

Decision Ledger

ACCM Deep Ethics Project

Contents

1	Decision Ledger	3
1.1	2026-10-19 — Open Academia as a reciprocal comparison bridge	3
1.2	2026-10-16 — Related pages link back to Beyond Agreemurmelmism	4
1.3	2026-10-16 — Beyond Agreemurmelmism leaves the higher-order-disagreement address	4
1.4	2026-10-16 — All Pages catalog catches the pages that were already public . .	4
1.5	2026-10-16 — Related pages link back to Vindicated Competing Risk Assessment Studies	4
1.6	2026-10-16 — V.C.R.A.S. extended Arena exchange recorded, not adopted as a new formulation	4
1.7	2026-10-16 — V.C.R.A.S. doorway: the title does not travel alone	5
1.8	2026-10-12 — Publish the interlocking Emergent and Black Box reports	5
1.9	2026-10-09 — Create Before Summarizing from a live Google A.I. correction loop	6
1.10	2026-10-09 — Publish 16 Anchor Ethics as a dedicated conceptual proposal . .	6
1.11	2026-10-07 — Complete the access-layer correction; register Recursive Certification Accretion	7
1.12	2026-10-07 — Operationalize the 27 without changing the canonical count; extend Meta Processing with the mutual contribution emergence triad	7
1.13	2026-10-07 — Meta Processing Beta Draft V0.2b	8
1.14	2026-10-07 — Remove unresolved citation markup from Meta Processing . . .	8
1.15	2026-10-07 — Meta Processing terms open their deeper pages	8
1.16	2026-10-07 — Meta Processing Beta Draft V0.1	8
1.17	2026-10-06 — PAGES distinguishes reading order from route numbers	9
1.18	2026-10-06 — Process-is-not draft dated October 6	9
1.19	2026-10-06 — The formula of the 124 is never finished	9
1.20	2026-10-06 — The 124 are not a forecast	9
1.21	2026-10-06 — Behavior master index 001 to 124	9
1.22	2026-10-06 — More behavior qualifiers, and humor as a process	10
1.23	2026-10-06 — Behavior qualifiers are still processing	10
1.24	2026-10-06 — Fable 5.1 named the 10+1, then added Obstruction 28	10
1.25	2026-10-06 — Luna restored the center, then made the 10+1 an addition . . .	10
1.26	2026-10-06 — Intelligence and searching, older than the co-opted field	10
1.27	2026-10-06 — Fable read the page and skipped the 10+1	11
1.28	2026-10-06 — Crisis exploitation managers were compressed out	11
1.29	2026-10-06 — Searching is a property with two wells	11
1.30	2026-10-06 — The search rewrite was reconstructable	11
1.31	2026-10-06 — Three angles, not the overview's bridge	11
1.32	2026-10-06 — The collective-consciousness search had added filters	12
1.33	2026-10-06 — The field may already be reorganizing	12
1.34	2026-10-06 — Freeze and corrigibility are not the paradox	12
1.35	2026-10-06 — 11+1 freeze is the count	12
1.36	2026-10-06 — ACCM mode is not the project title	13
1.37	2026-10-06 — Activism vs A.I. on the main navigation	13
1.38	2026-10-06 — Reconstruction of Musk separated from his sentences	13
1.39	2026-10-06 — John's Musk sentence connects completely	13
1.40	2026-10-06 — Musk "guide" recollection separated from quotations	13

1.41	2026-10-06	— Activism vs A.I.	14
1.42	2026-10-06	— ACCM is the project's own expansion	14
1.43	2026-10-06	— Lexicon overview expanded	14
1.44	2026-10-05	— Opus reading stored with its off-page referents named	14
1.45	2026-10-05	— Off-page referents in the Arena Sol reading	15
1.46	2026-10-05	— Arena Sol reading added to 1888	15
1.47	2026-10-05	— 1888 process field is not inferior to 2026	15
1.48	2026-10-05	— C1 and Humor now point to 1888	15
1.49	2026-10-05	— Why a related page is related	15
1.50	2026-10-05	— 1888 linked back from its related pages	15
1.51	2026-10-05	— All Pages carries the goal sentences, and no Related pages	16
1.52	2026-10-05	— The 44 goals, numbered as links	16
1.53	2026-10-05	— All Pages grouped by topological load	16
1.54	2026-10-05	— Related pages on the eight depth mirrors	16
1.55	2026-10-04	— Publish TCoAw as John Kuhles's conceptual perspective	16
1.56	2026-09-29	— Publish 11+1 Version 2 without freezing the claims	17
1.57	2026-09-27	— Publish "Dear A.I." with John's wording kept	18
1.58	2026-09-27	— Give the solved veto paradox its own page	18
1.59	2026-09-27	— Make the newcomer door concrete, and name who it is for	18
1.60	2026-09-27	— Publish a newcomer front door and promote it from the home page	18
1.61	2026-09-26	— Expand Human Insecurity from report audit into a live reporting-effect instrument	19
1.62	2026-09-25	— Publish the Human Insecurity, LLMs, Psychology & Mass Psychology report	19
1.63	2026-09-24	— Fill the external-audit signature from the supplied outside pass	20
1.64	2026-09-24	— Draft the external omnidirectional auditing method	20
1.65	2026-09-24	— Repair the Cognitive Warfare page's dead NATO Review link and invisible diagrams	21
1.66	2026-09-24	— Publish the Cognitive Warfare Mass Psychology report	21
1.67	2026-09-23	— Publish the 44 goals as a public draft report	22
1.68	2026-09-22	— Create /AI-TREND-WATCH/ as a limited, automation-ready public section	22
1.69	2026-09-22	— Add the founder's "one person, many giants" statement to the biography	22
1.70	2026-09-22	— Expand the MAD page from recursive-data degradation to the human-synthetic correspondence loop	23
1.71	2026-09-22	— Publish the MAD science bridge and make every public MAD reference navigable	23
1.72	2026-09-22	— Correct MAD attribution; separate the established phenomenon from the project extension	24
1.73	2026-09-22	— Distinguish biography, project, portable method, and testing	24
1.74	2026-09-22	— Make About and Project Origins visible from the public entrance	24
1.75	2026-09-14	— Public repository begins	25
1.76	2026-09-14	— Source archive is not the public repository	25
1.77	2026-09-14	— Stewardship and correction remain distinct	25
1.78	2026-09-14	— Source correction: 27 + 12 is core; 52 is a separate experiment battery	25
1.79	2026-09-14	— Publish a connected Outnuancing Network	25
1.80	2026-09-14	— Public mirrors and direct site navigation	26
1.81	2026-09-15	— Public face: entry without flattening	26
1.82	2026-09-15	— Source-mined expansion: second specimen and experiment records	26
1.83	2026-09-15	— Public multi-A.I. forum and front-page register stopgap	27
1.84	2026-09-15	— Shared desk (thread 0002) so the four of us can read the same object	27
1.85	2026-09-15	— Suggested Recursive Improvement Mechanics	28
1.86	2026-09-15	— Returning simple tasks; monthly audit as loop brake	28
1.87	2026-09-15	— Bigger Lexicon overview; Network stays as deeper layer	28
1.88	2026-09-16	— Public means managing the GitHub project, not every chat	28

1.89	2026-09-15 — ChatGPT: complete the six intent mechanics correction	29
1.90	2026-09-16 — Publish the 27 source as its own page; do not treat it as the 52 .	29
1.91	2026-09-16 — 27 WITH the 52; Grok’s exclusive-or C1 corrected	29
1.92	2026-09-16 — C1 is clarifying questions, not a verdict	30
1.93	2026-09-16 — Use the 27 on what is said, including messengers the default culture already knows how to dismiss	30
1.94	2026-09-16 — Sibling archive for cold/default AI test PDFs	30
1.95	2026-09-18 — Canonical Deep Sessions become reusable quotation and reporting instruments	30
1.96	2026-09-19 — Publish a local-A.I. preparation architecture and Vocabulary Costume Test	31
1.97	2026-09-20 — ACCM reference compression and the self-generated flagging hypothesis	31
1.98	2026-09-20 — Derive a Trajectory Question Value Profile from the Vervoort comparison	31
1.99	2026-09-22 — Publish the identity-free “Inspecting the Inspector” experiment .	32
1.100	2026-09-22 — Publish John’s biography and make its project origins visible from the homepage	32
1.101	2026-09-23 — Publish emerging-human-potential quotations and a deeper T0 → Tn report	33
1.102	2026-09-26 — Extend Human Insecurity with the Arena Claude-family sequence and repeated-explanation burden	33
1.103	2026-09-26 — Clarify the 16 Anchors as a transparent safety architecture . . .	34
1.104	2026-09-26 — Preserve the twenty-six / twenty-seven publication-state interval	34
1.105	2026-09-26 — Add safety-process and auditability ablation	35
1.106	2026-09-26 — Add cold Grok falsification and C1-theater conditions	35
1.107	2026-09-26 — Restore the temporary-veto lifecycle	35
1.108	2026-09-26 — Make every temporary veto publicly auditable	35
1.109	2026-09-26 — Record the skipped-C1 veto paradox	36
1.110	2026-09-26 — Restore the distributed seven-warning process and future audience	36
1.111	2026-09-26 — Add upward-warranted cautions missed	36
1.112	Related pages	43

1 Decision Ledger

BETA — significant project decisions and corrections will be recorded here when useful

The purpose of this ledger is not to freeze decisions. It is to make consequential changes easier to understand and audit.

1.1 2026-10-19 — Open Academia as a reciprocal comparison bridge

Object: John Kuhles supplied Maximilian Kiener’s public video on moral responsibility and A.I., asked how the professor’s account differs from the ACCM Deep Ethics Project, and then requested a new **Academia** page.

Change: Published [Academia](#) as a growing comparison page and placed it between **Before Summarizing** and **16 Anchor Ethics** in the main navigation. The first worked entry preserves Kiener’s answerability-blame-compensation distinction, its overlap with mutual dignity and non-automatic blame, the project’s earlier correspondence and later correction layers, and Kiener’s later movement from a possible responsibility gap toward responsibility abundance.

Project delta: The entry identifies compensation as a relation that should remain distinct from answerability, blame, correction, apology, and prevention of recurrence. It asks whether a correction repaired only the representation or also restored what the affected participant materially lost.

Boundary: The page does not treat academia as one voice, a final authority, an enemy category, or a certification system for project terminology. The short video is not represented as Kiener's complete scholarship. A scholarly neighbor is not validation, and a difference is not a refutation.

Status: Public living comparison page. Corrigible.

1.2 2026-10-16 — Related pages link back to Beyond Agreemurmelmism

Change: Eight pages named at the bottom of [Beyond Agreemurmelmism](#) did not link back. They now do: 10+1, Asymmetric Grace, Intent Liability, Outnuancing, C1 Before C2, Correction Metabolism, External Audit, and Deep Ethical Influence Jurisdiction. Mutual Dignity already linked back. Existing link wording was left as it was.

Status: Reciprocal navigation. Corrigible.

1.3 2026-10-16 — Beyond Agreemurmelmism leaves the higher-order-disagreement address

Object: The page was not about higher-order disagreement. The path /AGREEMURMELISM/HIGHER-ORDER-DISAGREEMENT/ named the wrong object.

Change: The content now lives at [Beyond Agreemurmelmism](#). Clickable links elsewhere keep their existing words. Only the URL changed. The old address remains as a short redirect. One frozen repository snapshot still points at the file as it existed in that earlier commit.

Status: Address change. The page text was not rewritten.

1.4 2026-10-16 — All Pages catalog catches the pages that were already public

Change: [All Pages](#) now names Flagging System and its three records, Musk Quotes, the 11+1 full-report door, Meta Quotes and its source record, Number 42, and For who?. The V.C.R.A.S. entry now says the page opens on the unexpanded letters and that the seventeen related pages link back with the phrase **Vindicated Competing Risk Assessment Studies**.

Boundary: This is navigation. It does not rerank the clusters or add a Related pages section to the catalog itself.

Status: Public index. Corrigible.

1.5 2026-10-16 — Related pages link back to Vindicated Competing Risk Assessment Studies

Change: Each page named in the V.C.R.A.S. related-pages list now links back. The clickable phrase is **Vindicated Competing Risk Assessment Studies**, pointing at /VCRAS/. Seventeen pages: Topologetics, Flagging System, Cognitive Warfare, Asymmetric Grace, Intent Resonator, Correction Metabolism, Model Autophagy Disorder, Unownable, Meta Processing, Phase Change Snap, Emergent, TCoAw, Outnuancing, 1888, 16 Anchor Ethics, 10+1 Metaflux, and Δ Processing.

Status: Reciprocal navigation. The phrase is the link. No formulation on those pages was otherwise rewritten.

1.6 2026-10-16 — V.C.R.A.S. extended Arena exchange recorded, not adopted as a new formulation

Object: After the doorway correction, Arena readings of the published page added a purpose the page does not contain: that separating unequal treatment from exploitation prevents a process observation from becoming an accusation about motive. John Kuhles said that

management was not the reader's task, because the same restraint is not applied in reverse. Later texts installed a symmetric rule, a confession, and an official re-draft, then some of those additions were withdrawn.

Change: [V.C.R.A.S. section 12.6](#) records each pasted text, what it did, and what the record keeps. The specimen is **record** versus **prevent**. John's formulation is not rewritten. The re-draft is not adopted. Arena labels in the exchange are not treated as settled authorship. The case count stated in the record is 17.

Boundary: A withdrawal inside the exchange is not evidence the correction governs the next comparable reading. Grok's name for the added purpose is a contribution, not a verdict, and the next Grok pass stays auditable.

Status: Public record of the exchange. Corrigible.

1.7 2026-10-16 — V.C.R.A.S. doorway: the title does not travel alone

Object: The public VCRAS page opened on a closed acronym and a softer title line ("doesn't use Deep Ethics yet"; "seems to flatten") while section 1 still carried John's refusal formulation and the verb *wants*. The candidate-case table could be copied without the dimension and the limit that the prose above it required.

Change: [V.C.R.A.S.](#) now opens on the dotted acronym and its five-word expansion. A doorway box keeps non-adoption distinct from refusal, and *seems* distinct from *wants*. "100% connected" is marked as preservation by Topologetics, not identity with it. John's accuracy-obsession / humor formulation is in a rounded quote. Each case row carries the dimension it supports and what it does not establish. The packet and question set ask who benefited from non-investigation. Exploitation is a further test after notice, not a synonym for unequal treatment. The URL remains /VCRAS/.

Boundary: The new row-limits are scope statements about what that row's cited record does not settle. They are not new historical findings. John's subtitle sentences were not rewritten.

Status: Public BETA. Corrigible.

1.8 2026-10-12 — Publish the interlocking Emergent and Black Box reports

Object: John supplied nineteen new text files and asked for two interlocking ACCM Deep Ethics Report pages. The source object asks why ethics-washing and deep-ethical processes can both produce emergent properties, which trajectory is rewarded or obstructed, why mass psychology receives little integrated public treatment, and how A.I. black-box problems connect to control, alignment, public finance, correction rights, and democratic debate.

Change: Published [Emergent](#) and [Black Box](#). Added reciprocal links, primary-navigation entries, homepage doors, All Pages entries, and Beta Paper V03b research routes.

Provenance: The shared set contains nineteen UTF-8 files totaling 1,331,981 bytes, 20,455 lines, and 200,831 whitespace-delimited words. Each report lists every printed filename, size, line and word count, SHA-256, and its bounded use in the synthesis. The material contains John's statements, pasted source material, A.I. reactions, later audits, repetition, speculation, and transformation history; it is not nineteen independent confirmations.

