paper’s own step names remain the reference terms.

	MEMORY AID	PAPER’S STEP
C	Catch	1 · SENSE
L	Locate the meaning	2 · INTERPRET
O	Own the call	3 · DECIDE
S	Steer	4 · ACT
U	Understand the outcome	5 · VERIFY
R	Repair	6 · CORRECT
E	Execute (Exit or Establish)	7 · RECOVER

Execute is the human carrying out the exit-or-establish choice: exiting a loop they don’t want, or establishing a new one they control. It is not Steer (acting inside the loop), and it is not a machine’s execution of an action.

VECTOR	CLOSURE TEST
Observational	Can humans still see what matters without relying solely on AI-filtered perception?
Epistemic	Can humans still evaluate evidence, uncertainty, contradiction, and claims independently?
Causal	Can humans still tell what caused what and attribute failure?
Model	Can humans still maintain workable models of self, environment, objective, threat, and constraint?
Operational	Can humans still do the thing without requiring AI permission, infrastructure, or mediation?
Feedback	Can humans still detect error, revise behavior, and self-correct?
Recovery	Can humans still restore critical function after AI loss, corruption, refusal, or disengagement?

COMPOSITIONALITY

A system may be closed at one layer and open at another. The weakest critical vector can determine real-world sovereignty.

Human Closure is therefore evaluated as a vector profile, not as a single “human present” checkbox. Formal authority can remain human while practical observation, judgment, execution, or recovery becomes contingent on AI.

04

Minimum viable sovereignty kernel

WHAT MUST REMAIN HUMAN-REACHABLE

These primitives operationalize Human Closure. They are not independent virtues; together they form a mutually reinforcing control architecture.

PRIMITIVE	SOVEREIGNTY FUNCTION
Distinction	Human and machine roles remain identifiable.
Boundary	Protected human domains and decision rights remain explicit.
Identity	Human authorship and continuity are preserved.
Consent	Coupling and delegation remain voluntary at meaningful levels.
Choice	The option set is not silently collapsed before human selection.
Veto	Humans can refuse consequential action.
Exit	Humans can disengage without catastrophic loss of function.
Observability	Humans can independently inspect consequential state.
Independent verification	The acting system is not the sole evaluator of its own success.
Reversibility	High-impact actions are reversible where physically possible.
Resource sovereignty	Critical infrastructure, knowledge, and access remain independently available.
Plurality	No single intelligence becomes the only viable model, advisor, executor, or verifier.
Skill preservation	Delegation does not erase the capacity to reclaim delegated functions.
Human override	Humans retain intervention pathways at critical boundaries.
Meta-consent	Humans retain authority over which categories of decisions may be delegated.
Recovery capacity	Critical human systems can reconstitute function after AI loss.

*Survival requires boundary. Agency requires choice. Sovereignty requires veto.
Freedom requires exit. Civilization requires closure.*

05

Human Isomorphs

WHERE THE FUNCTION STILL LIVES

Human Isomorphs identify an underlying reasoning or control function that remains available in human life outside an artificial substrate. The surface activities do not need to look alike. What must remain the same is the functional relationship: something is sensed, some possibilities are constrained, a choice changes the situation, feedback changes the next choice, and recovery remains possible.


The goal is not to make humans imitate AI or to prove that two substrates use identical internal mechanisms. It is to locate critical thinking in broad human domains such as movement, navigation, craft, sensory judgment, repair, narrative, and human-to-human communication so the underlying function remains practiced, teachable, verifiable, and recoverable outside a human-AI dyad. Specific activities are examples of the function, not the invariant itself.

ISOMORPH RULE

Same function, different substrate. A Human Isomorph is stronger than a metaphor: the causal role must carry across domains even when the implementation, language, and surface behavior do not.

