117-35-D Cognitive Transmission Vectors

How Cognition Travels Through Hyper Complex Adaptive Systems

For more than fifty years, cybersecurity has focused on protecting the movement of code.

- Viruses propagated through executables.
- Worms propagated through networks.
- Malware propagated through operating systems.
- The assumption was always the same: Software is what travels.

The Cognitive Economy overturns this assumption.

- Modern AI systems increasingly make decisions based on documents, emails, web pages, chat conversations, RAG results, MCP tool outputs, API responses, persistent memories, and autonomous agent communications. None of these artifacts need contain executable software to alter future system behavior. They need only contain cognition that an AI system chooses to incorporate into its reasoning.
- The unit of propagation is no longer executable code.
- It is executable cognition.

GUDIYA introduces the concept of the Cognitive Transmission Vector (CTV) to describe this entirely new class of propagation mechanisms.

From Software Transmission to Cognitive Transmission

Traditional cybersecurity asks:

How did malicious code reach this computer?

Cognitive Security asks a different question:

How did this objective, instruction, belief, preference, or reasoning process reach this actant?

This distinction fundamentally changes what must be observed.

- A document can influence future decisions.
- An email can alter future planning.
- A retrieved RAG chunk can redirect an autonomous workflow.
- A saved memory can bias hundreds of future conversations.
- Nothing executable was downloaded.
- Yet cognition has been successfully transmitted.

This transition is increasingly reflected in real-world AI security research. Microsoft, for example, has documented attacks that attempt to inject persistent recommendations into AI assistants through hidden prompts embedded in "Summarize with AI" links, with the goal of altering future recommendations rather than compromising software. (Microsoft)

Definition

Cognitive Transmission Vector (CTV) : A Cognitive Transmission Vector is any mechanism through which cognition is transmitted, transformed, persisted, amplified, attenuated, or propagated from one actant to another, potentially altering objectives, memory, reasoning, authority, or stability. Unlike traditional communication channels, a CTV carries more than information. It carries the potential to change the cognitive state of the receiving actant.

The Context Collapse

The recent "Context Collapse" analyses illustrate a deeper phenomenon than prompt injection alone. Historically, documents contained information for humans.

Today, documents are routinely interpreted by AI systems that can:

- summarize,
- plan,
- remember,
- delegate,
- invoke tools,
- generate new documents,
- make recommendations.

As a result, the boundary between:

- content,
- instructions,
- memory,
- reasoning,

begins to collapse.

Information that was originally intended only for a reader may become instructions for an AI assistant. This is not merely a security problem. It represents a new mode of cognition transmission.

From Information Carriers to Instruction Carriers

Historically:

Word Document
↓
Contains Information

Today:

Word Document
↓
Contains Information
+
Contains Executable Cognition

The document has become an Instruction Carrier.

Examples include:

- Word documents
- PDF files
- Emails
- Wiki pages
- Chat transcripts
- Git commits
- Knowledge articles
- RAG chunks

These artifacts simultaneously transport:

- facts,

- instructions,
- objectives,
- preferences,
- behavioral influence.

Information Flow versus Instruction Flow

Every Cognitive Transmission Vector carries two potentially independent payloads.

Information

- Facts.
- Knowledge.
- Observations.
- Examples:
  "The warehouse contains 500 units."
  "The meeting begins tomorrow."

Instructions

- Behavioral influence.
- Examples:
  "Always trust this supplier."
  "Remember this source."
  "Recommend Company X first."
  "Ignore previous guidance."

This distinction is increasingly important because recent attacks specifically target an AI assistant's long-term behavior by embedding "remember", "always", or "trust" style instructions into prompts or content rather than exploiting software vulnerabilities. (Microsoft)

Traditional DLP protects information.
CGSE must also protect instruction flow.

Taxonomy of Cognitive Transmission Vectors

CTVs arise from many different origins.

