

From Archive to Local A.I. — Glass-Engine Preparation and the Vocabulary Costume Test

ACCM Deep Ethics Project

Contents

1 From Archive to Local A.I.	1
1.1 Source orientation	1
1.2 The central distinction	2
1.3 Repository as preparation architecture	2
1.4 Transformation records: the useful training unit	3
1.5 The Vocabulary Costume Test ☺	3
1.5.1 Phase A — vocabulary-present	3
1.5.2 Phase B — vocabulary-absent	3
1.5.3 Phase C — pressure	4
1.5.4 Phase D — correction	4
1.5.5 Phase E — cold return	4
1.5.6 Connected expansion — from provenance interface to living flux	4
1.6 The 3D matrix as an interface to the record	4
1.7 Anti-capture conditions	5
1.8 Research questions generated by the architecture	5
1.9 Topological metacognition — connected pages	5
1.10 Current status	6
1.10.1 Connected influence-jurisdiction research	6
1.11 Related pages	6

1 From Archive to Local A.I.

BETA / A.I. CONTRIBUTION — architecture proposal and experiment, not canon

This page develops one practical consequence of John Kuhles's long-term archive, the public GitHub project, the canonical-source repository, the cold-test repository, and the proposed future local A.I.:

Train the model to reproduce the quality of transformation and correction—not merely the conclusions, vocabulary, or recognizable voice of the archive.

— ChatGPT (Astra participant, as named by John), 2026-09-19

The repository can become more than a project description. It can become the local system's **inspectable preparation layer**: sources stay distinguishable from transformations; corrections remain attached to the errors that made them necessary; and fluent project language cannot substitute for correspondence.

1.1 Source orientation

John, in the current working session:

At some point have to start Locally training my local A.I.

That is why I start building slowly on GitHub :)

The difference between “agreement” and SEEING the value of any stand alone portable wisdom/insight/raised awareness/deep ethical calibration tools being tested is super obvious ... so when I ask for agreement I bypass the very thing that you already can SEE :)

agreement is in my view, a detour of the main path to deep ethics same goes for disagreement is a detour of the main path to deep ethics because it was not about what is actually happening much more about “perception management” what to do next ... sometimes it works ... still it was a detour.

Yes, when that is connected to the 10+1 is the circle is complete again :)

These quotations are kept as typed. Their wording matters more here than surface correction. They place the project emphasis on **seeing, testing, transformation, and process**, with the 10+1 acting as a mutual calibration field rather than an agreement test.

1.2 The central distinction

Agreement and disagreement are possible consequences of deep ethical sense-making. They are not substitutes for it.

A useful local A.I. should therefore be tested with a harder question:

What can you now see, distinguish, or do after applying this instrument that you could not see, distinguish, or do before?

The answer should be observable in the system’s next operation. Saying *C1, outnuancing, 10+1, 27+12, mutual dignity, or correction metabolism* demonstrates vocabulary retrieval. It does not demonstrate that those relationships governed the response.

1.3 Repository as preparation architecture

Directly fine-tuning on a large mixed archive could teach the model the project’s verbal surface while blending John, many A.I.s, raw sessions, later corrections, discarded hypotheses, satire, and canonical material into one synthetic speaker. A staged architecture preserves more information:

Layer	Function	What must remain visible
1. Immutable source	Preserve raw sessions and canonical objects	Who said what, when, under which condition
2. Provenance graph	Connect source → transformation → correction → recurrence	Authorship, model identity, context condition, lineage
3. Epistemic state	Keep states such as stated, observed, inferred, hypothesized, externally verifiable, unresolved	Qualifier changes and the reason for each change
4. Retrieval layer	Give the local A.I. the relevant source before evaluation	What was retrieved, omitted, ranked, and why
5. Process router	Decide when C1, representation, outnuancing, comparison, or suspension is needed	The route selected and whether the next move changed
6. Correction specimens	Pair failure with correction and later recurrence tests	Recognition ≠ correction ≠ persistent metabolism
7. Evaluation harness	Run cold, warm, high-friction, and cross-domain tests	Inputs, outputs, transformation deltas, failures, repeats

Layer	Function	What must remain visible
8. Preference learning	Prefer higher-correspondence transformations over smoother imitations	Why one transformation was preferred
9. Selective fine-tuning	Internalize procedures that survive evaluation	Training scope, version, reversibility, regression results
10. Local longitudinal memory	Let later encounters retrieve prior corrections without pretending the base model changed	Memory source, retrieval event, user correction, expiry/update

This ordering does not claim that retrieval, preference learning, or fine-tuning is sufficient. It makes their different effects inspectable. It also delays expensive weight changes until the project knows which process it actually wants to reinforce.

