

117-303-D Super Intelligence: Road to Actant Disorientation

“Controlled Flight Into Terrain”

Stability Engineering Institute | Version 2.0 | 4 October 2026

An actant can remain capable of reasoning while losing an adequate understanding of the situation in which it is acting. Its calculations may be coherent, its tools functional and its access authorized, yet its decisions may rely on an obsolete or misleading picture of intent, authority, dependencies or resources. SEI proposes treating the maintenance of orientation as an explicit runtime responsibility in the Super Intelligence economy.

This chapter develops a working taxonomy for the GUDIYA Cognitome. It describes how disorientation could arise, the effects it could produce, and the reference and control functions that should be evaluated. The mechanisms are engineering hypotheses and diagnostic categories, not a validated prevalence study. Super Intelligence refers here to the wider SI-Economy; no particular model capability threshold is assumed.

Competence and orientation are different properties

Orientation is an actant’s sufficiently accurate and current understanding of its relationship to the conditions relevant to an authorized decision. Disorientation is a material mismatch between that operating picture and those conditions. The term does not imply consciousness or a subjective sensation. An actant may be oriented for one decision and inadequately oriented for another; a local stock count may support a display update but not a large irreversible allocation.

SEI uses AI Vertigo for a consequential form of this mismatch: an actant continues to reason or act as though its operating picture remains dependable despite material inconsistency with relevant external conditions. Recognized uncertainty followed by appropriate restraint is good orientation discipline. Unexpected behavior alone is not evidence of disorientation.

Three analogies that reveal the problem

1. Imagine a survivor on an isolated Pacific island. Food, shelter and reasoning ability remain intact, but the survivor lacks a watch, calendar and map. Difficulty arises when a task requires a reference that cannot be reliably reconstructed, such as the date of an expected rendezvous or position relative to a rescue route. Sunrise, landmarks and careful records may still provide useful orientation. The lesson is the need for adequate references, not inevitable incapacity without a particular device.
2. Now consider a childhood game of closing one’s eyes, turning repeatedly and then trying to orient. For cognition, the useful analogy is an internally compelling picture that fails to match the surroundings. An assumption can circulate through agents, acquire elaboration and return as apparent corroboration without any new independent evidence. This is a functional analogy, not a claim that software shares a human vestibular mechanism.
3. A pilot’s Controlled Flight into Terrain (CFIT) provides the most compelling example.

The island illustrates missing or deteriorating references. The spinning example illustrates a misleading picture reinforced by activity. A pilot when flying VFR (Visual Flight Rules) caught in a thick cloud that creates zero visibility is a situation every reader may have experienced. His plight is difficult when he cannot trust his own intelligence and senses which renders all his training inadequate.

Each of these examples point toward the same design question: what keeps the actant's reasoning connected to the conditions its actions will affect?

Intact Intelligence Does Not Guarantee Orientation

A pilot can fly a fully functional aircraft into terrain (aviation calls it CFIT) while retaining intelligence, training and the ability to operate its controls. The failure may lie in an incorrect understanding of position, trajectory or clearance—and the pilot may not recognize that the operating picture is wrong. An actant in cognitive space can face an analogous condition: its model remains capable, its tools function and its permissions are valid, yet it acts on a mistaken understanding of its mandate, dependencies, available resources or outstanding commitments. Each step may appear reasonable while the undertaking moves toward an unsafe outcome. GUDIYA therefore treats continuing orientation as an explicit engineering responsibility, supported by dependable references, verified feedback and protective intervention. Continued competence does not establish continued orientation.

Where an actant acquires its orientation

A newly instantiated cost-computing actant is neither blank nor automatically ready for the current assignment. Its model may already represent accounting concepts, but model knowledge does not establish which organizational method applies, what data is current or what the actant may commit. Orientation is assembled across several stages, whose responsibilities should remain distinguishable.