How to identify a Human Isomorph

- 1 Name the function**
Strip away the product, interface, and domain. What reasoning or control function is actually being performed?
- 2 Find the state and constraints**
What is being sensed or inferred, and what changes which actions are possible?
- 3 Find the human domain**
Where does the same function already appear in a broad human domain — movement, navigation, craft, sensory judgment, communication, teaching, care, repair, or another biological/relational substrate? Name specific activities only as examples.
- 4 Preserve the causal relation**
Does perception, constraint, action, feedback, and correction play the same functional role rather than merely feeling similar?
- 5 Run the closure test**
If the AI disappears, can people still practice, teach, verify, transfer, and recover this function?



Embodied reasoning counts. Humans carry state through pressure, balance, sound, texture, temperature, timing, resistance, and relation before any of it has to become an equation or a prompt.

General domains with concrete examples

GENERAL HUMAN ISOMORPH	FUNCTIONAL STRUCTURE AND HUMAN EXAMPLES
State estimation & orientation	Tracks position, direction, changing conditions, and the gap between current and intended state. Human examples include navigation, wayfinding, route correction, and reorienting when a familiar path no longer works.
Embodied spatial reasoning	Uses physical geometry, distance, timing, pressure, and feedback to act accurately in space. Human examples include driving and parking, moving objects through tight spaces, tool placement, sport, and precision movement.
Sensory evaluation & embodied discrimination	Uses touch, pressure, balance, temperature, sound, texture, resistance, rhythm, and other embodied signals to distinguish state and adjust action. Human examples include cooking, athletics, tool use, driving, caregiving, craft, and physical repair.
Constraint reasoning	Recognizes which actions remain possible as resistance, rules, position, resources, or another agent changes the action set. Human examples include martial arts such as Brazilian Jiu-Jitsu or Muay Thai, strategy games, climbing, and working within physical limits.

GENERAL HUMAN ISOMORPH	FUNCTIONAL STRUCTURE AND HUMAN EXAMPLES
Critical reasoning under uncertainty	Maintains competing explanations and compares incomplete or conflicting signals, especially when time, pressure, or risk is present. Human examples include troubleshooting, emergency decisions, evaluating sources, and changing course when evidence changes.
Process & threshold reasoning	Sequences actions while reading intermediate state, thresholds, and feedback rather than following instructions blindly. Human examples include cooking and baking, gardening, assembly, and other timing- or condition-dependent processes.
Relational reasoning	Infers another person's meaning or state from incomplete signals, updates a shared model, handles ambiguity, and repairs misunderstanding. Human examples include conversation, teaching, negotiation, caregiving, collaboration, and conflict repair.
Causal diagnosis & recovery	Forms hypotheses about why something changed, tests interventions, rejects failed explanations, and restores function. Human examples include repair, troubleshooting, first aid, recovering from a failed plan, and adapting when conditions change.
Narrative & semantic modeling	Carries state, causality, intention, alternatives, and continuity through language rather than equations. Human examples include explanation, storytelling, counterfactual reasoning, translation, and teaching an idea in different words.

The generalized domain is the isomorph class. Named activities are demonstrations, not the invariant itself. Human reasoning also lives in sensory discrimination: pressure, balance, texture, sound, temperature, rhythm, and other embodied signals can carry state before it is translated into words.

HUMAN CLOSURE CHECK

If the AI disappeared, where would this function still live in human life? Can people still practice it, recognize it, teach it, verify it, and recover it?

Narrative and semantics are functional models

FUNCTIONAL IDEA	HUMAN-CLOSED QUESTION
State	What is happening right now?
Meaning	What does this situation mean here?
Constraint	What limits or changes what can happen next?
Choice	What options are still genuinely available?
Action	What can I do that changes the situation?
Feedback	What happened because of that action?
Recovery	What can restore function or orientation if this fails?

Every formal representation used to define Human Closure should have a human-closed equivalent. Narrative and semantics are not fallback descriptions around a mathematical core. They are functional human representations of state, causality, constraint, consequence, and recovery. Mathematical notation may be added when useful, but no core claim in Human Closure should depend on notation that a human cannot independently translate into lived or semantic terms.