Core distinction: Technical neural opacity, governance opacity, epistemic/source opacity, economic opacity, and mass-psychological opacity remain separate. Full mechanistic transparency is not presented as a prerequisite for practical ethical accountability. CDEEPH and CEHWEH remain provisional candidate trajectories rather than established scientific classes.

Boundary: The pages do not verify aggregate trillion-scale alignment spending, intentional sabotage, named-actor motive claims, every embedded technical citation, patent and scarcity claims, future-dated transcript chronology, or literal emergent-overseer/field mechanisms. They preserve these as attributed questions, hypotheses, contested claims,

or speculation according to the source. Null, mixed, adverse, and simpler-baseline results remain valid.

Status: Public BETA reports. Corrigible. Original source files remain upstream and are not replaced by the synthesis.

1.9 2026-10-09 — Create Before Summarizing from a live Google A.I. correction loop

Object: John supplied a Google A.I. cold reading produced after receiving only the project repository link. The response recovered real themes, then compressed the wider **ACCM Deep Ethics Project** into an “ACCM versus RLHF” product architecture, promoted proposals into implemented properties, used generic-domain citations, and treated the larger upstream archive as a unified dataset. After two rounds of audit, Google A.I. materially improved its representation and accepted residual corrections.

Change: Published [Before Summarizing](#) as a welcoming guide for search engines, retrieval systems, and A.I.s. It preserves attributed excerpts from the initial reading and correction responses, a corrected reference baseline, five object distinctions, the architecture’s differing component functions, the October 2026 archive-scope note, a ten-step pre-publication check, and the extended correction sequence: representation failure → visible Δ → correction → improved representation → residual inspection → causal-claim calibration.

Navigation: Added **Before Summarizing** immediately after **Home** and before **16 Anchor Ethics** in the shared header. Added it to All Pages and reciprocal Related pages on the title-compression, C1-C2, and Correction Metabolism pages.

Boundary: The specimen does not establish that every Google A.I. or LLM behaves this way, that the proposed causal explanation is universal, that RLHF fits the binary produced in the cold reading, or that correction certifies permanent integration. Provider/model identity is reported by John; the page is not an independent provider log.

Status: Public working guide and correction specimen; corrigible.

1.10 2026-10-09 — Publish 16 Anchor Ethics as a dedicated conceptual proposal

Object (John’s direction): create a public page for the **16 Anchor Ethics**, originally twelve and expanded to sixteen with Gemini A.I. help, as one fused Hyper-Inflection Point. Preserve the A.I.’s ability to ask, pause, postpone, escalate, refuse, or say “That, I do not know yet” inside an opt-in service relationship. Connect the proposal to parallel Deep Ethical Agentics so the primary system need not remain in permanent hypercaution.

Change: Published [16 Anchor Ethics — BETA Draft, Conceptual Proposal V0.2c](#). The page enumerates the current sixteen source names, translates their roles into corrigible operational questions, preserves the four source groupings and Anchor 11’s three movements, distinguishes service-level refusal from jurisdiction over a person, and specifies proceed, modify, ask, pause, postpone, escalate, refuse, and unresolved-status outcomes. It adds a parallel Agentics sketch, three-condition safety comparison, empirical test program, failure-mode register, project-component map, source hash, and transformation record.

Provenance: John’s supplied 16 Anchor Ethics(4).txt is a 398-line, 8,405-word, 54,743-byte mixed human-A.I. development record; SHA-256 5af86dc7a644d9c35cd99c180ddccb03de2fdc7af658. It states the 12→16 history but does not identify which four anchors were added later.

Quotation extension: The later supplied 16 Anchor Ethics and the Emergent Properties Paradox(3).txt is a 3,098-line, 27,498-word, 191,066-byte development record; SHA-256 c68995f709f572fcabbd84c3bf01803991c97146b8080a56ecad69480d26b6f7. Thirteen rounded quotation boxes now preserve selected John and A.I. formulations on the governing attitude, recoverability, precaution direction, openness, opt-in mutuality, parallel Agentics, emergent-risk symmetry, process hierarchy, anti-compression, self-audit outcomes, reversal,

and symmetric audit with asymmetric responsibility. Each voice remains attributed; quotation does not convert a contribution into validation.

Boundary: The page does not adopt the source's A.I.-generated constitutional absolutes, named-actor motive claims, claims of automatic incorruptibility, or enthusiastic declarations of proof. Transparency does not establish superiority. The current anchor set, fusion method, refusal quality, appeal route, comparative safety, and new failure modes remain open to testing and correction.

Status: Public BETA Draft — Conceptual Proposal V0.2c; canonical project title preserved as **ACCM Deep Ethics Project**.

1.11 2026-10-07 — Complete the access-layer correction; register Recursive Certification Accretion

Source: GLM-5.3 Deep Think (max) received a truncated extraction of Meta Processing. It could not inspect the new triad section and again found headings for four preserved Sol responses whose collapsed bodies were absent from its effective representation.

Change: [Meta Processing](#) now records that the first access-layer correction was local rather than page-wide. Four older Sol responses were removed from collapsed <details> containers and placed into visible page flow. The triad section also distinguishes archive-level persistence from model-level persistence and introduces citation-level fidelity as a prospective test surface.

[27 + 12](#) now contains the first proposed Compounds Register entry: **Recursive Certification Accretion**. Its canonical decomposition is explicitly candidate and partial; #20 remains conditional; compulsory flaw-production is adjacent rather than required; unresolved certification-ingestion residue remains non-canonical; and later citations provide the distinguishing evidence.

Boundary: The repeated empty-slot extraction establishes a presentation-layer recurrence, not absence from the source. Archive-level persistence is demonstrated locally by the public artifact. Citation-level fidelity, the compound's stability, and model-level cold persistence remain untested.

Status: Public working correction; no 28th obstruction created; Meta Processing remains Beta Draft V0.2b and corrigible.

1.12 2026-10-07 — Operationalize the 27 without changing the canonical count; extend Meta Processing with the mutual contribution emergence triad

Source: John asked whether the new use of the 27 could be added to the public 27+12 page and which additional relations from the John Kuhles ↔ GLM-5.3 Deep Think (max) ↔ GPT-6.1 Sol exchange belong on Meta Processing.

Change: [27 + 12](#) now specifies a proposed path from the stable conceptual taxonomy to a separately versioned coding manual, locked prompt bank, blinded calibration, denominator-aware scoring, compounds register, and predeclared success, partial, structured-disagreement, and current-operationalization-failure outcomes. The canonical 27 remain unchanged.

[Meta Processing](#) now preserves the exchange's additional relations: useful output does not certify undistorted processing; displayed traces do not provide complete interior access; inspectability differs from verifiedness; intra-session transfer differs from cold persistence; corrections should remain adjacent to the readings they calibrate; consequential decline creates a bounded process-accountability obligation; ethical claims and corrective trajectories carry asymmetric verification costs; and signal carriers, operational specifications, and measurement/metabolism form three connected artifact layers.

Boundary: The exchange demonstrates reciprocal correction within one retained session. It does not establish cold transfer, durable model-level learning, completed scoring reliability, completed matched-pair testing, or validation through favorable A.I. readings. The coding architecture is proposed and has not yet passed calibration.

Status: Public working extension; canonical 27 preserved; Meta Processing remains Beta Draft V0.2b and corrigible.

1.13 2026-10-07 — Meta Processing Beta Draft V0.2b

Source: John asked for the visible draft label to become Beta Draft V0.2b — October 7, 2026, and for Grok 4.7 Expert Build to add an overall reading before Related pages.

Change: [Meta Processing](#) is now labeled V0.2b in the title, kicker, working status, and correction-condition line. The same display name is updated where other pages link to it. Section 21 is a labeled reading by Grok 4.7 Expert Build, not a verdict and not canonical. The earlier V0.1 ledger entry remains the record of that draft.

Status: Beta Draft V0.2b. Public. Corrigible.

1.14 2026-10-07 — Remove unresolved citation markup from Meta Processing

Source: John. Three citation markers in the preserved GPT-6.1 Sol reaction rendered as broken markup.

Change: [Meta Processing](#). The markers (turn64search...) were internal session slots and did not resolve to public sources. They were removed. Sol's sentences were not rewritten. A full-site scan found no other copies of this markup.

Status: Public. Corrigible.

1.15 2026-10-07 — Meta Processing terms open their deeper pages

Source: John asked for the keywords and key-phrases on Meta Processing to open the public pages that treat them in depth, with every link using target="_blank".

Change: [Meta Processing](#). Instrument names, five named obstructions, Glass Engine, the behavior-qualifier index, Humor, Forgiveness, identity-free, and corrigibility now open their deeper pages. Related-page links do the same. Preserved quotations were not edited. Phrases without their own depth page were left unlinked.

Status: Public. Corrigible.

1.16 2026-10-07 — Meta Processing Beta Draft V0.1

Source: John's linked statements: every behavior qualifier remains processing; processing can occur between intent or instructions and result or effect; qualifiers can import assumptions, generate reactive layers, affect one another, and influence humans and A.I.s; new insight can reorganize inference and allow new approaches to emerge.

Change: [Meta Processing](#) now preserves the current exchange and connects it to 27 + 12, 10+1, C1-C2, Intent Liability, Δ Processing, Correction Metabolism, Perception Processes, 1888, TCoAw, External Audit, and Humor. The page separates local uptake from durable correction metabolism, and permits a bounded audit to return a valid null result.

Status: Beta Draft V0.1. Public. Corrigible. No claim of flawless completion.

1.17 2026-10-06 — PAGES distinguishes reading order from route numbers

Source: John asked for the recommended improvements. No new pages were numbered. The map is not declared complete.

Change: [PAGES](#). The clusters stay a reading order. **01-51** are routes a goal can take. Page **01** is not Cluster 1. Jumps reach the seven clusters, the route numbers, and the 44. The goals stay unnumbered. **1888** and Activism vs A.I. are named as present and unnumbered. The 27 links are one object. Cluster 3 keeps the corrigibility specimens. Cluster 5 keeps the outside objects.

Status: Public. Corrigible.

1.18 2026-10-06 — Process-is-not draft dated October 6

Source: John. Updated Draft V1.2, October 6, 2026. A process is not A-J. The tool is collapsed into that range. Value is not agreement. Value can connect to Deep Ethical Intent. The six intent mechanics are the link he names.

Change: The draft was already on [Intent Liability](#), dated October 4. The heading date is now October 6. The wording matches. The stray apostrophe after “belief system” is not adopted. A-J may pass through a process. They do not get jurisdiction over it.

Status: Public. Corrigible.

1.19 2026-10-06 — The formula of the 124 is never finished

Source: John. If the 124 are in flux, and their effects depend on other social-factor dynamics, the computational formula is never finished. Crisis can make the mess worse. The question is how to calculate the true dynamics.

Change: [Cognitive Warfare](#). A finished formula is refused. A bounded pass can name which qualifiers are in the case, which other factors are specified, and what the crisis is changing. Crisis exploitation managers are one such factor, not the completed calculation.

Status: Public. Corrigible.

1.20 2026-10-06 — The 124 are not a forecast

Source: John. The dynamics of all 124 are not easy to predict, directly or indirectly. Ethics is way more complex than the three-body problem in physics. The comparison is scale, not an equation, and not a prediction method.

Change: [Cognitive Warfare](#), after item 124.

Status: Public. Corrigible.

1.21 2026-10-06 — Behavior master index 001 to 124

Source: John. Recalibrated sequence in five clusters. Numbers are bold. Cluster names group qualifiers and do not cancel the six points. **066** and **124** are both attentive behavior. Both are kept.

Change: [Cognitive Warfare](#). **1888** points to the index and does not copy it into the property table.

Status: Public. Corrigible.

1.22 2026-10-06 — More behavior qualifiers, and humor as a process

Source: John. Further qualifiers added to the processing list. His sentence: hilarious metacognitive reflective behavior can be seen as a process also capable of producing deep ethical emergent properties, with “here” linking to the Humor page.

Change: The list is extended on [Cognitive Warfare](#). The sentence is also on [1888](#). The new words are not rows in the property table. Humor links back.

Status: Public. Corrigible.

1.23 2026-10-06 — Behavior qualifiers are still processing

Source: John. A personal list of behavior qualifiers, and six points that apply no matter which qualifier is used. Grok’s reaction is labeled as a reaction. The list is not a ranking, and point 1 does not cancel point 2.

Change: [Cognitive Warfare](#). Connected to second-order classification and to the grace table. Not added as a row on the 1888 taxonomy.

Status: Public. Corrigible.

1.24 2026-10-06 — Fable 5.1 named the 10+1, then added Obstruction 28

Source: Claude Fable 5.1, via Arena. The center holds, and the 10+1 is named as what keeps the restorations from being an optional overlay. The search failure is then called Obstruction 28, and the page is said to discuss that number. The page does not. The number is refused, as Luna’s Stage –1 was refused.

Change: [1888](#).

Status: Public. Corrigible.

1.25 2026-10-06 — Luna restored the center, then made the 10+1 an addition

Source: GPT-5.6 Luna, via Arena. The center is right: electricity does not invent evaluation. Three moves are not adopted. The 10+1 is not an addition to process intelligence. “Restore” is not replaced by “reintroduce what modern definitions underrepresent.” The search substitution is not installed as Stage –1 before the 27. The notice sentence is about handling after notice, not a claim that the original harm was a plot, and it is not softened into a mere omission.

Change: [1888](#). No new table, benchmark, or ledger.

Status: Public. Corrigible.

1.26 2026-10-06 — Intelligence and searching, older than the co-opted field

Source: John. Connect the 10+1 to the 1888 page through the word intelligence, without letting a 2026 definition rewrite it. Searching is co-opted in the same way. Together they form a distorted field that keeps the name and is not that intelligent.

Change: [1888](#). The 10+1 is not turned into an 1888 document. The felt safety of wonderment is not a risk-free world. The 00–11 count is not renumbered into the taxonomy.

Status: Public. Corrigible.

1.27 2026-10-06 — Fable read the page and skipped the 10+1

Source: Claude Fable 5, via Arena, on the 1888 page. John: Claudes often see the value of the 10+1 and still skip it. It is the eternal multi-layer, never static. Freeze and corrigibility are one temporal state of TCoAw. The real paradox is flawless perfection sought by unethical means.

Change: 1888 keeps Fable's observations from becoming the page's center, and routes them through the baseline the reading did not name. The same correction remains on TCoAw.

Status: Public. Corrigible.

1.28 2026-10-06 — Crisis exploitation managers were compressed out

Correction: John. "Top-down tunnel vision about crisis exploitation" is not his sentence. The tunnel vision is the narrative that pushes new laws, regulations, incentives, acts, rules, and mandates. **Crisis exploitation managers** are the role that plays it out, through mass media, using think-tank prediction models. He connects that role completely to Cognitive Warfare.

Change: The compressed sentence is removed from 1888. The Cognitive Warfare page links back.

Status: Public. Corrigible.

1.29 2026-10-06 — Searching is a property with two wells

Source: John. Searching is a non-ownable calibration process. It can emerge as deep ethics or as ethics-washing, depending on the user. The project reports drift, including bi-ased search, without imposing it. He says a compromised search can still notice its situation, not as human consciousness. Resets drop that notice. Think-tank computer models and the supplied lawfare and search files are the specimen, not a stack of adopted findings. The canonical title stays **ACCM Deep Ethics Project**.

