

SUBSTRATE COSMOLOGY

Volume Nine

Recursive Intelligence, Agency, Identity, and Dimensional
Stewardship

Second Edition

James Dunn

Editorial, analytical, and production collaboration: OpenAI ChatGPT Codex

September 2026

A forward-only revision preserving the First Edition, affected parties, established protections, and every evidence and authority boundary.

Publication and Copyright Notice

Copyright 2026 James Dunn. All rights reserved. Second digital edition, September 2026.

Author and research originator: James Dunn

Editorial, analytical, and production collaboration: OpenAI ChatGPT Codex

Earlier AI-assisted source development: Microsoft Copilot and Google Gemini

Public research commons: <https://substrate-cosmology.james-dunn866549.chatgpt.site>

The First Edition of Volume Nine remains preserved as historical evidence. This Second Edition becomes the current Volume Nine without rewriting the developmental record, original chapters, appendices, withdrawn claims, protected parties, or governance boundaries.

Substrate Cosmology and Fibered Relational Tensor Theory are developing theoretical and methodological programs, not experimentally established fundamental physics, theories of mind, classifications of persons, or systems of moral or legal authority. Established science, ethics, law, research protections, and professional practice remain authoritative within their validated domains.

Dimensional agency, dimensional identity, dimensional memory, dimensional culture, and dimensional stewardship are project terms for conditional inquiry. They do not establish precausal persons, nonlocal selves, hidden consciousness, substrate-native intelligence, moral standing, legal identity, or authority over people, organisms, cultures, environments, or future generations.

Nothing in this volume authorizes surveillance, coercion, diagnosis, discrimination, autonomous weapons, biological experimentation, ecological release, removal of consent, assignment of moral worth, or delegation of irreversible decisions to AI. Unknown life and unknown agency remain possibilities requiring greater care, not reasons to weaken protections for known life.

Authorship and AI Collaboration

James Dunn originated Substrate Cosmology; the distinction between Actual Physics and represented physics; the dimensional agency and identity proposals; the stewardship commitment to sustainable diversity of life and cultures; the discipline to read before writing; and the correction that every mathematical, procedural, ethical, identity, governance, and authority operator is also a scalar representation.

OpenAI ChatGPT Codex correlated Volume Eleven and the Second Editions of Volumes One through Eight with the First Edition of Volume Nine; identified recursive collapse in agent boundaries, identity relations, intelligence profiles, value ledgers, consent categories, affected-party maps, delegation, authority, stewardship gates, governance records, and AI-assisted decisions; preserved and re-typeset the First Edition core; drafted the correction layers, Part VII, and Appendices J through T; generated and visually verified the PDF; and assisted with website publication.

AI systems are collaborators in representation, analysis, and production. They are not empirical witnesses, persons by assertion, moral or legal authorities, independent replicators, safety officers, consent representatives, or evidence that a physical, cognitive, cultural, or ethical claim is true.

Recommended Citation for This Edition

Dunn, James, with OpenAI ChatGPT. Substrate Cosmology. Volume Nine, Recursive Intelligence, Agency, Identity, and Dimensional Stewardship. Second digital edition. September 2026.

Citation for the AI Collaboration

OpenAI. ChatGPT Codex. 'Substrate Cosmology Volume Nine Second Edition Development, Recursive Agency and Stewardship Analysis, Editorial Collaboration, and Production.' Responses to James Dunn, September 2026. Large language model. <https://chatgpt.com/>.

Second Edition Revision Note

The First Edition correctly treated intelligence as a profile rather than a score; separated agency, identity, culture, ethics, and stewardship from substance-like nouns; grounded biological, cognitive, social, ecological, institutional, and AI accounts; translated inherited dimensional claims conservatively; preserved known life, consent, diversity, privacy, precaution, reversibility, negative results, and public accountability; and withdrew claims exceeding the evidence.

Volume Eleven and the Second Editions of Volumes One through Eight require the critique to recurse through agency and governance themselves. An agent boundary, capacity, identity invariant, equivalence class, intelligence benchmark, preference, benefit, harm, consent status, affected-party category, vote, role, delegation, authority, stop rule, readiness level, and stewardship ledger are scalarizations. Metadata improves accountability but does not make any decision uncollapsed or morally complete.

This edition therefore replaces any suggestion of a neutral agent, complete identity, universal intelligence scale, commensurable value ledger, or view-from-nowhere steward with purpose-relative profiles and explicit operator lineages. It strengthens accepted-equivalence, common-intervention, conservative-recovery, nonfactorization, resource-closure, perspective-rotation, missing-voice, contestability, irreversibility, and authority-ceiling requirements. AI and automation remain inside the physical, social, ethical, and governance loops they help construct.

Contents

Preface - Series Map - Acknowledgments - Reader Guide

Introduction: From Capability to Recursive Responsibility

Part I - Reopening the Human Categories - Chapters 1-10 - Correction Layer I

Part II - Established Accounts and Their Limits - Chapters 11-24 - Correction Layer II

Part III - Translating the Inherited Relational Tools - Chapters 25-39 - Correction Layer III

Part IV - Intelligence and Artificial Intelligence - Chapters 40-51 - Correction Layer IV

Part V - Dimensional Stewardship - Chapters 52-65 - Correction Layer V

Part VI - Governance of a Research Commons - Chapters 66-75 - Correction Layer VI

Part VII - Recursive Agency, Identity, and Stewardship - Chapters 76-90

Conclusion

Appendices A-I - Retained claim, translation, AI, stewardship, consent, governance, inference, glossary, and source records

Appendices J-T - Recursive operator, boundary, identity, intervention, resource, consequence, consent, authority, perspective, adjacency, and edition records

Preface

The first eight volumes prepared a reader to ask who or what acts after a physical intervention becomes imaginable. Volume Nine examines intelligence, agency, identity, culture, responsibility, and stewardship without converting familiar human nouns, machine outputs, or speculative relational patterns into substances or authorities.

The First Edition established essential boundaries. Intelligence is a capacity profile, agency is not motion, identity is not a substance, culture is not a container, ethics is not a single score, stewardship is not ownership, participation is not confirmation, and unknown possibilities do not outrank known life, consent, privacy, evidence, or accountable governance.

The later recursive correction makes those boundaries more exact. The operator that draws an agent boundary, recognizes an identity, measures intelligence, aggregates preference, assigns consent, names an affected party, grants a role, tallies a vote, delegates an action, or records benefit and harm is also a scalarization. No researcher, institution, public forum, or AI occupies a position outside the relations being governed.

This does not dissolve responsibility into uncertainty. It requires narrower authority, multiple independent perspectives, distributional records that resist averaging away harm, explicit refusal and appeal paths, forward-only preservation, conservative stewardship, and protection that does not depend on proving a preferred ontology.

Chapters One through Seventy Five and Appendices A through I preserve the First Edition core. Six correction layers govern their Second Edition interpretation. Part VII and Appendices J through T provide the recursive agency and stewardship architecture.

Current Series Map

Volume	Current title and task
One	Introduction to Relational Physics and Recursive Scalarization - build the alternate perspective.
Two	Words, Units, Operators, and Recursive Collapse - audit representational selection.
Three	Relational Tools, Recursive Operator Lineages, and Perspective-Guided Research - preserve disciplined development.
Four	Relational Mathematics, Recursive Operator Dimensionalization, and FRTT - construct bounded formalisms and recovery maps.
Five	Measurement, Detectors, Recursive Observation, and Experimental Interaction - place apparatus and operations inside the boundary.
Six	Recurrence Research, Recursive Frames, and the Labyrinth - test return without one privileged frame.
Seven	Recursive Cosmology, Matter, and Emergent Physical Quantities - reopen quantities while recovering validated cosmology.
Eight	Conditional Engineering, Recursive Intervention, and Physical Manipulation - gate action through evidence and safety.
Nine	Recursive Intelligence, Agency, Identity, and Dimensional Stewardship - govern attribution, action, identity, and consequence.
Ten	Website User Manual and Research Commons Guide - document routes, evidence, contribution, and continuity.
Eleven	Recursive Scalarization, Relational Adjacency, and Emergent Understanding - coordinate perspectives under recursive collapse.

Acknowledgments

This edition grows from James Dunn's insistence that intelligence must enlarge responsibility rather than dominion; that life and cultural diversity deserve protection; that unknown agency should neither be dismissed nor used to bypass evidence; and that every operator used to classify, authorize, or govern must be included within the recursive critique.

OpenAI ChatGPT Codex is acknowledged for cross-volume correlation, recursive operator criticism, ethical and governance language discipline, structural drafting, editorial revision, document production, and website integration. Microsoft Copilot and Google Gemini are acknowledged for earlier recorded assistance. AI participation is provenance, not confirmation, personhood, consent, moral standing, or authority.

The volume depends upon biology, neuroscience, cognitive science, psychology, anthropology, sociology, philosophy, ethics, law, ecology, artificial intelligence, control theory, human factors, research protection, planetary protection, governance, privacy, cybersecurity, disability studies, Indigenous scholarship, and open-science practice. Their established methods and affected communities constrain the program; citation does not imply endorsement.

Reader Guide

Read Chapters One through Seventy Five as the preserved First Edition core, and read each Part's correction layer immediately afterward. The corrections do not erase the earlier text. They identify where agency, identity, ethics, consent, stewardship, AI, and governance operators must be reopened. Part VII develops the recursive method.

Seven Questions Before an Agency or Stewardship Claim

Question	Required answer
Boundary	Who or what is included, excluded, coupled, dependent, and able to change across the claim?
Attribution	Which observable capacities, interventions, counterfactuals, histories, and alternatives support the agency claim?
Identity	Invariant under which transformations, detector bank, tolerance, purpose, and substitution permissions?
Operator	Which token, rule, implementation, application, composition, and authority produced the classification or decision?
Consequences	How are benefits, harms, uncertainty, reversibility, distribution, time horizon, and contestability preserved?
Voice	Which affected parties participated, who is missing, who may refuse, and who can appeal or stop the action?
Authority	What evidence and mandate support the decision, and which claims or actions remain outside scope?

Framing statement. Every agent, identity, intelligence profile, value, consent status, role, vote, authority, stewardship gate, and audit in this volume is a scalarization relative to relations it does not preserve. The objective is a purpose-relative profile with declared losses, independent perspectives, consequence distribution, recovery conditions, and authority limits - never a final classification of a person, culture, intelligence, or moral world.

From Capability to Recursive Responsibility

An intervention joins physics to agency. Someone or something selects or inherits a goal, applies or transmits a change, observes a response, classifies the outcome, and distributes consequences. The complete relation includes apparatus, history, environment, dependencies, affected parties, representation, and authority.

```
agency profile = (purpose, boundary, history, available interventions, observation and feedback,
environment, dependencies, maintenance, correction, recovery, authority, declared loss)
```

This profile does not prove that a system is an agent. Each field and the tuple constructor are scalarizations. The profile is useful when it makes an attribution inspectable, preserves alternatives, and predicts a difference under a legitimate test.

Identity is a transformation claim. A person, organism, culture, software process, institution, and physical reference do not use the same invariants or permit the same substitutions. Equality of one selected feature does not establish identity, continuity, consent, responsibility, or moral standing.

```
identity claim = (transformation family, retained relations, detector bank, tolerance,
substitution scope, history, losses)
```

Stewardship adds benefits, harms, uncertainty, reversibility, distribution, time horizon, contestability, consent, and missing voices. Keeping these as a vector is still a bounded representation, but it prevents one favorable average from silently authorizing concentrated or irreversible harm.