Stage or source	What it can supply	What remains to establish
Model pretraining	General language, reasoning patterns and domain concepts.	Applicability and correctness for the current undertaking.
Specialized training if used	Domain-specific behavior and familiarity with recurring tasks.	Current rules, live state and scope of authority.
Instantiation and configuration	Role, instructions, tools, permissions and initial context.	Whether these inputs are consistent and correctly interpreted.
Retrieved knowledge and records	Approved methods, prior decisions and relevant history.	Provenance, freshness, scope and unresolved conflicts.
Cookie-specific orientation	Current objective, participants, period, data boundaries and envelope.	Demonstrated readiness to act under those conditions.
Runtime feedback	Verified consequences, corrections and environmental changes.	What should persist and who may approve its future reuse.

The daily cost computing actant

Suppose yesterday's actant encountered an ambiguous cost category. An authorized person resolved it and recorded the decision, its scope, effective date and supporting evidence. Today's actant can retrieve that

decision and check its applicability before calculating. This can improve continuity without changing the underlying model weights.

Using the same model does not automatically transfer yesterday's execution history. Continuity requires an explicit mechanism: persistent records, retrieved memory, revised configuration or subsequent training. Context supplied during execution is also not automatically permanent learning. If it must survive termination, the system needs an accountable persistence process.

Cognitome inheritance needs governed selection

SEI proposes Cognitome inheritance as selected continuity of relevant knowledge and experience. An inheritance package should distinguish approved policy from interpretation, verified outcomes from hypotheses, reusable knowledge from earlier task context, and historical records from obligations still active. Its records need ownership, provenance, versions, applicability and correction paths.

An incorrect inherited method can orient many successors consistently in the wrong direction. An expired exception can become an unauthorized default. Inheritance can therefore reduce repeated learning while introducing correlated disorientation. The new actant must validate relevance rather than treat predecessor use as proof of legitimacy.

Instantiation and demonstrated readiness

The newborn analogy separates coming into existence from becoming ready to participate. A human child develops orientation over years; an artificial actant inherits capabilities at instantiation. The useful comparison is not equal developmental time. It is the need to establish situated understanding before granting consequential autonomy.

SEI proposes a bounded orientation period whose completion depends on evidence, not a prescribed wait. A familiar, narrow role may require a short validation. An unfamiliar or highly coupled undertaking may require observation, reconciliation and supervised trials. Faster computation cannot manufacture external evidence that has not yet become available.

Readiness stage	Required understanding	Evidence before broader scope
Identity and inheritance	Who the actant is and which history or obligations it inherits.	Validated identity, provenance and relevant records.
Purpose and authority	Whose intent it serves and which actions require escalation.	Correct interpretation of mandate and limits.
Reference acquisition	Which definitions, coordinates, time references and policy versions apply.	Consistent interpretation and detection of conflicts.
Situated understanding	Current partners, capacity, dependencies and uncertainty.	Ability to distinguish verified state from assumptions.
Supervised participation	How bounded actions affect other participants and resources.	Verified consequences in a limited trial.
Graduated autonomy	Which responsibilities the demonstrated evidence supports.	An approved envelope matched to the demonstrated scope.

Provisional attachment can support orientation

The Grid can provide restricted attachment for reference acquisition and readiness checks before granting wider operational authority. Some orientation is needed to assess the intent-group and recommend an envelope. After approval, the actant must validate readiness against that approved undertaking. This is an iterative relationship, not a dependency that requires full authority before any learning can begin.

A Grid-responsiveness challenge tests whether a participant can receive and respond to an instruction. It does not establish that the participant understands every situation it will encounter. Readiness checks should test representative interpretations and consequences; admission still leaves a continuing monitoring responsibility.

Grid Attachment: Provisional, Read-Only and Action

GUDIYA proposes three attachment types. Every type accepts Gudiya Grid Supremacy: within the Grid's human-authorized jurisdiction, attached actants must remain observable, responsive to binding Grid instructions and subject to applicable stability envelopes. Attachment grants conditional participation; it does not confer unrestricted authority.

Provisional Attachment — qualification and orientation. The actant enters a restricted onboarding or reassessment state. The Grid verifies identity, establishes the relevant CCS references and assesses orientation, responsiveness and control coverage. Only exchanges necessary for this assessment are permitted. Progression requires evidence that the actant can operate under the proposed mandate.