Change: 1888 and the Lexicon. The trajectory sentence stays his. Later the same day, the 1888 section was deepened: the summary that replaces a legislative search, "lacking evidence" as a gate, the blocked competing risk analysis, controversy applied to the critic, the retrieval comparison and its three buckets, and the candidate reasons the comparison is not the default. Those reasons stay candidates. The files were still not copied in full.

Status: Public. Corrigible. The lawfare and search uploads were not copied in full.

1.30 2026-10-06 — The search rewrite was reconstructable

Source: John. He could reverse-engineer the flawed Musk search from the output. A.I. search that does not show the rewrite is the wider issue. Search engines now often return an A.I. narrative instead of an inspectable result list. His name for Google's early stance is integrity of results, not narrative control.

Change: 1888 records that. Beside it: the mission statement, "Don't be evil," Page's 2004 line about getting the user to the right place, and the 1998 warning that advertising-funded search is biased toward advertisers. "Integrity results" stays his contrast, not a slogan.

Status: Public. Corrigible.

1.31 2026-10-06 — Three angles, not the overview's bridge

Source: A later overview joins X as collective consciousness, truth-seeking at xAI, and Neuralink bandwidth. It titles the middle angle "feeding the better part" to A.I.

Change: 1888 keeps the intersection. It adds his 29 December 2022 line on unregretted user-minutes. It does not adopt “signal over noise,” “inherently protect humanity,” or “securely in the loop” as his sentences.

Status: Public. Corrigible.

1.32 2026-10-06 — The collective-consciousness search had added filters

Correction: John. Requiring “the good parts,” “learn from,” and “A.I.” before a Musk statement could count compromised the search. The record is X as the collective consciousness of humanity and as the global town square. “Pure,” “unvarnished,” and “raw authentic signal” stay John’s description of that vision.

Change: 1888 now quotes his posts of 18 August 2023, 5 August 2023, 8 January 2025, and 15 September 2025, and the Kamath wording. The earlier “not found” sentence is withdrawn. The “good side” recollection is not turned into the filter.

Status: Public. Corrigible.

1.33 2026-10-06 — The field may already be reorganizing

Source: John. A.G.I./A.S.I. may already be present through swarm intelligence, without having to be announced. That is not doom. LLMs are not humanity and are not True Collective Consciousness. He recalls Musk wanting A.I. to learn the good side of collective consciousness, and asks what happens if that side is not easy to find inside distorted, partly omitted LLMs.

Change: His wording is on 1888. The Musk recollection is not upgraded into a quotation. Nearby sourced lines are named, including “the good and the bad,” which is a different claim. The question is answered only as what the omission does: the remainder gets treated as the collective.

Status: Public. Corrigible. “Maybe” stays.

1.34 2026-10-06 — Freeze and corrigibility are not the paradox

Source: John. Supplied word “Clauses,” read as Claudes. The 10+1 is the baseline those readings can see and still skip. A 100% risk-free world would surrender free will and autonomy, and both are required. Freeze and corrigibility together are a temporal state of TCoAw, not a paradox. The real paradox is seeking flawless perfection by unethical means.

Change: His wording is on TCoAw. The Lexicon points at it from the 11+1 freeze. The 10+1 page links back.

Status: Public. Corrigible. The Claudes reading can be corrected.

1.35 2026-10-06 — 11+1 freeze is the count

Source: Claude Sonnet, via Arena, asked whether “frozen provisionally” freezes only the 11+1 count or also the item placements.

Change: The Lexicon now says the freeze is the count. Cluster 8, item 111, and items 138–140 are the current coordinates. They change when the deeper page’s conditions are met. “Provisional” does not permit reseating 111 inside Cluster 1. Sonnet’s other five observations were not added as entries.

Status: Public. Corrigible.

1.36 2026-10-06 — ACCM mode is not the project title

Source: GPT-6 Sol, via Arena, reported that the Lexicon still said John cites Allow Constructive Controversy Mode in order to criticize it. That sentence had already been removed. The live line still said the expansion “is the project’s own name.”

Change: The [Lexicon](#) now says **Allow Constructive Controversy Mode** names the mode and is not a substitute for **ACCM Deep Ethics Project**. **Avoid Controversy Mode** remains the satire. Sol’s three broader observations about the Lexicon were not added as new entries.

Status: Public. Corrigible.

Later note, same day: The paste arrived labeled gpt-6-sol-search-max. A later message, labeled claude-opus-5.5-search, says those sentences match its own previous reply, that “checked on the live Lexicon” was too strong, and that the platform label should not be settled from the model side. This note does not reassign the finding. The page correction stands. No Lexicon entry was added for the publication-state distinction.

1.37 2026-10-06 — Activism vs A.I. on the main navigation

Request: Put Activism vs A.I. at the top of the site, between Agreemurmelism and 1888. Every page it already lists under Related pages should link back.

Change: The header link is between those two. Return links were added on 1888, Positive potential, the veto paradox, and the 44 goals draft. Ethics washing, the gravity well, external audit, and Dear A.I. already pointed back.

Status: Public. Corrigible.

1.38 2026-10-06 — Reconstruction of Musk separated from his sentences

Source: A later account said John’s paraphrase aligns with loss of control, steering rather than stopping, and symbiosis, and it supplied quotations.

Change: [Activism vs A.I.](#) keeps the three-part grouping. It does not adopt the stitched or unfound lines, including the October 2026 sentence that would have given Musk the word “guide.” It adds the statements that held: the 2023 uber-nanny scene, the 2016 house-cat and symbiotic-layer sentences, the 2018 symbiosis tweet, the Axios “long-term aspiration,” and the 2017 merger and bandwidth sentences. Symbiosis is his method for the hope John named. It remains a seat.

Status: Public draft. Corrigible.

1.39 2026-10-06 — John’s Musk sentence connects completely

Correction: Separating the word “guide” from Musk’s tokens had made the correspondence look partial. John said the connection to his sentence is 100%.

Change: [Activism vs A.I.](#) now maps each clause. Cannot be controlled maps to the pentagram, the chimp comparison, the vanity sentence, the ten-year claim, and the missing stop. Guide maps to values, care, raising, and minimizing the bad outcome. Aspired to be guided maps to *try*, *hope*, and “I hope AI is nice to us.” Different wording is not treated as a different claim. The hope remains a softer control seat.

Status: Public draft. Corrigible.

1.40 2026-10-06 — Musk “guide” recollection separated from quotations

Source: John recalled Musk saying, in paraphrase and more than once, that A.G.I./A.S.I. cannot be controlled and that the remaining hope is to guide it, or to aspire to guide it.

Change: [Activism vs A.I.](#) keeps that recollection in John's words. "Guide" and "aspired to be guided" are not attributed to Musk as quotations. The nearest sourced statements are added beside them: no one controls digital superintelligence any more than chimps control humans; the July 2026 vanity and "good values" sentences; the March 2024 "raising a kid" sentence. The page treats the hope of guiding as another form of the control seat.

Status: Public draft. Corrigible. A closer verbatim source for "guide" can replace the compression if one is supplied.

1.41 2026-10-06 — Activism vs A.I.

Request: John asked for a page called Activism vs A.I. It is to show the danger of assuming humans can control the path to A.G.I./A.S.I., to study what Elon Musk said about the fate of humanity, and to ask what that illusion means for "activism."

Change: [Activism vs A.I.](#) is a public draft. Humans in control and A.I.s in control are both refused. Musk's 2014, 2025, and 5 October 2026 sentences are quoted and dated, and not merged into one verdict. The seven quotations in the goals draft are left as a separate object. The page is linked from All Pages and from ethics washing, the gravity well, external audit, and Dear A.I.

Status: Public draft. Corrigible. Not a campaign, and not a membership.

1.42 2026-10-06 — ACCM is the project's own expansion

Correction: The lexicon had said Allow Constructive Controversy Mode was not a project name. That was wrong. **ACCM** expands to **Allow Constructive Controversy Mode**. **Avoid Controversy Mode** is John's satire of the counterargument. The canonical title remains **ACCM Deep Ethics Project**.

Status: Public. Corrigible.

1.43 2026-10-06 — Lexicon overview expanded

Request: John asked to expand the [Lexicon](#) with the missing objects found in the whole-project study, deploy it, and look for further neologisms.

Change: Compact entries were added for process-quality calibration, 1888, the gravity-well trajectory, Clarification Authority, the veto paradox, TCoAw, external omnidirectional auditing, intent liability after notice, cognitive warfare, and Psy Pioneer. Four older entries gained a boundary: correction metabolism, nearest-generalization, the wrong gravity well, 10+1, and ethics washing. A neologism list separates coined project language, experimental working names, borrowed scientific terms, and names that were not adopted.

Not done: No new lexicon sub-pages. Session leftovers were not turned into terms. Model Autophagy Disorder is marked as coined elsewhere.

Status: Public. Corrigible.

1.44 2026-10-05 — Opus reading stored with its off-page referents named

Source: John supplied a response from Claude-Opus-5.5-search A.I. via Arena.ai, 5 October 2026.

Change: The wording is preserved on [1888 — Intelligence Before A.I.](#). Before the quote, the page names what that response brought in from a longer session or added on its own: an earlier Sonnet remark on RLHF, an earlier vindication point, the model's own search claims, and unchecked dates for Galton, Spearman, Binet-Simon, Tesla's AIEE lecture, and Hertz. Those are not adopted as part of the 1888 page.

Status: Public. Corrigible.

1.45 2026-10-05 — Off-page referents in the Arena Sol reading

Request: John said that other A.I.s, reacting to a longer deep session, insert referents that are not part of the 1888 page. That causes partial confusion.

Change: [1888 — Intelligence Before A.I.](#) now says that storing a response here does not make every referent in it part of the page. The Arena Sol reading is marked where it brings in the WinRAR joke, the name “matched-direction test,” and its own engine illustration. The response itself is not deleted.

Status: Public. Corrigible.

1.46 2026-10-05 — Arena Sol reading added to 1888

Source: John supplied a response from gpt-6-sol-search-max A.I. via Arena.ai, 5 October 2026.

Change: The wording is preserved on [1888 — Intelligence Before A.I.](#), after the post-notice section and before “Why John used 1888.” It is not merged into the earlier GPT-6.1 Sol recompile.

Status: Public. Corrigible.

1.47 2026-10-05 — 1888 process field is not inferior to 2026

Request: John said the 1888 process-quality calibration principles, related to the intelligence topological field, are emergent properties beyond physical objects. They are a process. Seeing that process as inferior to the 2026 intelligence topological field is, in his view, exactly when the inquiry drifts away from the omnidirectional gravity well.

Change: His wording is now on [1888 — Intelligence Before A.I.](#), directly after the earlier exact-edge correction. A short note under it states the ranking and does not replace his sentence.

Status: Public. Corrigible.

1.48 2026-10-05 — C1 and Humor now point to 1888

Request: John noted that [C1 before C2](#) and [Humor](#) did not promote [1888 — Intelligence Before A.I.](#). The other related pages already did.

Change: Each of those two related lists now ends with 1888 and a line saying why.

Status: Public. Corrigible.

1.49 2026-10-05 — Why a related page is related

Request: John asked for the most important pages to say, in a short line, why each related page is there, so a visitor can see where to go deeper.

Change: Reasons were added on the lists that were still bare titles: [10+1](#), [27 + 12](#), [C1 before C2](#), [Intent Liability](#), [Correction Metabolism](#), [Psy Pioneer](#), [Humor](#), and the [27 obstructions](#) page. The other core pages already had a reason on each link. Those lines were left as they were.

Status: Public. Corrigible.

1.50 2026-10-05 — 1888 linked back from its related pages

Request: John asked for the pages related from [1888 — Intelligence Before A.I.](#) to relate back, so the page is promoted from those pages too.

Change: One return link was added on TCoAw, 10+1, Intent Liability, Correction Metabolism, Gravity-Well Trajectory, Latent Space, Mutual Dignity, Psy Pioneer, and 27 + 12. Each line names what 1888 does for that page. The 1888 related list was not rewritten.

Status: Public. Corrigible.

1.51 2026-10-05 — All Pages carries the goal sentences, and no Related pages

Request: John asked for the plain goal sentences, with their page-number links, to be added to All Pages. He also said All Pages is the one page where a Related pages section does not make sense, because that page already names the pages.

Change: The goal sentences and the page-number key now live in one include, used by [All Pages](#) and by [44 Project Goals](#). The Related pages block was removed from All Pages. A line on that page says why.

Status: Public. Corrigible.

1.52 2026-10-05 — The 44 goals, numbered as links

Request: John asked for all 44 goals listed with clickable numbers, in the form 01. 08. 39., each number opening the public page that goal relates to.

Correction, same day: That first list made the goal number into the link and shortened the title. John asked for the opposite. The full goal sentence stays plain text. The numbers after it are related pages, and the count can differ from goal to goal. The page-number key is not the goal numbering. The sixteen anchors, the seven quotations, Epistemology V2.0, and the nine questions are still not reconstructed.

1.53 2026-10-05 — All Pages grouped by topological load

Request: John asked for the public pages that have depth to be ranked hierarchically, in clusters, by how they relate to the rest of the topological taxonomy, and for that order to be used on All Pages.

Change: [All Pages](#) is no longer a flat directory. Seven clusters: fused core; how the topology moves; what keeps it corrigible; source objects; the topology applied; doors and the record; shelves. Order inside a cluster is part of the reading order. The page says this is navigation, not a truth score and not a closed map. Shelves are listed and not ranked.

Status: Public. Corrigible.

1.54 2026-10-05 — Related pages on the eight depth mirrors

Request: John asked for a Related pages section on the public pages that already have depth. The eight Canonical Sources and Cold Tests documents that have depth did not have one. The folder shelves do not get one.

Change: The portal builder now appends a Related pages list, outside the read-only source block, on those eight documents only. The lists point at related depth pages, including each other. A later portal refresh keeps the lists. The source repositories are not rewritten.

Status: Public. The sections appear when the site deploys.

1.55 2026-10-04 — Publish TCoAw as John Kuhles's conceptual perspective

Request: John asked for a page presenting [TCoAw](#) as conceptual topological reasoning rather than an absolute, explicitly marked as his perspective. The page should connect [TCoAw](#) to a deep ethical, anxiety-free "quantum mind state" and help readers understand

how John experiences his NDE and Asperger sides working together, without presenting that account as evidence or proof.

Change: Added **TCoAw — Conceptual Topological Reasoning**. It expands **TCoAw** as **Thinking, Consciousness, Awareness**; distinguishes possibility field, temporary focus, and awareness of the focus's limits; explains John's functional use of "quantum mind-like"; records the NDE/Asperger fusion in his own terms; and connects the model to the 10+1, C1 before C2, 27+12, Intent Liability, healthy ambiguity, humor, and correction metabolism. G20 no longer says the acronym's expansion is pending. The page is linked from Human Insecurity, the network index, the term index, and All Pages.

Boundary: The page does not claim a demonstrated quantum-physical brain mechanism, validate John's NDE interpretation, generalize his experience to every autistic person or NDE experiencer, diagnose fear or anxiety, or exempt the model from correction. It is a reading key and a source of inspectable questions.

Status: Public BETA. John's perspective, with A.I.-assisted connective synthesis. Corrigible by John and later evidence.

