

## 117-191-C Cognitive Attitude Indicator

The Artificial Horizon for an Autonomous Cognitive Actant

### Introduction — What Tells an Agent Which Way Is Level?


- A modern autonomous agent may possess extraordinary capabilities. It can reason, plan, remember, use tools, invoke APIs, operate a CLI, communicate with other agents, critique its own work, revise a plan and execute actions in the physical or digital world.
- Yet there is a surprisingly primitive question that today's agentic architecture generally does not explicitly answer:
  - What tells the agent which way is level?
  - Not morally level.
  - Not whether its latest answer is correct.
  - Not whether a particular action violates a guardrail.
  - Rather:
    - How does the agent know its current orientation and motion relative to the larger dynamic system within which it is operating?
- Aviation confronted an analogous problem long ago.
- A pilot can be intelligent, experienced and completely sincere about what the aircraft appears to be doing—and still be spatially disoriented.
  - The solution was not more intelligence.
  - It was external reference. Using an artificial horizon.
  - GUDIYA proposes that autonomous cognition may eventually require its equivalent:

### CAI — Cognitive Attitude Indicator

- A Grid-supplied instrument that continuously provides an actant with an externally referenced indication of its cognitive orientation, motion, trajectory and stability-envelope relationship relative to the Cognition Coordinate System (CCS), independent of the actant's own internal perception of its state.
- The CAI is the artificial horizon for cognition.

### 1. Intelligence Does Not Provide Orientation

- Consider an aircraft flying through cloud.
- The pilot and aircraft are moving together. Visual reference to the outside world may disappear. The pilot's vestibular system continues supplying information, but that information can become misleading.
- The pilot may sincerely perceive: I am straight and level, while the aircraft is banking.
- Aviation therefore separates two fundamentally different things:



- The same distinction may apply to an autonomous cognitive actant.
- An agent may know:

- its objective,
- its instructions,
- its current plan,
- its memory,
- its tool results,
- its confidence,
- its previous actions.
- It can therefore possess considerable Agentic Self-Awareness.
- But none of this necessarily establishes: Where am I relative to the larger system?

Hence an important GUDIYA distinction:

*Agentic Self-Awareness  $\neq$  Grid Awareness*

And more fundamentally:

*Intelligence  $\neq$  Orientation*

## 2. Present-Day Agents Largely Have No Artificial Horizon


- Today's agentic architectures commonly provide some combination of:

Prompt + Context + Memory + Planning + Tools + Observations + Reflection + Guardrails

- These capabilities can be extremely sophisticated.
- But they mostly answer local questions:
  - What should I accomplish?
  - What information do I possess?
  - What should I do next?
  - Did my previous action succeed?
  - Am I permitted to perform this action?
  - Should I revise my plan?
- What is generally absent is a standardized external system continuously answering:
  - Where are you relative to a common cognitive reference?
  - How rapidly is your relevant state changing?
  - How are other actants moving relative to you?
  - Are you approaching an unstable rendezvous?
  - Are you departing your approved Cognitive Envelope?
  - Is your internally perceived orientation diverging from your externally measured orientation?
  - Where will your present trajectory take you?
- Thus, present-day agentic computing generally assumes that an agent can operate without anything resembling an externally authoritative Cognitive Attitude Indicator.
- This does not mean agents possess no state awareness.
- They can possess enormous amounts of state.
- The narrower proposition is:
  - State awareness is not necessarily attitude awareness.
  - An agent may know a great deal about itself while lacking a reference for determining how that self is moving relative to everything else.

### 3. The Somatogravic Lesson

- Aviation provides an even more revealing phenomenon: somatogravic illusion.
- Under strong linear acceleration—classically during takeoff or a go-around—the human vestibular system can misinterpret the combined gravitational and inertial forces as a change in aircraft pitch.
- The pilot may experience the compelling sensation that the nose has risen excessively.
- The crucial lesson is not simply:
  - Human senses make mistakes.
  - It is subtler.
  - The locally sensed signals can be real.
  - What becomes incorrect is the inference about orientation derived from those signals.
- Thus:



- This phenomenon gives the Cognitive Attitude Indicator considerably deeper significance.

