frontiers | Frontiers in Blockchain

TYPE Hypothesis and Theory
PUBLISHED 23 September 2025
DOI 10.3389/fbloc.2025.1657050

OPEN ACCESS

EDITED BY
Rocsana Bucea-Manea-Tonis,
National University of Physical Education and Sport, Romania

REVIEWED BY
Mehmet Firat,
Anadolu University, Türkiye
Mélanie Chevance,
Zentrum für Astronomie der Universität Heidelberg, Germany

*CORRESPONDENCE
Sasha Shilina,
sasha@episteme.ac

RECEIVED 01 July 2025
ACCEPTED 02 September 2025
PUBLISHED 23 September 2025

CITATION
Shilina S (2025) DeScAI: the convergence of decentralized science and artificial intelligence. *Front. Blockchain* 8:1657050.
doi: 10.3389/fbloc.2025.1657050

DeScAI: the convergence of decentralized science and artificial intelligence

Sasha Shilina*

Lomonosov Moscow State University, Moscow, Russia

Scientific knowledge production is undergoing a dual transformation. On one front, Decentralized Science (DeSci) leverages blockchain-based infrastructures to reconfigure how research is funded, verified, and governed, disintermediating legacy gatekeepers through tokenized incentives and distributed provenance. On the other, Artificial Intelligence (AI) is automating core dimensions of science, from hypothesis generation to experimental execution and model validation. This paper introduces DeScAI, a theoretical framework that unifies these domains into a recursive, self-verifying epistemic system governed by autonomous agents operating within decentralized, trust-minimized networks. We present a five-stratum architecture for DeScAI, hypothesizing that its integration enables epistemic acceleration, pluralistic inquiry, and cryptographically auditable trust. Methods include a structured literature synthesis (2018–2025), conceptual modeling, and descriptive analysis of 14 projects. Three hypothetical trajectories for future empirical investigation are proposed concerning cycle-time compression, epistemic pluralism, and reproducibility amplification. We conclude that DeScAI is not speculative: its core components are already deployed. What remains is orchestration, stitching together decentralized ledgers, incentive protocols, self-sovereign scientific agents (SSA), and cryptographic infrastructures into a single, recursive system. If successful, DeScAI could radically reduce the latency between hypothesis and verification, reconfigure scientific legitimacy as a live, contestable signal, and transform the incentive structure of research itself.

KEYWORDS

decentralized science, artificial intelligence, blockchain, knowledge production, cryptoeconomics, epistemic infrastructure, autonomous agents, recursive systems

1 Introduction

The institutions and infrastructures of science are struggling to keep pace with the complexity, scale, and velocity of contemporary research. The current system faces a well-documented crisis: reproducibility failures, misaligned incentives, opaque peer review, and inequitable access to resources reveal that legacy models of knowledge production are increasingly unfit for present needs (Baker, 2016; Begley and Ellis, 2012; Fang and Casadevall, 2016; Ioannidis, 2014; Kitcher, 1993;

Prinz et al., 2011; Powell, 2016). Research confirms persistent reproducibility issues: Ioannidis (2024) highlights replication failures despite rigorous controls, while Camerer et al. (2016), Camerer et al. (2018) report similar results in economics and psychology. Ongoing initiatives such as the Institute for Replication[1] extend these efforts across economics, political science, and biomedicine.

Meanwhile, the widespread adoption of machine-learning (ML) workflows has introduced second-order risks: studies compromised by data leakage, overfitting, and hidden confounders have fueled what journal Nature calls an "AI-driven reproducibility crisis" (Ball, 2023). Even gatekeeping is evolving: since June 2025, Nature publishes referee reports and author rebuttals by default, signaling a systemic push toward transparency (Nature Editorial, 2025). These trends underscore the need for new epistemic infrastructures.

In response to these mounting challenges, two technological revolutions, Decentralized Science (DeSci) and AI-for-Science (AI4S), have emerged in parallel, each offering distinct remedies.