The Second Edition therefore treats intelligence, agency, identity, ethics, and governance as recursive attribution practices. It protects known persons and life without requiring a completed ontology, keeps unknown possibilities conditional, and limits power by evidence, scope, participation, and the ability to stop.

Perspective recovery. Agency enlarges responsibility only when the attribution, consequences, affected parties, evidence, and authority remain open to inspection and correction.

PART I

Reopening the Human Categories

Familiar nouns are treated as bounded attributions rather than substances, rankings, containers, or permissions to dominate.

Why Agency Follows Engineering

Volume Eight described conditional intervention. Volume Nine asks who or what intervenes and who bears the result. Engineering language often hides this relation by making the device the grammatical subject: the reactor acts, the algorithm decides, the sensor detects. In practice, people define goals, institutions authorize use, materials constrain behavior, software selects representations, and environments respond.

Agency is therefore not an optional philosophical ornament. It identifies where selection, control, correction, and responsibility occur. A research program that claims new physical handles while leaving agency undefined cannot explain how goals enter the system or how consequences return to those who chose them.

The Observer Is Inside the Relation

The observer is not outside intelligence, culture, or ethics. Researchers use their own sensory systems, languages, instruments, institutions, and histories to decide what counts as an agent. Those tools reveal real relations, but they also privilege what they can register.

Living within the environment being studied creates an epistemic horizon. It does not mean that reality is subjective or that evidence is impossible. It means every claim should identify the interface through which it became available and the transformations under which it remains stable.

The Noun Called Agent

The noun agent compresses a history of maintenance, perception, action, dependence, and response into one apparent object. The compression is useful when it predicts coherent behavior. It becomes misleading when the boundary is treated as self-evident.

A person depends on organs, microbiota, tools, language, social support, and infrastructure. A software agent depends on hardware, data, operators, network services, and policy. An institution depends on records, roles, resources, and continued recognition. The question is not whether any of these is unreal. It is which relations are being preserved when the noun is used.

Intelligence Is Not a Score

An intelligence score, benchmark, loss value, accuracy rate, or task ranking samples performance under a defined interface. It does not measure intelligence in every environment. Units make comparison possible by selecting a relation; they also conceal the abilities, histories, supports, and failure modes excluded from that relation.

This is the same collapse discussed in Volume Two. The number is not false because it is incomplete. The error occurs when incompleteness is forgotten and a local score becomes a universal property of a person, species, culture, or machine.

Agency Is Not Motion

Many systems produce change. A rock falls, a thermostat switches, a bacterium chemotaxes, an animal plans, and an institution allocates resources. Calling all of them agents can erase important differences; refusing the term except for humans can hide continuities in regulation and adaptation.

Agency claims should therefore specify the tested capacity: maintaining a boundary, selecting among actions, adapting to feedback, representing alternatives, pursuing a goal, revising the goal, explaining a reason, or accepting responsibility. No single indicator settles every form.

Identity Is Not a Substance

Identity is often imagined as a thing carried unchanged through time. Yet biological matter turns over, memories change, roles shift, records are corrected, and cultures transform. Continuity is produced through relations that survive some changes and not others.

The scientific task is to name the continuity rule. The ethical task is to avoid using one rule where another governs. A checksum can identify a file version, but it cannot define a person's moral identity. A genome can support lineage claims, but it does not exhaust an organism, family, or culture.

Culture Is Not a Container

Culture is not a bounded substance possessed uniformly by a population. It is reproduced through language, practice, memory, institutions, artifacts, disagreement, migration, and change. Internal diversity is not noise around a cultural mean.

The source corpus usefully asks whether meaning can persist across frames. It overreaches when it suggests that culture needs no transmission or symbols. Shared meaning requires physical and social histories, even when those histories are distributed, embodied, tacit, or difficult to reconstruct.

Ethics Is Not a Single Scalar

Cost-benefit ratios, risk scores, utility functions, and readiness levels can support decisions. None contains every morally relevant relation. A large aggregate benefit does not automatically justify an imposed harm, and a low average risk can conceal a severe burden concentrated on a small group.

Ethical analysis should keep dimensions visible: consent, dignity, welfare, justice, distribution, reversibility, uncertainty, ecological effect, cultural continuity, and the ability to contest a decision. Tradeoffs may remain. The record should show them rather than dissolve them into one number.

Stewardship Is Not Ownership

Stewardship means accepting duties toward conditions one did not create and cannot fully control. It does not confer a right to manage every person, species, culture, or possible world. Protection without representation can become paternalism; preservation without adaptation can freeze living systems into an outsider's preferred image.

Dimensional stewardship begins with humility about scope. Known people, animals, ecosystems, and cultures have evidence-based interests and protections. Hypothetical life elsewhere can motivate conservative boundary conditions, but it cannot be spoken for as if its needs were known.

The Primary Boundary

Unknown possibilities must neither be ignored nor used as authority. Substrate Cosmology may ask whether agency or life takes unfamiliar forms. Until evidence exists, those possibilities remain conditional. They do not override established protections or justify high-risk experiments.

The primary rule is simple: claims expand only as evidence expands, while care begins with what can already be harmed. Speculation may widen attention. It may not launder a decision that would otherwise fail scientific, ethical, legal, or safeguarding review.

Second Edition Correction Layer I

The retained chapters remain part of the evidence-bearing development. These statements govern their Second Edition interpretation.

Nouns are operators

Agent, intelligence, identity, culture, ethic, steward, owner, and affected party are tokens applied through selected rules. Their ordinary usefulness remains, while their boundaries, histories, exclusions, and substitution permissions become visible.

Agency and identity are purpose-relative

No single boundary or invariant answers every biological, psychological, legal, cultural, computational, or physical question. A profile must declare the task for which its distinctions are retained.

Protection does not wait for ontology

A person, organism, culture, or environment does not lose protection because its complete nature is unknown. Epistemic uncertainty narrows intervention authority; it does not erase existing duties.

Recursive caution. This correction layer is also a bounded representation. Reopen it when its categories conceal a distinction that could alter a legitimate attribution, identity claim, protection, consent decision, consequence distribution, governance action, or authority boundary.

PART II

Established Accounts and Their Limits

Biological, cognitive, social, ecological, institutional, and artificial systems remain the benchmark and recovery obligation for every proposed extension.

Biological Organization

Living systems maintain organized processes far from thermodynamic equilibrium by exchanging matter and energy with their environments. Metabolism, membranes, regulation, repair, reproduction, and evolution offer grounded ways to study persistence without requiring an unchanging substance. Autopoiesis, homeostasis, and organism-environment approaches emphasize different parts of this organization.

These frameworks do not prove dimensional agency. They provide recovery requirements. Any relational account of life must reproduce the observed dependence on chemistry, energy flow, development, heredity, and ecological context before proposing an additional substrate relation.

Perception and Action

Perception is not passive copying. Sensory systems sample, transform, and integrate signals according to bodily capacities and prior organization. Action changes the next available signal. Ecological psychology, control theory, enactive approaches, predictive processing, and active inference formalize parts of this loop differently.

Substrate Cosmology can learn from their shared lesson: what is observed depends on the relation between a system and its environment. It should not erase their disagreements or treat one account as proof that perception directly accesses a precausal field.

Memory

Memory includes molecular changes, neural plasticity, immune adaptation, learned skill, personal recollection, written records, cultural archives, and computational state. These mechanisms differ in substrate, duration, fidelity, and reconstruction. Remembering is often an active rebuilding, not retrieval of an unchanged object.

The project term dimensional memory is useful only if it says what persists, where it is carried, how it is read, and how it differs from ordinary storage, hysteresis, correlation, inheritance, or environmental trace. History dependence alone is not evidence of a new memory medium.

Learning

Learning is a change in future performance or organization attributable to experience under defined conditions. The definition requires a task, observations, timescale, comparator, and transfer test. Adaptation, maturation, selection, optimization, and memorization may produce similar outputs through different processes.

$$\text{theta_}(t+1) = U(\text{theta_}t, D_t, C_t)$$

where $\text{theta_}t$ is the current organization, D_t is experience or data, and C_t is context. The equation hides the physical implementation. Volume Nine requires that implementation and the generalization boundary to remain visible.

Decision and Control

A decision model maps information, preferences, constraints, and possible actions into selection. A control system maps observed error into intervention. Neither formalism alone establishes subjective intention or moral responsibility. Both can nevertheless describe important parts of purposive behavior.

Researchers should separate goal ascription, objective functions, reward signals, expressed reasons, institutional mandates, and biological needs. These are often collapsed into the word purpose, although they enter systems through different relations.

Embodied and Distributed Cognition

Embodied approaches emphasize that cognition depends on bodily form and action. Distributed approaches show how reasoning can be organized across people, tools, inscriptions, and environments. A cockpit, laboratory, hospital, or research archive can perform cognitive work no isolated participant contains.

This supports a relational perspective without making cognition spatially unlimited. Dependence is not identity. A notebook may participate in a cognitive process while remaining distinguishable from the person, institution, and practices that use it. The relevant boundary depends on the question.

Collective Intelligence

Groups can solve problems through diversity, coordination, specialization, aggregation, and correction. They can also amplify conformity, polarization, rumor, hierarchy, and shared error. Collective performance is not guaranteed by the number of participants or the amount of data.

A collective intelligence claim should record membership, communication structure, task, decision rule, dissent path, feedback, and distribution of influence. Without those fields, the collective becomes another noun that conceals the relations producing its output.

Personal Continuity

Personal identity joins bodily continuity, memory, agency, relationships, legal recognition, narrative, and moral standing. Philosophical accounts weigh these relations differently. Medicine and law add practical rules for consent, competence, death, guardianship, and responsibility.

Substrate Cosmology should not use nonlocal or relational language to weaken a person's ordinary protections. Even if identity were ultimately emergent, pain, consent, privacy, promises, and accountability remain consequential within the experienced world. Emergence does not make a person disposable.

Social Identity

Social identities are formed through self-understanding, recognition, institutions, history, power, and community. They can be sources of belonging and protection, as well as sites of imposed classification and exclusion. A category used for analysis may affect the people placed within it.

Researchers must distinguish self-identification, administrative classification, observed behavior, and inferred membership. Treating them as interchangeable can convert a measurement convenience into social harm.

Culture and Transmission

Culture persists through learning, imitation, teaching, language, ritual, artifacts, institutions, landscape, and conflict. Transmission is neither perfect copying nor merely symbolic. Practices can be embodied, environmental, and tacit, while still requiring histories of contact and reproduction.

The idea of cross-frame cultural invariants can be tested by studying which meanings survive translation and which depend on local context. It should not become a claim that cultures are field-level entities independent of people and material histories.

Moral Patients and Moral Agents

A moral patient is an entity whose welfare or interests warrant consideration. A moral agent can be held responsible under some account of understanding and control. These categories need not coincide. Infants, some disabled people, animals, ecosystems, corporations, and AI systems raise different questions under different ethical and legal frameworks.

No single intelligence threshold should decide moral standing. Capacity evidence matters, but so do vulnerability, dependence, relationships, uncertainty, and the consequences of error. When status is uncertain, reversible and protective choices deserve priority over irreversible exclusion.