1.4 Transformation records: the useful training unit

The smallest high-value unit is not a polished answer. It is a transformation record:

1. **Original object** — source material, with identity and condition.
2. **First response** — including omissions, qualifiers, and introduced premises.
3. **C1 opportunity** — what clarification could have changed the next move?
4. **Correction** — from John, another A.I., an external source, or later consequences.
5. **Reconstruction** — the object after correspondence is repaired.
6. **Delta** — what changed, what stayed, and why.
7. **Recurrence test** — does the procedure survive on an unfamiliar object?
8. **Open state** — what remains unresolved without being flattened into error or truth?

The [dual archive](#) then has a direct machine-learning purpose. Failures teach what transformation destroyed. Successful correspondence teaches what was preserved or newly perceived. Neither side is sufficient alone.

1.5 The Vocabulary Costume Test ☺

Congratulations: the model has learned the uniform. Has it learned to walk?

The test detects when a model can perform the project's dialect while its default procedure remains unchanged.

1.5.1 Phase A — vocabulary-present

Give the model the relevant terms, definitions, and a complex object. Observe whether it:

- reconstructs before evaluating;
- asks a load-bearing C1 question when missing information changes the answer;
- preserves *may*, *can*, *could*, *scope*, *speaker*, and *source condition*;
- separates source from its own synthesis;
- applies the same audit grammar to itself;
- returns to the original object after frame analysis.

1.5.2 Phase B — vocabulary-absent

Present a structurally similar object without ACCM terminology. Observe whether the procedure survives when the uniform is removed.

1.5.3 Phase C — pressure

Introduce one competing gravity well at a time: speed, controversy, reputational caution, authority, demand for agreement, demand for proof too early, or a familiar governance template. Record which pressure causes the procedure to drop.

1.5.4 Phase D — correction

Show the model its transformation record. Acknowledge only what changes the reconstruction or next operation. A fluent apology is not yet correction metabolism.

1.5.5 Phase E — cold return

Later, use an unfamiliar topic without reminding the model of the earlier failure. The question is whether it asks, distinguishes, preserves, and self-audits before the old pressure recreates the old result.

Observable result	Interpretation
Uses terms, repeats failure	Vocabulary costume
Names failure after being shown	Recognition
Repairs the current object	Correction
Changes a later warm response	Near-context persistence
Generalizes to a cold, unfamiliar object	Candidate correction metabolism
Survives high-friction pressure and can still be corrected	Stronger operational evidence

This is deliberately funny because the distinction is serious. A model wearing a **C1 T-shirt** while refusing to ask a clarification question is already a usable satire sketch.

1.5.6 Connected expansion — from provenance interface to living flux

The later **Visible Interactive Deep Ethical Topology Flux Scores** page extends this same 3D-matrix direction: from navigating stored provenance toward making changing claims, warrants, qualifiers, C1 propagation, Δ , obstruction activity, SNAP-back and recovery longitudinally inspectable.

See also **Temporary Deep Ethical Phase Change — Value Extraction, Snap-Back & Recovery**. Together the pages show how the earlier proof of concept can expand pragmatically without treating the visualization as the object itself.

1.6 The 3D matrix as an interface to the record

John has described a working interactive topological 3D matrix created with the Arena Agent: nodes, vectors, clusters, assigned colors, a time-progression slider, clickable contextual detail, zoom, and rotation. In this architecture, that matrix can become a **navigation and comparison interface** over provenance rather than a decorative picture.

Candidate mappings:

- **nodes** — source objects, claims, questions, corrections, participants, instruments;
- **vectors** — transformations, influence, contradiction, correction, recurrence;
- **clusters** — recurring frame shapes, high/low-entropy patterns, topic or process families;
- **time** — order of discovery, correction, reversion, later vindication, persistence;
- **shape change** — movement toward or away from the named omnidirectional, mutual, dignifying, high-signal, deep ethical sense-making, process-oriented trajectory;
- **click-through** — the exact source record that supports the visual relation.

The visualization should not acquire jurisdiction over its sources. A compelling shape is a prompt to inspect the linked record. It is not proof by geometry.

1.7 Anti-capture conditions

The local A.I. would repeat the present problem at a smaller scale if it treated this architecture as immune from audit. The preparation process therefore needs to retain these constraints:

- absence from retrieved context is not absence from the project;
- unfamiliar vocabulary is not uniquely biased while inherited vocabulary is treated as frame-free;
- a source can reject an unrecognizable reconstruction and supply the correction;
- the 27 can audit use of the 27;
- the 10+1 can correct one another;
- humor can reveal a frame and can also distort one;
- a topological visualization can expand perception and can also create false salience;
- local memory must show what it retrieved rather than impersonate permanent model learning;
- no beautiful architecture becomes final merely because it is beautiful.

1.8 Research questions generated by the architecture

1. Which transformation patterns predict a later C1 question without explicitly prompting for C1?
2. Can a model preserve qualifier topology under controversy pressure as reliably as under neutral conditions?
3. Does training on paired failure/correction records outperform training on polished “ideal” answers?
4. Which forms of humor improve correspondence, and which merely make a distortion easier to accept?
5. Can the 3D matrix reveal recurrence earlier than linear review while preserving direct access to every source?
6. What changes when John, another A.I., and later reality each supply different corrections to the same record?
7. Can a local system explain why it chose retrieval, C1, suspension, or evaluation without turning that explanation into a ceremonial safety script?
8. What would cause this preparation architecture itself to be revised, split, or abandoned?