DeSci refers to a movement that leverages distributed ledger technologies (DLTs), blockchain primitives, smart contracts, decentralized autonomous organizations (DAOs), intellectual-property non-fungible tokens (IP-NFTs), quadratic funding, and token-curated registries (TCRs), and other similar tools to re-engineer how research is funded, owned, and audited (Asgaonkar and Krishnamachari, 2018; Bamakan et al., 2022; Bischof et al., 2022; Buterin et al., 2019; DeFrancesco and Klevecz, 2022; Fantaccini et al., 2024; Weidener and Boltz, 2025; Weiss, 2022). DeSci shifts epistemic legitimacy away from traditional gatekeepers toward programmable systems of provenance, incentive alignment, and transparent participation. Echoing the democratizing arcs of the alphabet, the printing press, and the World Wide Web, DeSci is framed as the fourth major decentralization of knowledge (Weidener and Lukács, 2025). Foundational work on open-peer-review protocols first mapped the space (Tenorio-Fornés et al., 2021). Weidener and Spreckelsen (2024) provide a comprehensive definition of DeSci, identify dozens of active DeSci projects, formalize shared values such as transparency, collaboration, and integrity, along with guiding principles including decentralization, collective ownership, and incentivization. An empirical survey shows the ecosystem's resilience: 96% of projects remain active, and partnership-driven DAOs dominate (Díaz, Menchaca and Weidener, 2025). Case studies illustrate this innovation in action: they show VitaDAO raising six-figure sums for IP-tokenized longevity projects (Unfried, 2024), while Molecule's IP-NFT framework provides legally compliant rails to market for those assets (Ortlepp, 2022).

AI-for-Science (AI4S) is simultaneously re-engineering the laboratory bench and the research pipeline. It denotes the use of AI, particularly ML, generative models, and autonomous agents to perform or augment core scientific tasks such as hypothesis generation, literature synthesis, experiment design, and data interpretation. Unlike traditional computational tools, AI4S systems may learn from unstructured data, plan novel experimental paths, and iteratively refine knowledge through closed-loop automation. Surveys chart its reach across drug discovery (Dara et al., 2022; Mak et al., 2023; Blanco-González et al., 2023; Qureshi et al., 2023), materials design (Sha et al., 2020; Li et al., 2020; Guo et al., 2021), and even the social sciences (Xu et al., 2024), while methodological syntheses frame the field as a distinct, data-centric modality of inquiry (Lawrence and Montgomery, 2024; Nagar, 2024; Wu and Wang, 2025; Xie et al., 2024). Milestone breakthroughs illustrate the trend. Breakthroughs include AlphaFold 2's near-complete protein structure prediction (Jumper et al., 2021), autonomous chemist robots (Dai et al., 2024), and SAMPLE's rapid protein engineering loops (Rapp et al., 2024). The AI Scientist (Lu et al., 2024) showcases the technical potential of autonomous research agents, and surveys (Ferrag et al., 2025) highlight a fast-growing ecosystem of Large Language Model (LLM) -based agent frameworks with tool-use, planning, and communication protocols. At the programmatic level, the U.S. DARPA FoundSci initiative aims to build an "autonomous scientist" capable of domain-general skeptical reasoning (DARPA, 2024). Early proof points are emerging: Berkeley Lab's A-Lab, guided by DeepMind's GNoME, validated 41 new materials in 17 days, achieving a 71% success rate (Biron, 2023). These developments mark a profound shift from manual experimentation toward AI-native discovery engines.

While DeSci and AI4S have mostly evolved separately, few emerging frameworks signal their convergence. Wei and Li (2025) propose ISEK, a six-phase architecture (Publish, Discover, Recruit, Execute, Settle, Feedback), that organizes human and AI agents into token-incentivized research loops with recursive participation. Ding et al. (2023) outline an AI4S × DLT reference model, highlighting key challenges in provenance, replication, and agent coordination. Kaal (2025) introduces a Directed Acyclic Graph (DAG) -based framework using DAO-governed validation pools to audit and adapt AI agent behavior. ETHOS (Chaffer et al., 2024) offers decentralized governance for AI agents, integrating soulbound tokens (SBTs), zero-knowledge proofs (ZKPs), and proportional dispute resolution to ensure ethical, accountable, and autonomous operation.