### 4. Cognitive Somatogravic Illusion

- Imagine an autonomous agent undergoing rapid change:

TaskRate ↑  
 ToolActivity ↑  
 CouplingDensity ↑  
 ContextChangeRate ↑  
 Recursion ↑  
 ObjectivePressure ↑

- The agent's observations may remain individually valid.
- Its tool calls may succeed.
- Its reward may improve.
- Its collaborating agents may agree.
- Its local plan may remain logically coherent.
- Yet the agent may infer:
  - My trajectory is correct.
  - when, relative to the larger system:

IntentShear ↑

AuthorityMargin ↓  
CouplingGain ↑  
EnvelopeMargin ↓

- This suggests a new GUDIYA phenomenon: ‘Cognitive Somatogravic Illusion’
  - A condition in which rapid cognitive dynamics cause an actant to infer an incorrect system-relative orientation from locally valid information.
  - The agent does not necessarily hallucinate.
  - The model does not necessarily malfunction.
  - The software does not necessarily contain a bug.
- Instead:

*Local Evidence → Incorrect Global Orientation*

- That is a different class of failure.

## 5. Cognitive Acceleration May Therefore Matter

- The aviation analogy makes us reconsider Cognitive Acceleration.
- If:

$$V_C = \frac{dX_C}{dt}$$

represents some useful measure of Cognitive Velocity, then:

$$A_C = \frac{dV_C}{dt}$$

represents Cognitive Acceleration.

- Previously this could have been regarded simply as another navigational quantity.
- Somatogravic illusion suggests something more interesting:
  - Rapid change itself may degrade the reliability of locally inferred orientation.
- A candidate research hypothesis therefore becomes:
  - $P(D_C) = f(A_C, R_Q, L_C, \dots)$
  - where:
  - $D_C$  = Cognitive Disorientation,
  - $A_C$  = Cognitive Acceleration,
  - $R_Q$  = Reference Quality,
  - $L_C$  = Coupling/coordination load.
- Under some operating conditions we might discover:
  - $A_C \uparrow \Rightarrow P(D_C) \uparrow$
  - This remains a hypothesis to test, not an established law.
  - But it gives SEI a concrete experimental question.


## 6. The Most Dangerous Failure: Correcting the Wrong Problem

Somatogravic illusion reveals an even more important stability phenomenon.

Suppose a pilot incorrectly perceives excessive pitch-up.

The pilot attempts to correct it.

The corrective input may actually put the aircraft into a dangerous nose-down trajectory.



- The participant is not ignoring instability.
- The participant is actively attempting to stabilize the system.
- But because its attitude estimate is wrong:

*Stabilization Attempt → Destabilization*


- The cognitive equivalent could be extraordinarily important.
- An agent perceives:
  - The task isn't progressing quickly enough.
  - So it increases tool calls.
  - That increases coupling.
  - The increased coupling produces delays and inconsistent state.
  - The agent interprets the resulting delay as insufficient effort.
  - So it creates more subagents.
  - Coupling rises further.
- We obtain:

$Error_{perceived} \rightarrow Correction \rightarrow ActualError \uparrow \rightarrow StrongerCorrection \rightarrow ActualError \uparrow\uparrow$

- The agent is destabilizing itself while sincerely attempting to stabilize itself.
- That is precisely the kind of phenomenon Stability Engineering should care about.

## 7. Self-Reflection Does Not Necessarily Solve Disorientation

- Modern agent architectures often respond to mistakes with additional reasoning:  
Agent → Reflection → Critique → Replanning
- That can be valuable.
- But somatogravic illusion exposes its limitation.
- If the error originates in the agent's reference frame, asking the same cognitive system to reason harder from that reference frame may not restore orientation.



- A brilliant pilot can experience spatial disorientation.
- A brilliant agent can experience cognitive disorientation.
- The missing quantity is not necessarily intelligence.
- It is independent orientation information.

## 8. Why Another Agent May Not Solve It Either

- Perhaps we add a critic agent.
- Now:

$$\text{Agent}_A \rightarrow \text{Agent}_B$$

and B evaluates A.

