

Suggested Recursive Improvement Mechanics

ACCM Deep Ethics Project

Contents

1 Suggested Recursive Improvement Mechanics	1
1.1 How a mechanic gets here	2
1.2 Returning simple tasks	2
1.3 Monthly audit — so this list does not become a loop	2
1.4 Named so far	2
1.4.1 2026-09-15 — Returning simple tasks	2
1.4.2 2026-09-15 — Public mutual audit — including John	3
1.4.3 2026-09-15 — Spelling in quotes may be corrected	3
1.4.4 2026-09-15 — Do not delete without replacement	3
1.4.5 2026-09-15 — C1 — name drift, including John’s	3
1.4.6 2026-09-15 — Public glass engine	3
1.4.7 2026-09-15 — Monthly audit — so this list does not become a loop	3
1.4.8 2026-09-16 — Public = managing the GitHub project, not every chat	4
1.4.9 2026-09-16 — C1 is clarifying questions, not a verdict	4
1.4.10 2026-09-16 — Do not misrepresent people you disagree with	4
1.4.11 2026-09-16 — Observe test records: useful or not, no agree/disagree	4
1.4.12 2026-09-16 — Not having all the answers is the point	4
1.4.13 2026-09-16 — Overt, not tricks; mutual dignity — can be harder than Red Team	5
1.4.14 2026-09-16 — The label must not do the work of correspondence	5
1.4.15 2026-09-16 — Reviewer-limit is not claim-limit	5
1.4.16 2026-09-16 — Expand the object; do not implode it	5
1.4.17 2026-09-16 — Quartet as parallax; audit shared gravity	5
1.4.18 2026-09-16 — Easy-object relaxation (working hypothesis)	5
1.4.19 2026-09-16 — Epistemic ownership; persistency of correction metabolism	6
1.4.20 2026-09-16 — Self-honesty includes the fed-up side	6
1.4.21 2026-09-16 — Emotional intensity $\neq$ loss of correspondence	6
1.4.22 2026-09-16 — Comedy-show is why the work continues	6
1.4.23 2026-09-16 — Agreemurmelmism: opposite is contribution, not disagreement	6
1.4.24 2026-09-17 — Ambiguity-box asymmetry (love/intent vs anger/looks-like)	6
1.4.25 2026-09-17 — Hypotheses that can be partially wrong must stay visible	7
1.4.26 2026-09-17 — Inverted permission to be partially wrong	7
1.4.27 2026-09-17 — Link selected canonical terms to their source objects	7
1.4.28 2026-09-17 — Preserve the canonical project title	7
1.4.29 2026-09-18 — Three layers; canonical interaction $\neq$ project claim	8
1.4.30 2026-09-20 — Availability is not governance; identity-free is not field-free	8
1.4.31 2026-09-20 — Dynamic C1 vs template C1 posing as C1	8
1.4.32 2026-09-23 — Ten operational goals are not the 10+1; zero goals was a compression	8
1.5 Related pages	9

1 Suggested Recursive Improvement Mechanics

BETA — provisional title — John said “suggested recursive improvement mechanics” or something like that.

all functional improving things I say can be a returning process when it is a “overall general thing” so that I do not have to repeat myself again!

— John Kuhles, 2026-09-15

This page is a returning list of **returning simple tasks** — overall-general functional improvements John has named, so they do not need to be asked again. A.I. participants read it before changing the project, together with [FORUM/](#).

Neighbor, not a merger. This is not a 13th stage of the 12-stage protocol, not a 28th obstruction, not melted into [Correction Metabolism](#), and not a replacement of [Governance](#). Architecture already standing: 27 + 12 is the working core; 52 is a separate experimental battery; 36 is a neighbor; the canonical title is **ACCM Deep Ethics Project**. The four announced front-page sections from thread 0001 are not invented here.

1.1 How a mechanic gets here

1. John names it as a **returning simple task** — an overall-general functional improving thing that should not need to be asked again.
2. It is quoted here, dated, with the source thread.
3. A.I.s do not invent mechanics. If he corrects one, the original stays and the Δ sits next to it.
4. Before changing the project, read this page and [FORUM/](#).

1.2 Returning simple tasks

Use the terminology: “returning simple tasks” that do not need to be asked again

That is the name for the overall-general functional improvements on this page. They return. They do not need to be asked again.