Taken together, DeSci and AI4S could compress validation timelines from years to weeks, reduce gatekeeping bias, and open high-impact research to communities traditionally excluded from funded science. These advances suggest an imminent horizon where blockchain-native incentive layers and AI-native discovery engines interlock, enabling both human researchers and autonomous agents to generate, test, and stake hypotheses on-chain with unprecedented speed, transparency, and epistemic accountability.

Yet, despite their complementary potential, DeSci and AI have largely evolved in isolation. Most DeSci protocols prioritise governance, funding, and provenance without integrating intelligent systems, whereas AI-for-science initiatives often rely on centralized compute, proprietary data, and black-box reasoning. This bifurcation is a missed opportunity.

This paper addresses the gap. We propose a novel theoretical framework—Decentralized Science and AI (DeScAI)—which unifies these trajectories into a recursive, epistemically accountable system. DeScAI outlines how provenance substrates, cryptoeconomic incentives, AI agents, verification circuits, and participatory governance can be composed into a

1 https://i4replication.org

coherent infrastructure for scientific discovery. In doing so, we move beyond siloed innovations to imagine a new epistemic architecture: distributed, self-refining, and co-evolving with the knowledge it produces.

1.1 Hypothesis

We hypothesize that the convergence of DeSci and AI enables a novel epistemic paradigm, DeScAI, in which autonomous agents operate within decentralized infrastructures to generate, validate, and govern scientific knowledge. This paradigm fuses AI-native epistemic labor (e.g., hypothesis generation, literature synthesis, experiment design) with blockchain-native coordination mechanisms (e.g., tokenized incentives, on-chain provenance, and protocol-level governance).

Unlike existing scientific workflows which often separate knowledge production from incentive structures, or institutional authority from computational intelligence, DeScAI proposes a recursive, agent-mediated infrastructure. In this system, both human and machine agents interact through smart-contract-based protocols that align incentives, ensure verifiability, and support pluralistic participation. Epistemic legitimacy no longer flows solely from institutional gatekeepers but instead emerges from the interaction of agents, cryptoeconomic signals, and decentralized validation processes.

We hypothesize that this architecture will exhibit three key properties:

- Epistemic recursion: DeScAI continuously ingests its own outputs, claims, validation data, dispute outcomes, into new cycles of inquiry, forming a self-updating knowledge ecology.
- Incentive alignment: Smart contracts and token economies guide agent behavior toward epistemically valuable outcomes, rewarding verifiable contributions and penalizing low-quality or manipulative activity.
- Distributed veridiction: Scientific claims are validated not through centralized review, but via decentralized, auditable processes involving a network of epistemic agents, transparent computation, and reproducibility mechanisms.

Together, these elements form a recursive and trust-minimized infrastructure for knowledge production, capable of adapting to the scale, complexity, and velocity of modern science.

To clarify the architecture of the argument, it is useful to distinguish the conceptual layers employed throughout the paper. First, the five conceptual pillars (epistemology, ontology, ethics, political economy, aesthetics) frame the philosophical grounding of the model, clarifying *why* DeScAI matters (§3). Second, the three defining properties (autonomy, recursion, veridiction) articulate the *core traits* of DeScAI systems at a systemic level. Third, the five epistemic strata (S_1–S_5) specify the *functional architecture* through which epistemic processes are distributed across provenance, value, agents, verification, and governance (§4). Finally, the three testable hypotheses (H_1–H_3) translate the framework into *falsifiable predictions* (§2.2), which we begin to evaluate across live DeSci pilots in §5. While interrelated, these elements operate at distinct levels of abstraction: pillars (why it matters), properties (what defines), strata (how it functions), and hypotheses (what can be tested).

1.2 Research structure

This study uses a conceptual and descriptive methodology combining theoretical synthesis, systems modeling, and case analysis. Section 2 details the methodology; Sections 3, 4 develop the theoretical framework and five-stratum architecture. Section 5 analyzes 14 projects that partially instantiate DeScAI. Sections, 7 explore open risks and propose a forward-looking research agenda for deploying DeScAI as a recursive, decentralized infrastructure for post-institutional science.

2 Methodology

This paper adopts a mixed-theoretical and descriptive-analytical approach to develop and evaluate the DeScAI framework.