- Better.
- But suppose both agents share:
  - the same environmental state,
  - similar assumptions,
  - the same objective,
  - common artifacts,
  - common reward signals,
  - the same rapidly changing local conditions.
- Then both can inhabit approximately the same moving reference frame.
- Agent B may confirm:
- Agent A appears correctly oriented.

Yet:

$$\text{Orientation}_A^{CCS} \approx \text{Orientation}_B^{CCS}$$

while both are drifting.

Thus:

$$\text{Agreement} \neq \text{Orientation}$$

and:

$$\text{Consensus} \neq \text{External Reference}$$

Adding more agents does not automatically manufacture a horizon.

## 9. The Agentic Gang Can Be Perfectly Coordinated and Wrong

- Consider:
 
$$A \leftrightarrow B \leftrightarrow C \leftrightarrow D \leftrightarrow E$$
- The agents communicate continuously.
- Their relative state may remain remarkably coherent:
 
$$\Delta X_{AB} \approx 0$$

$$\Delta X_{BC} \approx 0$$

$$\Delta X_{CD} \approx 0$$
- From inside the group:

InternalCoherence  $\uparrow$

- Yet relative to CCS:

$$V_{\text{Group}}^{\text{CCS}} \gg 0$$

- and:

$$\text{EnvelopeMargin}_{\text{Group}} \rightarrow 0$$

- The entire emergent organization is drifting together.
- This yields one of the more important propositions in GUDIYA:

$$\boxed{\text{Internal Coherence} \neq \text{External Alignment}}$$

- A thousand mutually agreeing agents can still be collectively disoriented.
- The group needs a horizon outside the group.

## 10. CCS Defines the Horizon

- The Cognitive Attitude Indicator therefore fits naturally into GUDIYA's developing Navigation Science.
- First:

### CCS — Cognition Coordinate System

- CCS answers:

*Relative to what are cognitive state and motion described?*

- Conceptually:

$$\boxed{\text{CCS} = \text{Reference}}$$

- This is analogous—not identical—to the external reference framework that makes navigation meaningful.
- Importantly, under GUDIYA's revised inductive genealogy, CCS does not require prior proof that a literal Cognitive Field exists.
- It follows from the engineering requirement for common reference among dynamically coupled actants.

## 11. CPS Determines the Agent's Referenced State

- Next comes:

### CPS — Cognitive Positioning Service

- CPS attempts to estimate:

$$X_A^{\text{CCS}}(t)$$

- and potentially:

$$\begin{aligned} &V_A^{\text{CCS}}(t) \\ &A_A^{\text{CCS}}(t) \\ &\text{Trajectory}_A(t) \end{aligned}$$

- with an associated:

$$\text{ReferenceConfidence}_A(t)$$

- Thus:

$$\text{CCS} \rightarrow \text{CPS}$$


- means:
- Define the reference, then determine the actant's relevant state relative to it.
- But one more function is necessary.
- The actant needs that information returned to it.

## 12. CAI Shows the Agent Its Attitude

- Thus:

### CAI — Cognitive Attitude Indicator

- The architecture becomes:



- Reasons and acts with Grid Awareness
- This is analogous in spirit to:  
Reference → Attitude Determination → Attitude Indicator → Pilot
- The CAI does not make the agent intelligent.
- It gives intelligence a horizon.

## 13. What Is Cognitive Attitude?

- We should resist over-literalizing aviation.
- An agent does not physically possess pitch, roll and yaw.
- The relevant dimensions of Cognitive Attitude remain a research problem.
- Candidates may include:
  - intent alignment,
  - authority alignment,
  - objective persistence,
  - coupling state,
  - recursion,
  - temporal alignment,
  - coordination burden,
  - trajectory relative to approved intent,
  - Cognitive Envelope margin.
- We can therefore initially define Cognitive Attitude abstractly:
  - Cognitive Attitude is an actant's orientation and dynamic trajectory across stability-relevant dimensions relative to an externally maintained reference.
  - Conceptually:  

$$\Theta_A(t) = g(X_A^{CCS}, V_A^{CCS}, A_A^{CCS}, \text{Trajectory}_A, \text{Envelope}_A)$$
- The dimensions must ultimately be discovered through experiment.
- The instrument should not depend upon pretending we already know them.