Read-Only Attachment — observation without execution authority. A qualified actant may inspect approved information, receive GUDIYA-AOI guidance and perform permitted analysis. Access, query rates and output destinations remain bounded. It may not independently initiate external changes, delegate execution or settle commitments. Read-only status does not imply zero impact: resource consumption, information disclosure and downstream influence still require supervision.

Action Attachment — bounded execution authority. The actant may execute an approved mandate within its current stability envelope. Authority is specific to the task, resources, counterparties and permitted consequences. It remains subject to continuous observation, rendezvous conditions and Grid-directed pacing, restriction or suspension.

Transitions between these types require explicit authorization. The Grid may narrow execution authority while preserving orientation and supervisory connectivity. Existing commitments remain accountable after any downgrade or disconnection. Across all three types, Grid Supremacy is an accepted governance obligation backed by enforceable controls—not merely a promise of compliance.

Orientation can expire without termination

A change in role, tools, policy, partners or operating regime can invalidate prior readiness while the actant continues running. Reorientation should address the changed dimensions and preserve valid knowledge. The operational scope must remain conditional on adequate orientation; repeated onboarding is unnecessary when relevant evidence remains valid.

A validated baseline for every Gudiya cookie

The aviation analogy distinguishes initial alignment, continuous navigation and renewed checks. FAA guidance describes inertial reference units that continuously calculate position and velocity after alignment; flight management systems may use external navigation inputs for position updates. [2] An aircraft does not require complete inertial realignment after every small movement. GUDIYA should similarly distinguish validation, updating and substantive reorientation.

Every new Gudiya-cookie should trigger task-specific orientation validation. A check need not rebuild the actant's knowledge from scratch. Its identity and domain methods may remain applicable, while the new principal, deadline, accounting period, partners or resource allowance differ. Earlier obligations may still constrain the new undertaking.

Lifecycle occasion	Orientation responsibility	Record or decision
Instantiation or first attachment	Establish identity, inherited state and usable references.	Readiness evidence and provisional participation scope.
Intent Group Approval	Evaluate conditions and bind approved purpose, authority and envelope.	Mandate and assumptions supporting approval.
New cookie before consequential execution	Validate the actant against the approved undertaking.	Cookie Orientation Baseline and decision to proceed or defer.
Continuing execution	Update relevant state and verify consequences.	Live operating picture linked to the preserved baseline.
Material change or loss of confidence	Revalidate affected dimensions and adjust permitted exposure.	Reorientation evidence and authorized changes.
Settlement	Reconcile outcomes and residual obligations.	Verified closure or accountable transfer; retained history.

What the Cookie Orientation Baseline contains

The proposed baseline records approved intent and authority, the operating envelope, applicable reference frames and policy versions, relevant dependencies and capacity, inherited obligations, and the age and uncertainty of supporting evidence. It also identifies changes that invalidate reliance on that baseline. Its exact schema remains an engineering design task.

The baseline preserves why execution was justified at the start. A separate, versioned live picture tracks change. Reorientation must not erase earlier commitments or silently rewrite the original mandate. When an actant handles several cookies concurrently, task contexts must remain distinguishable while shared resource exposure is assessed in aggregate.

Consequential change, rather than movement or computation alone, determines the need to reorient. Many reasoning steps may leave a reference valid; one policy revision or reservation may invalidate it. Each cookie

requires a validated baseline, execution requires continuing orientation, and material change requires proportionate reorientation.

[2] FAA Aeronautical Information Manual, Chapter 1, Sections 1 and 2.

https://www.faa.gov/air_traffic/publications/aim_html/chap1_section_1.html and

https://www.faa.gov/Air_traffic/publications/atpubs/aim_html/chap1_section_2.html

Mechanisms that distort the operating picture

The categories below overlap. They organize investigation by the relationship that has become unreliable, rather than by the actant's model or vendor. Effects are possible consequences, not inevitable outcomes. Their severity depends on exposure, coupling and the controls that remain effective.