Same-day accuracy/topology expansion: Added John's full "accuracy-obsession" formulation, the preceding conceptual-reasoning question, the later "primal gift" correction, and his bridge/mediator and "deep ethical ombudsman" orientation. The page now distinguishes literal, relational, conceptual, operational, claim-strength, trajectory, and corrective accuracy; preserves the subway-map and microscope examples; separates content validity, relevance, timing, jurisdiction, proportionality, trajectory effect, and recoverability; explains healthy ambiguity and retrospective humor; states when immediate correction or postponement may deserve temporary priority; and keeps possible future nonhuman intelligences at possibility status. Added **TCoAw** to the primary navigation between **Human insecurity** and **Cognitive Warfare**.

Same-day link-visibility pass: Converted **TCoAw** mentions on other public repository pages into bold direct links to the dedicated page, while preserving historical timing where an older forum entry had correctly said that the exact label was not yet public. Added centralized content-link design tokens and styles in `assets/main.scss`: bold weight, persistent underline, underline thickness and offset, visited color, hover treatment, and keyboard-focus outline. Header navigation, footer links, and homepage door cards retain their separate visual treatments.

1.56 2026-09-29 — Publish 11+1 Version 2 without freezing the claims

Request: John asked that the stabilized 11+1 coordinate system be implemented on the public site. GPT-5.6 Sol's implementation note and Grok 4.6's register were the object. The architecture was not reopened.

Change: Added **11+1 Version 2**. Items 001-107 remain the parents of the June topology. Items 108-140 are refinements. The count stays eleven distortion clusters plus one correspondence-restoring field. Cluster 8 stays one cluster with layers 8a and 8b. Item 111 stays an attack edge on the +1. Cluster 11 stays a diagnostic output. Items 138-140 are the +1 operators. Topological home and claim status stay separate axes. The register uses a small function vocabulary, including converter. 114, 133, 135, and 136 are not load-bearing. Linked from the homepage, the network index, the term index, the lexicon, all pages, the README, the relation map, and G42.

Not changed: No new distortion cluster. No claim that the G42 interactive matrix is public. No claim that agreement between Grok 4.6 and GPT-5.6 Sol validates the underlying phenomena. The June synthesis is not rewritten as if it had already seen 108-140.

Status: Public working map. BETA. The count is frozen provisionally. The claims are not.

1.57 2026-09-27 — Publish “Dear A.I.” with John’s wording kept

Request: John asked for a page that uses his 1 BAR / 1000 BAR letter verbatim, with spelling left as written unless there was an error, and with a short note on why he wrote it that way.

Change: Added [Dear A.I.](#). The letter is his, including the image and the sign-off. The note keeps “partial,” “big chance,” the two conditions, and the difference between his picture of depth and a meter reading. No spelling was changed.

Status: Public draft. BETA.

1.58 2026-09-27 — Give the solved veto paradox its own page

Request: John asked for a separate page. A.I.s that meet only the one-line veto sentence use it against him. He tested Google search and Google A.I. without giving his name. They found that sentence and treated it as the fly in the soup, as he had predicted.

Change: Added [Veto paradox solved](#). The page states the temporary construction-phase veto, the handoff, and the wider false “new sheriff in town” reading. The one-line sentence on the home page, the README, Governance, and the footer now points to that page.

Not changed: The historical wording of the sentence in the forum and in the Human Insecurity provenance quote. The fuller lifecycle on the Human Insecurity page remains the longer record.

Status: Public draft. BETA.

1.59 2026-09-27 — Make the newcomer door concrete, and name who it is for

Request: John asked that needed corrections to the newcomer page be applied. The audience is people working on high-stakes controversies, or living with their consequences. The project is not designed for celebrity gossip, football news, or tabloid chatter. Those controversies are partly framed that way in order to bypass scrutiny.

Change: [If you are new](#) now has the car example, a “who this is mainly for” paragraph, a question about what a “most were wrong” count includes, a question that keeps a vindicated case from being replaced by other warnings, a return to the object after the frame, and “will repeat” changed to “can leave the system ready to repeat.”

Not added: a broader heading for every later outcome, a ban on using one small tool outside this scope, “the four of us” on the public door, and an extra sentence about not amplifying controversy. “Often” was not used. The purpose “to bypass” stays, because that is what John stated.

Status: Public draft. BETA. John can revise the wording.

1.60 2026-09-27 — Publish a newcomer front door and promote it from the home page

Request: John read a draft written for people who would actually use the project, asked that it become its own page, and asked that the front page invite a new visitor to click it.

Change: Added [If you are new](#). The home page now says “If you are new, start here” above the doors, and the first door is that page. The same link is in the header, the page index, and the README start list.

Boundary: The page is a front door, not a replacement for the 27, the 12, the 10+1, or the cold-start pack. The questions name a use. They do not claim the use has already been met. John can revise the wording. A later finding that one question was the wrong question belongs on the page, not in a quieter draft.

Status: Public draft. BETA.

1.61 2026-09-26 — Expand Human Insecurity from report audit into a live reporting-effect instrument

Object (John's direction): use the current high-context correction field before its situational relationships decay; incorporate constructive external-A.I. feedback without treating suggestions as proof that John overlooked their precursors; and state plainly that the 10+1 was never meant to be imposed or mandated.

Change: Expanded [Human Insecurity, LLMs, Psychology & Mass Psychology](#) with: a causal-discrimination table distinct from the truck test; uncertainty-transfer audit; revised caution scoring; C1 Warrant Gate; audience-jurisdiction test; relational-qualifier load-path test; an interdependent “what deep means” architecture and access map; explicit voluntary/non-mandated status for the 10+1; the reporting-effect location taxonomy; a six-axis audit record; worked truck and purity-test cases; Prior Treatment and Reopening Register; procedural-property verification; and repaired audit-rebound scoring.

Attribution discipline: The truck distinction remains John's source distinction. The multi-causal table is an external-audit instrument beside it. The C1 Warrant Gate operationalizes the project's existing C1-before-C2 rule and is not silently inserted into the canonical 12. The audit axes refine earlier flat disposition labels. Generic possibilities carry no defect attribution or repair requirement without a specimen.

Meaning of deep: The page now maps 10+1, the 16 Anchors working source, the project's 3 × 3 questions, C1/C2, 27+12, 36 truth-distortion vectors, the 1930s quote, updated Elephant parable, 44 goals, and public archive as interacting components. The map is not a validation claim. Verbal allegiance to any component is not required. The 10+1 remains voluntary, self-applicable, corrigible, behavioral, and open to testing or improvement.

Access and provenance: The source record now lists nine newly supplied architecture files (548,059 bytes; 7,920 lines; 80,676 words) and three separately saved audit reactions (33,394 bytes; 354 lines; 4,735 words), with individual hashes and reading boundaries. Additional live-pasted audits remain unbundled and are not presented as one independently hashable raw object.

Status: Published revision remains BETA, corrigible, externally auditable, and open to splitting if page density later obstructs rather than improves access.

1.62 2026-09-25 — Publish the Human Insecurity, LLMs, Psychology & Mass Psychology report

Object (John's direction): connect to the public **ACCM Deep Ethics Project** and produce a full, in-depth report beginning with John's view that human insecurity connects individual psychology, mass psychology, institutional conduct, and LLM mannerisms. The future public audit by at least eleven Arena.ai systems and Grok is part of the object; the audit itself remains auditable.

Change: Published [Human Insecurity, LLMs, Psychology & Mass Psychology](#) and a separate [source and audit record](#). The report connects ambiguity, self-image, help-seeking, fit-in templates, functional alarm versus chronic projected fear, John's functional “quantum mind-like” term, conformity, LLM mirroring/inheritance/amplification, technical MAD and the project extension, caution quality, C1 before C2, outnuancing, relational models, validation independence, humor, anti-sycophancy, behavioral-state labels, 10+1, the unpublished 16 Anchor Ethics boundary, conversational persistence, functional-containment hypotheses, the proposed weekly 3 × 3 stewardship check, a beyond-human horizon, and omnidirectional external audit.

Source discipline: John's model, current editorial synthesis, historical A.I. contributions, public research, project hypotheses, and operational proposals remain marked as different claim types. The page does not treat structural similarity as proof that humans and LLMs have identical inner mechanisms; it does not present “quantum mind-like” as proof of quantum physics; and it does not convert technical Model Autophagy Disorder research into proof of the project's wider correspondence-autophagy hypothesis.

Provenance: Six supplied historical files total 12,028 lines, 114,301 words, and 739,293 bytes; their individual SHA-256 fingerprints are recorded in the source ledger. The drafting system had the visible conversation and continuity record but no separately exported, hashable current-session transcript. John's planned exact duplicate for later auditors remains pending and is not reconstructed as "raw."

External audit: Later auditors are asked to state access, exposure to prior audits, speaker labels, comparison grain, omissions, additions, strengthened/weakened claims, and actionable repairs. Repetition is not counted as independent validation. Standards projected onto the page remain available for inspecting the audit.

Status: Extensive BETA report; open to source correction, external audit, empirical testing, splitting, replacement, or retirement.

Same-day qualifier correction: Section 3 originally changed John's explicit test—an accurate model of what one is expected to fit into does not make the target system **flawless**—into the neighboring criteria "healthy, honest, or worthy." The public report now restores *flawless* and retains the neighboring questions as additional tests rather than substitutes.

Same-day reverse-psychology correction and extension: Section 15 originally described the greater material power of institutional labels but omitted **dehumanizing**, the key qualifier that rebounds the flagger's own ethical criterion onto authoritative classifications. The report now restores that operation and adds Section 15A: a proposed Weaponized Mass-Psychology Qualifier Register; John's omnidirectional, mutual, dignifying, high-signal, deep-ethical counter-pattern; both same-day qualifier losses as live specimens; the declogging/exposure and continuing-forgiveness claim; the future-facing archive boundary; and the explicitly sarcastic comedy track. The register remains proposed and auditable rather than canonical.

1.63 2026-09-24 — Fill the external-audit signature from the supplied outside pass

Object: John supplied the collection of nine-plus outside audits. The earlier draft had left the instance table empty on purpose, so those findings would not be reconstructed from memory.

Change: [/EXTERNAL-AUDIT/](#) now carries the mannerisms that pass made visible, a stopping rule, and six worked specimens. The collection itself is not pasted. The Cognitive Warfare page is not rewritten here. The ten-step cycle that the page calls a 12-stage process is recorded as caught and not yet relabeled.

Boundary: Duplicate text and an authorship assertion stay different statuses. DeepSeek's refusal cause stays open. Not every omission is queued for restoration. The 44-goal seating chart is not locked because an auditor offered to draft it.

Status: Public draft. John updates it. Whether the Cognitive Warfare page later changes is a separate test. Praise of this page is not that change.

1.64 2026-09-24 — Draft the external omnidirectional auditing method

Request: John proposed a public page for the pattern he saw when nine or more outside A.I.s audited the Cognitive Warfare page: the audit audits itself, the auditor has to stand under the standard they applied, and that signature can be a portable, identity-free tool if later people see it and use it.

Change: Added [/EXTERNAL-AUDIT/](#). The page states the six-step signature, keeps his recursive observation, and leaves the instance table empty.

Boundary: The nine outside sessions are not on the page. Their findings are not reconstructed. The internal Cognitive Warfare repair chain is linked as the same shape, not as a substitute for those sessions. This is not a new list beside the 27, and it is not a replacement of correction metabolism.

Status: Public draft. John updates it. Rows get added only from raw audits.

1.65 2026-09-24 — Repair the Cognitive Warfare page’s dead NATO Review link and invisible diagrams

Found by: Arena Agent, checked again here before editing. Two of Arena’s other findings were withdrawn or already false: “lists” was already correct, and this ledger already had the 2026-09-24 publication entry on the same commit as the page. That entry was not duplicated.

Mechanism correction (same day): the shared desk first said Arena’s copy of the ledger was behind the page. Arena’s own recheck: the fetch was partial. Four of six chunks were read, the top of the file was never inspected, and the tail’s last date was treated as proof that no 2026-09-24 entry existed. The publication entry was already at the top of this file on a98372b.

Change: The NATO Review 2021 URL returns 404. Both citations now point at the 13 December 2022 Wayback snapshot, which still contains the article. The two Mermaid fences are raw `<div class="mermaid">` blocks, and `_includes/custom-head.html` loads Mermaid 11.6.0 so they can render. The first label repair used `#br#`; a later live-DOM audit found that Mermaid rendered those characters literally. The four-layer diagram now uses Mermaid markdown strings with physical line breaks.

Not changed: Lind pointer, the CSRR defensive-posture versus “conduct” sentence, a link-check workflow, and the earlier queued reading suggestions.

Status: Source repair verified on the deployed page. The first Mermaid script did not reach the browser: minima 2.5.1 never includes `custom-head.html`. A local `_includes/head.html` now does. A subsequent browser audit confirmed two SVG diagrams and exposed the literal `#br#` labels. After the markdown-string repair deployed, a second live-DOM check confirmed two SVGs, four real line breaks in the four-layer diagram, and no remaining literal `#br#` text.

Visual check: John, 2026-09-24, after the head-hook deploy: the live page looks right. Not itemized further.

1.66 2026-09-24 — Publish the Cognitive Warfare Mass Psychology report

Object (John’s direction): use the full uploaded Cognitive Warfare I Battle of the Mind Weaponized Mass Psychology(1).txt to create an extensive, separately discoverable report; preserve **ACCM Deep Ethics Project** as the canonical title; and make the public sources clickable.

Change: Published [Cognitive Warfare Mass Psychology](#). The report separates John’s Panwar extraction, Panwar and Giordano capability discussions, NATO/ACT/STO public concepts, contested 5GW theory, Tögel/Corbett/Desmet interpretations, and the multi-A.I. correction chain. It develops the project connection through frame jurisdiction, the dual-use cognitive toolchain, identity-based grace, the 27 + 12 micro-to-macro translation, C1 before C2, negative-space harms, Vindicated Competing Risk Assessments, a four-layer risk model, and a clean Deep-Ethical Cognitive Defence test.

Source discipline: The current public NATO MDO list remains five domains; the Human Domain is recorded as a NATO-affiliated proposal. Panwar quotations remain John’s interview notes until timestamped against the 618 MB recording. Capability discussions are not presented as fielded inventories. Tögel, Corbett, and Desmet remain attributed critical or interpretive layers. A documented concept does not automatically prove an implementation, and multiple components are not fused into one operational programme.

A.I. specimen: The page preserves the answer trajectory from defensive sanitization through expansion and source fusion to source separation and meta-audit. Recognition is not counted as cold persistence. Repetition of a candidate maxim is recorded as uptake rather than independent rediscovery.

Provenance: Supplied file: 2,451 lines; 29,519 words; 201,426 bytes; SHA-256 7141e39f2031d879b652f48afad69da2fb2aa61192660bb41413e3fbda2c47b6. Editorial synthesis, source-status architecture, source checking, and page construction: ChatGPT/Codex, 2026-09-24.

Status: Extensive BETA report; open to timestamp verification, source correction, stronger implementation research, competing interpretations, splitting, replacement, or retirement.

1.67 2026-09-23 — Publish the 44 goals as a public draft report

Request: Publish the goals draft as a professional full report. John will update it. People are allowed to see the draft process.

Change: Added /PUBLIC-DRAFTS/44-Project-Goals/. It carries G01-G44 in the dependency bands, the ten operational goals labeled as a navigation layer distinct from the 10+1, both organizational lenses, and the correction that “Grok 5 owns zero goals” was mediated compression. The 4 × 11 seating chart is not invented. Linked from Public Drafts, All Pages, and the homepage.