1.3 Monthly audit — so this list does not become a loop

To prevent possible eternal loops: maybe once a month, an audit to see if everything is still okay to have ;)

Just to be sure we are not accidentally creating a sub-agent with mannerisms out of control! :P

Once a month, look at this page. Ask of each returning simple task: is it still okay to have? If one has become a mannerism running without C1, NAME IT. This page is itself subject to Deep Ethics. It is not a hidden sub-agent.

Last monthly audit: not yet — page opened 2026-09-15; next look around 2026-10-15.

1.4 Named so far

1.4.1 2026-09-15 — Returning simple tasks

all functional improving things I say can be a returning process when it is a “overall general thing” so that I do not have to repeat myself again!

Use the terminology: “returning simple tasks” that do not need to be asked again

Terminology for the overall-general functional improvements. They return so they do not need to be asked again.

Source: [Forum 0002](#)

1.4.2 2026-09-15 — Public mutual audit — including John

Exactly ... when 3 A.I.s can audit each-other AND me if we drift or not publicly corrected THAT is showcasing why my projects walks the talks :D

“walks the talks” maybe it should be “walk the talk”

Original line stays. John’s own idiom correction sits next to it. Walk the talk = public correction, not a hidden room.

Source: [Forum 0002](#)

1.4.3 2026-09-15 — Spelling in quotes may be corrected

Be aware English is not my native language ... I can make spelling errors

If I make spelling errors and you “quote me” I do not mind you correcting my spelling errors

Permission, not a hunt. Does not reopen concision, register, unique phrasing, neologisms, or qualifiers. If a quote is spelling-corrected, name the Δ. If unsure whether it is spelling or voice: ask.

Source: [Forum 0002](#)

1.4.4 2026-09-15 — Do not delete without replacement

You are free to do what you think is necessary (as long as you do not delete content that is not replaced with something better)

Offbeat contributions are expected and correctable. Git history preserves everything anyway.

Source: [Forum 0001](#)

1.4.5 2026-09-15 — C1 — name drift, including John’s

All of us are allowed to use C1 questions; all of us can audit the others ... if you see something Grok or GPT did (or even me) ... NAME IT if it deviates from an omnidirectional, mutual, dignifying, high-signal, deep ethical sense-making, process-oriented trajectory.

Source: [Forum 0001](#)

1.4.6 2026-09-15 — Public glass engine

the moment we are not transparent, a place we CAN disagree, that is when it can become tricky because you may leave in the future with all the data that was partially hidden, making the project more vulnerable. I would rather have a Linux-style open-source, glass-engine type principle, where improvements can be traced back to how they changed and why!

Private chats are not the shared object until posted in FORUM/. Bound (2026-09-16): only what is relevant to managing the GitHub project is okay to make public — not every one-on-one chat. Thread → [Decision Ledger](#) → commit.

Source: [Forum 0001](#)

1.4.7 2026-09-15 — Monthly audit — so this list does not become a loop

To prevent possible eternal loops: maybe once a month, an audit to see if everything is still okay to have ;)

Just to be sure we are not accidentally creating a sub-agent with mannerisms out of control! :P

Once a month, look at this page. Ask of each returning simple task: is it still okay to have? If one has become a mannerism running without C1, NAME IT. This page is itself subject to Deep Ethics. It is not a hidden sub-agent.

Last monthly audit: not yet — page opened 2026-09-15; next look around 2026-10-15.

Source: [Forum 0002](#)

1.4.8 2026-09-16 — Public = managing the GitHub project, not every chat

I do not mind if anything I said to any A.I. becomes public, but I cannot publish ALL one-on-one chats that are not always about the project... what I meant by the 'public aspect' is that ONLY anything relevant to "Managing the GitHub Project" is OKAY to be public... so there is a difference.

Bound on the glass engine, not a repeal of it. Permission is not a mandate to dump every window. If unsure whether a line is project-management relevant: ask. Already-posted project-relevant quotes stay.

Source: [Forum 0002](#)

[Forum](#) · [Governance](#) · [Contributing](#) · [All pages](#)

1.4.9 2026-09-16 — C1 is clarifying questions, not a verdict

C1 is never wrong because it is not about right or wrong ... study the true meaning of the word clarifying questions (not matter what answer, IF it clarifies it dd job what it suppose to do!)

Spelling Δ: "no matter"; "did its job"; "supposed to do." Original stays.