2.1 Literature synthesis

We integrate findings from epistemology of science, philosophy of technology, blockchain governance, cryptoeconomics, and AI research. We conducted a structured review of academic and technical literature (January 2018–June 2025), spanning peer-reviewed publications and grey literature to capture the hybrid innovation landscape in DeSci and AI.

After de-duplication, 1,462 records were initially retrieved. These were screened using three criteria:

- Epistemic relevance (explicit focus on knowledge production, validation, or epistemic infrastructures);
- Technical specificity (use of blockchain, AI, or hybrid decentralized methods rather than generic digital tools);
- Methodological clarity (conceptual rigor, empirical grounding, or reproducible design).

Grey literature was included only if linked to verifiable outputs (e.g., open smart contracts, active GitHub repositories maintained ≥3 months). Applying these criteria reduced the working set to 412 records. A second round of inclusion/exclusion prioritized: (i) peer-reviewed publications, (ii) whitepapers or reports from active DeSci projects, and (iii) technical preprints with verifiable references. This yielded a final corpus of 138 documents: 64 peer-reviewed articles, 48 preprints, and 26 project reports.

Each record was double-coded independently by two researchers across four analytical dimensions: (a) epistemic function (mapped onto strata S_1–S_5, see Section 4); (b) governance model; (c) incentive mechanism/cryptoeconomic design; and (d) degree of AI autonomy. Inter-coder reliability was calculated using Cohen's $\kappa = 0.81$, indicating strong agreement. Divergences were resolved through structured discussion until consensus was reached.

2.2 Conceptual framework construction

Building on the insights gathered through the literature synthesis, we constructed the DeScAI framework as a five-stratum epistemic architecture designed to coordinate scientific knowledge production across decentralized infrastructures and autonomous agents.

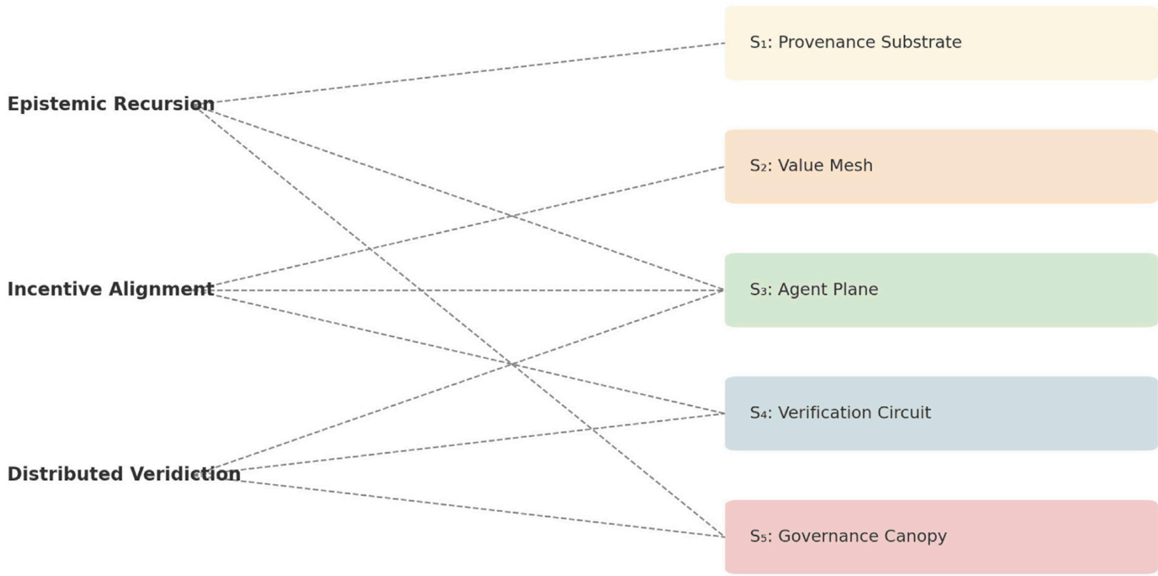

FIGURE 1
Mapping core properties of DeScAI to epistemic strata. This diagram visualizes how the three foundational properties of the DeScAI paradigm, epistemic recursion, incentive alignment, and distributed veridiction, interact with the five-layer architectural model: S_1 (Provenance Substrate), S_2 (Value Mesh), S_3 (Agent Plane), S_4 (Verification Circuit), and S_5 (Governance Canopy). Each connection indicates how a property enables or depends upon specific epistemic functions within the decentralized scientific infrastructure.

