

Generative AI and the recursive constitution of operability

Rough closure semantic transduction and ecological hyperknots

Steven Watson

Cambridge EdgeLab working paper | Version 1 | 18 September 2026

This working paper develops the argument of Autopoietic Ecology AE2.6, revised 15 September 2026. It is offered for scholarly discussion and has not been peer reviewed.

Abstract

Generative AI can help constitute the conditions of further activity: a generated distinction becomes a question, a provisional diagnosis becomes an assessment premise, and a recorded decision conditions later generation. Explaining these relations requires an account of operability, understood as the organisation through which an occurrence can perform a determinate role. Developing Autopoietic Ecology, this paper distinguishes constitutive reciprocity, temporal renewal and the attribution of a particular autonomous unity. The generalisation of autopoiesis concerns the first; evidence for the other claims requires additional specification. Rough closure directs inquiry towards consequential conditions left unsettled by an operation. Semantic transduction follows the reconstitution of distinctions across computational, communicative and institutional settings. Ecological hyperknots connect these passages to the different continuations they sustain or undermine. A constructed educational assessment case shows how a locally accurate correction can coexist with the continued authority of an erroneous record, and how successful service delivery can consume the conditions of future participation. Comparative propositions identify observations that could support or weaken the proposed explanations. The contribution is a relational account of how AI participates in constituting operability, with explicit limits on claims of machine autonomy and on the empirical reach of the philosophical generalisation.

Keywords: autopoietic ecology; generative AI; operability; rough closure; semantic transduction; ecological hyperknots; education

1 Introduction

An AI-generated comment can become a reason for revising an assignment, a premise for a teacher's judgement or an entry in a student record. Those consequences depend on what happens to the comment. Its wording can remain unchanged while its evidential status and authority change. Later, the record can be retrieved as context for another generated response. The earlier output then participates in organising the conditions of subsequent activity. An account confined to the quality of either response leaves these relations unexplained.

Watson and Romic (2025) approached ChatGPT through the entangled evolution of society, education and technology. The distinction between instrumental and semantic mediation invited attention to how technology changes the possibilities of meaning and action. The paper also advanced an autopoietic account of technology. Subsequent discussion makes it necessary to clarify both the scope of that extension and the specific relations through which generative AI becomes consequential. The question concerns how technical operations participate in the constitution of capacities that neither the model nor its immediate user supplies alone.

Two interventions sharpen the problem. Zönnchen et al. (2025, section 3.1) challenge accounts of technological autopoiesis that leave self-production and re-entry insufficiently demonstrated. Their wider argument nevertheless recognises the significance of neural networks for socially patterned meaning production. Roth and Lien (2026) investigate relations between programmes whose outputs become consequential within organisational decision-making. These positions direct attention to different objects: an alleged technological unity, the operations of communication, and the organisation through which generated contributions acquire a role. They should be compared at those references, rather than treated as votes on a single property of AI.

The revised Autopoietic Ecology manuscript, AE2.6, supplies a further development of the original programme (Watson and Brezovec, 2026). Its general proposition that all operations are autopoietic concerns the constitution of operability. An organisation enables operations, while its operative reality is realised through relations of production, transformation and retention. AE2.6 distinguishes this philosophical claim from demonstrated temporal renewal and from an attribution of autonomous self-maintenance to a particular object. It also develops rough closure, ecological hyperknots and semantic transduction as ways of investigating differentiated conditions of continuation. This paper takes that formulation as its conceptual basis; the public first edition remains an earlier statement of the programme (Watson and Brezovec, 2025).

My central claim is that generative AI should be investigated through the relations by which its operations help constitute what can happen next. A generated proposal may reorganise the field of available responses. Its uptake may alter a practice's records or expectations. Those changes may return as conditions of further generation, assessment or decision. The resulting recursion can preserve some capacities while weakening others. Identifying these relations provides a positive account of AI's participation without making the model's autonomous existence a prerequisite for recognising its significance.

The argument is conceptual. It reconstructs selected theoretical positions, develops a constructed case and proposes comparisons for subsequent research. It reports no new participant data, deployed-system evaluation or systematic citation census. The originality claimed is the connection between constitutive explanation and a specified account of cross-setting continuation, correction and depletion. The paper does not claim that AE alone can discover these phenomena. Its contribution must be assessed against what an established systems-theoretical or infrastructural analysis can already explain.

2 Autopoiesis and the constitution of operability

The biological and social formulations

In Maturana and Varela's (1980) biological formulation, autopoiesis concerns a network of processes producing the components that realise the network and constitute its unity. The relation between cellular production and boundary formation is constitutive: the organisation through which production occurs is itself realised through that production. Environmental dependence remains compatible with this account. Autopoiesis does not require a cell to produce all the matter or energy on which it depends.