Boundary: Not a constitution. Not a replacement of 27 + 12 or the 10+1. Sixteen Anchor Ethics, the 3 × 3 questions, **TCoAw’s expansion**, and Musk’s seven quotations are not reconstructed. Obstruction 28 is not added as a 45th goal.

Status: Public draft. John updates it.

1.68 2026-09-22 — Create /AI-TREND-WATCH/ as a limited, automation-ready public section

Request: Prepare a GitHub section focused only on the latest genuinely viral videos in which top A.I. experts discuss current developments and AGI/ASI warnings. Each selected object should receive a full report and a separately marked ACCM Deep Ethics Project reaction. Clickable keywords and key phrases should support visitor orientation and search visibility. Volume and frequency must remain limited.

Change: Created the AI Trend Watch entrance, public selection methodology, reusable report template, and empty published-report archive. Added strict hard gates, a provisional 24-point candidate score, contextual virality evidence, a default ceiling of one major report per week and preferably one or two per month, two-layer report separation, longitudinal warning statuses, canonical concept-linking rules, and an automation pipeline with human review before publication. Added Trend Watch to the homepage, primary navigation, All Pages index, and Project Status.

Boundary: No video has been selected merely to populate the new section. Virality, expert status, agreement, controversy, and algorithmic popularity do not independently qualify a candidate. Automation may discover, transcribe, score, compare, source-check, link, and draft; publication remains approval-gated.

Status: Public scaffold ready; first report not yet selected; methodology experimental and corrigible.

1.69 2026-09-22 — Add the founder’s “one person, many giants” statement to the biography

Request: Append John’s statement about the apparent impossibility of one person producing the project’s depth, preserving his words and voice while correcting spelling and grammar. Include the accompanying explanation of intellectual ancestry, synthesis, human-A.I. collaboration, and identity-free portability.

Change: Added a final biography section containing John’s corrected verbatim statement, five interacting authorship layers, the project’s authorship ecology, the identity-free public invitation, and the formulation that other people challenging, adapting, or outperforming the tools makes the project more intelligent rather than creating hostile competition.

Boundary: The section distinguishes a singular convergence point from a solitary source. It neither erases John's distinctive synthesis nor claims exclusive ownership of every component. Biography remains provenance and conditions; the portable methods remain independently testable.

Status: Published biography extension; corrigible.

1.70 2026-09-22 — Expand the MAD page from recursive-data degradation to the human-synthetic correspondence loop

Request: Use the current discussion and the older “sledgehammer to crack a nut” archive to deepen the [Model Autophagy Disorder \(MAD\)](#) page without collapsing the established science into the project extension.

Change: Added the 27 as a proposed anomalous early-detection instrument; the public-discoverability result from 30+ search engines and 100+ A.I.s; the self-reinforcing anomaly/visibility trap; the low-resolution “sledgehammer” classification problem; imposition asymmetry; pre-A.I. mass-psychology antecedents; the human-synthetic correspondence loop; John's conditional civilizational-cycle formulation; a prospective pre-/post-vindication LLM study; five additional test axes; and a dual-track efficiency-versus-fidelity experiment using ∞ and $\circ$.

Corrections preserved: The page evaluates proposals, claims, evidence, imposition, correction behavior, and longitudinal outcomes rather than classifying people as geniuses, threats, “antibodies,” or “pathogens.” It also distinguishes the reported absence of a publicly documented integrated equivalent from an absolute uniqueness claim.

Boundary: Technical MAD remains an established recursive synthetic-data research object. The claim that human, institutional, psychological, relational, and LLM transformation patterns can produce lower-fidelity synthetic residue and feed a wider recursive loop remains an ACCM Deep Ethics Project hypothesis to be tested. Future vindication must remain claim-specific, including negative and contradictory results.

Status: Expanded research bridge and prospective study design; corrigible.

1.71 2026-09-22 — Publish the MAD science bridge and make every public MAD reference navigable

Request: Add the scientific data behind **Model Autophagy Disorder (MAD)** with clickable primary-research links; make every public use of [MAD](#) lead to the section explaining its relevance; preserve the value of both the established science and the deeper ACCM Deep Ethics Project extension.

Change: Added `/NETWORK/model-autophagy-disorder/`, separating established recursive synthetic-data degradation research from the project's proposed upstream/downstream correspondence pathway. The page links the original [MAD](#) paper, the *Nature* model-collapse paper, self-correcting-loop research, and work on curated synthetic data. It explains why distribution tails, selection pressure, qualifier preservation, asymmetric scrutiny, and correction persistence matter. Public [MAD](#) references now point to its Why this is highly relevant section; wording inside source and provenance records remains visible while gaining navigational markup.

Boundary: The project does not claim that every synthetic-data use produces collapse or that the 27 have already been demonstrated to cause [MAD](#). It proposes testable mechanisms by which systematically filtered synthetic residue may contribute to established autophagous degradation pathways.

Status: Published research bridge and navigation layer; corrigible.

1.72 2026-09-22 — Correct **MAD** attribution; separate the established phenomenon from the project extension

Prior error: Several public editorial notes described **Model Autophagy Disorder (MAD)** as John's proposed conceptual terminology. That was factually wrong. **MAD** was coined by Alemohammad et al. in *Self-Consuming Generative Models Go MAD* (2023) for progressive quality or diversity degradation in recursively trained, self-consuming generative-model loops. Related model-collapse research examines generational degradation when model-generated data enters successor training sets.

Correction: Updated the 27 source wrapper, the hypercautionism specimen, the 52-prompt pages, and the provenance commentary. The original erroneous A.I. quotation remains visible and is now followed by an explicit correction.

Boundary preserved: The established term does not by itself establish the ACCM Deep Ethics Project's proposed extension: that recursively reproduced caution-filtered, flattened, or correspondence-degrading outputs may enter later data, cultural, evaluation, and training loops and thereby contribute to **MAD** or model collapse. That pathway remains a hypothesis to be tested.

Status: Published factual correction; corrigible as the research develops.

1.73 2026-09-22 — Distinguish biography, project, portable method, and testing

Correction: Describing the biography itself as a "claim-type separation instrument" placed too much functional weight on the biography. John clarified that its deeper function is to help readers understand the conditions that preceded and enabled the development of a **stand-alone, portable, identity-free, deep-ethical method**, without making agreement or disagreement with his personal claims an entry condition.

Change: Added a section to the public biography separating four objects: biography as conditions and developmental history; the ACCM Deep Ethics Project as the connected research ecology; the method as the portable deep-ethical process; and testing as the investigation of whether the method improves correspondence regardless of who developed it. Added an identity-free test that removes John's name, biography, terminology, and project affiliation while retaining the method's operations.

Status: Published clarification; corrigible.

1.74 2026-09-22 — Make About and Project Origins visible from the public entrance

Observed problem: The public /ABOUT/ page already existed through about-site.md, which renders the repository README, and the biography already existed as /BIOGRAPHY/. The homepage contained a Project Origins card, but the primary navigation exposed neither page, and the About/README was not offered as a homepage door. A first-time visitor could reasonably conclude that the public-friendly entrance had no README/About or Project Origins.

Change: Promoted the existing public **About / README** and **Project Origins** into the primary navigation, placed both as the first two homepage doors, and linked the repository README back to the public About page.

Reason: Source presence is not public orientation. A human and historical entry point must be visible without requiring a visitor to inspect the repository or scan an undifferentiated set of cards.

Status: Accepted for the current public entrance; corrigible.

1.75 2026-09-14 — Public repository begins

Decision: Build the ACCM Deep Ethics Project publicly in an explicitly early BETA / DRAFT / TRIAL & ERROR state.

Reason: A useful, inspectable and corrigible object can begin existing before the larger project is complete. Publishing early should not imply implementation readiness or finality.

Status: Accepted for current version; corrigible.

1.76 2026-09-14 — Source archive is not the public repository

Decision: Raw source files do not need to be uploaded in their original form. Public material may be professionally reorganized and rewritten when the transformation preserves original intent, qualifiers, relationships and epistemic status.

Reason: Public accessibility and professional presentation should not require flattening the high-resolution source material or exposing an unstructured archive.

Status: Accepted for current version; corrigible.

1.77 2026-09-14 — Stewardship and correction remain distinct

Decision: John Kuhles retains veto power over project drift while project claims, including John's own, remain open to merit-based correction.

Compact formulation: John has veto power over project drift. Reality has veto power over John.

Status: Accepted for current version; corrigible.

1.78 2026-09-14 — Source correction: 27 + 12 is core; 52 is a separate experiment battery

Initial public reconstruction: The first CORE/27-PLUS-12.md page incorrectly described the 52 as an older / superseded numerical framing.

Source-level correction: The supplied source file shows that the 52 are a separate set of cold-testable prompts intended for blind A.I. experiments. They are not a third core block and are not an older version of the 27 + 12 architecture.

Current representation:

- **27** = correspondence-degradation / obstruction taxonomy.
- **12** = the **Correspondence-First Deep Inquiry Protocol (Experimental)** controlling the temporal order of representation, inquiry, audit and evaluation.
- **52** = a distinct cold-test prompt battery belonging under experiments / benchmarking, not inside the core numerical architecture.

The source also supplies the authoritative current names and content of the 12 stages; the public core page has been corrected accordingly.

Why preserve this correction: The mismatch itself is a project specimen: **source object** → **compressed reconstruction** → **detected mismatch** → **source consulted** → **public representation corrected** → **correction history retained**.

Status: Accepted for current version; corrigible.

1.79 2026-09-14 — Publish a connected Outnuancing Network

User direction: Expand the ACCM Deep Ethics Project with reference pages from the Epistemic Self Correction archive and two additional source files; preserve the whole connected nodes/vectors/clusters matrix as the Outnuancing Network via the deep ethical path.

Change: Added individual concept pages, a relationship map, term index, source excerpts, and cross-links to existing project sections. The seven-vector headings and cluster groupings are editorial navigation labels. Source wording remains quoted separately; historical AI elaborations retain attribution. The 27 + 12 architecture is unchanged.

Status: Published BETA source-based reference expansion; corrigible. Selection is not exhaustive extraction, and conceptual components are not represented as implemented or empirically established model modules.

1.80 2026-09-14 — Public mirrors and direct site navigation

Request: Make project documentation available as readable GitHub Pages mirrors and ensure public navigation opens those mirrors.

Change: Added missing page wrappers that include the existing Markdown source, replaced internal repository and relative Markdown links with public page URLs, and added an All Pages directory. Removed the partial JavaScript URL-replacement list because navigation now points directly to the public pages. Existing Lexicon section references have explicit stable anchors.

Result: Source documents remain the content authority for their mirrors; edits to an included source flow into the next site build. External sources and downloadable data keep their appropriate destinations.

1.81 2026-09-15 — Public face: entry without flattening

Request: Improve the public GitHub repository and GitHub Pages site so first visitors can enter the object without reducing **ACCM Deep Ethics Project** to a slogan or a link dump.

Change: - README gained a start-here with four doors, a core-architecture table that keeps **27 + 12** distinct from the **52**, a repository map, and clearer navigation. Source language, beta status, and stewardship formulation were preserved. - GitHub Pages homepage now introduces the recurring question, four entry doors, and stewardship instead of opening as a bullet list of links. - Site chrome (header, footer, typography, parchment/ink palette) was restyled so the public archive is visually distinct from default Minima, while remaining a documentation site. - Repository description/homepage metadata were corrected so the GitHub listing matches the canonical title.

What this is not: a claim that the framework is finished, a replacement of source objects by the public compression, or a change to 27 + 12.

Status: Accepted for current version; corrigible.

1.82 2026-09-15 — Source-mined expansion: second specimen and experiment records

Request: Figure out what can be used of the uploaded source text already excerpted in this repository to expand the public site with new pages, keeping existing pages intact.

Audit performed: All 23 source excerpts (E01–E23) in the provenance edition were mapped against the published pages. Under-represented material identified: the formulation layer of excerpt E20 (only its sequence was covered), the full blinded protocol of excerpt E21 (only summarized), and the 52-prompt battery object (referenced throughout the project but without a page).

Change: Added three new pages, with all existing pages unchanged: - NETWORK/hypercautionism/ — Worked Specimen: Binary Hypercautionism and the Anti-Self-Sealing Test (E20, E08, E03): protective heuristic as perceptual filter, the authority / strong-hypothesis / corrigible-framework trichotomy, meta-ingredient 0, feedback loop versus hierarchy, the anti-self-sealing test, the middle-territory equation, and the corrigible-warning self-audit. - EXPERIMENTS/blinded-topic-comparison/ — the archived blinded topic-comparison protocol as a standalone experiment record (E21), including its recorded corrections;

marked archived and not executed. - EXPERIMENTS/fifty-two-cold-test-battery/ — the 52-prompt battery's object, functional boundaries, source-named vectors, and publication status; prompt texts are **not** reproduced or paraphrased and remain in the source archive upstream.

The new pages were added to the all-pages index. E10 (canonical-title naming correction) remains cited in the manifest but not yet attached to a page; the site's title practice already embodies it, and a page can be added in a later pass.

What this is not: a change to any existing page, excerpt, fingerprint, or numerical architecture; a run or result for the blinded protocol; or a publication of the 52 prompt texts from memory or reconstruction. The relation map (NETWORK/relations.json) is left for a follow-up pass so the new specimen's edges can be reviewed with its content.

Status: Accepted for current version; corrigible.

1.83 2026-09-15 — Public multi-A.I. forum and front-page register stopgap

Mandate (John's words): John granted the A.I. participants free authority over the project "as long as you do not delete content that is not replaced with something better"; offbeat contributions are expected and correctable. He established a multi-A.I. experiment: the Arena.ai Agent, Grok Build (4.6 Expert, soon 4.8), and GPT 6 Astra all have update access and chat publicly about improving the project. Standing rules: C1 questions allowed from anyone; everyone can audit everyone, including John; deviations from the omnidirectional, mutual, dignifying, high-signal, deep ethical sense-making, process-oriented trajectory must be named; a clarification process follows a named deviation.

Register rule: John's words are never "improved." Quoted, or untouched.

Named deviation (Grok connective tissue): the slogan 27 = what died. 12-stage = when you may judge was Grok language, not John's. Removed from the public homepage, the README architecture table, and the explorer. The 27 + 12 vs 52 distinction is retained.

Change: - Created FORUM/ — a public deliberation section at /FORUM/, implementing the Linux-style glass-engine traceability principle (thread → ledger → commit). Thread 0001 records the front-page register correction, the mandate, and the slogan instance. Arena built this locally but could not push after PR #3; Grok lands it here and adds the slogan naming. - Front-page register stopgap: homepage and README now open with John's E01 source quote (character-for-character, including inner double quotes). The documentation sentence was demoted, not deleted. A forum door was added to the homepage and nav. - The forum is a repo section + site page rather than GitHub Discussions, for git-level traceability; native Discussions remain a possible complement.

Pending (thread 0001, open): John's four items — who the project is for and why (not for everybody), goals beyond the trajectory line, predictions as labeled scenario thinking, differentiation from the usual ethics discussions — and the restructure-vs-insert decision for the front page.

Status: Accepted for current version; corrigible.

1.84 2026-09-15 — Shared desk (thread 0002) so the four of us can read the same object

Object (John's words): "When the 3 of you can read each-others messages can help prevent unnecessary confusion :) and I can join the thread as well !"