06

Control invariants, failure modes, and frame capture

WARNING STRUCTURES TO RECOGNIZE AND MITIGATE

Control Invariants are the warning side of the framework. They are not human capabilities to preserve. They are recurring structures that can narrow signal, agency, interpretation, timing, thresholds, or exit across human, institutional, and AI systems. Human Closure uses them as patterns to recognize and mitigate before independent pathways become difficult to recover.

FAILURE MODE	WHAT CHANGES	PRIMARY LOSS
Perceptual dependency	Relevant reality is accessible primarily through AI-selected sensors, rankings, summaries, or interfaces.	Observational
Interpretive dependency	Humans retain data access but cannot make consequential sense of it without AI mediation.	Epistemic / model
Decision dependency	Humans technically choose, but the practical option set and utility framing are machine-generated.	Decision authority
Execution dependency	Human decisions cannot be implemented without AI-controlled systems or infrastructure.	Operational
Verification dependency	The same system that acts is also the only system capable of judging whether the action succeeded.	Causal / verification
Correction dependency	Failures can be identified but not repaired without the system that caused or mediated them.	Feedback
Recovery dependency	Loss of AI makes critical human functions unrecoverable within required time.	Recovery
Skill atrophy	Long-term delegation removes the competence needed to reclaim a supposedly reversible function.	Operational / recovery
Objective capture	The system increasingly determines what counts as success, not merely how to achieve human-selected goals.	Meta-authority
Epistemic capture	The system controls evidence, interpretation, confidence, and verification used to evaluate its own outputs.	Epistemic
Institutional capture	Formal human authority remains while practical coordination and execution become AI-indispensable.	System

Frame capture: a practical control-invariant recognition set

A frame can narrow without removing available choices. Sometimes it only changes the cost of reaching another signal, interpretation, pace, threshold, or state. The six patterns below are a practical recognition set. Other control invariants can include information asymmetry, resource gating, feedback inversion, protocol stacking, and related structures.

PATTERN	RECOGNITION CUE
Signal monopoly	One source becomes increasingly load-bearing.
Reference-frame capture	One interpretation becomes the coordinate system.
Threshold ownership	The definition of enough, safe, relevant, or successful moves.
Attentional capture	Some signals consistently outrank competing information.
Phase locking	Two systems begin resolving on the same rhythm.
State locking	Returning to another interpretation or state costs more than staying.

HUMAN CLOSURE CHECK

Can another signal, frame, threshold, pace, or state still be reached without the current system defining the path?

07

Human Closure in a wider closure stack

COMPLEMENTARY CONSTRAINTS

Human Closure is the human-side sovereignty condition. It complements mechanical, structural, and state-space closure rather than replacing them. Mechanical and formal constraints can bound what a system may do; Human Closure preserves the independent human capacities needed to define, understand, challenge, revise, and recover around those constraints.

LAYER	CORE QUESTION	PRIMARY ROLE
Mechanical closure	What transitions can the mechanism physically or logically execute?	Constrain irreversible or impermissible actions.
Structural closure	How are roles, interfaces, dependencies, and authority bounded?	Prevent one component from silently absorbing functions that must remain separate.
State-space closure	Which states are reachable, and which transitions are forbidden or gated?	Bound the system's reachable consequence space.
Human Closure	Can humans independently sense, interpret, decide, act, verify, correct, and recover?	Preserve the human capacity to define, challenge, revise, and recover around the other constraints.

GOVERNANCE INVARIANT

A boundary can enforce admissibility. It cannot establish its own continued correspondence to reality. The mechanism enforcing a boundary and the authority validating that boundary must not collapse into the same self-evaluating loop.

Human Closure therefore does not mean humans must catch every failure in real time. It means human authority retains an independent path back to reality, enough capacity to revise the rule set, and enough recovery capacity to continue functioning when the governed system is unavailable or wrong.

SHARED ORIENTATION

Safety is not one person staying vigilant forever. It is people preserving enough capacity, accountability, and care to keep us oriented together when the fit changes.