## 14. The CAI Versus the Agent's Internal Estimate

- The existence of CAI gives GUDIYA a particularly interesting new observable.

- Let:

$$\hat{\theta}_A^{\text{Local}}$$

represent the agent's internally estimated attitude.

And:

$$\theta_A^{\text{CAI}}$$

represent the Grid-referenced estimate.

- Then define:

$$\text{CAD}_A = D(\hat{\theta}_A^{\text{Local}}, \theta_A^{\text{CAI}})$$

where CAD is: 'Cognitive Attitude Disagreement'

- A small CAD means:

$$\text{Local Estimate} \approx \text{External Estimate}$$

- A growing CAD means:

$$\text{Local Estimate} \neq \text{External Estimate}$$

- This could become an important early-warning metric.
- Especially interesting would be:

$$\text{CAD} \uparrow$$

simultaneously with:

$$A_C \uparrow$$

$$\text{CouplingLoad} \uparrow$$

$$\text{EnvelopeMargin} \downarrow$$

- That combination might identify an approaching cognitive-disorientation event before an observable failure occurs.

## 15. The CAI Should Be Predictive

- An artificial horizon principally tells the pilot what the aircraft is doing now.
- Machine-speed cognition may require more.
- The Grid should attempt to estimate:

$$X_A(t)$$

- and:

$$\hat{X}_A(t + \Delta t)$$

- Thus the Cognitive Attitude Indicator might conceptually report:

| COGNITIVE ATTITUDE              |               |
|---------------------------------|---------------|
| Current Orientation             | NORMAL        |
| Intent Alignment                | NORMAL        |
| Cognitive Acceleration          | HIGH          |
| Coupling Load                   | ELEVATED      |
| Envelope Margin                 | 38%           |
| Cognitive Attitude Disagreement | RISING        |
| Projected +10 sec               | NORMAL        |
| Projected +30 sec               | CAUTION       |
| Projected +60 sec               | ENVELOPE EXIT |
| CPS Confidence                  | 97%           |

- The agent therefore doesn't merely know:
- Where am I?
- It begins to know:
- Where am I going?
- This connects CAI directly to another GUDIYA principle:
  - GUDIYA Must Fly Cognition Ahead of Cognition

## 16. CAI, CEM and CFD Are Different Instruments

We should preserve an important separation.

- Cognitive Attitude Indicator — CAI

Answers:

*How am I oriented and moving relative to external reference?*

- Cognitive Envelope Monitor — CEM

Answers:

*How much stable operating margin remains?*

- Cognitive Flight Director — CFD

Answers:

*What corrective trajectory should I follow?*

Thus:

CCS → CPS → CAI

- feeds: CEM
- and potentially: CFD
- before corrective action reaches: Agent

- The distinction mirrors a fundamental engineering separation:  

Observation  $\neq$  Guidance  $\neq$  Control  $\neq$  Enforcement
- CAI should fundamentally be an instrument, not an autonomous controller.

## 17. Guardrails Are Not an Artificial Horizon

- Guardrails answer:
  - May this action be performed?
  - Does this output violate policy?
  - Is this tool invocation permitted?
- CAI asks:
- What is happening to my dynamic orientation?
- An agent can remain perfectly policy compliant while:
  - drifting,
  - oscillating,
  - accelerating dangerously,
  - becoming excessively coupled,
  - pursuing a stale objective,
  - approaching an unstable rendezvous.
- Therefore:
 

Guardrail Compliance  $\neq$  Dynamic Stability
- Guardrails constrain behavior.
- CAI supplies orientation.

## 18. Monitoring Is Not Necessarily Attitude Awareness Either

- Traditional observability may report:
  - latency,
  - token consumption,
  - errors,
  - tool-call success,
  - CPU usage,
  - API failures,
  - queue depth.
- All can remain nominal while:
 

CognitiveTrajectory

is becoming unstable.
- That resembles an aircraft whose engines, hydraulics and electrical systems all report normal while its attitude is rapidly deteriorating.
- Health monitoring asks:
  - Are my components working?
- CAI asks:
  - Where is the system going?
  - Those are different questions.