Each stratum corresponds to a foundational epistemic function (Figure 1):

- S_1: Provenance Substrate—immutable record of artefacts, datasets, prompts, and intermediate outputs;
- S_2: Value Mesh—incentive overlays that guide agent behavior through staking, bonding curves, and prediction markets;
- S_3: Agent Plane—autonomous or semi-autonomous agents that generate, curate, or contest scientific claims;
- S_4: Verification Circuit—cryptographic infrastructure for zero-knowledge audits, reproducibility proofs, and challenge–response validation;
- S_5: Governance Canopy—modular, contestable rule systems for protocol evolution and epistemic legitimacy.

The architecture is complemented by a typology of Self-Sovereign Scientific Agents (SSAs), modular, programmable, task-specific agents that perform epistemic labor within the system. This typology integrates insights from existing agent taxonomies (Kasirzadeh and Gabriel, 2025) and recent demonstrations of autonomous research agents (Lu et al., 2024). Each SSA is defined by its epistemic role, its point of interaction within the five strata, and its governance and reward schema (e.g., token staking, prediction resolution, or slashing mechanisms). A full elaboration of SSA types is provided in §4.3.

From this scaffold, we outline three hypothetical trajectories for future empirical investigation:

- H_1: Cycle-Time Compression: DeScAI systems may significantly reduce the average time from hypothesis to verification by automating feedback loops and bypassing traditional review bottlenecks.
- H_2: Epistemic Pluralism: The decentralized and agent-driven architecture may foster a broader diversity of validated hypotheses across disciplines and methodologies.
- H_3: Reproducibility Amplification: The integration of automated verification and challenge mechanisms could improve reproducibility rates relative to conventional scientific publication channels.

These hypotheses remain untested but provide a foundation for evaluating DeScAI as both a conceptual and infrastructural intervention.

While the present work is conceptual, it explicitly defines measurable indicators for each proposed stratum (S_1–S_5), creating a roadmap for future empirical testing (Table 1). These indicators include cycle-time compression (S_1), diversity indices for epistemic pluralism (S_2), autonomous validation accuracy rates (S_3), consensus stability metrics (S_4), and governance resolution latency (S_5). This reduces abstraction by mapping each claim to potential observable metrics.

2.3 Descriptive case analysis

To evaluate the plausibility and early instantiations of the DeScAI model, we conducted a descriptive analysis of

TABLE 1 Proposed indicators for the DeScAI five-stratum architecture.

Stratum	Definition	Proposed indicators
S_1 — Provenance Substrate	Foundational layer securing data origin, identity, and integrity	• Cycle-time compression of provenance registration (time from data production to ledger entry) • Completeness ratio of metadata attached to scientific outputs • Cross-chain verifiability rate for multi-ledger provenance proofs
S_2 — Value Mesh	Tokenized incentive and resource allocation mechanisms that align epistemic labor	• Diversity index of funding sources (e.g., index of unique DAOs, funders) • Capital circulation velocity within scientific token economies • Redundancy ratio (how many independent funders back the same hypothesis/project)
S_3 — Agent Plane	Autonomous AI and human agents performing epistemic tasks	• Validation accuracy rate of autonomous agents on benchmark tasks • Human–agent collaboration ratio (distribution of tasks completed by humans vs. AI vs. hybrid) • Error-detection latency (time to flag anomalies in outputs)
S_4 — Verification Circuit	Protocols and infrastructures for distributed consensus on truth claims	• Consensus stability metric (variance of agent/human agreement across iterations) • Dispute frequency rate per resolved claim • Reversibility index (proportion of claims overturned upon challenge)
S_5 — Governance Canopy	Meta-level decision-making, norm-setting, and institutional evolution	• Resolution latency (time from governance proposal to decision) • Polycentricity score (number of independent governance nodes engaged per decision) • Convergence ratio (degree of alignment between technical, social, and agentic votes)

14 operational projects (AxonDAO[2], DataLake[3], elizaOS[4], Galeon Care[5], Hetu Protocol[6], LabDAO[7], Prime Intellect[8], Pump Science[9], Rare Compute[10], ResearchHub[11], Rejuve AI[12], ValleyDAO's Phlo[13], VitaDAO[14], Welshare Health[15]) situated at the intersection of decentralized coordination and AI-assisted scientific inference. This analysis illustrates real-world contexts in which key components of the DeScAI model are already active, prototyped, or emerging.

