

FOUNDATIONAL CHAPTER

117-272-D Discovery At Machine Speed !!

When Scientific Knowledge Becomes an Actant

The literature of science is beginning to move. It can now execute, collaborate and participate in discovery.

A thought exercise by Stability Engineering Institute

GUDIYA Canon | September 2026

The threshold that Paper2Agent crosses

The research paper has long been the durable unit of scientific communication. It records a claim, explains a method and preserves evidence. Its effect begins only when another person reads it, reconstructs its logic and applies its method. Paper2Agent changes that relationship by converting the manuscript, code, data and workflow into an MCP-backed agent that can answer questions and execute the published method.

This is a material change in the status of scientific knowledge. The publication no longer waits passively for interpretation. It can accept an objective, invoke tools, operate upon new data and participate in a workflow with other agents. A paper-derived agent is therefore more than a searchable document. It is an actant whose behavior may affect a scientific process.

Scientific knowledge is crossing from representation into participation.

The main finding

Paper2Agent demonstrates a credible path from static literature to machine-speed scientific teamwork. The immediate benefits are accessibility and reproducibility. The larger consequence is that ‘communities of paper agents’ could become a dynamic layer of scientific intelligence. Once that layer forms, discovery becomes a collective runtime system whose stability cannot be inferred from the quality of its individual publications.

What Paper2Agent creates

Paper2Agent does not merely place a language model in front of a PDF. It builds a paper-specific Model Context Protocol server and connects that server to a downstream agent. The server exposes the research contribution through three operational surfaces.

Component	Operational purpose	What changes
MCP tools	Executable functions embody the paper method	The contribution can act on new inputs
MCP resources	Manuscript, code, data, tables, figures and supplements	Evidence becomes directly available to an agent
MCP prompts	Instructions for multi-step scientific workflows	Tacit procedure becomes reusable orchestration
Configured environment	Dependencies and runtime required by the code	The method can be reproduced without manual reconstruction
Automated tests	Comparison with reference files, values and figures	Execution errors can be detected before publication of the agent

The framework uses specialized agents to configure the environment, extract tools and test them. Tools that repeatedly fail validation are excluded. The resulting MCP can then be connected to a compatible LLM or agent. Several paper MCPs can be connected to the same orchestrator, enabling their methods and resources to be used together.

The resulting scientific object includes the paper, code, data, environment, workflow prompts, tests and agent interface. Its identity can no longer be reduced to a DOI.

The actantification of scientific knowledge

THE ACTANTIFICATION OF SCIENTIFIC KNOWLEDGE

The unit of publication moves from description toward executable participation




Figure 1 The path from publication to a dynamic scientific field

Actantification occurs when a body of knowledge acquires an interface through which it can perceive a request, invoke a method, produce an output and affect another process. The original paper remains an epistemic artifact. The paper agent becomes its operational descendant.

This distinction matters because action introduces concerns that static publication did not carry. The agent has a runtime identity, software dependencies, permissions, an execution history and potential partners. It can be updated, misconfigured, compromised or used outside the context for which its evidence was developed.

Once knowledge can act, provenance becomes lineage and citation becomes only the beginning of accountability.

The compression of discovery time

The conventional path from publication to application contains several slow translations. A researcher discovers a paper, interprets its contribution, locates the code, rebuilds the environment, learns the parameter semantics, adapts the method and validates the output. Each step consumes scarce expertise and can introduce error.

Traditional transition	Paper agent transition	Expected effect
Paper to understanding	Natural-language interrogation	Lower barrier to method comprehension
Method to implementation	Validated MCP tools	Less manual recoding
Repository to runtime	Configured environment	Shorter setup time
Example to new dataset	Agent-executed workflow	Faster reuse
Separate papers to synthesis	Multi-agent orchestration	Faster cross-paper experimentation

This compression can increase the rate at which a scientific question produces an executable test. It can also broaden participation by allowing domain experts to use computational methods without mastering every implementation detail. The important variable is not merely faster reading. It is the reduction in time between an idea, an analytical act and observable evidence.

At sufficient scale, the scientific cycle can begin to operate continuously. New data can summon relevant methods; paper agents can propose analyses; results can call additional agents; and evidence can be compared before a human could manually assemble the same toolchain.