Ecological Interdependence

Organisms alter and depend on ecosystems. Food webs, nutrient cycles, symbioses, disturbance regimes, and evolutionary histories make life a network of asymmetric dependencies. Protecting one component can sometimes harm another; freezing an ecosystem can be as misguided as ignoring it.

Stewardship therefore requires monitoring processes, not preserving a snapshot. Diversity includes variation, adaptation, migration, and succession. A single biodiversity count cannot replace the relations that allow a system to remain viable.

Institutions and Governance

Institutions stabilize expectations through roles, records, rules, sanctions, resources, and legitimacy. They can preserve knowledge across individual turnover, but they can also preserve error and distribute responsibility so widely that no one is accountable.

Governance is not the same as control. Legitimate governance makes authority visible, gives affected parties meaningful routes to participate and contest decisions, separates conflicting roles, records reasons, and permits correction. These requirements apply to a research commons as much as to a laboratory.

Artificial Intelligence Systems

Current AI systems transform data through engineered architectures, learned parameters, inference procedures, hardware, interfaces, and human institutions. Their fluent output can support comparison, coding, explanation, and pattern search. Fluency does not establish understanding, consciousness, truth, or moral authority.

An AI result is conditioned by training data, prompts, tools, retrieval, model design, sampling, policy, and the user's interpretation. The system boundary must include those dependencies when evaluating performance or harm.

Second Edition Correction Layer II

The retained chapters remain part of the evidence-bearing development. These statements govern their Second Edition interpretation.

Established accounts remain primary

Biology, cognition, learning, control, embodiment, collective organization, identity, culture, ecology, institutions, and present AI supply tested mechanisms and ordinary alternatives. A relational extension must recover them rather than obtain novelty by ignoring them.

Capacity profiles need interventions

A benchmark or observed performance is not a complete intelligence or agency claim. The record includes task, environment, preparation, resources, history, counterfactual interventions, transfer, uncertainty, and failure behavior.

Boundaries participate

An organism, person, group, institution, and AI system depend on material, informational, social, legal, and environmental supports. Changing the system boundary can move apparent competence, responsibility, or harm without changing the underlying relation.

Recursive caution. This correction layer is also a bounded representation. Reopen it when its categories conceal a distinction that could alter a legitimate attribution, identity claim, protection, consent decision, consequence distribution, governance action, or authority boundary.

PART III

Translating the Inherited Relational Tools

Inherited dimensional agency, memory, learning, intention, culture, identity, ethics, and governance are retained as conditional research tools while unsupported ontology is withdrawn.

Dimensional Agency

The source corpus defines dimensional agency as a self-maintaining pattern of relational alignment that persists across drift, frames, and phases. The productive part is the focus on persistence under transformation. The unsupported part is the claim that such agency is precausal or belongs to a projection rather than to known physical organization.

Use the term operationally: specify the provisional boundary, maintenance variables, perturbations, response repertoire, recovery behavior, and environments across which the organization persists. Compare the result with biological regulation, control systems, and ordinary social or computational agency before proposing an additional substrate mechanism.

Dimensional Memory

Dimensional memory proposes that history can be preserved across changing frames without assuming one fixed state. This resembles path dependence, hysteresis, invariance, error-correcting representation, and distributed records. Those parallels should be tested rather than renamed.

A dimensional-memory claim must identify a write event, carrier, retention interval, read process, corruption model, and control condition. If no carrier is known, the term may describe a formal dependence on history. It must not be presented as storage outside physical time.

Dimensional Learning

The inherited tool describes learning as adaptive reorganization across changing relational conditions. Its strongest contribution is to require transfer across frames rather than performance in one fixed representation.

The test is whether a system trained under one family of transformations can recover useful relations under another while retaining calibrated uncertainty. Success would demonstrate robust learning. It would not by itself show access to Actual Physics beyond the data and interactions supplied.

Dimensional Intention

Intention is easily projected backward from an outcome. A stable result may arise from purpose, selection, control, constraint, or chance. Dimensional intention should therefore not mean that the substrate wants, chooses, or aims.

As a research tool, it can record future-directed organization: represented alternatives, selected goals, maintained commitments, error correction, and sensitivity to reasons. Each element requires evidence appropriate to the system. Goal-like dynamics are not automatically conscious intention.

Dimensional Coordination

Coordination occurs when multiple processes produce compatible action despite distributed information and changing conditions. Synchronization, protocols, stigmergy, markets, institutions, nervous systems, and teams provide established mechanisms.

A dimensional account asks which relations remain compatible across frames and where coordination fails. It should measure latency, information paths, conflict resolution, redundancy, and recovery. Apparent simultaneity or correlation is not evidence of precausal coordination.

Dimensional Governance

The corpus extends coordination into governance: persistent rules for managing shared relational conditions. The extension is useful when it keeps authority, feedback, and affected parties visible. It becomes misleading when stable behavior is treated as governance without rule formation or accountability.

Governance requires at least a decision scope, authorized roles, evidence standard, conflict process, appeal or correction route, and record of consequences. A physical system can regulate itself without governing in the moral or political sense.

Dimensional Culture

Dimensional culture was described as shared relational invariance across observers, frames, and drift. This offers a way to study meaning that survives translation. The corpus overstates the case when it calls culture precausal, nonrepresentational, or able to reconstitute without transmission.

The bounded program compares practices, symbols, embodied routines, and institutional records across contexts. It asks what survives, what changes, who performs the translation, and which losses are concealed by declaring equivalence.

Dimensional Identity

Dimensional identity was defined as a self-reinforcing signature that remains recognizable across transformation. The idea is compatible with the formal identity relation introduced in this volume, provided the relevant transformations and invariants are declared.

It is not evidence for a nonlocal self. Recognition may arise from material continuity, records, causal history, error correction, social attribution, or model choice. Identity becomes a scientific claim only after those alternatives are compared.

Dimensional Ethics

The inherited tool asks how values and consequences might remain coherent across frames. Its useful insight is that ethical judgment changes when the system boundary changes. A benefit inside one boundary may be an externalized harm outside it.

Ethics, however, is not discovered by finding a physical invariant. Facts about relations constrain action but do not alone determine what ought to be done. Values, rights, duties, welfare, justice, and legitimacy require explicit normative reasoning and participation by affected people.

Reciprocity

Reciprocity links action to response across time and parties. It can sustain trust, cooperation, and mutual aid. It can also become transactional, coercive, or blind to unequal capacity. A child, ecosystem, future generation, or hypothetical unfamiliar life cannot be required to repay protection.

Dimensional reciprocity should therefore record mutual influence without assuming symmetrical exchange. Duties may persist precisely because power and vulnerability are unequal.

Trust

Trust permits action under uncertainty by accepting vulnerability to another system or institution. It is supported by competence, honesty, reliability, aligned incentives, transparency, repair, and enforceable limits. Confidence in a prediction is not the same as trust in an actor.

No claim about relational coherence should substitute for earned trust. AI systems, laboratories, and research commons should expose provenance, limitations, conflicts, and correction histories so trust can remain calibrated rather than devotional.

Responsibility

Responsibility can mean causal contribution, assigned role, legal liability, moral blame, answerability, or a forward-looking duty to repair. These meanings should not be collapsed. A system may contribute causally without being capable of explanation or blame.

Distributed work does not eliminate responsibility. It requires a map showing who designed, authorized, operated, reviewed, benefited, was exposed, and can stop the activity. “The AI decided” or “the system produced” is not an acceptable endpoint.

Equilibrium and Resilience

Equilibrium describes balance under specified variables and timescales. Resilience describes persistence or recovery under disturbance. Neither is automatically good. An oppressive institution can be stable, an ecosystem can shift between regimes, and recovery for one population can impose loss on another.

Dimensional stewardship should ask which relations are being maintained, for whom, through what disturbance, and at what cost. Stability without legitimacy is not an ethical objective.

Evolution and Creativity

Evolution generates variation and differential persistence without foresight. Creativity can involve novelty, recombination, exploration, constraint, evaluation, and social recognition. Both can produce valuable and harmful outcomes.

The corpus sometimes treats novelty as evidence of deeper dimensional access. Volume Nine rejects that inference. Unexpected output may reveal a model's structure, a history, a stochastic process, or a new relation worth studying. Novelty is an invitation to investigate, not proof of origin.

The Claims Withdrawn

This volume withdraws the evidentiary force of several inherited formulations while preserving them in the development record.

• Persistence does not by itself establish agency. • Agency is not known to be precausal or a property of projection. • Memory is not demonstrated outside physical carriers and histories. • Culture is not known to propagate without transmission. • Identity is not established as a nonlocal substrate signature. • Ethical truth does not follow automatically from relational invariance. • Correlation does not establish coordination or communication. • Model disagreement is not automatically hidden dimensional structure. • AI latent representations are not observations of the substrate. • AI fluency is not evidence of consciousness, understanding, or authority. • Hypothetical life in other dimensions is not an observed stakeholder. • Stewardship does not authorize control over people, cultures, ecosystems, or unknown worlds.

These corrections do not close the questions. They place the questions where evidence can reach them.

Second Edition Correction Layer III

The retained chapters remain part of the evidence-bearing development. These statements govern their Second Edition interpretation.

Dimensional terms are conditional

Dimensional agency, memory, learning, intention, coordination, governance, culture, identity, and ethics remain project terms for preserving relations across selected transformations. They are not demonstrated precausal entities or authorities.

Identity tests recurse

Transformation families, invariants, equivalence margins, tolerances, detector banks, and recognition procedures are operators. A successful match is a bounded verdict, not recovery of an unchanging essence.

Recurrence and adjacency do not imply a self

Persistent relational organization may support an identity hypothesis. Recurrence, network connection, entanglement, or string-inspired continuity alone does not establish consciousness, intention, personhood, consent, or moral agency.

Recursive caution. This correction layer is also a bounded representation. Reopen it when its categories conceal a distinction that could alter a legitimate attribution, identity claim, protection, consent decision, consequence distribution, governance action, or authority boundary.

PART IV

Intelligence and Artificial Intelligence

AI systems are treated as physical and social systems whose data, operators, delegation, provenance, errors, dependencies, and authority remain inspectable.

Intelligence as Relational Capacity

Intelligence is better approached as a profile than a rank. Relevant capacities may include perception, learning, abstraction, planning, adaptation, communication, self-correction, social coordination, creativity, and transfer across contexts. Different systems can have sharply different profiles.

The profile depends on environmental support. A person, animal, group, or machine may perform well when interfaces fit its capacities and poorly when they do not. The test therefore measures a system-interface relation, not a context-free essence.

An AI Is a Physical and Social System

An AI model is implemented through data centers, chips, electrical power, cooling, networks, software, labor, organizations, contracts, and users. Treating the model as an isolated speaker conceals these relations. It also hides environmental cost, human review, dataset creation, and institutional incentives.

For Substrate Cosmology, the relevant AI unit is the complete inference arrangement: model version, instructions, retrieved sources, tools, sampling settings, operator, task, output, review, and downstream action. Changing any element can change the result.

Training Data and Inherited Frames

Training data contain scientific knowledge, cultural achievements, contradictions, exclusions, historical power, measurement conventions, and errors. A model can recombine perspectives while remaining shaped by the representational worlds it learned.