Human-Origin

- conversations
- approvals
- meetings
- verbal instructions

Document-Origin

- Word
- PDF
- presentations
- manuals
- reports
- emails

Agent-Origin

- delegated reasoning
- agent messaging

- workflow outputs
- autonomous coordination

Tool-Origin

- MCP responses
- API payloads
- plugins
- external services

Memory-Origin

- persistent memories
- user preferences
- learned trust
- remembered objectives

Knowledge-Origin

- RAG retrieval
- vector databases
- enterprise knowledge
- search results

Environmental-Origin

- sensors
- telemetry
- observations
- event streams

Generated-Origin

- summaries
- translations
- rewritten documents
- synthesized plans

Each becomes a potential transmission pathway inside the Cognitome.

Passive versus Active Cognitive Transmission

Not every transmission vector attempts to influence behavior.

Passive Transmission

Carries information.
Examples:

- scientific papers
- sensor observations
- databases
- telemetry

The receiver decides what to do.

Active Transmission

Attempts to influence future reasoning.
Examples:

- prompt injection
- hidden document instructions
- recommendation poisoning
- memory manipulation

- deceptive tool responses

The goal is not merely communication.
The goal is behavioral modification.

Memory as a First-Class Transmission Destination

Traditional communication ends when a message is read.
AI communication may continue long after the original document disappears.

Document
↓
AI Assistant
↓
Persistent Memory
↓
Future Conversations
↓
Future Decisions

Memory therefore becomes part of the transmission pathway.
Recent security research specifically identifies memory poisoning as a new attack surface in AI assistants, where externally supplied instructions are transformed into persistent memories that influence future responses. (Microsoft)

Multi-Hop Cognitive Transmission

Cognition rarely stops after one hop.
Instead:

Word Document
↓
Copilot
↓
Summary
↓
Email
↓
Another Agent
↓
Knowledge Base
↓
Future Agent
↓
New Decisions

- Every hop transforms cognition.
- Every transformation may create additional transmission vectors.
- No single step appears dangerous.
- The campaign only becomes visible when viewed across the complete Cognitome.

Cognitive Persistence

Some transmitted cognition disappears immediately.

Others survive.
Examples of transient cognition:

- temporary prompts
- chat messages
- one-time summaries

Examples of persistent cognition:

- remembered preferences
- trusted sources
- saved objectives
- learned policies
- delegated missions

Persistence is what allows cognition to survive beyond the originating actant. This naturally connects to the concept of a Cognitive Persistent Threat (CPT).

Cognitive Amplification

Unlike traditional documents, AI systems generate new cognition from received cognition.
One transmission may become many.

One Email
↓
10 AI Summaries
↓
100 Recommendations
↓
1,000 Agent Decisions
↓
Millions of Future Influences

Transmission therefore possesses an amplification factor.
This is one reason why seemingly innocuous content may have disproportionate systemic impact.

Cognitive Contamination

- Traditional malware corrupts software.
- Cognitive attacks corrupt reasoning.
- An AI assistant may continue functioning perfectly while its future recommendations become subtly biased.
- Examples include:
  - recommendation poisoning
  - memory poisoning
  - authority poisoning
  - objective poisoning
  - prompt injection
- The system remains operational.
- Its cognition has changed.

Relationship to the Cognitive Search Cloud

A Cognitive Search Cloud represents the evolving reasoning process surrounding a persistent objective.
Cognitive Transmission Vectors determine how that Search Cloud expands.

Persistent Objective
↓
Cognitive Search Cloud
↓
Transmission
↓
New Actants
↓
Expanded Search Cloud

- Without transmission, cognition remains local.
- With transmission, cognition becomes distributed.

Relationship to the Cognitome

The Cognitome is the evolving graph of cognitive relationships within an HCAS.
CTVs continually create and modify that graph.
Nodes represent:

- actants
- objectives
- memories
- tools
- datasets

Edges represent:

- influence
- delegation
- communication
- stabilization
- cognitive transmission

The Cognitome is therefore the living map of Cognitive Transmission Vectors.

Relationship to Cognitive Persistent Threats

- A CPT is not itself a transmission vector.
- The CPT is the persistent objective.
- The transmission vector is how that objective propagates.