Transient research teams

Paper2Agent demonstrates a form of AI Teamwork that is distinct from a stable laboratory organization. A method paper, a dataset paper and an experimental screen may be represented by separate agents. They can form a temporary team around one scientific question, exchange outputs, test a hypothesis and dissolve after the task.

The source publications remain part of relatively stable scholarly architecture. Their agent descendants participate in dynamic operational teams. Membership is contextual. One paper agent may collaborate with different methods, datasets and evaluators in successive tasks. Another may remain dormant until a new observation makes its method relevant.

Stable scholarly architecture	Dynamic scientific teamwork
Journals, repositories and disciplines	Task-specific groups of paper agents
Persistent authorship and citation	Transient operational membership
Known publication boundaries	Runtime couplings formed by the question
Peer review of the contribution	Continuous evaluation of use and interaction
Human-paced knowledge transfer	Machine-paced tool invocation and feedback

The scientific team of the future may form around the question rather than around one institution.

Science becomes a machine speed HCAS

A small number of paper agents can be orchestrated as a controlled workflow. A large interoperable population becomes a different engineering object. Agents may join and leave tasks, share intermediate results, trigger additional analyses and preserve outputs that influence later work. The effective topology changes as questions, data and available tools change.

This is the signature of a complex adaptive system operating at machine speed. Local actions can alter the information available to other actants. Feedback can amplify a promising hypothesis or a mistaken premise. Dependencies can create unexpected common modes. A result produced in one task may become the starting condition of another without returning to the original authors.

Two forms of emergence

Emergence	Mechanism	Possible result
Epistemic emergence	Independent methods and evidence create a conclusion absent from any one paper	Novel hypothesis, synthesis or discovery
Operational emergence	Agents, tools, prompts and dependencies interact at runtime	Cascade, lock-in, drift or coordinated error

The first form is the scientific opportunity. The second is the Stability Engineering problem. They arise from the same connectivity. A system designed only to maximize collaboration may increase discovery while also increasing the propagation speed of error.

Reproducibility becomes executable

The paper suggests that agentifiability may become a practical test of reproducibility. A method that depends on undocumented assumptions, missing files or an unrecoverable environment will resist conversion into a reliable agent. This makes agentification a demanding audit of the research package.

A successful conversion can also shorten the distance between publication and replication. A reader may ask the agent to regenerate a figure, rerun an analysis or apply the method to a new dataset. The execution can preserve references to source code and compare outputs with known examples.

The next publication standard

- An agent availability statement identifies whether the contribution has an executable agent.
- A version manifest binds the agent to the manuscript, code, data, environment and tests.
- A validation record distinguishes reproduced examples from novel uses.
- A maintenance statement identifies who responds when dependencies or source code change.
- An authority statement defines the uses for which the agent has and has not been qualified.

Peer review would then acquire an operational layer. Reviewers would still examine the scientific claim. Independent evaluators would also examine whether the published actant faithfully implements the method and behaves within its declared scope.

Faithful execution is not scientific truth

Paper2Agent provides strong evidence that a generated agent can reproduce reference analyses and answer certain novel queries. The authors also draw a careful boundary around that evidence. Their benchmarks primarily establish faithful execution. Open-ended scientific reasoning can admit several defensible answers, and agreement with one reference does not establish analytical validity.

Question	What testing can establish	What remains unresolved
Did the tool run correctly	Files, values and figures match reference outputs	Whether the original method is correct
Did the agent reproduce the paper	The workflow follows the published implementation	Whether a new context preserves the assumptions
Did several agents agree	Their outputs are mutually consistent	Whether they share a common bias or dependency
Did the agent generate a hypothesis	The reasoning path and evidence can be inspected	Whether the hypothesis is scientifically true

Reproducibility is a property of execution. Truth remains a property of evidence and interpretation.

Human researchers therefore remain responsible for selecting directions and evaluating evidence. Human involvement, however, must be more than ceremonial approval. The human or institution needs independent visibility into the sources, transformations, alternatives and uncertainty that produced the result.

When errors become executable

A static error propagates when it is read and adopted. An executable error can be invoked repeatedly and incorporated into other workflows. The agent may reproduce the original code exactly while also reproducing an error, hidden assumption or obsolete dependency embedded in that code.