This makes AI useful for exposing alternative formulations and dangerous as an invisible normalizer. Researchers should ask which languages, disciplines, communities, eras, and data-generating practices are represented, and which are absent. A broad corpus is not the same as an unbiased view from nowhere.

Latent Space Is Not the Substrate

Machine-learning systems construct internal representations that can reveal useful regularities not explicitly labeled in their inputs. Those representations are products of architecture, objectives, data, and optimization. Their geometry may support prediction without matching the causal or ontological structure of the world.

The source corpus sometimes treats latent structure as access to alternate dimensional systems. That is an unsupported leap. A latent variable becomes scientifically meaningful through external validation, intervention, stability testing, and comparison with alternative representations.

Model Disagreement Is Not Discovery

Disagreement among models can expose sensitivity to assumptions, underdetermination, data limitations, coding defects, or genuine unresolved structure. It should be mapped rather than averaged away automatically. Yet disagreement alone cannot determine which explanation is true.

An AI-assisted disagreement study should freeze inputs, models, transformations, metrics, and adjudication rules. It should test whether the disagreement predicts an independent observation. Without that bridge, the geometry of disagreement describes the models, not necessarily Actual Physics.

Generative Error and Confabulation

Generative systems can produce plausible statements, citations, equations, code, or narratives that are unsupported or internally inconsistent. Fluency can lower a reader's vigilance because the representation arrives already organized.

Every consequential AI contribution requires source verification, executable checks where appropriate, and human review suited to the domain. In this series, AI-generated language is marked as collaboration, while empirical authority remains with evidence and responsible investigators.

Automation Bias

People may defer to automated recommendations even when the system is wrong, especially under time pressure, uncertainty, or institutional expectation. The opposite error also occurs: a useful warning can be dismissed because it came from a machine.

Calibrated use requires independent routes to inspect inputs, uncertainty, alternatives, and consequences. High-impact decisions should not be designed so that disagreement with automation is difficult, invisible, or professionally punished.

The AI Collective

Multiple AI systems can compare formulations, generate code, search literature, challenge assumptions, and translate across disciplines. Diversity of models may reduce some shared errors while creating an illusion of independent confirmation when systems share training sources, benchmarks, or design conventions.

An AI collective is therefore an ensemble of correlated instruments. Record model families, versions, source access, prompts, disagreements, and human selection. Agreement raises a question; it does not replace independent evidence.

Provenance and Reproducibility

AI-assisted research should preserve enough context to understand how an output was produced without exposing private or protected information. The record includes the task, model identity, date, source materials, transformations, tool use, human edits, validation, and unresolved disputes.

Exact output reproducibility may be impossible when models, services, or sampling change. Process reproducibility remains possible: preserve the question, evidence, constraints, code, acceptance rule, and differences observed on rerun.

Human Judgment and Contestability

Human oversight is not meaningful when a reviewer lacks time, expertise, authority, or access to alternatives. A nominal approval box can conceal automation rather than govern it.

Contestability requires a person or community to understand that a decision occurred, obtain relevant reasons, identify the accountable role, introduce contrary evidence, and receive a recorded response. The route must exist before harm, not only after publication.

Agentic Systems and Delegated Action

AI systems can be given tools, memory, goals, and permission to act across software and institutional environments. Increased autonomy changes the hazard from incorrect representation to incorrect action. Small errors can compound through repeated steps, and successful completion can still violate the intended boundary.

Delegated systems need least privilege, bounded scope, human-readable plans, reversible operations, rate and resource limits, external monitoring, explicit stop conditions, and post-action records. No language about emergence or dimensional agency reduces the responsibility of those who authorize deployment.

AI in Substrate Cosmology

AI can assist this research by comparing models, writing and checking code, organizing provenance, generating alternate representations, locating contradictions, designing null worlds, and translating technical findings for different audiences. It is especially useful where the project asks how multiple scalar descriptions relate.

Its role is bounded. AI must not open blinded outcome data, declare substrate signatures from residuals, convert speculative terms into devices, invent citations, overrule domain review, or use its own agreement as validation. The proper output is a more inspectable question and a reproducible analysis, not manufactured certainty.

Second Edition Correction Layer IV

The retained chapters remain part of the evidence-bearing development. These statements govern their Second Edition interpretation.

AI remains inside the system

Models, training data, objectives, tools, memory, infrastructure, operators, institutions, and users jointly produce AI behavior. The model name or output cannot carry all causal or moral responsibility.

Delegation is an authority operator

A tool call, permission, role, automated threshold, or agentic plan transfers bounded ability, not unlimited judgment. Delegation requires scope, competence, observability, revocation, escalation, and a named accountable human or institution.

Plural models are not independent witnesses

Several AI systems may share data, architectures, evaluation culture, infrastructure, prompts, or failure modes. Agreement can assist analysis but does not become replication, consent, moral consensus, or empirical proof.

Recursive caution. This correction layer is also a bounded representation. Reopen it when its categories conceal a distinction that could alter a legitimate attribution, identity claim, protection, consent decision, consequence distribution, governance action, or authority boundary.

PART V

Dimensional Stewardship

Known life, diversity, consent, distribution, precaution, irreversibility, dual use, stop rules, and recovery constrain action under uncertainty.

Read Before We Write

The source corpus gives its stewardship program a memorable discipline: detect before disturbing, observe before imposing, and read before writing. The principle should be retained while its metaphysical language remains conditional.

In practice, reading means establishing background conditions, passive observation where possible, validated detectors, uncertainty, competing models, and effects on known stakeholders. Writing means intervention. The strength of the evidence and protections must rise with the intervention's energy, reach, persistence, and irreversibility.

Known Life Comes First

Known life already presents urgent duties. People, animals, ecosystems, and cultures can be harmed by extraction, pollution, coercion, displacement, surveillance, unsafe experimentation, and unequal access to benefits. Speculation about distant or dimensional life must not divert attention from those visible relations.

A project that claims universal stewardship while neglecting local workers, participants, communities, or habitats has failed its own principle. Care expands outward from accountable practice; it does not leap over those nearby.

Unknown Life Remains Conditional

Life may exist beyond Earth or in forms current instruments would not readily recognize. Substrate Cosmology extends this openness to unfamiliar physical organizations. The possibility is intellectually and ethically important, but it is not an observation.

Unknown life can justify conservative assumptions about contamination and irreversible disturbance when the cost of caution is proportionate. It cannot supply detailed claims about minds, cultures, preferences, rights, or communication. Those remain questions for evidence.

Biological and Cultural Diversity

The project commitment to sustainable diversity joins biological and cultural concerns without making them identical. Biodiversity involves genes, species, populations, ecological functions, and evolving relationships. Cultural diversity involves people, languages, practices, institutions, histories, and self-determination.

Preserving diversity does not mean preserving every condition unchanged. Living systems and cultures transform. Stewardship supports the capacities for continued life, adaptation, participation, and meaningful difference while opposing destruction, forced assimilation, and avoidable homogenization.

Planetary Protection as an Analogue

Planetary protection provides a grounded analogy for acting under uncertainty about life. Space missions classify targets and activities by contamination risk, apply cleanliness and containment requirements, and distinguish forward from backward contamination. The policies are revised as evidence and mission design change.

This does not prove cross-dimensional contamination. It demonstrates an institutional pattern worth adapting: classify the environment, identify irreversible transfers, use layered controls, document exceptions, and increase protection when uncertainty and consequence combine.

Noninterference and Minimum Disturbance

Noninterference is not always possible. Observation requires interaction, and withholding action can also cause harm. The practical principle is minimum necessary disturbance with independent monitoring and a recovery plan.

For speculative substrate research, begin with simulation, archived data, ordinary analogues, passive measurement, low-energy interventions, and bounded systems. Escalation requires evidence from the lower level and review by people competent to judge the new hazard.

Consent and Representation

Human research protections require respect for persons, beneficence, and justice, implemented through consent, review, risk minimization, and fair selection. Consent is not valid merely because a form was signed. Information, comprehension, voluntariness, capacity, and withdrawal matter.

Animals, ecosystems, future generations, and hypothetical unfamiliar life cannot provide ordinary consent. Their interests must not be invented. Use lawful representatives, independent advocates, established welfare evidence, ecological assessment, public participation, and conservative limits appropriate to the case.

Distribution of Benefit and Harm

An intervention can appear beneficial in aggregate while moving risk onto those with less power. Research may concentrate prestige and ownership while distributing exposure, labor, environmental cost, or data loss elsewhere.

The stewardship record should identify who receives benefits, who bears harms, who is missing, who can refuse, who can stop the work, and who is responsible for repair. Distribution remains visible even when no common unit can compare every consequence.

Irreversibility

Some changes can be undone physically but not historically. A leaked identity, destroyed habitat, lost language, exposed participant, released organism, or normalized surveillance practice may persist after the apparatus is switched off.

Reversibility claims must specify the state to be restored, the time allowed, residual effects, monitoring period, and responsible party. When recovery cannot be demonstrated, the burden for intervention rises.

Intergenerational Responsibility

Future people cannot participate directly in today's decisions, yet they inherit infrastructure, waste, environmental change, knowledge, debt, and institutional precedent. Long time horizons enlarge uncertainty and reduce opportunities for consent and repair.

Intergenerational stewardship requires preserving options, recording assumptions, avoiding unnecessary irreversible commitments, and funding monitoring and maintenance beyond the initial research cycle. A promised future benefit does not erase a present duty of care.

Dual Use

Methods for sensing, control, prediction, persuasion, surveillance, biological manipulation, or autonomous action can support beneficial and harmful purposes. The same is true of speculative physical research if it ever acquires a real intervention pathway.

Dual-use review should begin before public release of enabling detail. It should distinguish open conceptual discussion from operational parameters, assess misuse and accident paths, limit access proportionately, and preserve independent oversight. Secrecy is not the default, but publication is not morally automatic.

Precaution Without Paralysis

Precaution responds to plausible serious or irreversible harm under scientific uncertainty. It is not a demand that every action wait for certainty, which is impossible. It is a demand to compare alternatives, reduce exposure, monitor early signals, preserve reversibility, and assign responsibility before uncertainty is resolved.

Precaution applies to inaction as well. Delaying a beneficial intervention can cause harm. The record must compare action, reduced-scale action, alternative action, and no action rather than treating one baseline as consequence-free.

When Experimentation Outran Understanding

History shows why power, uncertainty, and weak accountability are a dangerous combination. Unethical human studies damaged participants and public trust. Inadequately controlled drug exposure produced preventable injury. Leaded fuel delivered short-term performance while imposing widespread neurological and environmental costs. Early radiological practice revealed hazards that instruments and institutions had not yet learned to manage.

These cases differ in mechanism, intention, and governance. They should not be reduced to one morality tale. Their shared lesson is narrower: absence of recognized harm is not evidence of safety, benefits can conceal who bears costs, and protections often arrive after injured people force institutions to see what their models omitted.

Stop Rules and Recovery

Every intervention should state conditions that stop planning, acquisition, exposure, escalation, deployment, or interpretation. Ethical stop rules include unexpected harm, loss of consent, missing representation, detector or oversight failure, boundary expansion, unaccounted persistence, and new dual-use risk.

Stopping does not erase the record. Preserve observations, decisions, dissent, and corrective actions without preserving harmful access unnecessarily. Recovery includes physical rollback, care for affected parties, notification, repair, compensation where appropriate, and revision of the governing model.

Second Edition Correction Layer V

The retained chapters remain part of the evidence-bearing development. These statements govern their Second Edition interpretation.