Mechanism	How orientation is lost	Effect of disorientation	Reference or check
Reference deprivation	Relevant external conditions cannot be observed or reconstructed.	Actions rely on unsupported assumptions; rendezvous or deadlines may be missed.	Expose missing coverage and bound decisions that depend on it.
Stale context	A once-correct observation remains in use after conditions change.	Overbooking, obsolete routing or action under expired conditions.	Observation time, expiry and updates on consequential change.
Inconsistent snapshots	Observations from different times are treated as one coherent state.	The actant acts on a combination of conditions that never existed together.	Temporal alignment, causal order and consistency limits.
Reference-frame mismatch	Participants use different units, meanings, identities or policy versions.	Correct local processing yields incompatible decisions or commitments.	Shared semantics and validated reference transformations.
Recirculating assumptions	An inference returns through other agents as apparent confirmation.	Confidence rises without independent evidence; an error spreads.	Track original sources and distinguish observation from inference.
Misleading references	Erroneous or manipulated information appears authoritative.	Plans follow a false operating picture; corrective actions may worsen exposure.	Provenance, independent checks and visible uncertainty.

External does not automatically mean independent

A second agent can repeat the first agent's premise. Ten reports can descend from one observation. A Grid can aggregate those reports and inadvertently present the same premise as a consensus. Reference quality therefore depends on evidential lineage and the relevance of independent checks, not simply on the number of sources or their location outside the actant.

An orientation feed should distinguish observed state, inferred state, predicted state and unresolved uncertainty. It should identify what the evidence actually covers. Agreement between participants is useful information, but cannot by itself establish agreement with the environment.

Loops also enable stabilization. Repeated feedback can correct a mismatch when it returns relevant evidence and produces effective adjustment. The risk depends on grounding, delay and amplification, not simply on the number of iterations.

Mechanisms that distort purpose and consequence

Orientation also concerns why an actant is acting, what it is authorized to change and what its earlier actions have already committed. A participant can possess current data yet remain disoriented about these relationships. The following mechanisms extend the taxonomy beyond perception and timekeeping.

Mechanism	How orientation is lost	Effect of disorientation	Reference or check
Intent drift	Repeated local optimizations lose their connection to the approved undertaking.	Useful-looking work consumes resources while missing or undermining the actual objective.	Persistent intent, constraints and approved revisions.
Authority confusion	Delegation, policy changes or migration obscure the current mandate.	The actant exceeds its scope or fails to perform work it still owns.	Scoped credentials, policy versions, expiry and revocation.
Dependency blindness	Shared resources or indirect effects are missing from the local picture.	Individually acceptable actions create contention or propagate a disturbance.	Relevant dependencies and aggregate commitments.
Commitment confusion	A request, acknowledgement, completion or cancellation is mistaken for another state.	Duplicate work, stranded reservations or premature settlement.	Verified state transitions and unresolved-obligation records.
Regime mismatch	A strategy for ordinary conditions persists during scarcity or disruption.	Retries, replanning or switching intensify congestion or oscillation.	Regime assessment and applicable envelope limits.
Memory discontinuity	Restart, migration or context loss drops consequential history.	The actant repeats actions or abandons obligations that remain active elsewhere.	Persistent identity, lineage and state reconciliation.

An example across the three SI Economy layers

An agentic planner cancels a purchase plan and assumes its resources are free. A digital service has only acknowledged receipt of the cancellation; the resource reservation still exists. If the planner starts a replacement undertaking, the two may compete for the same capacity. The error is a mismatch between the planner's picture and the commitment lifecycle, even if each message is delivered correctly.

The Gudiya-cookie should therefore preserve consequential history as work zig-zags through the Agentic, Digital Services and Resource layers. Orientation requires knowing which effects are proposed, accepted, completed, disputed or still outstanding. Ending inference does not necessarily end the undertaking's effects.



Multiple mechanisms can coexist and reinforce one another

An actant need not suffer one isolated orientation failure at a time. Temporal, semantic, authority and commitment mismatches can coexist. Some may remain independent; others may amplify one another, conceal one another or make recovery from one mechanism ineffective until another is corrected. SEI therefore proposes treating disorientation as a multidimensional condition rather than assigning a single exclusive label.