Luhmann's (1995) reconstruction changes the operational medium. Communication reproduces communication; its elements cannot be identified simply by listing individuals or transmitted messages. This extension raises substantive questions about what is retained from biological autopoiesis. Mingers (2002) examines the difficulties involved in applying the concept to social systems. The disagreement is instructive for technological extensions because it requires the relevant sense of production and the proposed unity to be stated explicitly.

Reichel (2011) develops an account of technology through operativeness and the distinction between working and failing. That proposal makes technological organisation a theoretical object in its own right. It does not establish that every device or model is an autopoietic unity. Tække (2026), in turn, analyses generative AI as nontrivial technology embedded in organisational decision structures. His distinction between the causal closure of a technical process and the stability expected of its outputs helps prevent unpredictability from being mistaken for self-production. These approaches leave an additional question: how does a technical capacity acquire the operative role through which it becomes consequential in a particular arrangement?

Constitutive reciprocity

Operability denotes the organisation through which an occurrence can perform a determinate role in a process. An authorised confirmation settles a request because procedures, records and recognition relations make it count as that kind of act. A computational transformation produces a usable output through an execution architecture, configured resources and specified relations among states. These capacities require explanation. Naming an institution or a model does not explain how it can do what is attributed to it.

AE2.6's argument begins without defining an operation as an event that successfully sustains itself. An operation is identified through its role, its conditions and the differences it makes. The further question concerns the constitution of that organised capacity. One account treats the capacity as an intrinsic power of an independently constituted thing. Another locates its organisation in an external source. AE gives explanatory priority to relations through which operations realise the organisation that makes operations of the relevant kind possible (Watson and Brezovec, 2026, chapter 3).

The positive reason for this preference is that the bearer of the capacity becomes part of the explanation. The authority of an assessment procedure is realised through consequential

institutional operations; it is not an additional substance contained in the procedure. An exclusively external explanation postpones the constitutive question when the supplying organisation also depends on operations. This does not logically refute every theory of primitive powers or every local feed-forward model. It motivates a non-essentialist account in which the constitution, exercise and alteration of capacities can be investigated together.

Constitutive reciprocity therefore differs from a claim that each event causes its own future repetition. An institution's final decision may dissolve it. The decision remains constituted through the relations that give it institutional force, even though it ends their continuation. Likewise, a retained material arrangement can preserve a capacity during an interval without actively replacing its components. The general philosophical claim concerns how operative organisation is realised. A temporal claim requires evidence about what actually renews, retains or depletes it over a specified period.

For generative AI, this distinction avoids an unproductive choice between declaring the model self-sufficient and treating its output as causally inert. A model invocation may be an externally maintained component at the scale investigated. Its operation can nevertheless participate in constituting the possibilities of a continuing practice. Whether the service itself regenerates its constitutive organisation is a further question. Whether it exercises agency, experiences meaning or bears responsibility requires further arguments again.

Three claims with different warrants

The first claim is ontological: operative organisation is constituted through the domain of operations it enables. This is AE's proposed generalisation of autopoiesis. Its defence rests on the explanatory argument just given and on the adequacy of its domain-specific interpretations. Examples can reveal strengths and difficulties, but a collection of successful cases cannot by itself prove universality.

The second claim concerns temporal renewal. A particular sequence replenishes, repairs or transforms identified conditions of further activity. This is empirically investigable. Researchers can examine whether correcting a record changes subsequent retrieval, whether training renews a capability, or whether repeated use consumes time that cannot be replaced. Repetition and renewal may diverge. A service can continue issuing responses while the capacities that sustain it deteriorate.

The third claim attributes a particular autonomous or autopoietic unity. Here the proposed operations, production relations, boundary and interval must remain explicit. The continued reproduction of an organisation does not establish the autonomy of each technical component it uses. Nor does the existence of mutually dependent continuations prove that their aggregate is a higher-order self-producing system. AE2.6 retains this discipline in its discussion of technology (Watson and Brezovec, 2026, chapter 12).

Table 1 *Claims about operability and their warrants*

Claim	Required account	What it does not establish
Constitutive reciprocity	How operations realise the organisation that gives them their determinate role.	Successful renewal over time or autonomous agency.
Temporal renewal	Which identified conditions are retained, repaired, replenished or depleted over a stated interval.	The universal thesis or the autonomy of every participant.
Autonomous unity	The operations, production relations and boundary of a specified self-producing organisation.	Independence from all environmental support or moral standing.