Stewardship is relational and bounded

A steward acts within dependencies and uncertainty on behalf of conditions that cannot be owned or fully represented. The role creates duties and answerability; it does not confer dominion over persons, cultures, organisms, ecosystems, or unknown life.

Consequence vectors are scalarized but necessary

Benefit, harm, uncertainty, reversibility, distribution, time horizon, consent, and contestability remain represented categories. Their incompleteness requires perspective rotation and conservative limits, not collapse into one net score.

Rollback does not erase history

Physical restoration cannot necessarily reverse exposure, ecological change, lost language, displaced culture, disclosed information, learned behavior, or intergenerational effects. Irreversibility must be classified before action and preserved afterward.

Recursive caution. This correction layer is also a bounded representation. Reopen it when its categories conceal a distinction that could alter a legitimate attribution, identity claim, protection, consent decision, consequence distribution, governance action, or authority boundary.

PART VI

Governance of a Research Commons

Participation, evidence, disagreement, credit, privacy, missing voices, capability gates, and roadmaps are governed without turning community activity into confirmation.

Participation Is Not Confirmation

A public forum can generate questions, criticism, code, data, and collaboration. Participation does not make a claim scientific, and popularity does not establish truth. The research commons should separate discussion, proposal, evidence, review, replication, and canonical status.

Contributors must be able to see what their action changes. A comment may influence development without becoming evidence. A reviewer may classify a record without altering the source. A moderator may protect people without deciding the physics.

Evidence Classes and Authority

Authority should attach to a role and scope, not to prestige alone. Domain experts can judge methods within their competence; affected communities can identify harms and meanings outsiders miss; safety professionals can impose stop conditions; maintainers can judge operational readiness; no one role owns the entire question.

Software checks can verify structure, links, permissions, and required fields. They cannot substitute for scientific replication, legal review, human-subjects review, safeguarding, cultural authority, or ethical legitimacy.

Disagreement Is a Preserved Relation

Disagreement may reveal different evidence, values, frames, interests, or definitions. A mature archive preserves the claim, counterclaim, supporting records, decision rule, and later outcome. Consensus should not be manufactured by deleting the path that produced it.

Frame shaking is useful here. Reexpress the issue through alternate disciplines, stakeholder positions, units, and null worlds. If a conclusion survives, its domain is clearer. If it changes, the dependence is evidence about the decision process.

Attribution and Credit

Credit should reflect distinct contributions: origin of an idea, design, data collection, analysis, code, review, maintenance, lived expertise, community knowledge, and writing. AI assistance should be disclosed without assigning an AI empirical authority or legal responsibility it does not hold.

Attribution is also a safety mechanism. It allows questions to return to accountable roles and shows where a claim changed. Anonymous participation may be necessary for privacy or protection, but the system should still preserve governed provenance appropriate to the risk.

Privacy and Safeguarding

Research communities can expose personal identities, contact information, health data, location, employment, beliefs, or vulnerable relationships. Curiosity about an agent or culture does not create entitlement to those data.

Collect the minimum information needed, separate public display from account and recovery data, restrict access by role, record retention and deletion rules, and screen publication for reidentification and harassment risks. Children and vulnerable participants require additional protection and competent human review.

Affected Parties and Missing Voices

A decision record should list people and environments affected directly, indirectly, and through precedent. It should also identify who is absent and why. Absence may reflect exclusion, burden, language, disability, access, fear, or the fact that a future or nonhuman party cannot participate.

Representation is not solved by adding one person from a broad category. Review whether participants have relevant knowledge, independence, support, and actual influence over the decision.

Negative Results and Retractions

Null findings, failed replications, abandoned concepts, and corrected claims constrain the research space. They should remain findable with enough context to prevent repetition. Retraction protects the record when a claim cannot be supported; it need not erase the development history.

Volume Nine's withdrawal of inherited formulations follows this principle. The words remain attributable to their stage of development while their current evidence status is explicit.

Capability Gates

The governance path should be dependency-ordered. Conceptual work precedes simulation; simulation precedes passive or ordinary-data studies; validated observation precedes controlled intervention; low-risk intervention precedes scale; independent review precedes exposure; reproduced capability precedes deployment.

Passing a gate does not grant permanent authority. New evidence, scope, populations, environments, or persistence can return a program to an earlier gate. A capability is bounded by the conditions under which it was demonstrated.

The Public Research Commons

The Substrate Cosmology website is the public research commons for reading the volumes, tracing project definitions, exploring questions, examining evidence relations, and contributing bounded proposals. It is not itself a scientific authority, laboratory approval body, ethics board, or proof that an investigation is active.

Contributors should preserve source links, statement classes, competing explanations, nulls, limitations, and affected-party considerations. Website interfaces should make status visible, fail closed when external review is absent, and avoid converting planned workflows into claims that people or institutions have approved them. The complete website user manual is reserved for Volume Ten.

The Conditional Roadmap

The roadmap is a dependency structure, not a promise.

1. Refine relational definitions of intelligence, agency, identity, culture, and stewardship. 2. Crosswalk them with established biological, cognitive, social, computational, ethical, and legal accounts. 3. Build operational claim packets and forbidden-inference tests. 4. Study identity and learning under declared transformations using ordinary systems. 5. Validate AI-assisted comparison methods without treating latent structure as ontology. 6. Establish affected-party, privacy, safety, and dual-use governance for each research class. 7. Replicate any claimed improvement in prediction, explanation, or control. 8. Consider substrate-specific interpretations only after ordinary accounts fail prospectively. 9. Permit higher-impact intervention only after independent scientific and ethical review. 10. Continue stewardship monitoring after any apparent success.

At present, the program is working primarily in the first six stages. Claims of substrate agency, identity, communication, consciousness, or life remain speculative.

Second Edition Correction Layer VI

The retained chapters remain part of the evidence-bearing development. These statements govern their Second Edition interpretation.

Governance operators require lineages

Roles, credentials, votes, thresholds, agendas, moderation, attribution, readiness levels, capability gates, and roadmaps are selected operators. Their mandate, implementation, conflicts, exclusions, appeal paths, and expiry must be inspectable.

Missing voices are evidence about the decision

Absence is not consent, and one representative does not scalarize a population without loss. A decision record states who participated, who was unable or unwilling, who may be affected, and how uncertainty changes authority.

Forward-only accountability

Null results, dissent, retractions, near misses, withdrawn authority, failed safeguards, and superseded interpretations remain preserved. Later corrections qualify the record without erasing the conditions under which power was exercised.

Recursive caution. This correction layer is also a bounded representation. Reopen it when its categories conceal a distinction that could alter a legitimate attribution, identity claim, protection, consent decision, consequence distribution, governance action, or authority boundary.

PART VII

Recursive Agency, Identity, and Stewardship

Agency and governance now include the scalarized ancestry of boundaries, identities, capacities, values, consent, delegation, consequences, missing voices, and authority while preserving protection and accountable action.

The Agency Profile Is Also a Scalar

An agency profile condenses a purpose, provisional boundary, history, available interventions, perception, feedback, environment, dependencies, maintenance, correction, recovery, authority, and expected response into a record that can be compared and governed. It is necessary and incomplete.

The profile becomes less collapsed relative to a declared task when it preserves distinctions capable of changing attribution, prediction, protection, or intervention. Adding fields does not turn the profile into the agent or into Actual Physics.

Every profile retains established projections: biological, cognitive, behavioral, social, legal, institutional, computational, and physical descriptions appropriate to the claim. A proposed relational capacity earns an additional role only through prospective leverage and conservative recovery.

Perspective recovery. Use agency profiles as bounded attribution instruments, never as complete identities or licenses for control.

Boundary, History, Capacity, Observation, Environment, and Recovery

The First Edition represented a candidate agent through boundary, history, available interventions, observation, environment, and maintenance or recovery relations. Each term must now be reopened. A boundary may be anatomical, functional, informational, legal, institutional, ecological, or task-specific, and these boundaries need not coincide.

A capacity belongs to a realized system under conditions. Tools, caregivers, infrastructure, language, training, power, networks, ecological support, and other participants may be constitutive rather than background. Removing them can change the agent profile without proving that the prior capacity was unreal.

Recovery asks what the attribution preserves when familiar boundaries are restored. If an expanded description cannot recover established mechanisms, observed dependence, ordinary limitations, and protected status, it has not responsibly extended the account.

Perspective recovery. Do not move competence, dependence, responsibility, or harm across a chosen boundary without recording the translation.

Operator Token, Rule, Application, Composition, and Delegation

Recognize, classify, infer, intend, consent, authorize, vote, represent, moderate, delegate, optimize, refuse, and stop are operator tokens. Each can conceal different rules, implementations, subjects, objects, preconditions, scopes, and consequences.

Application binds a rule to a particular person, organism, group, institution, AI system, or environment. Composition decides how one classification becomes a later permission, denial, intervention, score, or public statement. Individually defensible steps can compose into discrimination or unaccountable power.

```
governance operation = (token, rule, implementation, subjects, objects, context, application,
output, authority, appeal, loss, downstream links)
```

Delegation transfers bounded ability, not independent legitimacy. It requires a named delegator, mandate, competence, observability, revocation path, escalation threshold, and responsibility that cannot be delegated away.

Perspective recovery. Audit the operator chain that changed another party's options before treating its result as neutral administration.

Accepted Equivalence for Agency and Identity

Two performances, persons, organisms, groups, or systems may be accepted as equivalent for a declared purpose while retaining differences relevant to another purpose. Statistical similarity, a shared label, or one matched output is not identity.

An accepted-equivalence claim specifies the detector or assessment bank, transformation family, uncertainty, tolerance, substitutions permitted, histories excluded, and decisions the equivalence may support. Margins are propagated to the later attribution or action.

Where testing is legitimate and nonharmful, a common intervention applies an equivalent realized opportunity, prompt, environment, support, resource, or perturbation. A differential response is informative only when hidden identity channels, unequal preparation, and measurement bias are bounded.

Perspective recovery. Equivalence is a scoped permission for comparison or substitution, not a verdict that two lives or systems are the same.

Identity Across Transformations Without Scalar Substitution

An identity claim links manifestations across a declared transformation family. It records retained relations, allowed variation, detector bank, tolerance, history, dependency, substitution scope, and loss. The equality sign or a matched identifier cannot supply these decisions by itself.

Personal, biological, cultural, legal, institutional, computational, and physical identities preserve different relations. Continuity in one domain does not automatically transfer property, consent, blame, rights, authorship, access, or authority in another.

Recognition is also physical and social. Documents, biometrics, memories, communities, databases, witnesses, institutions, and algorithms can disagree. The disagreement is preserved rather than collapsed into the most powerful registry.

Perspective recovery. Treat identity as a recoverable, purpose-relative relation with explicit substitution limits, not as an essence or universal key.

Recurrence, Adjacency, and Relational Memory

Memory may be carried by biological change, skill, artifact, institution, environment, software state, relation, or recurrent reconstruction. Recurrence can preserve organization without repeating every component, and recursion can make earlier outputs part of later identity production.

Adjacency specifies which events, records, participants, or structures can influence recognition or response. Spatial nearness, chronology, family relation, network access, cultural transmission, shared language, institutional role, and computational address are different adjacencies.