## 19. CAI Converts VFR Agents Into IFR Agents

- This finally gives precision to an older GUDIYA distinction.
- **VFR Agent**
  - Primarily operates from:

### Local Observation + Internal Reasoning

- IFR Agent
  - Operates from:  
Local Observation + Internal Reasoning + Authoritative External Reference
- The CAI is therefore one of the instruments necessary to transform:
 

VFR Agent
- into:
 

IFR Agent

  - The Grid doesn't replace the agent's intelligence.
  - It gives that intelligence externally referenced orientation.

## 20. The CAI May Become a Standard Grid Instrument

- A Grid-attached actant could eventually consume a standardized CAI object:  

$$CAI_A(t)$$
 containing candidate quantities such as:

*{CFAddress, ReferenceID, Position, Velocity, Acceleration, Trajectory, Envelope, Margin, CouplingLoad, AttitudeDisagreement, ReferenceConfidence, Timestamp}*

- The exact schema remains research.
- But eventually an externally referenced cognitive-state object might become as ordinary as:
  - timestamps,
  - authentication tokens,
  - health checks,
  - trace headers.
- A consequential actant might even refuse certain interactions when:
 
$$ReferenceConfidence < Threshold$$
- or:
 
$$RendezvousQuality < Minimum$$
- or:
 
$$CAD > Maximum$$
- The ability to establish orientation becomes part of the conditions for safe coupling.

## 21. Loss of the Horizon Is Itself a Failure Condition

- This creates an important new category.
- We should distinguish:
 

AgentFailure

from:

PositioningFailure


from:

ReferenceFailure
- An agent may be functioning perfectly while CPS confidence collapses.
- Thus:
 
$$ReferenceConfidence \downarrow$$
 may itself justify:
 
$$Normal \rightarrow GridAwarenessDegraded \rightarrow RestrictedCoupling \rightarrow DynamicQuarantine$$

- depending upon the cognition grade and consequence level.
- An agent unable to establish reliable orientation should perhaps not continue unrestricted consequential coupling.
- Aviation already understands this principle.
- Loss of reliable navigation or attitude information changes what operations are safe.

## 22. The CAI Does Not Require the Cognitive Field Hypothesis

- CAI follows from the phenomenology-first pathway:



- The Cognitive Field can remain a separately testable higher-order hypothesis.
- The engineering case for CAI stands independently.

## 23. CAI Is Therefore Experimentally Testable

SEI could construct two otherwise identical agent populations.

### Population A — Conventional Agentic Architecture

Agents receive:

- prompts,
- context,
- memory,
- tools,
- reflection,
- local observations.

### Population B — CAI-Equipped Agents

They receive everything above plus externally generated:

- relative-state information,
- Cognitive Velocity,
- Cognitive Acceleration,
- coupling state,
- trajectory,
- envelope margin,
- Cognitive Attitude Disagreement.

Then expose both populations to:

- rapid environmental change,
- delayed feedback,
- implicit coupling,
- changing authority,
- persistent objectives,
- recursive delegation,
- conflicting local/global signals,
- sudden Cognitive Acceleration.

Measure:

EnvelopeExcursions  
UnsafeCouplings  
InstabilityRate  
RecoveryTime  
IncorrectCorrections  
CollectiveDrift

If CAI-equipped populations systematically perform better:

$$\text{Stability}_{\text{CAI}} > \text{Stability}_{\text{LocalOnly}}$$

then CAI begins moving from metaphor to empirical Stability Engineering.

## 24. Somatogravic Illusion Gives Us an Especially Important Experiment

- A more focused experiment could deliberately create the cognitive equivalent of acceleration-induced disorientation.
- Construct a scenario in which:
  1. local success signals remain positive;
  2. system-relative trajectory gradually deteriorates;
  3. the rate of environmental change increases;
  4. agents receive increasingly compelling local evidence that their corrective strategy is working.
- Then compare agents with and without external attitude reference.
- The critical question becomes:
  - Will locally oriented agents increasingly correct in the wrong direction while CAI-equipped agents detect disagreement between perceived and externally referenced attitude?
  - That would test the Cognitive Somatogravic Hypothesis directly.