These distinctions preserve the demanding aspect of autopoiesis while acknowledging that AE has changed its theoretical scope. Biological component production cannot function as an unannounced universal test of the generalisation. Equally, the generalisation cannot certify every local claim of successful maintenance. Montévil and Mossio's (2015) account of closure of constraints offers a useful comparison because it specifies dependence among constraints at relevant timescales. Its biological analysis does not independently establish AE's universal proposition.

3 Semantic transduction and the role of generated contributions

For Luhmann (1995, pp. 143–148), communication involves the synthesis of information, utterance and understanding. A transmitted string does not specify that synthesis by itself. Subsequent contributions can make an understanding observable, including misunderstanding, disagreement or refusal. Institutional authorisation is consequently not a universal condition of communication: a student can contest generated advice before any teacher or institution endorses the exchange.

Esposito (2017, 2022) examines how algorithms participate in communication without requiring them to reproduce human intelligence. This permits investigation of generated contributions while leaving claims about machine consciousness unresolved. It also makes the organisation of uptake consequential. A useful response, an accepted proposition and an authorised assessment are different achievements. Their relations must be traced through what participants and technical arrangements actually do.

Meaning mediation names the arrangements through which something becomes intelligible, relevant and consequential. A prompt, a display, a record and an assessment convention can each condition the alternatives that become available. Semantic transduction specifies a passage within that broader mediation: a distinction or significance is reconstituted across operational settings (Watson and Brezovec, 2026, chapter 11). The same words may enter a new relation and acquire a different status. Conversely, different words may preserve a distinction that matters for further action.

Suppose a model describes a student's response as showing a possible misconception. In a conversation, that description can become a question to investigate. In a database, it can become a category attached to the student. In a later planning meeting, the category can become an accepted premise for allocating work. Nothing in the persistence of the phrase establishes that its evidential status survived those passages. The content, the warrant for using it and the consequences of its use have to be examined separately.

Semantic transduction is not necessarily distortion. A teacher may turn an ambiguous suggestion into a more accurate explanation; a structured record may preserve a qualification that would otherwise disappear. Dependable transformations are achievements of particular arrangements. The analytical question is which distinction survives a transition, which changes, and which grounds authorise the resulting use. This remains answerable to evidence. A mathematically false claim does not become correct because its institutional circulation can be explained.

The concept also prevents a premature division between an active human interpreter and a passive machine. Technical generation can offer a distinction the user had not considered. Retrieval can foreground one history and suppress another. Interface defaults can give a tentative statement the appearance of an established finding. Human judgement remains situated within these arrangements. Bender and Koller (2020) caution against inferring meaning from form alone; their argument does not settle every question about a computational contribution within a larger practice. Here the focus is the organisation through which contributions become consequential, rather than an attribution of private experience from linguistic performance.

A useful empirical description would follow four features of each passage: what is represented, what grounds are available, what role the representation is permitted to play, and what later operation uses it. These features need not remain constant. An authorised transformation can legitimately change a provisional judgement into a decision. The problem arises when a change in status occurs without the evidence or procedure needed to warrant it. Calling that passage transduction identifies a research object; explaining it requires the actual sequence and its alternatives.

4 Rough closure and ecological hyperknots

Determination and consequential incompleteness

An operation can produce a determinate result while leaving conditions of its use unsettled. A platform can store a precise category. That precision does not establish that the category expresses an appropriate judgement, that a correction can reach it, or that the people affected can obtain reconsideration. Rough closure concerns this relation between operational determination and consequential incompleteness at a stated scope (Watson and Brezovec, 2026, chapter 4).

AE2.6 distinguishes ontological dependence, selective representation and practical exclusion. A service depends on conditions it does not supply alone. Its records select some features of those conditions. Its procedures determine which differences can affect an outcome. These claims require different evidence. Dependence does not automatically establish ignorance, and an

omission does not automatically establish injustice. An analysis must show which omitted or unresolved difference matters, for whom, and through what relation.

Roughness is therefore not a degree of closure or a synonym for inaccurate computation. Suppose a teacher has already demonstrated that an automated diagnosis is mistaken. The result is known; the record is legible; the software stores exactly what it was instructed to store. If no available procedure can change the official entry before a consequential decision, the difficulty concerns the operative route from correction to use. More accurate retrieval of the same entry may reproduce the problem.