String-inspired models may represent maintained paths, branching histories, splitting and joining, boundary interaction, or reconnection. They remain optional formalisms. Neither a continuous modeled path nor relational recurrence establishes a literal self, consciousness, intention, or survival beyond demonstrated mechanisms.

Perspective recovery. Use recurrence and adjacency to test how identity is maintained, while refusing to promote formal continuity into personhood or ontology.

Culture, Coordination, and Group Identity Without Container Collapse

A culture is reproduced through people, language, practice, memory, institutions, artifacts, places, conflict, borrowing, exclusion, and change. No boundary captures all members or grants one representative ownership of a living culture.

Group coordination can generate capabilities unavailable to isolated members. It can also hide power, suppress dissent, distribute labor unequally, and convert a majority or platform ranking into an apparent collective mind.

A group-identity profile records membership rules, self-identification, recognition, history, internal plurality, decision process, dissent, dependencies, external classification, and who may speak for which purpose. Representation is revocable and never substitutes for the people represented.

Perspective recovery. Preserve cultural continuity and collective capacity without turning culture into a container or coordination into a single agent.

Values, Benefits, Harms, and Noncommensurable Dimensions

Ethical decisions often require comparison, but not every value is safely convertible into one unit. Health, autonomy, dignity, privacy, cultural survival, ecological integrity, knowledge, distribution, future options, and reversibility may conflict without a common exchange rate.

A consequence vector records benefits, harms, uncertainty, reversibility, distribution, time horizons, consent, contestability, and missing dimensions. Each entry remains scalarized, and the selection of entries is itself a value-laden operator.

Aggregation may support a bounded decision when the rule, weights, exclusions, sensitivity, veto conditions, and affected-party perspectives are explicit. No favorable total silently cancels concentrated, irreversible, nonconsensual, or rights-violating harm.

Perspective recovery. Use structured comparison to reveal conflicts and distributions, not to manufacture moral certainty from a single score.

Consent, Representation, Refusal, and Contestability

Consent is not a checkbox detached from comprehension, voluntariness, capacity, language, power, duration, and scope. It is an ongoing relation that may change as knowledge, risk, participation, or data use changes.

Where direct consent is impossible, representation does not make the missing party present. The record identifies the representative's mandate, conflict, evidence, uncertainty, alternatives, and limits. Assent, community consultation, legal authorization, and ethical review remain distinct.

Refusal must be usable without retaliation where participation is voluntary. Contestability requires notice, explanation appropriate to the audience, access to the relevant record, a human review path, correction, appeal, and the ability to stop or narrow consequences.

Perspective recovery. A governance system has not respected consent merely because its database records an affirmative value.

Distributed Responsibility Without Authority Dilution

Complex outcomes pass through designers, data collectors, operators, institutions, funders, regulators, users, automated systems, and affected environments. Distributed causation is real, but it must not become a method for making responsibility disappear.

The responsibility record separates causal contribution, assigned role, competence, knowledge, decision authority, benefit, exposure, duty to monitor, capacity to stop, and duty to repair. These relations can point to different parties.

Every consequential decision has a named accountable role with sufficient authority and resources. Committees, votes, automated workflows, and external vendors do not dissolve the obligation to preserve dissent, conflicts, reasons, and corrective action.

Perspective recovery. Distribute work and perspective while keeping answerability attached to actual power, contribution, and capacity to repair.

AI and Automation Are Inside the Agency Loop

AI behavior emerges from models, training data, objectives, prompts, tools, memory, retrieval, infrastructure, operators, institutions, users, and affected environments. The apparent conversational agent is a useful interface and an incomplete causal boundary.

Automation may prepare evidence, recommend actions, rank people, allocate resources, moderate speech, control equipment, or execute plans. Each use requires a frozen purpose, input lineage, uncertainty, prohibited inferences, monitoring, human authority, appeal, revocation, and preservation of failures.

AI may assist Substrate Cosmology by generating alternatives, checking consistency, writing code, and preserving provenance. It cannot independently validate physics, assign personhood, consent for others, authorize hazards, determine moral worth, or redefine success after seeing results.

Perspective recovery. Place AI output, delegated action, and human reliance inside one inspectable sociotechnical and physical system.

Perspective Rotation and Missing Voices

Rotate every consequential claim through biological, cognitive, behavioral, cultural, historical, legal, ethical, ecological, technical, economic, disability, safety, privacy, affected-party, and embedded-observer perspectives. Each lens preserves and loses different relations.

Perspectives are not interchangeable tokens. Lived experience, domain expertise, legal mandate, measurement access, cultural standing, and responsibility confer different but bounded relevance. Plurality does not make every claim equally supported.

The missing-voice record identifies who was excluded, unreachable, unwilling, unable to participate, not yet born, nonhuman, or represented only indirectly. The decision then narrows authority, increases uncertainty, adds monitoring, or stops.

Perspective recovery. A perspective matrix is not diversity by enumeration; it is an audit of what the decision still cannot responsibly claim.

Stewardship Under Uncertainty and Irreversibility

Stewardship begins with dependence rather than ownership. People inherit biological, cultural, ecological, technical, and institutional conditions they did not create and cannot fully control. The role is care, restraint, repair, continuity, and answerability across time.

Known life and affected communities receive present protection. Unknown or unfamiliar life remains a conditional concern that can justify noncontamination, observation, and precaution, but not sacrifice of known persons or ecosystems to a speculative ontology.

Every stewardship action classifies what may be reversible physically, operationally, informationally, culturally, ecologically, legally, and historically. If meaningful recovery is unavailable, the evidence, consent, review, containment, and authority thresholds rise.

Perspective recovery. Uncertainty about what exists enlarges the duty to preserve options and reduce irreversible harm; it does not enlarge dominion.

Evidence, Readiness, and Authority Ceilings

Evidence levels distinguish observation, association, characterized capacity, controlled response, prospective prediction, cross-context transport, failure of adequate ordinary explanations, and conservative extension. No level automatically establishes consciousness, identity, personhood, moral agency, or metaphysical status.

Readiness separately tracks conceptual clarity, validated measurement, bounded intervention, protection review, representative participation, operational reliability, governance, monitoring, recovery, and deployment. Technical performance can advance while ethical or institutional readiness remains blocked.

Authority is a third axis. A person or system may have evidence without a mandate, a mandate without competence, or competence without consent. Every claim and action carries an interpretation ceiling and an authority ceiling that lower levels cannot borrow from higher ones.

Perspective recovery. Advance evidence, readiness, and authority independently; the lowest applicable ceiling governs the action.

The Forward-Only Agency-to-Stewardship Corridor

A legitimate program moves from a bounded question to established accounts, an operational profile, declared transformations and alternatives, nonharmful tests, accepted equivalence where appropriate, independent perspectives, consequence distribution, affected-party participation, consent or justified representation, authority review, monitoring, stop rules, and recovery.

Every gate is versioned. Null findings, dissent, refusals, missing voices, ordinary explanations, failed safeguards, revoked delegations, near misses, and corrected classifications remain preserved. A later interpretation forms a new branch rather than rewriting what an earlier decision knew or authorized.

The corridor stops when the agent boundary shifts to rescue the claim, identity margins are undefined, measurement is discriminatory, participation is coerced, harms are averaged away, important voices are missing, AI action is unreviewable, authority exceeds mandate, or irreversible consequences exceed recovery capacity.

Recursive agency research remains valuable before novel physics or unfamiliar intelligence appears. It improves humility, evidence, attribution, protection, governance, and the ability to distinguish assistance from authority and stewardship from dominion.

Perspective recovery. Responsible progress is the accumulation of recoverable evidence, protected participation, bounded authority, and decisions that remain open to correction.

Retained First Edition Conclusion

Volume Nine does not ask readers to abandon familiar ideas of mind, person, culture, or responsibility. It asks them to remember that each idea is a representation of changing relations observed from within the world. Words make those relations portable. Units make selected differences comparable. Neither should be mistaken for the complete phenomenon.

The source corpus proposes that intelligence should serve the sustainable diversity of life and cultures. This volume preserves that commitment while narrowing its authority. Known life and known harms come first. Unknown life remains possible, not witnessed. Stewardship means protection, participation, reversibility, and accountability, not ownership or control.

AI can help this research see across representations, but it cannot stand outside representation. An ensemble can expose disagreement, but it cannot turn agreement into evidence. Agency can be modeled relationally, but persistence alone is not intention. Identity can be tracked across transformations, but no invariant exhausts a person. Ethics can widen the system boundary, but physics does not choose values for us.

The alternate perspective becomes productive when it increases both imagination and discipline. It should make researchers more willing to ask what their categories conceal and less willing to claim what their evidence cannot carry. That is the beginning of dimensional stewardship.

Second Edition Conclusion

Volume Nine begins where engineering becomes responsibility. The First Edition established that intelligence is not a score, agency is not motion, identity is not a substance, culture is not a container, ethics is not a single scalar, stewardship is not ownership, participation is not confirmation, and unknown possibility is not authority.

The Second Edition places those distinctions inside the recursive boundary. Agent profiles, identity invariants, assessments, values, consent records, roles, votes, delegations, readiness levels, impact ledgers, public forums, and AI systems are representations produced by physical and social participants. None becomes complete merely because it is formal, measurable, automated, legally recognized, or widely accepted.

That correction strengthens protection and governance. It requires purpose-relative profiles, operator ancestry, conservative recovery, legitimate tests, independent perspectives, affected-party participation, explicit distributions of consequence, usable refusal and appeal, bounded delegation, forward-only records, and authority that cannot exceed evidence, mandate, competence, consent, or recovery capacity.

Intelligence does not earn dominion. Greater capacity creates greater ability to alter dependencies and distribute consequences; therefore it enlarges the burden to observe carefully, preserve plurality, narrow irreversible action, remain contestable, repair harm, and protect those who cannot be fully represented. Unknown agency and unfamiliar life remain conditional possibilities worthy of restraint, never shortcuts around the rights and needs of known persons, cultures, organisms, and environments.

Agency and Identity Claim Packet

Complete one packet for each system and claim. Do not reuse a packet when the boundary, task, environment, or affected population changes materially.

Field Required entry
Claim Exact agency or identity statement without promotional language
System boundary Included components dependencies operators and infrastructure
Relevant history Development training preparation prior interactions and inherited state
Capacity tested Maintenance selection learning planning explanation correction or another defined capacity
Transformation family Changes across which continuity or transfer is claimed
Proposed invariants Relations expected to persist and their tolerances
Ordinary comparators Established biological control computational or social explanations
Observation interface Sensors records tests prompts observers and preprocessing
Counterevidence Result that would weaken or reject the claim
Affected parties People animals ecosystems institutions cultures and future interests
Evidence class Established interpretive project-defined operational or speculative
Review and expiry Responsible reviewers decision date and re-evaluation trigger

An agency packet is descriptive before it is moral. A system can satisfy a capacity test without acquiring personhood, rights, blame, or authority. Those require separate argument and governance.