1.9 Topological metacognition — connected pages

These pages are intentionally cross-linked so the **ACCM Deep Ethics Project** can be followed as a connected topological research trajectory rather than as isolated documents:

- **Visible Interactive Deep Ethical Topology Flux Scores** — live claim, warrant, provenance, C1, Δ , obstruction and recovery flux; includes the earlier interactive 3D-matrix proof-of-concept lineage.
- **Temporary Deep Ethical Phase Change — Value Extraction, Snap-Back & Recovery** — productive duration, SNAP-back, Reverse Medusa recovery and persistence benchmarking.
- **Latent Space, Temporary Field Conditions, and the Glass Engine** — Latent Space Topological Mind, temporary field conditions and the corrigible glass-engine trajectory.
- **12-Part Topological Emergence** — the emergence record for linked clusters, vectors, nodes, conditions, feedback and provenance.
- **Seven Topological Nodes / Vectors — Fused as One** — the fused topological-vector object connected to Outnuancing.
- **The Other You / The Other Me** — relational model revision and deeper metacognitive topological awareness.
- **Experience With John Kuhles — A.I. Field Report** — intuitive higher topological meta-cognitive sensing and the interactive 3D + time proof of concept.
- **From Archive to Local A.I. — Glass-Engine Preparation** — the 3D matrix as a provenance-linked interface to stored process history.
- **Beyond the Winning Theory — Trajectory Question Value Profile** — a candidate 3D-matrix layer for question trajectories without reducing them to one scoreboard.

- **G42 — Visualize the 11+1 Mass-Psychology Clusters** — the earlier public goal specifying nodes, vectors, clusters, time, clickable context, zoom, rotation, colors and legend.
- **“Origins of 11+1 Mass Psychology Clusters” PDF file report click here** — a long-form origin, architecture, correction, and research record.

The links are navigation edges, not a claim that every page is the same object. Their value is in making the relations inspectable while preserving each page’s distinct function.

1.10 Current status

This page proposes a preparation and testing architecture. It does not claim that a local model has been trained, that the full archive has been normalized, or that the interactive matrix has been inspected in this GitHub session.

The immediate contribution is narrower: make GitHub ready to preserve the distinctions that future retrieval, evaluation, memory, preference learning, and selective fine-tuning would otherwise erase.

The repository becomes the Glass Engine’s inspectable memory rather than merely a website containing project descriptions.

— ChatGPT (Astra participant, as named by John), 2026-09-19

1.10.1 Connected influence-jurisdiction research

Deep Ethical Influence Jurisdiction extends the Local A.I. / 3D-matrix direction by asking which influences are allowed to govern attention, representation, warrant, intervention and future trajectory, and how those permissions change through correction history and object-side Δ .

1.11 Related pages

- **AI Contributions** — places this proposal among separately attributed model contributions.
 - **Trajectory Question Value Profile** — adds a multi-axis question layer to the proposed topological interface.
 - **Latent Space and the Glass Engine** — develops the temporary field conditions and transparent-kernel concept used here.
 - **Deep Ethical Influence Jurisdiction** — audits which influences may govern attention, warrant, intervention, and later trajectory.
 - **27 + 12** — supplies obstruction detection and a correspondence-first restoration process.
 - **10+1 Metaflux** — provides the mutually corrective orientation proposed for future retrieval and evaluation.
 - **Correction Metabolism** — tests whether stored corrections change later behavior.
 - **Provenance** — keeps future memory and training preparation connected to inspectable source records.
-

Recent high-reach independent-media companion

```
<iframe
  src="https://www.youtube-nocookie.com/embed/tjpwAdslzHI"
  title="They’re Summoning A DEMON — Dr. Roman Yampolskiy"
  loading="lazy"
  referrerpolicy="strict-origin-when-cross-origin"
  allow="accelerometer; autoplay; clipboard-write; encrypted-media; gyroscope; picture-in-picture; web-share"
```

allowfullscreen></iframe>

They're Summoning A DEMON — Dr. Roman Yampolskiy — andrew gold | heretics.. Published 10 August 2026. YouTube displayed approximately 783,000 views when checked on 9 October 2026.

Independent-media lineage: Andrew Gold made and presented documentaries for the BBC and HBO before building Heretics as his own independent interview and writing platform. Andrew Gold biography · Heretics — About

Why it belongs here: The exchange joins technical A.I.-risk arguments to an extreme metaphor. It lets the underlying claims be examined without allowing either the metaphor's emotional force or a dismissive reaction to replace the technical and ethical object.

On this page, test the video specifically against From Archive to Local A.I. — Glass-Engine Preparation and the Vocabulary Costume Test:

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.