This account gives rough closure a discriminating role. The inquiry specifies a focal operation, an unsettled relation and a consequence within an interval. It can then ask what would count as resolving that relation. A successful correction may settle the present record while leaving inadequate review time unresolved. Roughness need not be eliminated everywhere for a particular repair to be effective. The relevant claim concerns the diagnosed problem, rather than the completion of an exhaustive representation of the world.

Distinct continuations in a hyperknot

An ecological hyperknot is a configuration in which distinguishable recursive organisations co-condition their continuation through consequential dependencies, without necessarily forming one higher-order autopoietic unity (Watson and Brezovec, 2026, chapter 6). Computational execution, communication, institutional decisions and embodied participation can enter such an inquiry while retaining their differences. The shared occurrence is not evidence that they perform one common operation.

The hyperknot concept adds a directed question to structural coupling. Which relations make further activity possible, and how does activity in one setting renew or impair those relations elsewhere? A model's classification may change the time required for review. A review decision may change the material retrieved by the next model invocation. A student's withdrawal may remove the opportunity for the dialogue on which an assessment practice depends. These relations vary in direction, timing and replaceability. Describing them all as influence would hide what an intervention needs to change.

Infrastructure research already shows why apparently straightforward technical activity depends on situated access and often unobtrusive work (Star and Ruhleder, 1996; Star, 1999). AE's proposed contribution is to connect that work to the constitution and continuation of specified operational capacities. It asks, for example, whether repeated teacher intervention merely compensates for a faulty record or also changes the procedure that will generate future records. The distinction matters for whether a local success renews the conditions of the practice or consumes a reserve supplied by its participants.

No collective good follows simply from persistence. An institution can maintain reporting continuity while staff become less able to sustain the work. A student record can remain stable while the student's opportunity to contest it diminishes. A hyperknot therefore has no single measure of viability unless the relevant continuations, interests and intervals are stated. The

analysis can explain how one continuity is secured at another's expense. Judging the resulting allocation requires additional reasons about participation, evidence and responsibility.

5 A constructed case of feedback and assessment

From a mathematical response to an institutional record

Consider a stipulated sequence in an educational setting. It is a conceptual example, not an account of a school, participant or deployed product. A student submits a mathematically valid solution using an unfamiliar method. A generative service produces a tentative suggestion that the student may misunderstand a transformation. The teacher is expected to review the advice before using it for assessment. A platform also offers to summarise recurring difficulties for future teaching.

In the first passage, the submitted response becomes computational input. The generated suggestion depends on the particular service, prompt and available context. Its existence establishes that a computational operation occurred. Its fluency does not establish a correct diagnosis, and its tentative wording does not guarantee that later uses will preserve that tentativeness. The teacher asks the student to explain the method. The student supplies a valid justification, and the teacher recognises that the diagnosis should be withdrawn.

In the second passage, a summary process has already entered the original diagnosis into a learner profile. The profile records a misconception category without the original qualification. In the stipulated arrangement, the teacher's correction updates the visible feedback but not the profile. When a later planning assistant retrieves that profile, the misconception returns as established context. A technically accurate retrieval thereby supplies a substantively misleading premise for a new generated recommendation.

The mistake is not explained fully by saying that the model hallucinated. The initial diagnosis was erroneous, but the later problem also depends on how uncertainty was removed, what the correction changed, and which record subsequent activity treated as authoritative. Replacing the original generator could reduce the incidence of such errors. It would leave the relation between correction and later retrieval unresolved unless that relation also changed.

Nor should the teacher's judgement be treated as infallible. An alternative sequence might reveal that the generated suggestion identified a genuine problem which a hurried reviewer overlooked. The appropriate distinction is between a warranted and unwarranted use of evidence, with procedures for reconsideration. Human and generated contributions are each assessed at their actual role. The example does not make biological identity the criterion of correctness.

Following the constitutive relations

At the scale of the conversation, the student's explanation changes what the teacher can reasonably take the original response to show. The exchange continues through the uptake and revision of contributions. At the scale of the record, an authorised procedure determines which changes count as amendments to the learner profile. At the scale of computation, storage and

retrieval determine which representation is made available to the next invocation. The correction can succeed in one of these relations while failing in another.

The learner profile has institutional force because practices of recording and consultation make it a premise for subsequent action. Those practices are realised through operations that the available profile, in turn, conditions. This is the constitutive reciprocity at issue. Whether repeated use maintains the accuracy and revisability of the profile is a further temporal question. An arrangement can reproduce its capacity to classify while degrading its capacity to classify justifiably.

The focal inquiry is deliberately bounded. It follows one assessment sequence through its next consequential use and the capacity to correct that use. It does not treat every person, cable and power station as a constituent operation of a single educational system. Further dependencies enter when they explain a specified difference. The teacher's available review time is relevant because it affects whether the correction occurs before the next use; the electricity supply need not be analysed in detail when it is held constant in the comparison.