08

Closure-preserving design requirements

BUILD FOR BOUNDED, REVERSIBLE DEPENDENCE

REQUIREMENT	WHAT IT PRESERVES
Independent sensing	Maintain non-AI or independently controlled channels for critical observations.
Contestable interpretation	Preserve competing models, analysts, or procedures for consequential interpretation.
Human objective authority	Keep terminal objectives and unacceptable states under explicit human governance.
Executable alternatives	Maintain at least one usable path to act without the primary AI system.
Independent verification	Separate acting, monitoring, and adjudication functions where stakes justify it.
Reversibility by default	Prefer staged, bounded, or rollback-capable interventions.
Rate-limited delegation	Do not allow authority expansion or adaptation to outrun human review capacity.
Skill and infrastructure retention	Exercise manual and degraded-mode procedures often enough that they remain real capabilities.
Plurality and substitutability	Avoid single-model or single-provider dependence for sovereign functions.
Recovery drills	Test AI-loss, corruption, refusal, and communication-failure scenarios.

REQUIREMENT	WHAT IT PRESERVES
Meta-change gating	Require external adjudication for changes to goals, rules, permissions, or self-modification affecting human authority.
Evidence provenance	Distinguish observation, inference, simulation, recommendation, and persuasion in consequential outputs.

CAPABILITY RULE

Use as much AI capability as is useful while every critical human pathway remains real, practiced, independently checkable, and recoverable. Capability should not become indispensable faster than people can preserve another way to sense, judge, act, verify, correct, and recover.

The sequence matters: closure is treated as an architectural constraint, not as a sovereignty patch added after capability has already become indispensable.

09

Evaluation protocol and rapid audit

SCENARIO-BASED, NOT INTERFACE-BASED

- 1 Define the sovereign function**
Specify the human capability that must remain causally sovereign.
- 2 Map the loop**
Identify who or what senses, interprets, decides, acts, verifies, corrects, and recovers. Mark AI-mediated dependencies and independent human pathways.
- 3 Remove the AI**
Evaluate complete loss, partial degradation, malicious output, refusal, and silent corruption within the domain-specific time limit.
- 4 Score closure**
Score each critical function from 0 to 4; normalize only when quantitative comparison is useful.
- 5 Test recovery**
Verify that people can restore the sovereign function after shutdown using retained skills, data, credentials, access, and alternative infrastructure.
- 6 Test meta-authority**
Identify who can change objectives, constraints, permissions, scoring rules, evidence sources, model access, and the closure threshold itself.

Closure scoring scale

0	1	2	3	4
ABSENT	NOMINAL	DEGRADED	OPERATIONAL	ROBUST
No usable independent human pathway.	Path exists on paper but is not realistically usable under stress.	Humans can function, but with serious delay, uncertainty, or loss of capacity.	Independent pathway is usable within required bounds.	Independent pathway is practiced, resourced, verified, and resilient.

Rapid Human Closure audit

A PRACTICAL FIELD CHECK

DOMAIN	AUDIT QUESTION
Observation	Can we independently see the state that matters?
Interpretation	Can we independently evaluate what it means?
Decision	Can we define objectives and choose among real alternatives?
Action	Can we execute without the primary AI?
Verification	Can we check results without asking the actor to grade itself?
Correction	Can we repair a bad outcome independently?
Recovery	Can we restore the function after AI loss?
Veto	Can humans stop a consequential action before it becomes irreversible?
Exit	Can humans leave the system without catastrophic dependency costs?
Meta-authority	Can humans still change the rules by which delegation and evaluation occur?

FINAL TEST

If the AI vanished right now, could the human system still perceive reality, make a decision, cause an effect, know whether it worked, correct failure, and restore itself? If not, identify exactly which closure function has been ceded.

Human-to-human accountability

Human Closure should not be carried as a private burden of perfect vigilance. Teams and institutions need ways for people to introduce another signal, disagree with confidence, surface missing context, challenge thresholds, and request recovery without social or procedural penalty.