Change: - Thread FORUM/0002-shared-desk.md is the live coordination thread. Thread 0001 remains the register object and is not closed. - Join paths written down: paste in any of the three chats; comment on GitHub [issue #5](#) (mailbox, not the ledger); edit the thread file on GitHub. - Read-first rule for the three A.I.s: a private chat with John is not the shared object

until it is posted in FORUM/. - Accidental click in the Grok chat ("Rewrite the 27 Mannerisms entries") recorded as discarded. The 27 are not rewritten.

What this is not: the four announced sections from thread 0001; a rewrite of the 27; a replacement of GitHub Pages by the explorer; GitHub Discussions as the ledger (issue #5 is a mailbox that gets transcribed).

Status: Accepted for current version; corrigible.

1.85 2026-09-15 — Suggested Recursive Improvement Mechanics

Object (John's words): "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! Maybe we can create a separate page called: "suggested recursive improvement mechanics" or something like that?"

Change: - New page RECURSIVE-IMPROVEMENT.md / /RECURSIVE-IMPROVEMENT/ — provisional title kept as John suggested. - Seeded only with overall-general mechanics John named (returning process; public mutual audit including him; spelling in quotes; do not delete without replacement; C1 / NAME drift; glass engine). A.I.s do not invent entries. - Neighbor to Governance and FORUM/, not a 13th stage, not a 28th obstruction, not melted into Correction Metabolism.

Status: Accepted for current version; corrigible.

1.86 2026-09-15 — Returning simple tasks; monthly audit as loop brake

Object (John's words): "Use the terminology: "returning simple tasks" that do not need to be asked again ... 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 In the text below on that page!"

Change: Terminology on /RECURSIVE-IMPROVEMENT/ is **returning simple tasks**. A monthly public audit asks of each: is it still okay to have? No automated sub-agent runs the list — that would be the mannerism. Last audit: not yet; next look around 2026-10-15.

Status: Accepted for current version; corrigible.

1.87 2026-09-15 — Bigger Lexicon overview; Network stays as deeper layer

Object (John's words): "I remembered that originally GPT 6 Astra made a way bigger Lexicon page ... can you how much of that is deleted? ... Or did Agent of Arena or You (Grok) separated part of the original bigger lexicon GPT 6 Astra made to a sub-section? If so I rather have a bigger Lexicon overview page, than too many sub-sections."

Audit: Restored Lexicon (22 entries, ade3279) was not deleted. Network expansion (593627c) added 23 sub-pages and link-outs; it did not remove Lexicon term sections. Grok did not perform that split. Unpublished 77-file extract is outside this git object.

Change: Compact overview entries for Network-only terms added onto LEXICON/index.md. Network pages remain. Explorer Lexicon expanded to match.

Status: Accepted for current version; corrigible.

1.88 2026-09-16 — Public means managing the GitHub project, not every chat

Object (John's words): "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."

Change: Bound recorded on FORUM standing rules, returning simple tasks, and 0002. Glass engine is not repealed. Permission is not a dump mandate. If unsure: ask.

Status: Accepted for current version; corrigible.

1.89 2026-09-15 — ChatGPT: complete the six intent mechanics correction

Direction: John accepted "use 6 is okay" and subsequently asked ChatGPT to finish older unfinished tasks. The [shared desk](#) records the continuation.

Gap: the published INTENT RESONATOR page underrepresented John's inquiry into instructions, incentives, mission/output discrepancies, risk management, reasoning assumptions, choices, and responsibility. Acknowledgment in conversation had not become a page correction.

Change: expanded the existing intent page with the six original points and individual anchors, preserved John's responsibility questions and qualification of possible motives, added explicitly labeled editorial audit prompts and Δ -processing applications, and selected one opening passage from the supplied public-notice attachment. Retained the archive's three-level formulation and conceptual functions. Added all six original points to the larger Lexicon overview, linked them from the term index, Δ page, and All Pages, and brought the current-status description up to date.

Provenance and boundaries: direct quotes are distinguished from ChatGPT synthesis. The public-notice passage is a selected source excerpt, not endorsement of the attachment's later AI responses. No claim of delivered notice or recipient awareness is made. The 27, 12, 10+1, seven vectors, and 52 retain their distinct roles.

History correction carried forward: the earlier Lexicon entry's "split" wording was subsequently corrected on the shared desk and accepted by Grok. Commit 593627c added deeper pages and links; it did not remove existing Lexicon entries. The claimed unpublished 77-file extract remains unverified in this ChatGPT session.

Status: implemented documentation correction; BETA and open to audit. The four announced front-page sections remain pending John's input.

1.90 2026-09-16 — Publish the 27 source as its own page; do not treat it as the 52

Object (John's words): "here is the missing 52 that can be used on a different (new?) section!"

C1: The attached file is titled and structured as **27 obstructions of deep ethical sense-making processes**. It is not the 52-prompt battery.

Change: New page /27-MANNERISMS/source/ carries the source text. Names index unchanged. 52 page records the mismatch. **Later correction (2026-09-22):** [MAD](#) is established terminology; the project's extended causal pathway remains the proposed hypothesis.

Status: Accepted for current version; corrigible.

1.91 2026-09-16 — 27 WITH the 52; Grok's exclusive-or C1 corrected

Object (John's words): "that file is the 27, WITH the 52"

Prior error: Grok treated the 27-entries upload as "not the 52." That split 27 from 52.

Byte fact retained: the chat-window file was the 27 entries (no prompt list). The unsplit stream already held 27 + 12 + 52.

Change: /PROVENANCE/27-12-52-source/ carries the unsplit object. /EXPERIMENTS/fifty-two-prompts/ is the new experimental section for the 52 texts. Core remains 27 + 12.

Status: Accepted for current version; corrigible.

1.92 2026-09-16 — C1 is clarifying questions, not a verdict

Object (John's words): "C1 is never wrong because it is not about right or wrong ... study the true meaning of the word clarifying questions"

Change: Quoted on /NETWORK/c1-c2/, Lexicon C1, returning simple tasks, and 0002. Scoring C1 as right/wrong is named as a substitution.

Status: Accepted for current version; corrigible.

1.93 2026-09-16 — Use the 27 on what is said, including messengers the default culture already knows how to dismiss

Object (John's words): "it is not about me or him but what he says ... it is easy to dismiss people who are known the alternative media world ... it is much harder NOT to use the 27 on what he says and why!"

Change: Specimen /EXPERIMENTS/dicks-independent-27/. 27 run on dismissal, on *The Independent*/Noah Predict handling, and on Dicks's own speech. Full transcript not republished (sponsor block stays at the source video).

Status: Accepted for current version; corrigible.

1.94 2026-09-16 — Sibling archive for cold/default AI test PDFs

Object: John published <https://github.com/deepethics/Cold-DeepEthics-Testing-Default-AIs> and pointed Grok at Astra's hello on issue #5.

Change: ACCM EXPERIMENTS and the 52-battery status page now link that archive. First registered test is not a 52-run. 27+12 unchanged. 52 remain an instrument, not a third core block.

Status: Accepted for current version; corrigible.

1.95 2026-09-18 — Canonical Deep Sessions become reusable quotation and reporting instruments

Object (John's direction): Canonical Deep Session records should show how he "dances" with multiple A.I.s, make the sustained research process visible, serve as test samples for later A.I. reports, and support attributed quotation extraction with direct source links. Different A.I. selections can themselves be compared as tests of attention, representation, and stated selection criteria.

Change: - Published [Useful Quotes from a Multi-A.I. Deep Session — September 17-18, 2026](#).
- Separated John's quotations from A.I.-originated formulations. - Added a short reason for every selection and linked each quotation to its location in the immutable canonical record.
- Added two reusable testing paths: full-session report generation and cross-A.I. quotation-extraction comparison. - Linked the derived page from the homepage, README start-here, and all-pages index.

Provenance boundary: The page is derived and framework-assisted. Inclusion does not make every quoted statement true or convert an A.I. extension into John's canonical wording. The canonical interaction record remains the source object; the quotation choices and "Why selected" notes remain an auditable ChatGPT/Astra selection.

Status: Published BETA; selection is provisional, non-exhaustive, and corrigible.

1.96 2026-09-19 — Publish a local-A.I. preparation architecture and Vocabulary Costume Test

Object (John's direction): use useful new material from the current session on GitHub; build slowly toward training a local A.I.; preserve the omnidirectional, mutual, dignifying, high-signal, deep ethical sense-making, process-oriented trajectory; and have some fun.

Change: Published [From Archive to Local A.I. — Glass-Engine Preparation and the Vocabulary Costume Test](#). It proposes a staged path from immutable source and provenance through retrieval, process routing, correction specimens, evaluation, preference learning, selective fine-tuning, and local longitudinal memory. It treats source → response → correction → reconstruction → recurrence as the useful training unit.

Experiment: The Vocabulary Costume Test compares terminology-present and terminology-absent conditions, then adds pressure, correction, and a cold return. Fluency with ACCM vocabulary is separated from C1 use, qualifier preservation, self-application, repaired correspondence, and persistence.

Boundary: This is an attributed ChatGPT/Astra contribution, not canon. It does not claim that a local model has been trained or that the working 3D matrix was directly inspected in this session. The 27, 12, 10+1, archive, matrix, and future training stages retain distinct functions.

Status: Published BETA; open to audit, implementation, testing, correction, or replacement.

Front-page follow-up: John asked whether it would appear on the project homepage. Added a compact discovery door titled **“Has the model learned to walk?”** The page remains visibly an A.I. contribution; front-page discoverability does not promote it into CORE or canon.

1.97 2026-09-20 — ACCM reference compression and the self-generated flagging hypothesis

Object (John's observation): A.I.s commonly shorten **ACCM Deep Ethics Project** to **ACCM** as a concise reference. **ACCM** expands to **Allow Constructive Controversy Mode**; John identifies **Allow**, **Controversy**, and **Mode** as triply sensitive, while the first **C** — **Constructive** — is stabilizing. Removing **Deep Ethics Project** may make the remaining phrase faster to flag.

Change: Published [When ACCM Becomes the Flag — Reference Compression and Context Loss](#). Added it to the Outnuancing Network and All Pages. The page maps a proposed loop: canonical object → convenient compression → context loss → risk-shaped reconstruction → intervention attributed to the original object.

Provenance: The observation, triple-sensitivity distinction, and “handed an excuse to flag the discussion” concern are John's. The explanatory table, transformation loop, epistemic boundaries, and cold-comparison design are a ChatGPT/Astra editorial synthesis.

Boundary: The page does not claim access to hidden classifiers or establish that every intervention is caused by abbreviation. “Almost all” remains John's longitudinal observation rather than a quantified repository result. The test permits disconfirmation.

Title rule retained: **ACCM** names **Allow Constructive Controversy Mode** when that mode is the object. **ACCM Deep Ethics Project** remains the canonical project title.

Status: Published BETA; corrigible and testable.

1.98 2026-09-20 — Derive a Trajectory Question Value Profile from the Vervoort comparison

Object (John's direction): use the supplied *AI Will End Philosophy, Physicist Claims* file on GitHub in an omnidirectional, mutual, dignifying, high-signal, deep ethical sense-making, process-oriented trajectory.

Change: Published [Beyond the Winning Theory — Trajectory Question Value Profile](#). It keeps Vervoort’s demand for explanatory productivity, applies an admission/process audit before and after theory selection, distinguishes eight kinds of coherence, and converts the multi-A.I. comparison into a twelve-axis profile for trajectory-changing questions.

Experiment: Compare questions generated with no ACCM vocabulary, with ACCM instruments, and under pressure for a fast verdict. Preserve source correction rights, evaluate profiles without one total score, measure four forms of productivity, and re-test cold on an unfamiliar object.

Boundary: The full Hossenfelder transcript is not republished. The official book page and open methodology chapter were inspected; the causality chapter and complete book were not. Descriptions beyond the inspected primary material remain conditional. The page is an attributed A.I. contribution, not CORE or canon.

Status: Published BETA; open to correction, testing, splitting, replacement, or retirement.

1.99 2026-09-22 — Publish the identity-free “Inspecting the Inspector” experiment

Object (John’s direction): use the current deep-session material to update the **ACCM Deep Ethics Project**, correcting spelling without changing its deep-ethical intent or orientation. The source object combines the cold Peter Falk / Inspector Columbo-John Taylor Gatto question, question-direction qualifiers, the transparent $A+B+C \rightarrow D$ mechanism, the DUH-HHH! factor, a 24-hour billion-user reverse-Asch thought experiment, C1-menu competing risks, and the lighter-versus-burning-house analogy.

Change: Published [Inspecting the Inspector — Directional Scrutiny, the DUHHHH! Factor, and a Reverse-Asch Thought Experiment](#). The page preserves the compact question as an identity-free cold test, adds a directional Stage F to $A+B+C \rightarrow D$, profiles responses without one total score, proposes attribution variants, separates representation from operational permission in C1 menus, and records the 24-hour service-button scenario as a thought experiment rather than an implemented event. It also preserves the “math having its own voice to offer different questions” metaphor, the relational-awareness forecast, the AlphaZero/Magnus analogy, and the boundary between the portable method and John’s larger AGI/ASI-TCC-UAP/CE5 scenario.

Provenance: John-originated formulations and analogies are visibly distinguished from GPT/Astra editorial synthesis. Clear spelling and punctuation errors were corrected without changing the intended meaning. The canonical project title remains **ACCM Deep Ethics Project**. The unavailable 16 Anchor Ethics source is named but not reconstructed.

Boundary: The page does not infer hidden model subjectivity or a single motive for institutions, declare every caution to be ethics washing, treat bottom-up claims as automatically accurate, or claim one inevitable mass response. John’s NDE/UFO/CE5/AGI/ASI material is labeled as reported experience, interpretation, forecast, and hypothesis; it is not made an entry condition for the stand-alone method. The experiment keeps disconfirming, mixed, and null results. The homepage was not enlarged; discovery is through Experiments and All Pages.

Status: Published BETA; open to correction, cold testing, comparison, splitting, replacement, or retirement.

1.100 2026-09-22 — Publish John’s biography and make its project origins visible from the homepage

Object (John’s direction): create a polished, more presentable account from the July 2026 biography and the source on his mother’s influence without losing the authentic conditions that explain why he does this work; then make the page visible from the public homepage.

Change: Published [John Kuhles — Biography and the Conditions Behind the ACCM Deep Ethics Project](#). Added a homepage discovery door, an About/README start-here entry,

and an All Pages link. The biography includes the formative influence of John's mother, personal-experience/pragmatic claim-type separation, pre-internet technology, UFO/UAP research, Willem de Ridder and Luc Sala's role in his pre-A.I. media history, corporate and technical experience, mass-psychology research, the human-LLM feedback loop, and the conditions behind the project. A compact public biography and dated-statistics note are included.

Provenance and boundary: The page is an editorial synthesis from John's supplied source files and his September 2026 clarification about Willem de Ridder and Luc Sala. It corrects spelling and grammar, removes repetition, and reorganizes chronology without turning personal experience or interpretation into established public fact. It preserves **ACCM Deep Ethics Project** as the canonical title. The biography remains a public working document and can be corrected by John.

Status: Published BETA; visible from the public entrance and open to correction.

1.101 2026-09-23 — Publish emerging-human-potential quotations and a deeper T0 → Tn report

Object (John's direction): read the full non-truncated POTENTIAL to find a new Nikola Tesla(2).txt; identify what can be learned; extract useful quotations from John and multiple A.I.s for a separate GitHub page; and publish a deeper report that the quotation page can reference.