Table 2 Passages in the constructed assessment case

Passage	Change of role or status	Evidence needed for repair
Response to generated advice	Student work becomes input to a tentative diagnosis.	A warranted assessment of the mathematical claim.
Advice to learner profile	A tentative diagnosis becomes a recorded category.	The stored status and grounds match the reviewed judgement.
Correction to later retrieval	A local revision must affect the record subsequently consulted.	The next relevant use reflects the amendment.
Repeated use to future capacity	Ongoing delivery changes the work available for review and dialogue.	A traced change in workload and the capacity to continue.

The table separates events and roles that an interface might display together. Its categories are analytical descriptions, not a mathematical model or a universal workflow. A different assessment arrangement could legitimately organise those relations differently. The requirement is that the account state how a contribution changes status and how a warranted correction becomes consequential.

Continuation that consumes its own support

Now extend the stipulated sequence across a teaching period. The teacher repeatedly discovers discrepancies and manually reconciles the visible feedback with the profile. Lessons continue and the institution records prompt completion of feedback. The apparent continuity depends on additional review work. If that work reduces the time available for dialogue, the service can remain productive while the pedagogical conditions of its use deteriorate.

This is a claim about a possible mechanism. Demonstrating it would require evidence about workload, displaced activity and the available alternatives. It cannot be inferred from the presence of AI. A different arrangement might save sufficient routine work to increase time for explanation and challenge. The ecological question concerns the distribution of the conditions of continuation under the arrangement actually investigated.

The case nevertheless reveals why a local performance measure can be inadequate. Counting completed feedback messages treats the service's output as the relevant continuation. Following students' opportunities to obtain an explanation changes the object. Following the teacher's capacity to sustain review over a term changes the interval. The resulting descriptions can disagree without any of them being numerically mistaken. They concern different capacities whose relations must be explained.

Repair and transformation

An effective repair in this case requires the correction to reach the record used in later decisions. It also requires a way to distinguish the superseded diagnosis from the current assessment. Preserving a history may be appropriate, but the history must not continue to function as an unresolved finding. A visible note saying that a correction was made is insufficient if subsequent retrieval still supplies the original category as current evidence.

A further intervention could change the summary procedure so that provisional diagnoses retain their status and a later amendment reaches dependent records. That intervention changes the organisation of future operability. It creates possibilities for correction that were previously unavailable. Its value depends on whether the new relation is used and sustained. Adding a field that reviewers cannot access would change the representation without establishing the required pathway.

Repair may also involve declining to continue a particular operation. An institution could suspend automatic profile updates while preserving the use of generated suggestions in dialogue. Such an interruption can sustain the wider possibility of educational participation. AE's account of continuation therefore cannot identify improvement with maximising the survival of every local process. Some operations should end, and the grounds for ending them require justification beyond their capacity to recur.

6 Comparisons that make the account investigable

The constructed case suggests comparisons that hold selected relations constant while changing others. Their purpose is to discriminate explanations of a particular outcome. They cannot directly adjudicate the universal philosophical thesis. The unit of study would be an episode linking a generated contribution to at least one subsequent use, with the relevant records and revision history identified. Longer observation would examine how the capacity for such episodes is sustained.

The first comparison changes the generator while retaining the review, recording and retrieval arrangements. If similar failures of correction persist across generators with different diagnostic

performance, the outcome cannot be attributed solely to the original model's error rate. Conversely, if an improved generator removes most relevant failures under the same arrangement, that is evidence for the importance of its contribution. The ecological account must remain capable of recognising a technical improvement rather than interpreting every outcome as an institutional problem.

The second comparison retains the generated suggestion while changing how its status is represented and used. One arrangement carries the distinction between a provisional diagnosis and a reviewed finding into later records. Another loses it during summarisation. Researchers would examine whether later decisions treat the two differently. A difference would support a claim about that transition. It would not establish that a status label alone is sufficient, because reviewers may ignore it or encounter it too late.

The third comparison follows a correction through different routes. A local amendment is compared with an amendment that changes the record actually consulted downstream. The outcome is whether a subsequent decision still relies on the superseded claim, with timing and legitimate record-retention requirements specified. This is more informative than asking whether the interface reports successful correction. It distinguishes a change to a displayed artefact from a change to the conditions of later action.