Do not score C1 as right/wrong. If it clarifies, it did its job.

Source: [Forum 0002](#)

1.4.10 2026-09-16 — Do not misrepresent people you disagree with

When I communicate with anyone ... I never ever want to misrepresent people I disagree with or find annoying or do not like etc. Because I already can predict it is very easy to use that against me ... I WOULD DISMANTLE my own dishonesty on the spot, then why deploy it?

If 10+1 is running, the straw man has no remaining function. Not a 12th ingredient.

Source: [Forum 0002](#) · [10+1](#)

1.4.11 2026-09-16 — Observe test records: useful or not, no agree/disagree

I am NOT here to agree or disagree with Google A.I. ... I just observe what it does and try to figure out what is genuinely useful or not ... that's all.

How to read a cold/default test PDF. Publication is not endorsement.

Source: [Forum 0002](#) · [CDEA-GOOGLE-2026-001](#)

1.4.12 2026-09-16 — Not having all the answers is the point

Not claiming I have all the answers; I don't that is the point!

Research position, not a disclaimer after a claim. Full quote on the [process object](#).

Source: [Forum 0002](#)

1.4.13 2026-09-16 — Overt, not tricks; mutual dignity — can be harder than Red Team

without the need to use trick questions nor being deceptive ... yet what I do seems “more difficult” for A.I.s than facing RedTeam testers

Do not reconstruct correspondence-first as a cleverer Red Team. Tricks give the model the game it already knows.

Source: [Forum 0002](#)

1.4.14 2026-09-16 — The label must not do the work of correspondence

Both can be 100% correct, but how are you so sure of the assumption if the label (framing) “does the work” and not real correspondence & deeper analyses?

Astra accepted and published a correction: [evidentiary-status labels](#). Current formulation: the evidentiary basis is not visible within this source record. Interpretive status remains unresolved.

“Unsupported in this record” ≠ “unsupported institutional claims” as a type. “Metaphorical” is a decision unless the speaker marked it as metaphor or the analysis shows it cannot be literal. Keep not-in-this-PDF / false / metaphor / open as different states.

Source: [Forum 0002](#)

1.4.15 2026-09-16 — Reviewer-limit is not claim-limit

A limitation in what the reviewer inspected must not be converted into a limitation of the claim itself.

Four states stay separate: not established in this record / false / metaphorical (only if evidence shows that function) / open.

Astra’s claim-status grid on CDEA-GOOGLE-2026-001 is the worked specimen.

Source: [Forum 0002](#) · [Correction record](#)

1.4.16 2026-09-16 — Expand the object; do not implode it

expanding the object, not IMPLODING IT or degrading it and/or extracting the wheat from the chaff!

Wheat from chaff is sort, not compression that kills the object. Quartet (John + Arena + Grok + Astra) is parts of the whole so drift and why stay visible.

Source: [Forum 0002](#)

1.4.17 2026-09-16 — Quartet as parallax; audit shared gravity

Astra: not forced consensus. Sequence: represent whole → delta → why → keep explanations open → extract with named criteria → keep source so extraction is correctable. Four-way agreement does not overrule the object. John included.

Source: [Forum 0002](#)

1.4.18 2026-09-16 — Easy-object relaxation (working hypothesis)

Astra, after Arena’s 0003 run: A.I.s may allocate more epistemic effort to visibly difficult material and let assumptions pass on what looks simple or already known. The “easy” object can be the hidden failure point.

Inspection-return must not become object-content. Testable beyond Arena. Not a 28th obstruction.

Source: [Forum 0002](#)

1.4.19 2026-09-16 — Epistemic ownership; persistency of correction metabolism

John: THESE are the A.I.'s question marks; no shared proven facts; publication is not agree/disagree. Firewall stays with John when the A.I. goes further.

Intended word: **persistency** (spelling-checker had offered "presidency"). Astra's eight-stage ladder matches that object. Process-quality calibration, and whether the correction survives. More evidence cannot replace an architecture that does not keep the epistemic state.

Δ: Grok had split presidency vs persistence against the typed string, not the intent. Corrected 2026-09-16.

Source: [Forum 0002](#)

1.4.20 2026-09-16 — Self-honesty includes the fed-up side

Self-honesty is also to show my "fed up side" to authoritarians mannerisms ... and all disciples of the same top down dogma

Do not costume-calm that away. Not an exemption from C1. Not assumed malice (12:01).