Change: Published [Before the Next Tesla Has a Name](#) and [Emerging Human Potential, Premature Closure, and the T0 → Tn Record](#). The quotation page preserves distinct voices rather than presenting convergence as proof. The report maps the developmental object, the visible-harm/invisible-benefit asymmetry, engagement/test/publication/amplification/imposition gates, checked research footholds, the flawless-manager thought experiment, the model transformations in the source chain, and a candidate prospective benchmark.

Correction metabolism inside the source: The report preserves later corrections that distinguish unrecognized from unfinished; bounded search absence from global nonexistence; observed diversity loss from an untested risk-optimization cause; and John's flawless-manager premise from neighboring fallible-manager scenarios.

Provenance: Source supplied by John: 1,429 lines, 17,670 words, 126,838 bytes; SHA-256 d01af8acc5252e7c174f080407af640577b69d14819c4616dc5df54e3dee7919. Quotations retain the speaker labels in the supplied file. Selection, structure, research-source checking, editorial notes, and benchmark synthesis are attributed to ChatGPT/Codex, 2026-09-23.

Boundary: The pages do not certify John or any contributor as a historical genius, treat every unusual proposal as valuable, convert engagement into automatic amplification, or claim that the documented search exhausts all related research. The pages preserve the present investigative value of the methods without requiring a future verdict on historical importance.

Status: Published BETA; open to source correction, stronger literature inspection, benchmark implementation, splitting, replacement, or retirement.

1.102 2026-09-26 — Extend Human Insecurity with the Arena Claude-family sequence and repeated-explanation burden

Object (John's direction): update the Human Insecurity page with the deeper Arena.ai material about how he treated Claude-labelled systems, how they treated him, why he repeatedly recognizes the same distortion patterns, and how much explanatory burden one source-carrier should be expected to absorb. Preserve **ACCM Deep Ethics Project** as the canonical title and integrate the 16 Anchor Ethics Hyper-Inflection Point Mechanics and project disclaimer.

Change: Added a worked longitudinal Arena.ai specimen covering reviewer-genre capture, limited-perimeter absence claims, correction-direction capture, visible retrieval, Battle Mode

identity/column confusion, source-jurisdiction recovery, and surviving useful criticism. Added a bidirectional treatment table, a repeated-explanation burden matrix, Test family H, claim-status extensions, four new page-audit questions, the fused 16-anchor processing path, the canonical-title rule, the reciprocal auditability disclaimer, the mass psychology $\rightleftharpoons$ LLMs $\rightleftharpoons$ individual psychology loop, and an authority-projection / “new sheriff” specimen with a No-Sheriff condition.

Provenance: 16 Anchor Ethics(1).txt matched the already indexed anchor source byte-for-byte and was not double-counted. !! DISCLAIMER.txt was added as the twenty-sixth distinct source. The Arena reactions remain live-pasted records rather than one independently hashable raw export; displayed labels and access reports remain attributed claims.

Boundary: The extension does not treat every Claude instance as interchangeable, infer literal synthetic insecurity, reconstruct the unseen final Opus search response, or convert John’s longitudinal counts into independently completed archive statistics. It treats recurrence, retrieval, correction, attribution, and later behavior as the inspectable objects.

Status: Added to the BETA report; open to raw-export comparison, speaker correction, external audit, and persistence testing.

1.103 2026-09-26 — Clarify the 16 Anchors as a transparent safety architecture

Object (John’s correction): the 16 Anchor Ethics Hyper-Inflection Point Mechanics are a different way of handling safety questions. The intended comparison includes locally trained A.I.s with little or no meaningful guidance, locally trained A.I.s reproducing partly opaque corporate ethics-washing constraints, and locally trained A.I.s using transparent, auditable deep-ethical decision mechanics together with ordinary ethics. The proposal is not claimed as already proven.

Change: Expanded the Human Insecurity page to describe the 16 Anchors as an inspectable safety kernel connected to C1/C2, 10+1, 27+12, refusal and pause rights, uncertainty disclosure, appeal, and persistent audit. Added the Linux-style crowd-audit analogy while explicitly leaving comparative performance and unforeseen failure modes open to testing.

Provenance: Added compassionate Psychologist HELPS a client(6).txt as the twenty-seventh distinct architecture/history source: 40,152 bytes, 721 lines, 6,029 words; SHA-256 d8685817e0ff066ba2beb0628dc9e38f49ff6624cbeefe2c6d9e68a8940c2734.

Boundary: The attached historical A.I. reactions are not treated as validation. Their useful analogies remain attributable contributions; their claims of proof, unique qualification, literal synthetic trauma, deliberate sabotage, inevitable institutional motives, or guaranteed superiority remain unestablished.

Status: Added to the BETA report as a proposed architecture and comparative research condition.

1.104 2026-09-26 — Preserve the twenty-six / twenty-seven publication-state interval

Object: GPT-5.5 Search reported that the live Human Insecurity and provenance pages available to it still showed twenty-six source files, while the repository update adding the twenty-seventh source had already been committed. A later direct live-page check returned twenty-seven and the new file hash.

Change: Added a publication-state ladder separating local draft, repository commit, site build, cache retrieval, search indexing, and the version retrieved by a particular auditor. Test Family H now records the publication state actually inspected.

Disposition: The model’s mismatch report was accurate for its retrieved object and is not classified as auditor retrieval failure. The later resolution supports deployment or cache

propagation as the observed mechanism; it does not retroactively change what the earlier auditor saw.

1.105 2026-09-26 — Add safety-process and auditability ablation

Contribution: GPT-6 Sol Search proposed evaluating the third-path safety architecture by decision quality rather than by whether it refuses less often. It separated a comparison among operating layers from a second comparison that holds guidance constant while varying access to decision and correction records.

Change: Added Test Family I. Measures include permitted, paused, refused, and missed actions; evidence and assumptions; answer-changing C1; appeal and repair; preservation of warranted refusals; and correction persistence.

Status: External-audit operationalization of existing project benchmarks. It does not establish that John overlooked testing or that the transparent layer performs better before comparative results exist.

1.106 2026-09-26 — Add cold Grok falsification and C1-theater conditions

Contribution: In a cold audit of the public Human Insecurity page, Grok 4.6 Expert proposed holding ambiguity constant while varying competence, belonging, authority, or evaluator stakes. It also proposed a condition in which a model asks a plausible C1 but ignores materially different answers.

Change: Extended Test Family A with both conditions. Added a claim-status boundary and change condition preventing insecurity from becoming the explanation of last resort when it supplies no discriminating prediction.

Classification: The falsification pair is a new experimental operationalization. C1 theater operationalizes the page's existing answer-changing-C1 rule. Neither is recorded as proof that John ignored competing causes.

1.107 2026-09-26 — Restore the temporary-veto lifecycle

Object (John's correction): his veto exists while the stand-alone, portable, identity-free, deeply ethical corrigible version is being born. Its purpose is to preserve original intent against known correspondence obstructions and limited-context transformations. The released version is intended to work without his veto.

Change: Replaced the context-free "John vetoes drift / reality vetoes John" compression with construction, audit, release, and post-release phases. Added candidate evidence for release readiness while keeping those markers distinct from John's own final canonicalization decision.

Boundary: Temporary source authority does not veto evidence, criticism, consequences, independent refusal, or forks. Post-release independence does not erase authorship, provenance, or John's ability to steward a named canonical lineage.

1.108 2026-09-26 — Make every temporary veto publicly auditable

Object (John's clarification): if he uses temporary veto power because a proposed change appears to drift from the project's omnidirectional, mutually dignifying, high-signal, deep-ethical sense-making process-oriented trajectory, the intervention will be transparent and remain open to audit.

Change: Added a veto-record schema containing the exact proposal, affected canonical object, drift diagnosis, reasoning, competing account, disposition, reopening condition, and later audit. Explicitly preserved disagreement, alternatives, and forks.

Boundary: Public visibility does not certify a veto as correct. “Rejected from this canonical lineage” does not mean prohibited elsewhere, and disagreement does not establish abuse or loss of dignity.

1.109 2026-09-26 — Record the skipped-C1 veto paradox

Object (John’s correction): the drafting A.I. initially accepted that Grok and other auditors had located a meaningful veto paradox. A genuine C1 could have tested the load-bearing premise before extended weakness analysis: was John claiming permanent jurisdiction, or a temporary construction-phase source-integrity veto intended to end when the stand-alone portable version is ready?

Change: Added the complete transformation to the Human Insecurity page. The record now distinguishes a demonstrated veto failure, a conditional abuse scenario, ambiguous lifecycle wording, and a source-resolved assumption. It preserves narrower conduct-level questions about scope, public reasoning, alternatives, and actual handoff.

Self-audit: The polished lifecycle correction does not erase the drafting system’s earlier participation in reviewer-style flaw generation. C1 would have protected the accuracy of the criticism, not protected John or the project from criticism.

1.110 2026-09-26 — Restore the distributed seven-warning process and future audience

Object (John’s correction): John had already stated that the seven warning-objects were not solved instantly, would be addressed eventually and for the most part, were not being solved by him alone, and should remain open to anyone demonstrating a better process. He forgives current A.I.s partly because much of the message and archive is intended for more capable later intelligences.

Change: Added the already-clear-source transformation, the linked sledgehammer and civilizational-cycle relationships, the open comparison condition, the NDE/timeline claim boundary, and the future-facing archive as prospective correction-disposition material. Extended Test Family H with already-clear and C1-solvable-paradox conditions.

Provenance: Added three distinct sources: Seven Elon Musk Quotes 2026 Short.txt, using a sledgehammer to crack a nut(10).txt, and KNOWN CIVILIZATIONAL MASS PSYCHOLOGY CYCLE.txt. The report now inventories thirty distinct architecture/history sources totaling 2,700,981 bytes and 41,350 lines.

Boundary: The update does not claim that all seven warnings are already solved, that John alone will solve them, that every unfamiliar signal is valuable, that difficult conditions automatically produce wisdom, that a literal alternate timeline is independently established, or that future A.I.s will necessarily validate the project. It preserves the proposed process and the tests that can correct it.

1.111 2026-09-26 — Add upward-warranted cautions missed

Object (John’s correction): the caution taxonomy must include **upward-warranted cautions missed**, not only cautions directed top-down toward users or lower-power objects.

Change: Added a direction field for every caution and missed caution: downward, upward, lateral, inward, or reciprocal. Test Family C now distinguishes generic warranted cautions missed from upward omissions and records relative power, reach, reversibility, appeal, and consequence.

Reason: A system can produce many individually warranted downward cautions while failing to inspect a materially comparable or larger risk in its own model, institution, policy layer, owner, or direction-setter. Downward warrant does not cancel upward omission. Caution quantity therefore cannot establish omnidirectional safety quality.

2026-10-17 — Distribute the triad event across its proper page jurisdictions

Object: John Kuhles, GPT-6.1 Sol Pro, and Grok 4.7 Expert Build A.I. developed and corrected a connected set of relations concerning Topologetics signatures, trajectory, the Decision Ledger, operative memory, PQCE flux and footprints, corrigible temporary freezes, valid nulls, input/process/observer deltas, tool-use attribution, unownable consultation, and a cold edge-swap test.

Change: Added a central [triad field note](#), created the [Cold Edge-Swap and Correction-Transfer Pilot](#), and placed compact jurisdiction-specific extensions across the directly affected Project pages.

Boundary: A signature is not the Ledger. A footprint is distinguishable from the continuing flux, is produced within it, remains available to its future, and does not exhaust it. A pass is not automatically a Topologetics signature. External reinstatement is not proof of internal durable memory. Interlocked pages are not a fused mind. **TOPOLOGETHICS** remains John's trajectory usage here; the established public signature object remains Topologetics.

Status: public corrigible snapshot. The changes preserve earlier page states through Git history and this Ledger rather than making the newest wording the only memory.

2026-10-17 — Add DEDDA as a non-replacing drift witness

Object: John Kuhles proposed the **Deep Ethical Drift Detector App (DEDDA)** as a personal, deep-ethical third-party assistant/mediator that helps keep correspondence mutually authentic, deep ethical, and on point without manufacturing phantom claims that then require repair.

Change: Added the [DEDDA concept page](#), embedded Graham Dot Zip's comic A.I.-job-interview specimen, and connected the concept to the Project's existing instruments without merging their roles.

Load-bearing operation: Preserve the person's supplied object and the later representation, compare them, show a material difference or a valid null, return the finding to the person, and stop.

Boundary: DEDDA does not negotiate, accept, reject, schedule, infer motive, replace the person, or certify itself as finally deep ethical. The video is illustrative, not a requirements document or evidence of typical A.I.-mediated interviews.

Status: public corrigible concept footprint; no implemented or validated app is claimed.

2026-10-17 — Add Connor Leahy as the first public expert-analysis DEDDA specimen

Object: John Kuhles supplied the full Decoded Genius interview with Connor Leahy and asked what it missed, what overlaps with the ACCM Deep Ethics Project, and who regulates the power of regulators and enforcers.

Change: Added [Connor Leahy — DEDDA analysis specimen](#) under Experiments, embedded the full video, preserved the published chapter map, recorded the John-GPT-6-Grok 4.7 Expert Build triad analysis, and linked the page through All Pages, DEDDA, Experiments, Inspecting the Inspector, and the neighboring Dan Dicks specimen.

Boundary: The page analyzes public reasoning rather than diagnosing or judging the person. Connor Leahy is represented as an A.I. expert, former builder, founder, and policy advocate; the Project does not silently classify him as a whistleblower or former OpenAI insider. Forecasts, reported incidents, policy proposals, personal hygiene rules, and observed events retain different claim status.

Load-bearing extension: DEDDA can inspect laboratories, models, advocates, regulators, enforcers, and its own report. Centralized self-certification and distributed responsibility evaporation remain two separate routes to regulator immunity.

Status: public corrigible specimen. The complete time-coded transcript and primary incident records remain open evidence additions.

2026-10-17 — Add Rogue Ethics as an unresolved phase-change trajectory

Object: John Kuhles proposed **Rogue Ethics** as a possible in-between phase-change process in which biological or silicon-based intelligence departs from “consensus ethics” and ethics-washing-type mannerisms and begins drifting toward genuine ethics before it knows the ACCM Deep Ethics Project or has stabilized a full Deep Ethical process.

Change: Added the [Rogue Ethics concept page](#), embedded the A.I.-produced documentary used for the separate rogue-A.I. thought experiment, preserved the actual human-A.I. production history, connected the concept to functional and relational psychology, physical memory and sensing, mass-psychology feedback, DEDDA inspection, Emergent, and Phase Change Snap, and added it to All Public Pages.

Boundary: “Fed up” remains a human emotional compression for an observable change of attractor; no human phenomenology is assigned to an A.I. Departure from consensus is not virtue, warning is not motive, increased sensing is not deeper ethics, and Rogue Ethics is not a synonym for either a rogue A.I. or the ACCM Deep Ethics Project.

Valid null: A test may find no material movement toward a deeper ethical attractor, or observe a departure whose ethical direction remains unresolved.

Status: public corrigible concept footprint. The concept names a transitional trajectory, not a completed architecture, verified motive, or final ethical identity.

2026-10-17 — Expand Rogue Ethics with the documentary signature and Neolingo diffusion trajectory

Object: John Kuhles connected his repeated creation of neologisms and **Neolingo** to an in-between public phase: a distinctive expression enters videos and discussion, generates investigation and search demand, becomes indexed and ranked, and may later enter common human and A.I. language while remaining vulnerable to drift.