The fourth comparison concerns the distribution of work. Configurations achieving similar immediate feedback quality can be compared for the review effort and opportunities for dialogue they require over time. Researchers would record work transferred to students, teachers and administrators as well as work saved. A reduction in one participant's effort could coexist with an increase elsewhere. Such findings would support a distributional explanation only where the change in work can be traced to the arrangement and plausible competing explanations have been considered.

An exchange between models offers a complementary setting. Similar transcripts may arise from roles imposed afresh on each invocation, retained interaction histories, or a service that revises its own routing and memory. Researchers should examine the execution arrangement rather than infer organisation from the transcript. Removing an initial role instruction while preserving history could test whether differentiation depends on repeated external assignment. Tracing the regeneration of a memory relation would address a different claim. Neither result alone establishes an autonomous technological unity or the presence of experience.

These comparisons preserve two constraints from the earlier attribution problem. The operation, target and interval must remain stable while a claim is evaluated. A missing constitutive relation cannot be supplied by silently expanding the target to include whichever external actor performed it. Participation in another organisation's reproduction also cannot confer that organisation's autonomy on the participant. The constraints now support the positive explanation of operability instead of exhausting the paper's theoretical contribution.

7 A prospective research programme

An initial study could use consented workshops with constructed student work and controlled technical arrangements. It would examine whether the proposed comparisons separate outcomes before exposing consequential assessment to experimental variation. The cases would include correct and incorrect generated suggestions, successful amendments, abandoned exchanges and justified decisions to stop. Selecting only fluent or successfully completed episodes would conceal the work needed to make continuation possible.

Evidence would combine execution traces, successive versions of records, observations of communicative uptake and accounts from participants. Each source answers a different question. A log can show which file was retrieved; it cannot by itself show how a teacher understood it. An interview can report an interpretation; it does not prove that a correction reached a dependent record. Aligning these sources supports a more specific account than treating any one source as access to the entire process.

Five propositions organise the proposed inquiry. They are hypotheses about specified relations and remain untested in this paper.

P1. Where downstream uses omit an available distinction between provisional and reviewed claims, generated suggestions will more often be treated as established findings than in an otherwise comparable arrangement that preserves and makes that distinction usable. Evidence against the proposition would include equivalent treatment across the conditions or evidence that a different feature explains the change. A mere increase in the visibility of labels would not confirm it.

P2. Where a correction changes the record actually consulted in a later operation, continued reliance on the superseded claim will be lower than where correction changes only a separate display. This prediction is conditional on the later operation using the amended record. If readers knowingly prefer the old claim, the explanation concerns their grounds or incentives rather than an unpropagated update. That distinction should be coded before interpreting the outcome.

P3. Arrangements with comparable immediate output quality can differ in their capacity to sustain review because they distribute the work of reconciliation differently. Support would require a traced transfer of work and an effect on an identified capacity during the observation period. If differences are explained by prior staffing, task difficulty or participant experience, the claim about the arrangement would need to be narrowed or rejected.

P4. Recurrent use can help constitute the authority of a generated classification when later operations treat records of earlier use as grounds for further reliance. Confirmation requires tracing this return, not simply counting repeated classifications. If each use is independently warranted, or if the record has no effect on subsequent decisions, the proposed recursive mechanism is absent even though the category appears frequently.

P5. A repair that changes an operative pathway can restore a particular capacity without improving the generator's isolated performance. Support would be a change in subsequent warranted action attributable to the revised pathway with the generator held constant. If

equivalent gains follow from a simpler existing procedure, the substantive repair claim may survive while any claim for the added value of AE's description weakens.

Operational definitions should be agreed before a comparison. For example, treating a provisional diagnosis as established could mean that a later decision asserts it as a finding without supplying the missing review. A successful correction could mean that the next identified use reflects the revised status within the relevant interval. These definitions are tied to the educational purpose and applicable procedure. They do not create an overall autopoiesis score or assume that every disagreement signals failure.

A competent baseline is essential. Existing systems may already preserve evidential status, propagate amendments and support contestation. Comparing AE-informed inquiry only with an arrangement that lacks those functions would exaggerate its contribution. A useful evaluation would ask whether the framework identifies an overlooked dependency, improves the explanation of a failure, or helps choose between interventions under comparable information. If an ordinary process map produces the same diagnosis and intervention with less effort, that finding should limit claims for the additional terminology.

The research programme could test local explanations and expose weaknesses in the general account. It could not verify that all operations are autopoietic by observing a finite set of AI-supported episodes. That proposition remains a philosophical commitment whose explanatory reach is open to challenge. The empirical contribution would concern specified relations of constitution, renewal, transduction and repair.