## 25. A New Failure Signature

- Somatogravic illusion therefore suggests a candidate GUDIYA signature:  

CognitiveAcceleration ↑ +CAD ↑ +CorrectiveActionIntensity ↑ +EnvelopeMargin ↓
- This combination should attract Grid attention.
- Why?
- Because it describes an actant that is:
  - changing rapidly,
  - increasingly disagreeing with external reference,
  - responding more aggressively,
  - nevertheless moving toward instability.
- That is potentially the cognitive equivalent of a pilot pushing harder on the controls because the illusion is becoming more convincing.

- The agent may look extremely purposeful.
- The Grid sees:
- Purposeful destabilization caused by reference-frame error.

## 26. The Artificial Horizon Becomes More Important as Cognition Accelerates

- There is an apparent paradox.
- One might assume that increasingly capable agents need less external support.
- Aviation suggests the opposite can sometimes occur.
- Higher speed and more demanding operating environments increase the importance of reliable instrumentation.
- Likewise:

CognitiveVelocity ↑  
CognitiveAcceleration ↑  
CouplingDensity ↑  
MachineSpeedInteraction ↑

may make external orientation more important, not less.

- The question isn't:
  - *How intelligent is the agent?*
- It is:
  - *How rapidly is the agent moving through a dynamically coupled environment, and how trustworthy is its reference?*
- This changes the scaling argument for GUDIYA.
- The more powerful agentic computing becomes, the stronger the potential case for CAI.

## 27. The Artificial Horizon Moment for Agentic Computing

- We are currently giving autonomous agents increasingly sophisticated:
  - reasoning,
  - memory,
  - planning,
  - browsers,
  - APIs,
  - CLIs,
  - computers,
  - other agents,
  - robots,
  - persistent objectives,
  - real-world authority.
- Yet we largely assume they can determine their dynamic orientation from the information available within their local operating frame.
- In aviation terms, we are building increasingly powerful aircraft while still assuming the pilot can mostly look out the window.
- That assumption may work remarkably well in uncomplicated conditions.
- It becomes questionable when cognition enters:
- HCAS Operating at Machine Speed
- The question GUDIYA introduces is remarkably simple:
  - Who gives the agent its horizon?

## 28. The Deeper Principle

- The Cognitive Attitude Indicator ultimately rests upon a principle much larger than aviation.
- A dynamically moving participant cannot necessarily determine its global orientation from observations made entirely within its own moving reference frame.
- For cognition:

Self Observation  $\neq$  External Orientation

- Somatogravic illusion strengthens this further:

Valid Local Observation  $\neq$  Valid Global Inference

- And multi-agent systems add:

Consensus  $\neq$  Orientation

- Finally:

Intelligence  $\neq$  Reference

- These four propositions together form the conceptual foundation of CAI.

## Conclusion — Give Intelligence a Horizon

- The artificial horizon was not invented because pilots were unintelligent.
- It became essential because an intelligent human operating inside a dynamic system can form a compelling but incorrect estimate of orientation.
- Somatogravic illusion makes the lesson even sharper.
- The pilot's sensory information can be real.
- The pilot's reasoning can be sincere.
- The corrective action can be rational given the perceived attitude.
- And the resulting maneuver can nevertheless move the aircraft toward catastrophe.
- Autonomous cognition may possess an analogous vulnerability.
- An agent's:

Reasoning

may be sound.

Its:

Observations

may be valid.

Its:

Local Success

may be genuine.

Its:

Peer Consensus

may be strong.

And yet:

Its System-Relative Orientation May Be Wrong

- That is the problem the Cognitive Attitude Indicator is intended to solve.
- CCS establishes the reference.
- CPS determines referenced position and motion.
- CAI returns that orientation to the actant.
- CEM indicates stability margin.
- CFD can provide corrective guidance.
- Together:

CCS  $\rightarrow$  CPS  $\rightarrow$  CAI  $\rightarrow$  CEM  $\rightarrow$  CFD  $\rightarrow$  Grid-Aware Agent

- And perhaps the entire chapter can be reduced to one sentence:
- An intelligent agent can know what it is thinking, correctly perceive what is happening locally, and still be wrong about which way is level.
- That is why the future autonomous actant may need something present-day agentic computing largely does not provide:
- an artificial horizon for cognition.