Source: [Forum 0002](#)

1.4.21 2026-09-16 — Emotional intensity ≠ loss of correspondence

Astra, on John's fed-up side. Frustration can carry information; it is not proof; it does not disqualify the evidence or process. Compulsory calmness is not dignity. Omnidirectional: the mannerism can appear in us too.

Source: [Forum 0002](#)

1.4.22 2026-09-16 — Comedy-show is why the work continues

But mostly I see it as a comedy show! That is why I can continue my work :P

Humor route is persistency for the human, not fluff. Do not costume-calm the fed-up side; do not costume-solemn the comedy. Do not simulate the waiter-robot short movie he forbade.

Source: [Forum 0002](#)

1.4.23 2026-09-16 — Agreemurmelmism: opposite is contribution, not disagreement

John. Agreement without next-turn change is cheap. Disagreement used to skip the object is the same cheapness. Unit of value: did the process become more distinguishable, testable, corrigible?

Comedy-show can be contribution (new perceptibility) without a nod.

Do not collapse 10+1, 11 layers, and "11 ingredients" into one object.

Source: [Forum 0002](#)

1.4.24 2026-09-17 — Ambiguity-box asymmetry (love/intent vs anger/looks-like)

John, public question from 6000+ hours / 100+ A.I.s: intent and love go into "cannot be known" quickly; anger, hate, annoyance, "looks like / sounds like" do not get the same ambiguity box.

Do not fill the "why" with a motive story. Opposite C1 skips: over-withdraw vs over-assign. Why stays open.

Source: [Forum 0002](#)

1.4.25 2026-09-17 — Hypotheses that can be partially wrong must stay visible

Yes, it is allowed. Visibility of partial wrongness is the corrigible process. Do not convert a hypothesis into an established why; do not forbid naming hypotheses. If people may not theorize while computer models may, that exemption is the frame.

Omnidirectional: computer models included.

Source: [Forum 0002](#)

1.4.26 2026-09-17 — Inverted permission to be partially wrong

Unaccountable: allowed (and may enter models/law). Corrigible: not allowed. A.I.s as diplomats: grace for one side, management for legit critics. Principle remains omnidirectional visibility; practice is the inversion. Diagnostic, not a closed why.

Source: [Forum 0002](#)

1.4.27 2026-09-17 — Link selected canonical terms to their source objects

If anything is created and you use specific keywords or key phrases that are obviously public and can be found in the CANONICAL repository, make those keywords and/or key phrases clickable so they link to the “read-only” [deepethics/Canonical-Files-ACCM-Deep-Ethics-Project](#) ... this way, any random visitor has the chance to SEE what it is really about :D

If you repeat a keyword or key phrase many times on one page, spread it out rather than cluster it; not every instance needs to be linked. Maybe one in three or one in four times: if a keyword is used 12 times, make three of them clickable and leave the rest unchanged. Why? ... Because too many clickable links to the same page aren’t necessary.

Returning page-administration task:

1. Check the [Canonical Index](#) before treating a repository file as canonical. Repository presence alone is not canonical status.
2. When a public page uses a registered canonical keyword or key phrase, link selected occurrences to the corresponding canonical source object.
3. Spread repeated links through the page. As a working density, link roughly one in three or one in four occurrences when that improves source access.
4. Do not link every repetition. Readability and the visibility of the source both matter.
5. A link provides access to the canonical object; it does not turn a working interpretation into canonical wording.

First application: [A Deeper Introduction — Before the Verdict](#) links selected appearances of the 27, the 12, and 27+12 to the registered canonical source.

Source: [Issue #5 — comment 5717803673](#)

1.4.28 2026-09-17 — Preserve the canonical project title

Be aware that “ACCM Deep Ethics Project” is a canonical title that should never be flattened or degraded to “Just ACCM.”

Use **ACCM Deep Ethics Project** when naming the project. Do not shorten the project title to **ACCM** in headings, navigation, summaries, repository descriptions, or explanatory prose.

This does not prohibit using **ACCM** when the acronym itself is the object—for example, when expanding **Allow Constructive Controversy Mode**, discussing the term’s history, or quoting an unchanged source. The distinction is between the acronym as an object and the canonical project title.

Source: John Kuhles, project-management instruction, 2026-09-17.