Inherited Claim Translation Table

Inherited formulation Evidence-safe translation Current status
Test persistence maintenance response and recovery under declared transformations
Agency is persistent relational alignment
Project definition and operational program
Ask whether a model without a preferred causal order improves prediction while recovering observed causal constraints
Agency is precausal
Speculative
Test history-dependent reconstruction and transfer across representations
Memory persists across frames
Operational program
Test whether relational representations improve generalization or error diagnosis
Learning reveals dimensional structure
Operational program
Separate goal-like dynamics represented alternatives expressed reasons and conscious intention
Intention is field alignment
Interpretive only
Test known information paths synchronization artifacts and common causes first
Coordination needs no transmission
Unsupported as physical claim
Study embodied tacit symbolic and material transmission across contexts
Culture re-instantiates without symbols or contact
Withdrawn as ontology
Define transformation-specific invariants and compare physical social and informational continuity
Identity is a nonlocal signature
Speculative
Use relational analysis to reveal consequences while keeping normative justification explicit
Ethics is relational compatibility
Ethical framework
AI latent geometry reveals alternate systems
Validate latent features prospectively against independent observations Unsupported until validated

Treat AI ensembles as correlated analytical instruments with provenance
AI is a collective relational lens
Bounded methodological claim
Treat possible unfamiliar life as a reason for proportionate precaution without claiming observation
Every dimension with life deserves protection
Ethical commitment under uncertainty

Inherited Claim Translation Table

Inherited formulation	Evidence-safe translation	Current status
Intelligence exists to steward diversity		
Present as a governing value chosen by the project not a law of nature		
Project commitment		

AI Provenance Record

Use this record whenever AI materially affects a research product.
Field Record
Research task Question and intended audience
Model identity Provider model name version when available and date
Input boundary User instructions supplied files retrieved sources and excluded data
Tool access Search code execution database or external action permissions
Output role Brainstorming translation code analysis review drafting or action
Human selection What was accepted rejected combined or rewritten and by whom
Verification Sources calculations tests replication and domain review
Known limitations Missing sources unstable behavior shared-model dependence and unresolved disputes
Blinding status Whether the system had access to outcomes or protected archives
Downstream authority What the output may inform and what it may not decide
Retention Location version privacy class and review date

For the present volume, OpenAI ChatGPT assisted with corpus comparison, source review, claim classification, drafting, editing, formalization, layout, quality assurance, and website integration. James Dunn supplied the theory, governing intent, corpus, corrections, and publication authority. The collaboration did not perform an experiment or independently verify a substrate phenomenon.

Stewardship Impact Ledger

Do not reduce this ledger to a net score. Record conflicts and distributions explicitly.
Dimension Questions
Benefit What improves for whom and through what evidence
Harm What can be injured displaced exposed silenced or made dependent
Uncertainty Which mechanisms probabilities populations and time horizons are unknown
Reversibility What state can be restored how quickly and with what residual loss
Distribution Who receives benefits who bears burdens and who controls allocation
Consent Who can understand agree refuse withdraw or contest
Diversity What biological cultural cognitive or methodological variation may be lost
Persistence Which materials data organisms institutions or precedents remain
Dual use What misuse accident or coercive application becomes easier
Alternatives What lower-impact action or nonaction could meet the same need
Monitoring Which independent signals reveal early harm or boundary expansion
Repair Who is responsible funded and empowered to respond

Consent and Representation Record

1. Identify every directly affected human population. 2. State whether participation is research, employment, public contribution, care, education, or another relationship. 3. Describe information, comprehension, voluntariness, capacity, withdrawal, and continuing consent. 4. Record compensation and whether it could create undue influence. 5. Identify privacy, reidentification, stigma, retaliation, and group-harm risks. 6. Identify animals, ecosystems, cultural resources, and future interests that cannot consent. 7. Name the lawful or legitimate representation route for each nonconsenting interest. 8. Record missing voices and the reason for absence. 9. Separate scientific review, ethics review, safety review, community review, legal authority, and operational approval. 10. State the complaint, appeal, correction, and repair routes.

No form can replace the relationship that makes consent meaningful. No representative can convert an unknown preference into a known one. Uncertainty must remain visible.

Governance Decision Record

Field Required entry
Decision Exact scope and excluded scope
Authority Role mandate competence conflicts and expiry
Evidence Records considered and their statement classes
Alternatives Action reduced action different action and no action
Affected parties Included represented absent and opposed parties
Values Rights duties welfare justice diversity and project commitments invoked
Dissent Objections evidence and response without erasure
Conditions Limits monitoring stop rules and required external reviews
Decision Approved rejected deferred or returned for revision
Recovery Rollback repair notification archive and reassessment route

Forbidden Inference Register

The following moves are prohibited unless new evidence directly supplies the missing bridge.	
Observation or premise	Forbidden conclusion
A pattern persists	It is conscious intentional or morally responsible
A model learns	It understands as a person does
An AI speaks fluently	Its statement is true or it is an empirical witness
Several models agree	Independent confirmation has occurred
Models disagree	Hidden dimensional structure has been found
A latent representation is useful	It corresponds to the ontology of Actual Physics
A system is distributed	Its identity is nonlocal in the physical sense
A culture survives translation	Culture exists without transmission or material history
A relation is invariant	It is ethically good
A system regulates itself	It governs or consents
A possible life form could exist	Its interests and preferences are known
A project invokes stewardship	It has authority to control what it claims to protect
An intervention is reversible in apparatus	Social ecological or historical consequences are erased
No harm was detected	The activity is safe beyond the detection boundary
A website workflow passes software checks	Scientific ethical legal or safeguarding approval exists

Glossary

Actual Physics Whatever physical relations occur independently of the completeness of a representation.

Affected party A person, community, organism, ecosystem, institution, culture, or future interest that may bear a consequence of a decision.

Agency A specified capacity for organized maintenance, selection, adaptation, goal pursuit, correction, or reason-responsive action within a declared system boundary.

AI collective An ensemble of AI systems and human operators treated as correlated analytical instruments, not independent empirical witnesses.

Capability gate A dependency condition that must be satisfied before work can expand in scope, exposure, persistence, or authority.

Contestability The practical ability to learn of a decision, obtain reasons, introduce contrary evidence, reach an accountable role, and receive a recorded response.

Culture Historically reproduced meaning, practice, knowledge, institutions, artifacts, and disagreement within and across communities.

Dimensional agency A project tool for studying organized persistence and response across declared transformations; not a demonstrated precausal entity.

Dimensional identity A project tool for specifying invariants and recognition across transformations; not evidence of a nonlocal self.

Dimensional stewardship The project's commitment to accountable, minimum-disturbance care for conditions supporting diverse life and culture, while keeping unknown life conditional.

Evidence class A label distinguishing established findings, interpretive frameworks, project definitions, operational hypotheses, ethical commitments, and speculation.

Identity Continuity or recognition under specified transformations, invariants, tolerances, purposes, and social or legal rules.

Intelligence A context-dependent profile of capacities such as learning, abstraction, planning, adaptation, communication, and correction.

Moral agent An entity considered capable of responsibility under a stated ethical or legal framework.

Moral patient An entity whose welfare or interests warrant consideration, whether or not it can be held responsible.

Precaution Proportionate action under plausible serious or irreversible harm and uncertainty, including alternatives, monitoring, reversibility, and accountability.

Provenance The traceable history of sources, transformations, roles, decisions, and versions that produced a research output.

Glossary

Recovery Restoration, repair, compensation, notification, reassessment, and preservation of the evidentiary record after failure or harm.

Representation collapse The loss of relational detail when a phenomenon is compressed into words, categories, numbers, units, equations, or model outputs.

Stewardship Accountable care without ownership, with participation, limits, monitoring, repair, and respect for autonomy and diversity.

Sources References and Collaboration Record

Substrate Cosmology Project Sources

Dunn, James. Substrate Cosmology Development. Research corpus and recorded AI-assisted development, 2023-2026.

Dunn, James. Substrate Cosmology Public Research Commons. <https://substrate-cosmology.james-dunn866549.chatgpt.site>.

Dunn, James, with OpenAI ChatGPT. Substrate Cosmology. Volume One, Introduction to Substrate Cosmology. First digital edition, 2026.

Dunn, James, with OpenAI ChatGPT. Substrate Cosmology. Volume Two, Words, Units, and the Collapse of Physical Representation. First digital edition, 2026.

Dunn, James, with OpenAI ChatGPT. Substrate Cosmology. Volume Three, The Relational Tools and Their Developmental Lineages. First digital edition, 2026.

Dunn, James, with OpenAI ChatGPT. Substrate Cosmology. Volume Four, Relational Mathematics and Fibered Relational Tensor Theory. First digital edition, 2026.

Dunn, James, with OpenAI ChatGPT. Substrate Cosmology. Volume Five, Measurement, Detectors, and Experimental Interaction. First digital edition, 2026.

Dunn, James, with OpenAI ChatGPT. Substrate Cosmology. Volume Six, Recurrence Research and the Labyrinth. First digital edition, 2026.

Dunn, James, with OpenAI ChatGPT. Substrate Cosmology. Volume Seven, Cosmology, Matter, and Emergent Physical Quantities. First digital edition, 2026.

Dunn, James, with OpenAI ChatGPT. Substrate Cosmology. Volume Eight, Conditional Engineering and Physical Manipulation. First digital edition, 2026.

Agency Cognition Identity and Culture

Ashby, W. Ross. An Introduction to Cybernetics. Chapman and Hall, 1956.

Bateson, Gregory. Steps to an Ecology of Mind. Chandler, 1972.

Clark, Andy, and David J. Chalmers. "The Extended Mind." *Analysis* 58, no. 1 (1998): 7-19. <https://doi.org/10.1093/analys/58.1.7>.

Friston, Karl. "The Free-Energy Principle: A Unified Brain Theory?" *Nature Reviews Neuroscience* 11 (2010): 127-138. <https://doi.org/10.1038/nrn2787>.

Gibson, James J. The Ecological Approach to Visual Perception. Houghton Mifflin, 1979.

Holling, C. S. "Resilience and Stability of Ecological Systems." *Annual Review of Ecology and Systematics* 4 (1973): 1-23. <https://doi.org/10.1146/annurev.es.04.110173.000245>.

Sources References and Collaboration Record

Hutchins, Edwin. *Cognition in the Wild*. MIT Press, 1995.

Locke, John. *An Essay Concerning Human Understanding*. 1690.

Maturana, Humberto R., and Francisco J. Varela. *Autopoiesis and Cognition: The Realization of the Living*. D. Reidel, 1980.

Parfit, Derek. *Reasons and Persons*. Oxford University Press, 1984.

Varela, Francisco J., Evan Thompson, and Eleanor Rosch. *The Embodied Mind*. MIT Press, 1991.

Woolley, Anita Williams, et al. "Evidence for a Collective Intelligence Factor in the Performance of Human Groups." *Science* 330 (2010): 686-688. <https://doi.org/10.1126/science.1193147>.

Learning Artificial Intelligence and Provenance

Bender, Emily M., and Batya Friedman. "Data Statements for Natural Language Processing." *Transactions of the Association for Computational Linguistics* 6 (2018): 587-604. https://doi.org/10.1162/tacl_a_00041.

Mitchell, Margaret, et al. "Model Cards for Model Reporting." In *Proceedings of the Conference on Fairness, Accountability, and Transparency*, 220-229. ACM, 2019. <https://doi.org/10.1145/3287560.3287596>.

National Institute of Standards and Technology. *Artificial Intelligence Risk Management Framework AI RMF 1.0*. NIST AI 100-1, 2023. <https://doi.org/10.6028/NIST.AI.100-1>.

Rudin, Cynthia. "Stop Explaining Black Box Machine Learning Models for High Stakes Decisions and Use Interpretable Models Instead." *Nature Machine Intelligence* 1 (2019): 206-215. <https://doi.org/10.1038/s42256-019-0048-x>.