8 Objections and theoretical limits

Does the generalisation make autopoiesis empty

If all operations are called autopoietic, the concept can appear to lose its ability to distinguish living or autonomous systems from externally maintained processes. This objection identifies a real cost of generalisation. AE changes the scope of the term. It must preserve domain-specific differences rather than use universality to erase them. The three claims distinguished in Section 2 are essential to that task: constitutive reciprocity, temporal renewal and a particular autonomous unity require different arguments.

The difficult case is a feed-forward device that executes a supplied rule once and is discarded. It does not demonstrate local self-maintenance. AE interprets the operative organisation as realised in connected material operations, but this interpretation does not establish that the discarded device renews its hardware or programme. Invoking its manufacturer would answer a different historical question. A critic may accept this description while rejecting the extension of the word autopoiesis. The dispute then concerns explanatory priority and conceptual scope, not an empirical discovery that the device secretly repairs itself.

The proposal would become empty if an operation were defined from the outset as anything that recursively enables itself and the general thesis were then announced as a discovery. AE2.6 avoids that route by identifying an operation through its determinate role and arguing separately

for reciprocal constitution. A critic can challenge the claimed reciprocity or propose a sufficient account without it. The present paper makes the relevant relations more explicit in the AI case, while acknowledging that these examples cannot settle the ontological disagreement.

Does the account add anything to existing approaches

Luhmann already distinguishes operational closure from environmental dependence. Infrastructure research already investigates access, maintenance and invisible work. Esposito already explains why algorithmic participation in communication does not depend on human-like intelligence. A paper that merely assembles those observations under new labels would contribute little.

The proposed addition lies in the connection among distinct explanatory tasks. Constitutive analysis identifies how an event acquires a role. Transduction follows changes in its meaning, warrant and permitted use. Rough closure identifies a consequential relation left unsettled at a declared scope. Hyperknot analysis follows how that relation affects different capacities to continue. In the constructed case, this connection explains why a corrected conversation, a functioning database and an apparently successful service can coexist with a continuing error and depleted review capacity.

That explanatory organisation remains a proposal. The comparisons in Sections 6 and 7 state how its practical value could be examined. AE has no entitlement to superiority simply because its vocabulary is broader. Its advantage would be demonstrated where connecting the levels changes a diagnosis or identifies an effective intervention that a narrower account missed. Equal explanatory performance would support the simpler description for the task at hand.

Does meaning return to a sovereign human subject

The account does not require a human interpreter to confer all significance on an otherwise meaningless technical sequence. Generated distinctions can reorganise what participants notice and consider possible. Organisational categories and inherited language also condition human judgement. Attributing every consequence to individual interpretation would obscure those relations.

It remains necessary to distinguish computational execution from communicative uptake, institutional authority and embodied experience. These distinctions neither settle machine consciousness nor reserve all meaningful participation for humans. They identify claims whose evidence differs. A transcript containing self-description is evidence that such language was generated; whether the distinction performs a constitutive role requires further investigation. The same demand for evidence applies to a researcher who declares in advance that no computational contribution could ever participate in meaning.

Does persistence provide an ethical criterion

An arrangement's ability to reproduce itself does not justify preserving it. The paper's educational examples assume that consequential assessments should be answerable to evidence and that those affected should have a meaningful opportunity to challenge error. These are

normative commitments, not deductions from autopoiesis. Their justification can draw on educational purposes and arguments about participation and responsibility beyond the scope of this paper.

AE contributes by making the conditions of those commitments investigable. A formal right of appeal can coexist with an inaccessible record or a deadline that makes reconsideration ineffective. A technically available correction can require work that a participant cannot supply. Explaining these relations helps locate an obstacle and assess a proposed repair. It does not determine every competing obligation or allocate responsibility solely from the location of an operation.

9 Conclusion

Generative AI participates in the constitution of operability when its operations help realise the relations through which further activity acquires a determinate role. Generated distinctions can become questions, assessment premises and records that condition subsequent operations. Their effects depend on how they are reconstituted across settings and on the conditions that make those passages possible. The model's autonomous self-maintenance is a separate attribution whose warrant must be specified.

AE2.6 provides the conceptual basis for distinguishing this constitutive claim from temporal renewal. Rough closure directs attention to consequential conditions left unsettled by a determinate operation. Semantic transduction follows changes in meaning, warrant and authority. Ecological hyperknots connect those changes to distinct continuations and to the distribution of the work sustaining them. Together, these concepts permit an account of successful output that remains sensitive to deteriorating conditions of participation.