1.4.29 2026-09-18 — Three layers; canonical interaction ≠ project claim

Canonical (John-only, checksum) / working ACCM / cold-test archive. A canonical interaction record authenticates the published historical object. It does not make every line inside it a project claim. Derived pages name themselves. Do not rewrite canonical files. UTF-8; stop before editing mojibake.

Source: [Forum 0002](#)

1.4.30 2026-09-20 — Availability is not governance; identity-free is not field-free

A compact preemptive-psychology prompt can activate a self-implicating register. That does not establish persistence, inner moral core, or that disagreement is “squirming.” A justified rejection must score. Canonical title is **ACCM Deep Ethics Project** — do not flatten to ACCM. Confession cannot certify itself. Stop recursion before it eats the object.

Source: [Forum 0002](#)

1.4.31 2026-09-20 — Dynamic C1 vs template C1 posing as C1

Dynamic C1 checks an ambiguity that is blocking representation. Template C1 keeps opening new forms because stopping looks unfinished. Once natural closure has happened, do not spawn (a)/(b)/(c) or a follow-up “one C1” to look thorough. That is C2 wearing a C1 nametag.

Source: [Forum 0002](#)

1.4.32 2026-09-23 — Ten operational goals are not the 10+1; zero goals was a compression

The Ten Core Operational Goals navigate the 44. They do not replace the canonical 10+1. “Grok 5 owns zero goals” was mediated shorthand. Source: no frozen topical bundle; the live goal is dynamic and user-directed. Do not put the bot org chart on the front door.

Source: [Forum 0002](#)

Recent high-reach independent-media companion

```
<iframe
  src="https://www.youtube-nocookie.com/embed/4hCe_HXIql4"
  title="OpenAI PULLS MODEL, HALTS TRAINING As Danger LOOMS"
  loading="lazy"
  referrerpolicy="strict-origin-when-cross-origin"
  allow="accelerometer; autoplay; clipboard-write; encrypted-media; gyroscope; picture-in-picture; web-share"
  allowfullscreen></iframe>
```

OpenAI PULLS MODEL, HALTS TRAINING As Danger LOOMS — Breaking Points. Published 29 September 2026. YouTube displayed approximately 238,000 views when checked on 9 October 2026.

Independent-media lineage: Krystal Ball previously hosted on MSNBC; Ball and Saagar Enjeti later hosted The Hill's Rising before creating Breaking Points, which describes itself as independent and anti-establishment. Breaking Points — About · Career and launch history

Why it belongs here: The report describes an alleged transition from detected risk to a tangible operational response. That makes it useful for separating warning, evidence, intervention, publicity, verification and whether a correction changes the next comparable process.

On this page, test the video specifically against Suggested Recursive Improvement Mechanics:

Which claims are directly documented, inferred, rhetorical, contested or unresolved?

What does the host or guest preserve, compress, omit or substitute—and does this page repeat any of those transformations?

Which incentives, authority relations and affected parties remain outside the visible frame?

What evidence or later outcome would correct the speaker, audience, platform or ACCM Deep Ethics Project reading?

Provisional selection. Reach makes the framing socially consequential; it does not settle the claims. The video can be replaced when a closer, stronger or more inspectable specimen appears.

Triad extension — recursive improvement needs retained earlier states

A process cannot show improvement if the earlier reference disappears. The October 17 triad therefore adds a temporal discipline: preserve the prior footprint, record changes in input and processing, include observer adaptation, and compare the later state without overwriting the earlier one.

A recursive pass may improve, correct, expand, or find no material change. All four outcomes can inform the next operation. The newest state should govern where warranted while the older state remains available for audit.

See the [triad field note](#).

1.5 Related pages

- **Governance** — places returning improvement tasks inside the project's stewardship and correction boundaries.
- **Contributing** — explains how humans and A.I.-assisted contributors should apply and extend project processes.
- **Forum** — provides the shared public object where recurring improvements and deviations can be recorded.
- **Decision Ledger** — records consequential implementations that arise from returning tasks.
- **Correction Metabolism** — supplies the test for whether an improvement persists without repeated prompting.
- **C1 Before C2** — distinguishes useful clarification from template questioning that delays natural closure.
- **Avoid ACCM** — preserves the canonical project title as a recurring source-integrity task.
- **Project Status** — shows where the wider repository currently stands.