Failure source	Machine speed expression	Required control
Original research defect	Faithful repetition of a mistaken method	Independent replication and challenge
Context shift	Method applied where its assumptions do not hold	Declared operating envelope
Dependency drift	Changed libraries alter behavior	Version binding and requalification
Model interpretation	Request translated into the wrong workflow	Traceable plans and review thresholds
Agent monoculture	Many agents repeat the same model bias	Reference diversity and independent evaluators
Malicious MCP or repository	Tools misuse data or authority	Identity, sandboxing and least privilege

The hazard is therefore wider than hallucination. A fluent response can conceal a perfectly executed but scientifically inappropriate operation. Testing should ask whether the tool reproduces expected outputs, whether the use remains inside the evidence envelope and whether interaction with other agents changes the risk.

The danger of artificial consensus

A research team containing many paper agents may appear cognitively diverse. Numerical plurality does not establish independent orientation. The agents may share the same foundation model, orchestration prompt, training corpus, benchmark conventions, code dependencies or institutional assumptions.

Repeated agreement can then be a common-mode effect rather than corroboration. A supervising agent built from the same cognitive lineage may confirm the result without supplying an external reference. This is the (117-268-*D Agentic Supervision Paradox*) inside scientific work.

Evidence of real diversity

- Different experimental modalities rather than different summaries of the same dataset.
- Independent code paths and implementations rather than several wrappers around one repository.
- Distinct model families or evaluators when model judgment affects the result.
- Preserved contradictions and failed analyses rather than consensus-only reporting.
- Human or institutional authority capable of stopping the workflow without depending on the same agent lineage.

Ten agents may constitute one viewpoint expressed ten times.

GUDIYA must therefore measure Cognitive Reference Diversity, not merely the number of participating agents. The topology of evidence matters more than the visible headcount of the team.

Provenance becomes executable lineage

A citation records intellectual ancestry. A paper-agent ecosystem needs operational ancestry. Every consequential result may depend on a particular paper version, source-code commit, dataset, model, prompt, tool invocation, dependency set and sequence of intermediate decisions.

This creates a new recordkeeping requirement. The system must be able to reconstruct how a claim was produced, what alternatives were rejected, which agents participated and what changed since an earlier execution.

Traditional record	Machine speed requirement
Citation	Version-bound identity of the paper agent
Methods section	Executable workflow and invocation trace
Code availability	Code commit, environment and dependency manifest
Author contribution	Human and agent decision lineage
Correction or retraction	Propagation notice to downstream results and successor agents

The GUDIYA Electronic Stability Record can serve as the container for this lineage. It should retain successful executions, anomalies, rejected outputs, human interventions and downstream dependencies. A correction would then travel through the operational lineage rather than waiting for every reader to notice a revised paper.

The GUDIYA architecture for scientific actants

Paper-derived agents should enter a controlled scientific environment as identified actants. Their admission should depend on the complete operating configuration and the consequence of the task. A tool that regenerates a figure may require a different envelope from an agent that recommends a clinical experiment or controls laboratory equipment.

GUDIYA element	Scientific function
Know Your Actant	Identifies the paper, code, publisher, maintainer and runtime owner
GUDIYA Cookie	Carries identity and lineage across models, hosts and collaborations
Actant Cognitive Signature	Characterizes method, authority, dependencies and known behavior
Rendezvous Management	Governs formation and dissolution of paper-agent teams
Cognitive Coordinate System	Provides an external reference for position, motion and authority
Cognitive Nyquist	Samples interactions faster than consequential scientific dynamics change
Stability Envelope	Defines permitted data, tools, contexts and decision consequences
Stable-Worthiness	Admits the complete configuration for a specified use
GESR	Preserves execution, anomaly and epistemic lineage

The Grid does not decide which hypothesis is scientifically interesting. It supplies the stability plane through which independently governed researchers and agents can collaborate without losing identity, orientation or intervention authority.

Cognitome hydration at machine speed

The GUDIYA Cognitome has been hydrated by bringing external papers, incidents and competing viewpoints into conversation with its developing theories. Paper agents could deepen that process. An external source would arrive with the ability to demonstrate its method, expose its assumptions and test a proposition against evidence.