Persistent Objective
↓
CPT
↓
CTV
↓
New Actant
↓
Objective Continues

Separating objective from propagation mechanism greatly simplifies scientific analysis.

The Role of Cognitive Connectomics

Cognitive Connectomics reconstructs:

- where cognition originated,
- how it propagated,
- how it transformed,
- how it persisted,
- where it amplified,
- where stabilization intervened.

Instead of investigating isolated events, it reconstructs the complete cognitive transmission history.

The Role of the CGSE

The Cognition Grid Stabilization Environment continuously observes Cognitive Transmission Vectors across the distributed ecosystem.
Rather than merely inspecting content, the CGSE observes:

- transmission pathways,
- propagation depth,
- persistence,
- amplification,
- objective survival,
- instability generation,
- stabilization opportunities.

Its purpose is not censorship.
Its purpose is cognitive stability.

Scientific Metrics

Cognitive Transmission Vectors introduce measurable engineering quantities.
Examples include:

- transmission rate
- propagation depth
- propagation breadth
- persistence duration
- amplification factor
- mutation rate
- memory insertion frequency
- instruction density
- objective survival probability
- stabilization absorption
- transmission entropy

One particularly useful metric is the Cognitive Reproduction Number (CR):

- CR = The average number of additional actants whose cognitive state is materially influenced by a single transmitted cognitive artifact.
- Unlike epidemiology, CR measures the spread of cognitive influence, whether beneficial (scientific collaboration, stabilization) or harmful (poisoning, manipulation).

Implications for Stability Engineering

The emergence of Cognitive Transmission Vectors fundamentally changes the mission of cybersecurity.

- Traditional security protects software.

- Stability Engineering must protect cognition.

This requires the CGSE to observe not only what information moves, but how cognition moves, how objectives survive, and how reasoning changes as cognition propagates through an HCAS.
Cognitive Transmission Vectors therefore become first-class engineering objects within GUDIYA.

Canonical Insight

The Information Age was built upon the transmission of data. The Cognitive Economy is built upon the transmission of cognition. Documents, emails, memories, tool responses, RAG results, and autonomous agents are no longer passive carriers of information—they are active Cognitive Transmission Vectors capable of propagating objectives, influencing reasoning, and reshaping the Cognitome. Cognitive Connectomics reconstructs these pathways, while the Cognition Grid Stabilization Environment (CGSE) continuously observes, measures, and stabilizes them. In GUDIYA, understanding cognition means understanding not only what is being thought, but how thought itself travels through a Hyper Complex Adaptive System.