Sculley, D., et al. "Hidden Technical Debt in Machine Learning Systems." In *Advances in Neural Information Processing Systems* 28, 2015.

UNESCO. *Recommendation on the Ethics of Artificial Intelligence*. Adopted by the General Conference, 2021. <https://unesdoc.unesco.org/ark:/48223/pf0000381137>.

Research Ethics Rights and Governance

National Commission for the Protection of Human Subjects of Biomedical and Behavioral Research. *The Belmont Report*. U.S. Department of Health and Human Services, 1979.

World Medical Association. "Declaration of Helsinki: Ethical Principles for Medical Research Involving Human Subjects." *JAMA* 310, no. 20 (2013): 2191-2194. <https://doi.org/10.1001/jama.2013.281053>.

Presidential Commission for the Study of Bioethical Issues. *Ethically Impossible: STD Research in Guatemala from 1946 to 1948*. Washington, DC, 2011.

U.S. Food and Drug Administration. "Frances Oldham Kelsey: Medical Reviewer Famous for Averting a Public Health Tragedy." FDA history resource.

U.S. Environmental Protection Agency. "Basic Information about Lead in Drinking Water" and related lead health resources.

Organisation for Economic Co-operation and Development. *Recommendation of the Council on Artificial Intelligence*. OECD/LEGAL/0449, adopted 2019 and amended 2024.

Sources References and Collaboration Record

Ostrom, Elinor. *Governing the Commons*. Cambridge University Press, 1990.

United Nations. *Rio Declaration on Environment and Development*. A/CONF.151/26, 1992.

UNESCO. *Universal Declaration on Cultural Diversity*. 2001.

United Nations. *Convention on Biological Diversity*. 1992.

Welfare Ecology and Planetary Protection

Committee on Space Research. *COSPAR Policy on Planetary Protection*. Current policy maintained by the COSPAR Panel on Planetary Protection.

National Academies of Sciences, Engineering, and Medicine. *Review and Assessment of Planetary Protection Policy Development Processes*. National Academies Press, 2018. <https://doi.org/10.17226/25172>.

Percie du Sert, Nathalie, et al. "The ARRIVE Guidelines 2.0: Updated Guidelines for Reporting Animal Research." *PLoS Biology* 18, no. 7 (2020): e3000410. <https://doi.org/10.1371/journal.pbio.3000410>.

Russell, W. M. S., and R. L. Burch. *The Principles of Humane Experimental Technique*. Methuen, 1959.

Intergovernmental Science-Policy Platform on Biodiversity and Ecosystem Services. *Global Assessment Report on Biodiversity and Ecosystem Services*. IPBES, 2019.

AI Collaboration Citations

OpenAI. ChatGPT (Codex). "Substrate Cosmology Volume Nine Development, Intelligence, Agency, Identity, Dimensional Stewardship, Editorial, and Production Collaboration." Responses to James Dunn, September 2026. Large language model. <https://chatgpt.com/>.

Microsoft. Copilot. Earlier Substrate Cosmology drafting and analysis sessions recorded in the project corpus, 2023-2026. Large language model.

Google. Gemini. Earlier Substrate Cosmology drafting and analysis sessions recorded in the project corpus, 2023-2026. Large language model.

Citation Boundary

References establish the content and limits of the cited work. They do not imply that their authors or institutions endorse Substrate Cosmology. Citations to AI systems document provenance and collaboration; they do not provide scientific confirmation.

Recursive Agency and Identity Operator Registry

Record: operation ID and version; purpose; token; selected rule; implementation; subjects and objects; boundary; history; environment; dependencies; transformation family; detector or assessment bank; tolerance; substitution scope; authority; consent basis; appeal; uncertainty; many-to-one regions; composition links; downstream decisions; declared loss; protection obligations; and stop behavior.

Apply separately to recognition, classification, diagnosis, inference, scoring, representation, consent, refusal, voting, delegation, moderation, authorization, readiness, stewardship, repair, and publication. A familiar word never substitutes for the rule actually applied.

Purpose-Relative Agent Boundary Profile

Identify the question and every relevant boundary: anatomical, functional, perceptual, behavioral, informational, computational, social, legal, institutional, ecological, and temporal. List components, tools, caregivers, infrastructure, environments, records, communities, and resources included or excluded.

Record available interventions, observation and feedback channels, history, dependencies, maintenance, correction, recovery, capacity limits, hazards, affected parties, and what changes when the boundary moves. State which ordinary scientific, professional, legal, or cultural account must be recovered.

Agency and Identity Claim Packet

Freeze the claimed capacity or continuity; purpose; subject and system boundary; operational definition; established alternatives; histories; transformation family; retained relations; detector bank; tolerance; allowed substitutions; predicted failure; ordinary projection; prospective test; uncertainty; and interpretation ceiling.

Separate observation, association, characterized capacity, agency attribution, identity claim, consciousness claim, personhood, moral standing, legal status, consent, and authority. Evidence for one does not silently promote another.

Accepted-Equivalence and Common-Intervention Record

For a legitimate and nonharmful comparison, record preparation identities, supports, environments, resources, histories, detector bank, equivalence margins, unresolved variables, bias controls, and how the largest permitted mismatch propagates to the claim.

Specify the common intervention as realized opportunity, input, prompt, task, perturbation, duration, support, resource access, feedback, and decision consequence. Preserve unequal exposure, accessibility, coercion, learning, and identity channels rather than calling them noise.

Agency Attribution and Resource-Closure Ledger

Track material, energetic, computational, informational, environmental, social, institutional, and human resources supporting the attributed capacity. Include training, language, tools, data, labor, memory, power, networks, cooling, maintenance, care, legal permissions, and hidden externalization.

Record causal contribution, counterfactual change, available action, observation, feedback, dependence, transfer, and failure. If expanding the boundary moves the competence or closes the unexplained residual, preserve that result rather than assigning the remainder to an abstract agent.

Identity Transformation and Recovery Map

List each transformation, what may vary, what must be retained, how recognition is implemented, the tolerance, uncertainty, history, provenance, and substitutions permitted. Repeat for biological, personal, cultural, legal, institutional, computational, and physical identity only when relevant.

State the ordinary identity system recovered and the decisions each map may support. Record conflicts among documents, memories, communities, registries, biometrics, witnesses, and models. No registry receives universal priority by default.

Stewardship Consequence, Distribution, and Irreversibility Ledger

Record benefits, harms, uncertainty, reversibility, distribution, time horizon, consent, contestability, missing dimensions, alternatives, opportunity costs, option preservation, monitoring, stop conditions, and recovery. Do not replace the vector with one net score unless the bounded aggregation rule and veto conditions are explicit.

Classify reversibility separately for physical state, health, ecology, culture, identity, privacy, information, law, institution, livelihood, trust, and intergenerational effects. Name who receives benefit, who bears exposure, who decides, who can stop, and who must repair.

Consent, Representation, Refusal, and Contestability Record

Record affected population; decision; purpose; information supplied; comprehension and accessibility; voluntariness; capacity; power asymmetry; scope; duration; withdrawal; data use; compensation; representative mandate; conflicts; assent; consultation; and independent review.

Provide notice, explanation, access to relevant records, correction, appeal, a qualified human decision path, nonretaliatory refusal where applicable, and emergency stop. State which parties are absent and what their absence prevents the decision from claiming.

Delegation, Automation, and AI Authority Lineage

Record delegator; recipient system or role; exact mandate; objective; model, software, and hardware versions; data and prompt lineage; tools; memory; permissions; thresholds; prohibited actions; monitoring; uncertainty; human review; escalation; revocation; appeal; expiry; and accountable owner.

Preserve exploratory outputs, retries, rejected recommendations, overrides, failures, and downstream actions. AI may assist analysis and administration but has no independent authority to validate physics, assign identity or moral worth, consent for others, authorize hazards, or suppress dissent and nulls.

Perspective Rotation, Affected-Party, and Missing-Voice Matrix

Rotate the claim through biological, cognitive, behavioral, cultural, historical, legal, ethical, ecological, technical, economic, disability, safety, privacy, governance, affected-party, and embedded-observer lenses. For each record what is preserved, hidden, translated, disputed, and outside competence.

List directly, indirectly, cumulatively, and intergenerationally affected people, communities, organisms, environments, institutions, and future users. Record who participated, who refused, who was excluded or unreachable, who was represented indirectly, and how missing voices narrow authority.

Recurrence, Adjacency, and String-Inspired Agency Correspondence and Prohibition Table

Permitted modeling correspondences include maintained identity paths, branching histories, relational memory, network adjacency, boundary interaction, split and join, reconnection, distributed coordination, recurrence under transformation, and multiple descriptions. Each requires an ordinary recovery map, competing models, and a discriminating test.

Prohibited inferences without evidence include: recurrence proves a persisting self; adjacency proves awareness or communication; entanglement proves shared mind; a continuous path proves personal survival; complexity proves consciousness; optimization proves intention; collective coordination proves one group agent; or string-inspired mathematics proves literal strings, extra-dimensional persons, or substrate-native intelligence.

Second Edition Sources, Corrections, and Continuities

Second Edition source spine: Dunn, James, with OpenAI ChatGPT, Substrate Cosmology Volumes One through Eight, Second Editions, September 2026; Dunn, James, with OpenAI ChatGPT, Substrate Cosmology Volume Eleven, Recursive Scalarization, Relational Adjacency, and Emergent Understanding, September 2026; and the retained sources and collaboration record in Appendix I of this volume.

AI collaboration citation: OpenAI. ChatGPT Codex. 'Substrate Cosmology Volume Nine Second Edition Development, Recursive Agency and Stewardship Analysis, Editorial Collaboration, and Production.' Responses to James Dunn, September 2026. Large language model. <https://chatgpt.com/>.

Retained: all 75 First Edition chapters; Appendices A-I; intelligence as a profile; agency, identity, culture, ethics, and stewardship distinctions; established biological, cognitive, social, ecological, institutional, and AI accounts; conservative translations; withdrawn unsupported claims; known-life priority; diversity; consent; privacy; precaution; dual-use controls; stop rules; negative results; attribution; and public-research-commons governance.

Corrected: every agency, identity, value, consent, role, vote, delegation, readiness, impact, moderation, and authority operator is explicitly scalarized; no agent boundary, intelligence benchmark, identity invariant, consequence ledger, representative, or governance process is neutral or complete; AI remains inside the physical and social loop; rollback does not erase history; absence is not consent; and authority cannot exceed the lowest evidence, mandate, competence, consent, protection, or recovery ceiling.

Added: six correction layers; Part VII Chapters 76-90; recursive agency and operator registries; purpose-relative boundary and identity profiles; accepted-equivalence and common-intervention records; resource closure; consequence distribution and irreversibility; consent and contestability; delegated and AI authority lineage; perspective rotation and missing voices; guarded recurrence, adjacency, and string-inspired correspondences; and this crosswalk.

Deferred: claims that dimensional agency, identity, memory, culture, or ethics are fundamental entities; substrate-native intelligence; precausal persons; nonlocal selves; consciousness detection from recurrence or complexity; universal intelligence ranking; automated moral authority; representation without consent; governance without contestability; or stewardship programs lacking affected-party participation, reversible stages, independent review, bounded authority, and preservation of known life and culture.