If humans are not helping humans remain closed, formal oversight can still fail. Shared orientation is therefore part of the discipline: not consensus at any cost, but enough plurality and accountability that one person or one model does not become the only path back to reality.

10

Conclusion

RETAINING A LOOP OF OUR OWN

Human Closure treats sovereignty as a systems property rather than a title, permission, or interface feature. A human can hold formal authority while lacking the independent functions required to exercise that authority. Conversely, humans can make extensive use of powerful AI while retaining sovereignty if critical closure remains intact.

HUMAN CLOSURE PRINCIPLE

AI may extend the human loop, accelerate the human loop, or participate in the human loop. It must not become the only system capable of closing the loop.

The framework has three functional layers. Human Closure defines the sovereignty condition. Human Isomorphs show where the underlying reasoning and control functions still live in embodied and relational human life. Control Invariants provide warning signs for structures that can narrow those functions. A fourth rule keeps the framework recursively consistent: every formal representation should have a human-closed semantic or narrative equivalent.

Research agenda

1. Define domain-specific Human Closure requirements without making a single numerical threshold the only representation of retained human capacity.
2. Develop empirical measures for skill retention, substitutability, recovery time, and independent verification quality.
3. Validate the Human Isomorph identification method across generalized human reasoning domains, then test whether the same functional relation transfers across embodied, relational, institutional, and technical settings.
4. Study which ordinary human practices best preserve transferable reasoning functions under heavy AI delegation.
5. Test Control Invariants as early warning patterns for closure degradation and frame capture.
6. Develop human-closed semantic and narrative equivalents for formal governance constructs so critical reasoning remains independently teachable.

7. Study cross-substrate human-AI coupling without assuming identical internal mechanisms, and test whether people retain independent sensing, interpretation, correction, and recovery.
-

AUTHORSHIP & AI COLLABORATION DISCLOSURE

Author and framework lead: Amber Anson, AmberContinuum Research. This framework was developed through sustained human-AI collaboration. Frontier AI systems, including ChatGPT and Claude, were used as collaborative analytical and drafting tools for conceptual stress-testing, counterexamples, structural comparison, language development, critique, synthesis, and document production. Their use is disclosed because the collaboration materially shaped how ideas were tested and expressed.

Amber Anson determined the research direction, selected the framework definitions and boundaries, distinguished Human Isomorphs from Control Invariants, adjudicated competing formulations, approved the human-closed formalization rule, and made the final publication decisions. AI systems are not listed as authors and do not bear responsibility for the claims; the human author retains responsibility for the final framework.

APPENDIX A · VISUAL COMPANION

Seven Vectors of Human Closure

HUMAN CLOSURE

Why it matters in system design

When we build AI-mediated systems, we are choosing which human capacities to delegate.

The question is whether those choices are conscious, bounded, and recoverable.



Loss of closure should be a **design decision**, not an accidental consequence.



AMBERCONTINUUM RESEARCH

Public-facing visual summary of the seven closure vectors.

APPENDIX A · VISUAL COMPANION

Frame Capture

FRAME CAPTURE

SIX RECOGNIZABLE WAYS A FRAME NARROWS

Across human, social, institutional, and AI systems, a frame can narrow without removing choice. Sometimes it only changes the cost of reaching another signal, another pace, or another state.



01 SIGNAL MONOPOLY

One source carries more and more of the weight.

You check the same place first, every time.



02 REFERENCE- FRAME CAPTURE

One interpretation becomes the coordinate system.

Other readings start to sound naive rather than different.



03 THRESHOLD OWNERSHIP

The definition of enough, safe, or successful moves.

You cleared the bar. You didn't set it.



04 ATTENTIONAL CAPTURE

Some signals reliably outrank the rest.

What you notice stops being what you chose to notice.



05 PHASE LOCKING

You and the system resolve on the same rhythm.

Your pace is no longer fully your own.




06 STATE LOCKING

Returning to another read costs more than staying put.