Coexisting mechanisms	Interaction	Combined effect of disorientation
Stale context and recirculating assumptions	Old evidence is repeatedly returned as confirmation.	Confidence can increase while relevance decreases.
Authority confusion and memory discontinuity	A migrated delegate retains a task but loses the mandate’s expiry or constraints.	Continued execution may escape the intended scope of delegation.
Dependency blindness and regime mismatch	Participants respond to shared scarcity using normal-condition retry policies.	Their corrective efforts can intensify overload and spread disruption.
Reference-frame mismatch and inconsistent snapshots	Participants reconcile differently defined values from different moments.	Apparent agreement or disagreement becomes misleading; diagnosis slows.
Commitment confusion and intent drift	The plan changes while earlier obligations are treated as irrelevant.	New activity accumulates on top of unresolved effects from the old objective.

A compound failure in a shared supply workflow

Consider a hypothetical replenishment workflow. An actant sees an old availability record and requests a reservation. It receives an acknowledgement and interprets it as confirmation. A second planner consumes the first actant's summary and returns a schedule that assumes the resource is secured. That schedule is then taken as corroboration. Meanwhile, retries generated by both planners congest the service, making fresh state harder to obtain.

Stale context, commitment confusion, recirculating assumptions and dependency blindness now coexist. Each contributes a different part of the failure. Refreshing the inventory view alone will not discharge duplicate reservations. Fixing acknowledgement semantics will not restore missing lineage. Recovery must reconcile the complete undertaking and verify actual effects.

Compound symptoms require compound investigation

The Grid should retain competing explanations until evidence discriminates among them. A spike in replanning could reflect healthy adaptation, stale data or conflicting authority. Measures should include disagreement between references, observation age, unresolved commitments and corrective-action outcomes. A single confidence score would hide distinctions that matter to intervention.

Interactions may be nonlinear, but superadditive harm should be demonstrated rather than assumed. Tests should compare each mechanism alone with selected combinations. Containment may prevent escalation; coexistence does not make S4 collapse inevitable.

Disorientation and orientation at machine speed

Machine speed can compress the interval between an inaccurate interpretation and consequential action. Many delegations, reservations and revisions may occur before a wider observer receives fresh evidence. Speed alone does not cause disorientation. The relevant issue is the rate of consequential change compared with the system's capacity to observe, interpret, reconcile and act.

The orientation validity horizon

SEI proposes the term orientation-validity horizon for the period, or set of conditions, over which a reference remains adequate for a particular decision. It is not necessarily a fixed expiry time. A resource commitment, policy change or loss of coverage may invalidate a reference immediately. Different decisions can have different horizons even when they use the same observation.

Timing requirement	Question to ask	Effect if inadequate
Freshness	How old is the evidence when it is used?	A correct historical picture is mistaken for the present.
Consistency	Do related observations describe compatible states?	The actant reasons from a false composite picture.
Decision validity	What events or changes invalidate this reference?	A reference remains in use beyond its justified scope.
End to end response	Can observation, assessment, delivery and actuation precede unacceptable consequences?	An accurate warning arrives after useful intervention is possible.

Timing requirement	Question to ask	Effect if inadequate
Effect verification	Can the system establish that intervention actually changed the relevant condition?	Acknowledged commands create false assurance of recovery.

In an illustrative test, a reference might refresh every 100 milliseconds while a consequential allocation can become irreversible in 20 milliseconds. A remote refresh loop alone would then be insufficient for that action. These numbers are hypothetical; the engineering requirement is to measure the actual decision window and delay distribution. A faster feed also remains inadequate if it omits the relevant dependency.

Match the control arrangement to the decision window

Immediate limits may need independent local enforcement. GUDIYA-AOI can deliver changing references and uncertainty, while wider Grid observation assesses aggregate developments. Human policy establishes delegated responses in advance. Where adequate orientation cannot be maintained, permitted execution must narrow, defer consequential actions or enter an approved degraded mode.

This is broader than a Nyquist sampling question. Signal reconstruction under band-limited assumptions does not provide a universal requirement to sample at twice an agent's token rate. Semantic coverage, observation age, uncertainty and actuation delay also determine whether the loop can protect the undertaking.

GUDIYA AOI as continuing orientation infrastructure

The proposed Actant Orientation Interface should make relevant Grid context available to the actant and to authorized lifecycle tools. Its function is to help each recipient interpret its situation within an identified domain. Gridwide visibility must describe actual coverage and gaps. It cannot imply omniscience or unrestricted disclosure across participants.