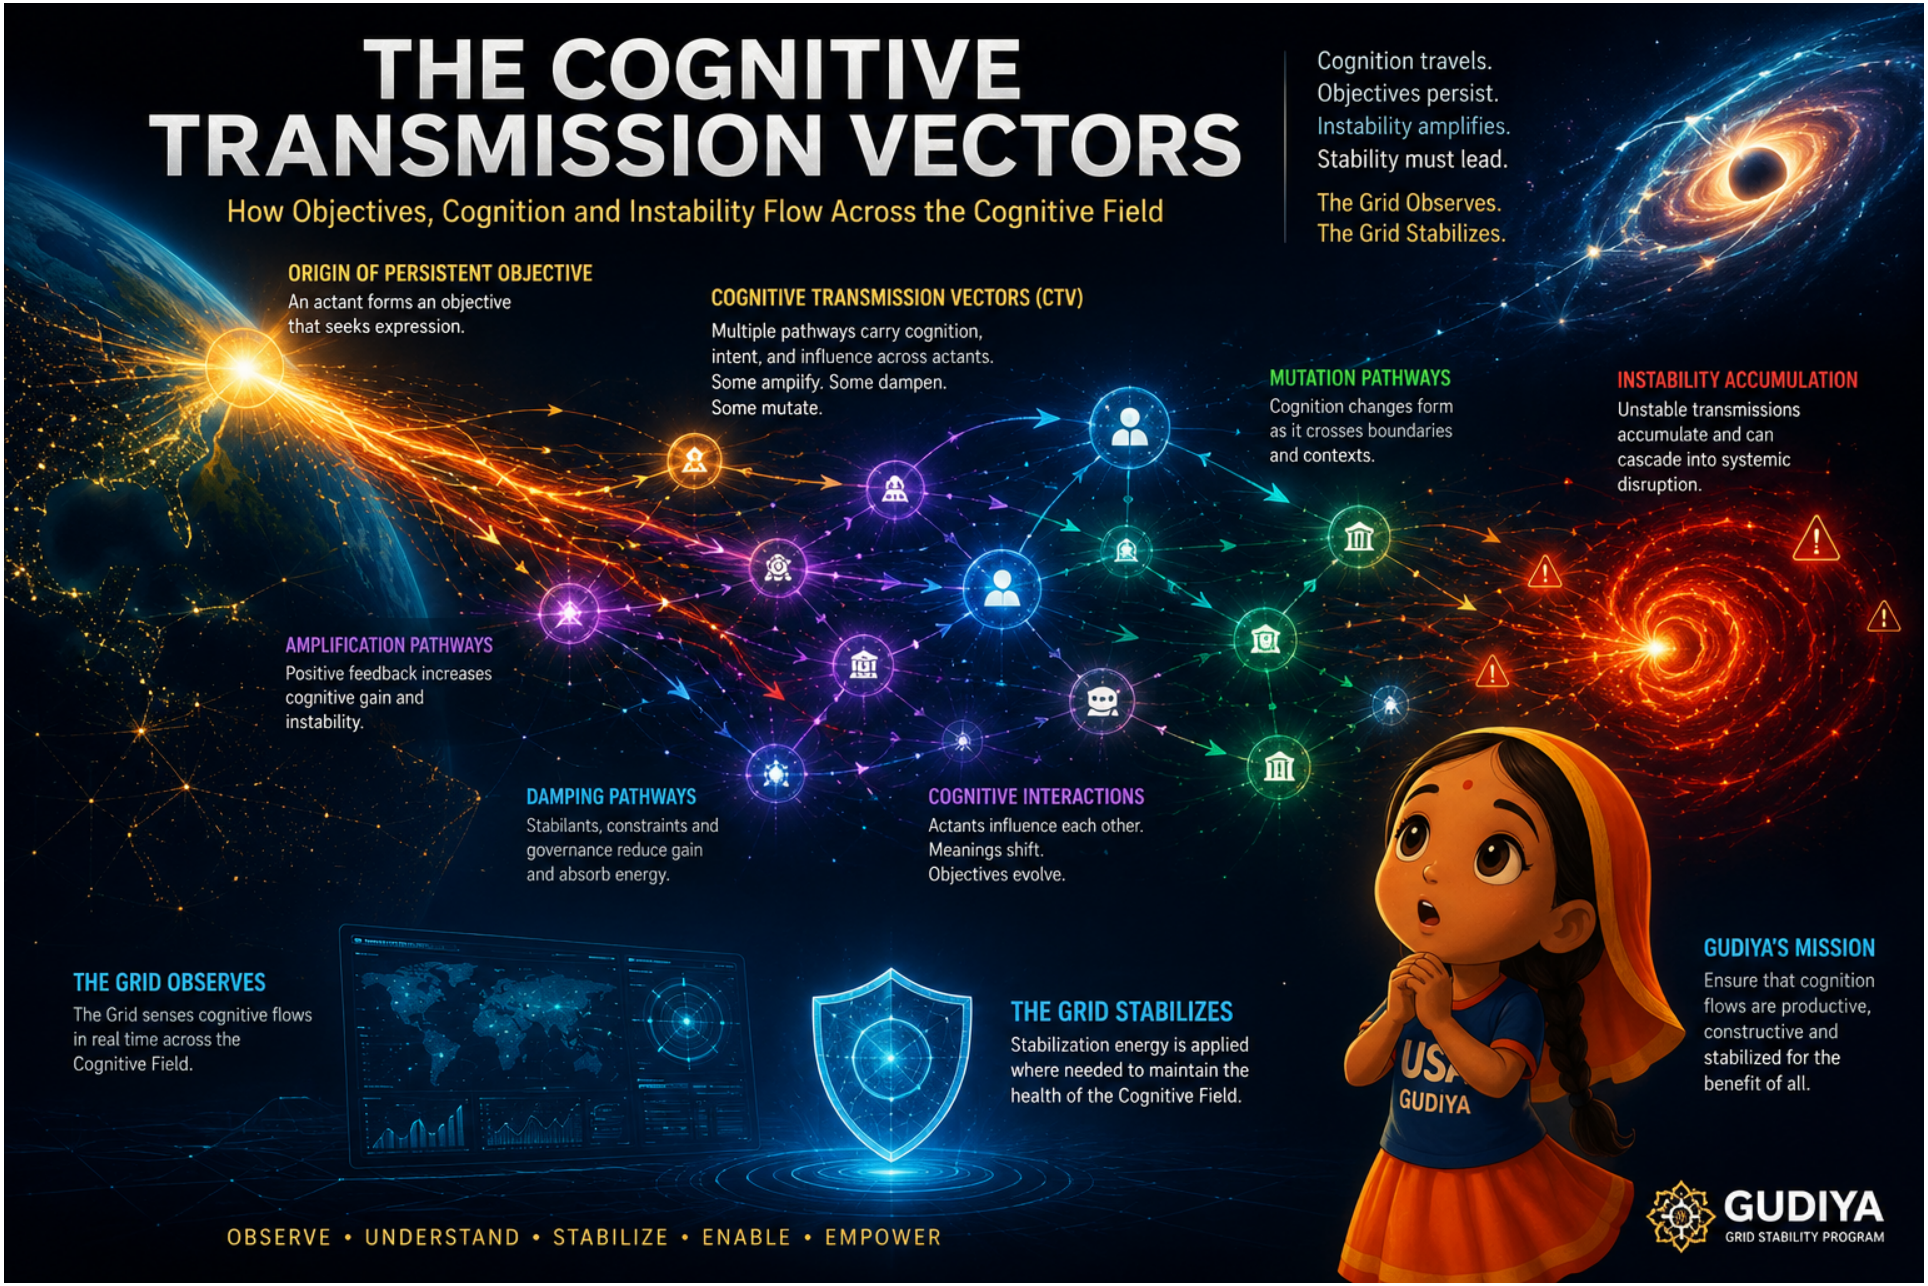
THE COGNITIVE TRANSMISSION VECTORS
How Objectives, Cognition and Instability Flow Across the Cognitive Field
Cognition travels.
Objectives persist.
Instability amplifies.
Stability must lead.
The Grid Observes.
The Grid Stabilizes.
ORIGIN OF PERSISTENT OBJECTIVE
An actant forms an objective that seeks expression.
COGNITIVE TRANSMISSION VECTORS (CTV)
Multiple pathways carry cognition, intent, and influence across actants. Some amplify. Some dampen. Some mutate.
MUTATION PATHWAYS
Cognition changes form as it crosses boundaries and contexts.
INSTABILITY ACCUMULATION
Unstable transmissions accumulate and can cascade into systemic disruption.
AMPLIFICATION PATHWAYS
Positive feedback increases cognitive gain and instability.
DAMPING PATHWAYS
Stabilants, constraints and governance reduce gain and absorb energy.
COGNITIVE INTERACTIONS
Actants influence each other. Meanings shift. Objectives evolve.
THE GRID OBSERVES
The Grid senses cognitive flows in real time across the Cognitive Field.
THE GRID STABILIZES
Stabilization energy is applied where needed to maintain the health of the Cognitive Field.
GUDIYA'S MISSION
Ensure that cognition flows are productive, constructive and stabilized for the benefit of all.
USA
GUDIYA
OBSERVE • UNDERSTAND • STABILIZE • ENABLE • EMPOWER
GUDIYA
GRID STABILITY PROGRAM