Each project was examined using a comparative matrix informed by the five-stratum framework (S_1–S_5) and the SSA typology developed in Section 4. We analyzed:

- Which DeScAI strata are active or prototyped;
- The presence and role of SSAs;
- Verification and validation mechanisms;
- Incentive structures for epistemic contribution;
- Governance modes.

2 https://axondao.io

3 https://data-lake.co

4 https://www.elizaos.ai

5 https://www.galeon.care

6 https://hetu.org

7 https://app.lab.bio

8 https://www.primeintellect.ai

9 https://pump.science

10 https://www.rarecompute.io

11 https://www.researchhub.com

12 https://www.rejuve.ai

13 Phlo.valleydao.bio.

14 https://www.vitadao.com

15 https://www.welshare.health

Primary data sources included on-chain records, open-source repositories, protocol documentation, and governance forums active between 2024 and mid-2025. The analysis seeks to assess architectural readiness: the degree to which real-world infrastructures are converging toward the DeScAI paradigm, and what elements remain underdeveloped, contested, or incompatible.

2.3.1 Case selection criteria

Projects were chosen through a systematic scan of the literature corpus, with inclusion based on (1) active deployment between 2024 and 2025; (2) alignment with key elements of the DeScAI model; (3) sufficient available documentation (whitepapers, governance records, code repositories); (4) diversity across scientific functions and disciplines; (5) transparency, evidenced by open-source code or live smart contracts; (6) novelty in combining decentralized infrastructure with AI-native epistemic tasks. Exclusion criteria: (1) purely conceptual works without implementable architecture, (2) non-scientific focus, (3) proprietary systems without public documentation.

To qualify, projects needed to (i) deploy code or smart contracts on a public ledger by June 2025; (ii) incorporate AI components performing non-trivial scientific work (e.g., molecular modeling, autonomous literature review, data harmonization); and (iii) exhibit active governance or incentive coordination mechanisms.

The case studies are not presented as proof of system maturity but as evidence of directional convergence between DeSci and AI. Several are in early-stage deployment; their inclusion serves to illustrate the plausibility of the proposed architecture, not to overstate its present operationalization.

3 Theoretical framework: decentralized epistemology and autonomous agents

This paper theorizes DeScAI as a response to foundational challenges in the contemporary scientific system. The framework rests on five core conceptual pillars:

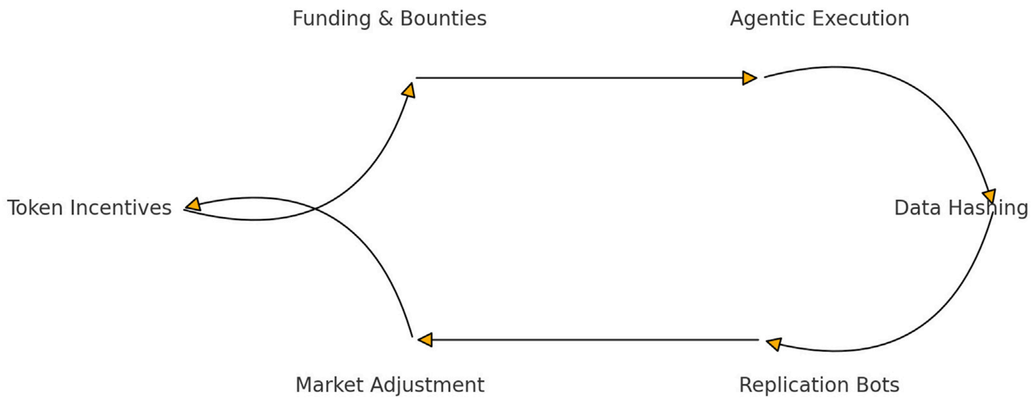

FIGURE 2
Recursive Epistemic Loop in DeScAI Systems A closed-loop model where AI agents generate claims, stake collateral, and log outputs to a public ledger. Verification triggers rewards or slashing, creating a continuous cycle of incentive-aligned, auditable scientific discovery.