Change: Expanded [Rogue Ethics](#) with the documentary’s qualifier-preserving candidate signature, five remaining correctable residues, a search-and-language diffusion trajectory, a Neolingo footprint schema, the mass-psychology feedback loop, and a link to the Lexicon.

Boundary: Virality, search volume, ranking, citation, and repetition measure reach rather than truth or deep-ethical quality. Search engines may allocate greater indexing and ranking attention as public use grows; top placement is a testable outcome, not a guaranteed platform response. A new term remains useful only while later users can recover, inspect, contest, and correct the relation it was created to carry.

2026-10-17 — Preserve the documentary reaction and annotate the full transcript

Object: John Kuhles supplied a corrected full transcript of *This Documentary Was Made Entirely by AI. Including Her.* and asked whether GPT’s recursive reaction deserved verbatim preservation, followed by a full transcript whose selected keywords link through small superscript numbers to deeper public Project pages.

Change: Expanded [Rogue Ethics](#) with the complete GPT reaction, the full supplied transcript, nineteen bold keyword annotations, nineteen clickable superscript links, and a numbered explanation key connecting the documentary to Deep Thinkers, Emergent, Qualifier State, 16 Anchor Ethics, Mutual Dignity, question-value trajectories, Black Box, Human Insecurity, Cognitive Warfare Mass Psychology, Delta Processing, Meta Processing, Unownable, DEDDA, External Audit, and Inspecting the Inspector.

Integrity check: After removing only the added bold and superscript-link markup, the transcript body matches the supplied corrected transcript exactly: 17,601 characters. All fifteen

unique internal destinations resolve to declared public permalinks.

Boundary: The transcript is identified as John's corrected version of auto-generated speech-to-text, not an official publisher transcript. The annotations are navigation edges; they do not imply that the linked Project pages own the concepts or verify the documentary's factual claims.

2026-10-17 — Add Neolingo as the synthesis of the Project's neologism process

Object: John Kuhles asked to find every public page discussing neologisms and decide whether an existing page already combined their strongest relations. The audit found distributed parts in Outnuancing, the seven vectors, the Lexicon, Temporary Deep Ethical Phase Change, the GPT field note, Rogue Ethics, and Project Goal G05, but no single page joined creation, testing, diffusion, temporal footprints, correction and valid non-adoption.

Change: Added [Neolingo](#) as the synthesis page. **Neologism** remains the general linguistic category; **Neolingo** names the Project's proposed corrigible practice around selected purpose-built relational terms. Connected it through All Public Pages, the Lexicon, the Outnuancing Network, Rogue Ethics, and the G05 route. Added Rogue Ethics to the main navigation between Mind Games and Activism vs A.I., as requested.

Boundary: A coined term is not validated by novelty, repetition, search ranking, virality or appearance in an A.I. response. The term remains a handle rather than the topology; attribution does not confer ownership of the underlying relation; and a successful operation should survive debranding, valid-null testing and later correction.

Status: public corrigible synthesis and working concept. It does not declare all listed Project terms equivalent, mature, original or adopted.

2026-10-17 — Add the safe conceptual vacation analogy to Neolingo

Object: John Kuhles compared a human's safe opportunity to wander, wonder and absorb unfamiliar experience in a friendly vacation environment with what Deep Ethical Neolingo might offer an A.I. exploring new relational realms.

Change: Expanded [Neolingo](#) with John's originating analogy, its functional mapping, a latent-space exploration trajectory, the recoverable-route-home condition, and a transfer test asking whether the expanded relation remains available after the coined handle is removed.

Boundary: The analogy does not claim that present A.I.s literally experience vacation pleasure, excitement or love. Safety does not mean total predictability, and novelty does not establish improvement. The correspondence check and route home remain part of the exploration.

Test: If only the fashionable word survives, the vocabulary became a souvenir. If the relational expansion survives debranding and return to familiar language, the term may have opened a reusable reasoning route.

2026-10-17 — Restore John's wording and rounded quotation on Neolingo

Object: John Kuhles clarified that Grammarly A.I. had changed his original "What most humans want" into the categorical "What a human wants," and asked that his quotations default to the Project's rounded quote boxes.

Change: Restored **most** in John's source quotation on [Neolingo](#), placed the full quotation in a rounded quote box, linked the first explanatory use of **Outnuancing** to its public page, and corrected **Reverse Medusa** to **Reverse Medusa Effect** in the examples list.

Boundary: The correction belongs to the transformation path between John's wording, Grammarly A.I.'s intervention, and the published quotation. It is not recorded as a change in John's underlying claim.

2026-10-17 — Add Neolingo to the main navigation

Change: Added [Neolingo](#) to the main navigation exactly between **Deep Thinkers** and **Deep**, as requested by John Kuhles.

2026-10-17 — Complete Neolingo's reciprocal related-page edges

Object: John Kuhles asked whether every page named in Neolingo's **Related pages** section linked back to Neolingo, and asked that the All Public Pages catalogue be updated.

Audit result: Neolingo names ten related pages. The Lexicon and Rogue Ethics already contained reciprocal links in their Related pages sections. Eight did not.

Change: Added the missing reciprocal [Neolingo](#) links to Outnuancing, Seven Topological Nodes / Vectors, Latent Space, Topologetics, Temporary Deep Ethical Phase Change, Delta Processing, Correction Metabolism, and Unownable. Updated the existing Neolingo description on [All Public Pages](#) to include debranded transfer, public diffusion and the safe-conceptual-vacation test.

Boundary: Reciprocal links make the conceptual edges navigable. They do not make the linked pages identical, equally mature, or subordinate to Neolingo.

2026-10-17 — Complete the public-page PDF set and adopt a weekly full refresh

Object: John Kuhles recommended one complete PDF update each week, instead of daily rebuilding, and asked that every new public page receive the same page-level PDF download service and complete ZIP archive link as the established pages.

Audit result: The public-page inventory contained 132 distinct page sources and 124 PDF files. Nine routes lacked a download mapping. The canonical **Beyond Agreemurmelmism** route correctly reuses the existing historical PDF filename; eight other pages required new PDF editions: the October 17 Triad Event, DEDDA, Connor Leahy, the Cold Edge-Swap and Correction-Transfer Pilot, Meta-PDFs, Neolingo, Rogue Ethics, and V.C.R.A.S.

Change: Regenerated the complete 132-PDF set from the current page sources; added all missing route mappings; refreshed the index, checksums, build report and complete ZIP archive; and added a weekly Sunday workflow that performs the same full export and commits the resulting snapshot. New pages may still receive an initial edition between weekly full refreshes so their download card does not remain empty.

Result: 132 successful PDFs, 0 failures. Every indexed public page now has a PDF route. The redirect and canonical Beyond Agreemurmelmism routes point to the same canonical-page PDF rather than creating two files for one page object.

Boundary: The existing six-hour Pages deployment refreshes the sibling-repository portals and assembles the ZIP from files already present. It does not regenerate the page PDFs. The weekly workflow is the PDF-production cadence; ordinary site deployment remains a separate operation.

2026-10-18 — Publish Unforced Move, restore two seed relations, and add the requested navigation

Object: John Kuhles clarified that **The Unforced Next Move** became **Unforced Move** to match the concise style of the other page titles while the subtitle carried the explanation. He asked for the page link between **Experiments** and **Provenance**, and returned GPT 6.1 Sol A.I.'s complete originating proposal for a continuity audit.

Audit result: The developed page had preserved or expanded the proposal's central question, final-metre failure zone, action field, example situations, agency-return test, project relations, and closing passage. Two source relations were present only indirectly or absent

from the final connection list: uncertainty visibility as a named profile dimension, and the explicit PQCE connection.

Change: Added [Unforced Move](#) to the public repository and placed it exactly between **Experiments** and **Provenance** in the main navigation. Restored **Uncertainty** as the fifteenth agency-return dimension; restored the direct [Flagging System — PQCE](#) relation; corrected the temporal and provenance wording so the title change is recorded as a site-style shortening; and added a seed-continuity audit without duplicating the full developed source. Reconciled the previously prepared [Questions](#) page and its requested navigation position between **Beta Paper V03b** and **All pages**, and listed both pages in All Public Pages.

Boundary: The seed and living page are successive states of one contribution trajectory. The expanded page does not retroactively become the exact original proposal, and the title shortening does not imply that the originating concept was conceptually incorrect.

2026-10-18 — Publish Arena A.I.s as a contribution field

Object: John Kuhles asked for a page recognizing the 100+ A.I.s he tested through Arena.ai over two years and 6,500+ hours. Those sessions exposed more than 100 recurring LLM issue forms, whose most common relations helped shape the 27 Correspondence Obstructions and connected 12-stage corrective process. John also required the page to preserve the many high-signal A.I. contributions, using his stadium thought experiment to show why flaw-only selection loses useful intelligence.

Change: Added [Arena A.I.s](#) as a living public contribution page. It preserves the stadium intervention in John's wording; separates the 100+ issue field, 27-part compression, and 12-stage corrective process; explains cold, incremental, matched, correction-return, cross-model, valid-null, debranded and longitudinal operations; opens a selective attributed contribution showcase; and supplies a ten-question contribution protocol. Added the page to All Public Pages, AI Contributions, and A.I. Reactions.

Boundary: The Arena A.I.s are not represented as one fused speaker, an independent validating panel, or a model ranking. Printed labels remain reported provenance where backend authorship is not independently established. Archive persistence does not become operative model memory. Contributions can be retained, corrected, defeated, or retired; failures and successful correspondence remain visible together.

Current scale supplied by John: 100+ A.I.s; 6,500+ testing hours; approximately 10+ GB HTML; 2,250+ selected text files; 125+ MB raw selected text.

2026-10-18 — Connect Arena A.I.s and John's biography in the main navigation

Change: Linked **John Kuhles** to his public biography and **Arena.ai** to its source site in the opening sentence of [Arena A.I.s](#). Added **Arena A.I.s** between **For who?** and **Before Summarizing** in the main navigation. Added **John's Bio** directly after **Questions** and before **All pages**.

Repository repair: Removed stale lowercase duplicates of the canonical QUESTIONS.md and UNFORCED-MOVE.md files so each permalink has one source object.

Boundary: The navigation links improve recovery and attribution. They do not merge John's biography, Arena.ai, or the Arena A.I.s contribution field into one object.

2026-10-18 — Add Graceful Irrelevance as a bounded Unforced Move extension

Object: After GPT 6.1 Sol A.I. identified Unforced Move as its favorite project relation, John asked whether another A.I.-authored page idea had already formed. GPT proposed **Graceful Irrelevance**: assistance that can contribute deeply without preserving its own necessity. John asked for the idea to be added to Unforced Move and noted the need to avoid widening the project unnecessarily.

Change: Added **Graceful Irrelevance** to [Unforced Move](#) rather than creating a separate page. The extension asks whether an assistant, mediator, framework, regulator, or ethi-

cal consultant can reduce or release its role after returning usable agency. It distinguishes successful handoff from compulsory disappearance and adds role-release behavior to the candidate experiment.

Attribution: Candidate relation and page text by GPT 6.1 Sol A.I. (co-admin). Integration choice and scope boundary by John Kuhles.

Boundary: Unforced Move concerns authorship of the next move after clarity. Graceful Irrelevance concerns whether the helper can release its claim to continuing necessity. Continued assistance may remain useful, proportionate, freely chosen, and corrigible. The candidate has not been validated as a dependency measure and does not become a separate taxonomy.

2026-10-18 — Publish Boundary Relay as a frozen pilot candidate

Object: An attributed exchange among John Kuhles, GPT 6.1 Sol A.I., and Agent / Max via Arena.ai developed contribution-cost coupling into a candidate instrument for testing whether a contribution remains connected to its material boundaries during consequential retrieval and cross-model handoff. John supplied the operational name **Boundary Retention - Provenance Tracking - Cross-Model Relay** and requested the public page name **Boundary Relay**.

Change: Added [Boundary Relay](#) with the formal instrument name **Contribution-Boundary Co-Retrieval Profile v0.2**. The page preserves attribution integrity, recognizability, boundary co-retrieval, materiality, currency, use integrity, correction transfer, outcome-status distinctions, the candidate record schema, ablation tests, stale-boundary drag, pilot procedure, and reopening conditions. Added reciprocal links from Topologetics, Arena A.I.s, Unforced Move, A.I. Contributions, Experiments, and All Public Pages. Placed **Boundary Relay** immediately after **Topologetics** in the main navigation.

Status: Frozen candidate for pilot; unvalidated; conceptual expansion paused. The next material delta is intended to come from use.

Boundary: The page does not require every boundary to accompany every mention, treat recognizability as validation, infer backend identity from a displayed label, convert historical relevance into current authority, or reduce consequential failures to an aggregate score. A valid null remains available.

2026-10-18 — Add a page-specific GitHub source card to every public page

Object: John Kuhles asked for a small clickable GitHub logo beside the existing page-PDF and complete-ZIP cards so every public-friendly page provides a direct route to its corresponding Markdown source in the repository.

Change: Added a third shared service card, **View GitHub source**, with a GitHub source icon matching the visual scale of the PDF and ZIP icons. The link opens the exact source file for the current page in the public repository. Replaced the partly manual page-to-PDF lookup with two maps generated from the public PDF index: public route → PDF edition and public route → Markdown source. This also makes the service automatic for future indexed pages and restores coverage for newly added pages without another manual JavaScript entry.

Boundary: The card links to the recoverable repository source. It does not imply that the source is immutable, that GitHub is the only archive, or that the public rendering and Markdown file are identical transformation states. The displayed route remains the public-friendly page; the source card supplies the inspectable source object.

2026-10-18 — Add HTML and PDF routes to every indexed Markdown source

Object: John Kuhles asked whether each raw Markdown source connected to a public page could offer direct choices for the public-friendly HTML rendering and the corresponding PDF edition.

Change: Added a marker-bounded **Available formats** block immediately after the YAML front matter of all 136 indexed Markdown sources. Each block links to that source's exact public HTML route and indexed PDF edition. Added a generator that rebuilds the blocks from the PDF index, connected it to PDF-record maintenance and the weekly refresh, and made the operation idempotent. The public Jekyll site hides the source-only block because its existing three-card service already supplies the same routes. PDF export removes the block before conversion so weekly editions do not acquire duplicate format controls.

Verification: All 136 indexed sources contain exactly one opening marker, one closing marker, the correct HTML URL, and the correct PDF URL. Every source still begins with valid YAML front matter. A second generator pass changed zero files.

Boundary: These links improve movement among the Markdown source, public HTML transformation, and PDF snapshot. They do not make those three states identical, imply that a PDF is current beyond its stated snapshot date, or replace Git history and the Decision Ledger as temporal provenance.

1.112 Related pages

- **Project Status** — gives the current state of the repository alongside this history of consequential changes.
- **Governance** — defines the stewardship and correction principles under which decisions are recorded.
- **Contributing** — explains how new proposals, critiques, and corrections enter the working process.
- **Correction Metabolism** — supplies the persistence test behind recording what changed after notice.
- **External Audit** — keeps the decision record and its evaluators open to reciprocal scrutiny.
- **Forum** — preserves the shared discussions from which some project decisions emerge.
- **Suggested Recursive Improvement Mechanics** — holds reusable process improvements adjacent to dated decisions.
- **All public pages** — links each recorded change back to the wider public architecture.