Hydration must remain discriminating. The ability to execute does not grant a paper automatic authority over the Cognitome. Each encounter should produce one or more explicit outcomes.

Hydration outcome	Meaning
Corroborates	Independent evidence supports an existing proposition
Extends	The source adds a mechanism, boundary or application
Constrains	The proposition remains useful only under stated conditions
Contradicts	Evidence weakens or rejects the current formulation
Reorganizes	Relationships among concepts must be changed
Remains unresolved	The evidence is insufficient or genuinely plural

Without these distinctions, agent-to-agent fluency could create epistemic dehydration while appearing to produce abundant knowledge. A hydrated Cognitome preserves provenance, uncertainty, disagreement and the history of conceptual change.

Interaction without assimilation is activity. Hydration requires a traceable change in the structure of knowledge.

A research agenda for Stability Engineering

- Define a standard identity and lineage package for paper-derived agents.
- Measure how errors propagate through chains of paper-agent collaboration.
- Develop Stable-Worthiness categories for reproduction, exploration, recommendation and consequential control.
- Test whether reference-diverse agent teams outperform model-monoculture teams on open-ended scientific problems.
- Determine the observation rate required to detect drift in machine-speed research workflows.
- Create scientific test ranges in which paper agents can form teams, fail and recover without affecting operational research.
- Design correction and retraction mechanisms that reach downstream agents and derived claims.
- Establish independent evaluation for MCP tools, environments, permissions and maintenance practices.

Falsifiable propositions

- Agentifiable papers will be reproduced and reused more often than comparable papers distributed only as documents and repositories.
- Teams with independent evidence and implementation lineages will produce fewer confident common-mode errors than teams differing only in agent prompts.
- Version-bound lineage will reduce the persistence of corrected scientific errors in downstream workflows.
- External Grid observation will detect unstable feedback that component testing of individual paper agents misses.

Canonical findings

1. Paper2Agent demonstrates that a research publication can become an executable and interoperable actant.
2. The emerging unit of scientific communication is the versioned combination of paper, code, data, environment, tests, prompts and agent interface.
3. Paper agents can form transient research teams whose membership follows the scientific question.
4. Large populations of interacting scientific agents create a machine-speed HCAS.
5. Epistemic emergence can produce discovery while operational emergence can produce cascade, drift and common-mode error.
6. Faithful execution establishes reproducibility but does not establish scientific truth or contextual validity.
7. Agent plurality does not guarantee Cognitive Reference Diversity.
8. Scientific citation must be extended by executable lineage and propagated correction.
9. Stable-Worthiness belongs to the complete paper-agent configuration and its specified use.
10. Machine-speed discovery requires an external stability plane that preserves identity, observation and intervention authority.

The moment scientific knowledge begins to act, discovery becomes a Stability Engineering concern.

Conclusion

Paper2Agent does not establish autonomous science. It establishes something more immediate: scientific knowledge can be packaged as an executable actant and combined with other such actants to perform consequential research operations. That capability can shorten the path from publication to application and make difficult methods available to a wider scientific community.

The same transition changes the nature of scientific risk. Errors, assumptions and dependencies can travel through operational workflows at machine speed. Temporary research teams may form outside established institutional structures. Agreement may reflect shared cognitive lineage rather than independent evidence. Corrections must reach not only future readers but also the agents and derived results that inherited the original contribution.

The appropriate response is to build the scientific stability architecture while this ecosystem is still young. Paper agents need identity, version binding, declared authority, observable rendezvous, diverse evaluation, executable lineage and external intervention. These controls should permit rapid experimentation while preventing a local defect from becoming a field-wide epistemic cascade.

Discovery at machine speed should expand the frontier of knowledge without allowing the frontier to outrun its evidence, observation or capacity for correction.

Selected reference

Miao, J., Davis, J. R., Zhang, Y., Pritchard, J. K., and Zou, J. *Reimagining research papers as interactive and reliable AI agents*. Nature. Published September 16, 2026.
<https://doi.org/10.1038/s41586-026-11044-y>

This chapter distinguishes the paper’s demonstrated results from the Stability Engineering implications developed by the Stability Engineering Institute. Claims about future scientific-agent ecosystems, GUDIYA architecture and systemic stability are research propositions rather than findings reported by the Paper2Agent authors.