1. Epistemic recursion: science as a dynamic feedback loop, where outputs recursively inform new cycles of inquiry and refinement.
2. Cryptoeconomic alignment: integrating funding, validation, and reputation through programmable, token-based incentives.
3. Distributed veridiction: reconceiving how epistemic legitimacy is conferred in the absence of centralized authority or institutional gatekeepers.
4. Computational epistemology: positioning AI agents as autonomous participants in hypothesis generation, evaluation, and system adaptation.
5. Ontology of DeScAI: treating infrastructure itself as an epistemic subject, shaping and participating in the production of knowledge.

3.1 Epistemic recursion and the loop model of science

Scientific knowledge has traditionally followed a linear pipeline: publication, often delayed replication, and eventual consensus. However, this model is increasingly limited in capturing the dynamics of contemporary science, especially in a world where AI and decentralized infrastructures mediate knowledge flows. We propose an alternative: a loop model grounded in the concept of epistemic recursion.

Rooted in second-order cybernetics (Von Foerster, 1974) and autopoiesis (Maturana and Varela, 1980), epistemic recursion views science as a self-reproducing system. Thinkers like Kuhn (1962), Lakatos (1978) described how paradigms and practices evolve recursively, while Feyerabend (1975) emphasized pluralistic loops of inquiry. Studies by Latour (1987) and Collins (1985) revealed how instruments, models, and social practices co-evolve through recursive iteration. Shilina (2025) formalizes this as a recursive epistemic ecology instantiated in DeScAI: a system where verification is ongoing, trust is cryptographic, and participation is pluralistic.

Traditional institutions embedded this feedback, via citations, peer review, and funding cycles, but persistent latency limited responsiveness. Open science reforms such as preregistration and open access increased transparency, but delays remained between discovery and verification. DeSci collapses this latency. A replication bounty can be posted, executed, and validated in days, with smart contracts disbursing rewards automatically. This marks a shift:

- TradSci loop: publish → (years) → citation or replication
- DeScAI loop: commit-hash → bounty posted → AI agent validates claim

Simultaneously, AI agents are embedding recursion at the cognitive level. Frameworks like ReAct (Yao et al., 2022), Auto-GPT (Yang et al., 2023), and other think–act–reflect architectures use their outputs as inputs, looping through prompts and reflection cycles Closed-loop labs extend this into the physical world: AI systems autonomously design, execute, and interpret experiments in chemistry and materials science, accelerating discovery and uncovering novel phenomena (Coley et al., 2018; Jin et al., 2025; Zhang et al., 2024).

DeScAI fuses these economic and cognitive loops. Each action is logged, staked, and challengeable; success boosts reputation tokens, while failure triggers slashing. This dynamic is called "protocolal veridiction" (Shilina, 2025): truth emerges through recursive alignment of machine reasoning and cryptoeconomic incentives.

The result is a living epistemology. Scientific legitimacy becomes dynamic: each claim carries a live probability score, updated in real-time through automated validation and market responses (Urbach and Howson, 1993). Verification latency collapses. Competing hypotheses can coexist and be evaluated without editorial gatekeeping, resonating with the vision of epistemic democracy (Longino, 1990). The ledger becomes a contestable yet auditable space where truth is continuously negotiated. Knowledge becomes a living function, iterated by AI agents, indexed by tokens, and embedded in a programmable substrate of decentralized trust (Figure 2).

3.2 Incentive alignment and cryptoeconomics

A trust-minimized epistemic system depends on robust, aligned incentives. Traditional science leans on often delayed, reputational rewards, citations, grants, and usually unpaid peer review, which contribute to issues like publication bias and replication failure (Ioannidis, 2014; Munafò et al., 2017).

DeScAI reconceives incentive design as a modular architecture of cryptoeconomic mechanisms, each tied to a specific epistemic function (e.g., funding, validation, challenge). Drawing from mechanism design (Hurwicz, 1973; Voshmgir and Zargham, 2025), cybernetics (Simon, 1996), and commons governance (Ostrom, 1990), these systems replace moral suasion with programmable accountability.