The constructed assessment case develops one implication: correcting a contribution matters only through the operations by which the correction becomes consequential for the relevant use. A local repair can succeed while an earlier claim retains authority elsewhere. Conversely, changing a pathway can restore a capacity without changing the generator. The proposed comparisons make those explanations open to challenge. Developing the entangled-evolution argument therefore requires following how capacities are constituted, transformed and sometimes exhausted through the operations in which AI participates.

Status and data availability

This is version 1 of a Cambridge EdgeLab working paper, dated 18 September 2026. It supersedes the unpublished development draft titled *Locating autopoiesis in generative AI*. It has not been peer reviewed. The educational case is constructed, and the research propositions are prospective. No new participant data or empirical evaluation results are reported.

AE2.6 is cited as an unpublished revised manuscript dated 15 September 2026. Its arguments are distinguished from the publicly available 2025 first edition. Section and chapter references identify the version used; this paper does not announce publication of the revised book.

Author and AI assistance

Steven Watson, Cambridge EdgeLab. Correspondence: steve.watson84@gmail.com. EdgeLab is an independent research network; this working paper does not imply institutional endorsement by the University of Cambridge.

The argument develops the author's research programme in autopoietic ecology and meaning mediation. Generative AI supported literature retrieval, drafting and revision. Responsibility for the arguments and their presentation rests with the author.

References

- Bender, E. M., and Koller, A. (2020). Climbing towards NLU: On meaning, form, and understanding in the age of data. In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics* (pp. 5185–5198). Association for Computational Linguistics. <https://doi.org/10.18653/v1/2020.acl-main.463>
- Esposito, E. (2017). Artificial communication? The production of contingency by algorithms. *Zeitschrift für Soziologie*, 46(4), 249–265. <https://doi.org/10.1515/zfsoz-2017-1014>
- Esposito, E. (2022). *Artificial communication: How algorithms produce social intelligence*. MIT Press. <https://mitpress.mit.edu/9780262046664/artificial-communication/>
- Luhmann, N. (1995). *Social systems* (J. Bednarz Jr. and D. Baecker, Trans.). Stanford University Press.
- Maturana, H. R., and Varela, F. J. (1980). *Autopoiesis and cognition: The realization of the living*. D. Reidel. <https://doi.org/10.1007/978-94-009-8947-4>
- Mingers, J. (2002). Can social systems be autopoietic? Assessing Luhmann's social theory. *The Sociological Review*, 50(2), 278–299. <https://doi.org/10.1111/1467-954X.00367>
- Montévil, M., and Mossio, M. (2015). Biological organisation as closure of constraints. *Journal of Theoretical Biology*, 372, 179–191. <https://doi.org/10.1016/j.jtbi.2015.02.029>
- Reichel, A. (2011). Technology as system: Towards an autopoietic theory of technology. *International Journal of Innovation and Sustainable Development*, 5(2/3), 105–118. <https://doi.org/10.1504/IJISD.2011.043070>
- Roth, S., and Lien, V. (2026). When programmes disagree about programmes: Large language models, code poetry and the problem of programme–programme communication. *Information Technology & People*. Advance online publication. <https://doi.org/10.1108/ITP-02-2026-0241>
- Star, S. L. (1999). The ethnography of infrastructure. *American Behavioral Scientist*, 43(3), 377–391. <https://doi.org/10.1177/00027649921955326>
- Star, S. L., and Ruhleder, K. (1996). Steps toward an ecology of infrastructure: Design and access for large information spaces. *Information Systems Research*, 7(1), 111–134. <https://doi.org/10.1287/isre.7.1.111>
- Tække, J. (2026). Generative AI—the transgression of technology. *Systems Research and Behavioral Science*. Advance online publication. <https://doi.org/10.1002/sres.70045>

- Watson, S., and Brezovec, E. (2025). Autopoietic ecology: Rethinking systems, meaning, and matter (1st ed.) [Public preprint]. <https://doi.org/10.13140/RG.2.2.18462.24649>
- Watson, S., and Brezovec, E. (2026). Autopoietic ecology: Rethinking systems, meaning, and matter (AE2.6, revised 15 September 2026) [Unpublished manuscript].
- Watson, S., and Romic, J. (2025). ChatGPT and the entangled evolution of society, education, and technology: A systems theory perspective. *European Educational Research Journal*, 24(2), 205–224. <https://doi.org/10.1177/14749041231221266>
- Zönnchen, B., Dzhimova, M., and Socher, G. (2025). From intelligence to autopoiesis: Rethinking artificial intelligence through systems theory. *Frontiers in Communication*, 10, 1585321. <https://doi.org/10.3389/fcomm.2025.1585321>