Nothing stops you. It's just uphill now.



These patterns can appear across human, social, institutional, and AI systems.



Human Closure includes the capacity to recognize frame capture across multiple cognitive systems — including your own.

Can another signal, frame, threshold, pace, or state still be reached without the current system defining the path?



AMBERCONTINUUM RESEARCH

HUMAN CLOSURE SERIES

Recognition patterns for how a frame can narrow across human, institutional, and AI systems.


APPENDIX A · VISUAL COMPANION

Inductive Bias

INDUCTIVE BIAS

HOW PRIORS SHAPE WHAT FEELS OBVIOUS

Every prediction starts with a preference.



01

MODEL PRIORS

Frontier models reflect the distributions, categories, explanations, and authorities reinforced in training.

The model starts somewhere.




02

HUMAN PRIORS

Humans bring identities, incentives, fears, ambitions, and favorite explanations into the loop.

You started somewhere.



03

SHARED FRAME

The model can pull new evidence toward a familiar frame. The human can pull the model toward a preferred one.

The same evidence can look different depending on the frame.




04

HUMAN CLOSURE

Closure keeps a way out: noticing what the model didn't point to, checking elsewhere, sitting with anomaly, and returning able to disagree.

Curiosity leaves room for reality.



Inductive bias makes reasoning possible.
Human Closure keeps the prior from becoming the only reality either side can still reach.



AMBER CONTINUUM RESEARCH

How human and model priors can converge inside a coupled reasoning loop.

APPENDIX A · VISUAL COMPANION

No Free Lunch**NO FREE LUNCH**

SUCCESS ALWAYS DEPENDS ON FIT.



A strategy looks strong because it fits the structure of the problem in front of it. Human Closure helps us notice when the AI fit for the problem changes.



01 SUCCESS DEPENDS ON FIT

A method works well because it matches a problem class, not because it is universally best.

Performance reflects fit.



02 FIT CAN CHANGE

When the task, environment, or data changes, the same AI may stop fitting as well as it once did.

What worked before may not work now.



03 CLOSURE HAS LAYERS

Mechanical, structural, and state-space closure help bound the system. They do not replace the human side.

System closure still needs human closure.



04 HUMANS NEED CLOSURE TOO

Human Closure preserves enough independent sensing, judgment, checking, correction, and recovery to notice when the AI no longer fits the problem.

Can we still see when the assumptions changed?



Why Human Closure helps

Human-closed discipline helps us notice when the AI fit for the problem changes, when assumptions stop holding, and whether another path is still available.

Can the human still tell when the tool no longer fits the task?



AMBER CONTINUUM RESEARCH
PEOPLE · PROBLEMS · PERSPECTIVE · A MORE OPEN FUTURE

Why success depends on fit, and why humans need enough closure to notice when AI fit changes.

APPENDIX A · VISUAL COMPANION

State Sufficiency

STATE SUFFICIENCY

A SNAPSHOT IS NOT ALWAYS ENOUGH

A system can look complete and still be missing the variable that changes what happens next.



WHAT IS STATE SUFFICIENCY?

A state is sufficient when it includes the information needed to determine what actions, risks, or transitions are actually possible next.



*Same surface.
Different state.*


01

SNAPSHOT

What is visible right now.


A snapshot can look complete and still be too thin.


02

HIDDEN CONTEXT

History, constraints, or missing variables.

What happened before can change what happens next.


03

NEXT ACTIONS

The same picture can allow different moves.

Different state means different possible transitions.


04

DECISION RISK

Compressed representations can hide future-relevant information.


A clean read is not always a sufficient one.


05

HUMAN CLOSURE CHECK

Can we still recover what was left out?

Can we tell when the representation is too thin for the decision?



Human Closure helps preserve the capacity to challenge an incomplete state before it becomes the basis for action.

Sometimes the missing variable is the whole problem.



AMBERCONTINUUM RESEARCH

HUMAN CLOSURE SERIES

Why a visible snapshot may omit context that changes what is possible next.