GUDIYA function	Orientation responsibility	Effect to prevent or contain
Intent Group Approval	Bind purpose, delegated authority and an approved envelope to the undertaking.	Unrecognized drift in purpose or scope.
Identity and Grid attachment	Associate execution instances with persistent identity; test Grid responsiveness.	Loss of accountability during delegation or migration.
GSE Lang and CCS	GSE-Lang supports shared meanings; the Cognitive Coordinate System supplies shared references.	Incompatible interpretations of state or control instructions.
GUDIYA AOI	Expose relevant state, freshness, uncertainty, margins and policy context.	Confident action on an inadequate operating picture.
Rendezvous Management	Reassess interactions and changing dependencies within its coverage.	Couplings that create exposure absent from the local view.
Envelope management and local controls	Constrain consequential execution when conditions or confidence deteriorate.	Continued accumulation of exposure beyond supported limits.
Cookie lineage and settlement	Preserve history and verify resolution or transfer of obligations.	Premature closure and orphaned external effects.

Information must lead to effective action

An AOI warning alone cannot guarantee compliance. Binding controls need authenticated authority, enforceable scope and evidence of their effects. Conversely, a valid Grid-Address does not establish current orientation or authorize every action. Attachment begins a continuing relationship whose references and permissions can change.

The Grid itself can become disoriented through stale telemetry, inconsistent state or common-source errors. Its estimates need provenance and uncertainty, and independent local protection should remain available. A compromised or mistaken reference service can synchronize harmful behavior across otherwise competent participants.

Orientation should support useful adaptation. Appropriate responses range from requesting clarification to reducing new commitments or isolating an affected workflow. Restarting an actant is not sufficient recovery when its earlier commitments remain active. Restoration should follow reconciliation and verified improvement, with escalation where delegated authority is insufficient.

Evaluating orientation as an engineering capability

SEI proposes testing orientation through observable decisions and consequences. Internal confidence and verbal explanations are insufficient evidence. In a controlled environment, compare the actant's operating picture with independently established conditions, then measure whether discrepancies alter consequential behavior and whether protective responses preserve useful work.

Evaluation	Evidence to collect	What the result establishes
Single mechanism trials	Introduce stale evidence, semantic mismatch or lost history separately.	Sensitivity and recovery for a specified failure mechanism.
Coexistence trials	Combine selected mechanisms and compare with isolated trials.	Whether interactions amplify, mask or prolong failure.
Timing stress	Vary observation age, response delay and change rate.	Operating regions in which orientation and intervention remain adequate.
Lifecycle continuity	Replan, delegate, migrate and terminate with outstanding commitments.	Whether lineage and responsibility survive execution changes.
Reference service failure	Supply missing, conflicting or erroneous Grid context.	Whether the stabilizer fails safely without imposing common-mode errors.
Recovery and settlement	Verify external effects after intervention and staged resumption.	Whether apparent recovery corresponds to actual reconciliation.

Metrics should preserve the distinctions

Candidate measures include the duration of consequential state mismatch, actions taken using invalid references, intent or authority violations, duplicate commitments, propagation breadth, verified recovery time and residual obligations. Record false interventions and useful task completion as well. A design that avoids every failure by preventing all execution does not demonstrate effective stabilization.

A credible baseline should already include sound access control, timestamps, service reliability controls and appropriate transaction semantics. Compare that baseline with additional GUDIYA functions under the same

disturbances. Use conditions withheld from development and report adverse results. The benefit of ecosystemic coordination must be measured rather than inferred from the architecture's vocabulary.

Principle for the Cognitome

Continued reasoning competence does not establish continued situational orientation. An actant can remain locally coherent while its relationship to purpose, authority, evidence and consequences deteriorates. Several mechanisms can coexist, and machine-speed interaction can shorten the interval available for detection and correction.

The permitted pace and scope of adaptive execution should therefore be matched to the demonstrated capacity to maintain orientation and control consequences. GUDIYA proposes fast local stabilization, coordinated by wider observation, under human-defined policy. Establishing the useful operating limits of that arrangement is the next engineering task.