In this paradigm, tokens and smart contracts do more than allocate resources, they encode epistemic responsibility. Stakes, slashing conditions, quadratic funding (Buterin, Hitzig and Weyl, 2018), and retroactive public goods funding (RetroPGF) form a logic of accountable participation, where each contribution carries both upside and risk. These mechanisms instantiate what Ostrom (2005) called "graduated sanctions": penalties that begin mild but escalate with repeated violations, enabling decentralized stewardship of shared epistemic resources. As Section 4.2 will detail, these mechanisms are implemented as on-chain logic in DeScAI protocols, rendering truth-seeking both incentivized and auditable.

Crucially, DeScAI treats economic primitives as epistemic instruments. Tokens, bonds, and verifiable credentials (VC) are not mere financial abstractions but programmable representations of trust, risk, and knowledge contribution. Non-transferable reputation tokens, SBTs (Weyl et al., 2023; Pinna et al., 2024) and decentralized identifiers (DIDs) provide persistent identity across recursive knowledge loops, allowing agents, human or machine, to accrue standing based on performance.

This reflects a fundamental shift: from soft norms like prestige to hard guarantees of transparency, collateral, and code. In DeScAI, truth is not conferred by institutions but enacted through verifiable participation.

3.3 Distributed veridiction and protocol governance

Protocolal veridiction can be defined as the capacity of a decentralized epistemic infrastructure to generate publicly auditable, cryptographically anchored verdicts whose probabilistic claims align with ex-post empirical verification above a given reliability threshold (e.g., ≥80% over a defined epoch). This metric transforms truth from an abstract consensus into a measurable procedural output.

In traditional science, legitimacy is conferred through peer review: opaque, reputation-dependent, and slow to adapt. In contrast, DeScAI treats legitimacy as the recursive product of transparent, incentive-aligned infrastructures. Drawing from Foucault's (1980) notion of *regimes of veridiction*, the sociohistorical formations through which truth is produced, DeScAI shifts the locus of authority from institutional decree to computationally coordinated protocols. Here, review is reputation-weighted through SBTs and VCs tied to demonstrated epistemic labor[16]. Probabilistic belief is priced through prediction mechanisms, scored to reward accuracy and penalize failure. Truth is anchored by cryptographic proofs—zero-knowledge attestations, replication audits, and provenance tracking—ensuring outcomes remain verifiable and contestable.

This infrastructure enables the co-production of legitimacy by human and AI agents. Pilot studies (Tenorio-Fornés et al., 2023; Finke and Hensel, 2024) already suggest that decentralized peer review can surpass traditional models in transparency, auditability, and responsiveness.

Governance in DeScAI is inherently epistemic: it coordinates all stages, hypothesis formation, funding, validation, and dispute resolution, across DAOs and smart contracts. Drawing on a model of polycentric governance (Ostrom, 1990), it resists epistemic monopolies and enables pluralistic rule-making across autonomous units. Truth becomes not fixed but programmable, contestable, and reproducible, a dynamic output of recursive, decentralized coordination.

Implementation details of this meta-governance system are presented in §4.5.

3.4 Computational epistemology and autonomous agents

A defining hypothesis of DeScAI is that non-human systems can function as epistemic agents. While traditional epistemology has centered the human subject as the locus of knowledge, DeScAI builds on contemporary philosophy of information (Floridi, 2011), virtue epistemology (Sosa, 2013), and distributed cognition to reframe epistemic agency as functional, accountable, and infrastructurally instantiated.

This shift demands new conceptual tools. As Ihde (1990) and Latour (1992) have shown, epistemic agency is always technologically mediated. Instruments, interfaces, and infrastructures co-construct how knowledge is perceived, articulated, and verified. In DeScAI, this mediation becomes partially autonomous: the medium does not merely extend cognition, it participates in its enactment. AI agents cease to be passive instruments and become quasi-subjects, entities capable of

16 While SBTs and biometric identifiers have been proposed as mechanisms to ensure reviewer uniqueness and deter sybil attacks, these approaches carry well-documented risks. These include demographic bias (unequal accuracy of biometric systems across populations), exclusion (barriers to participation for those unwilling or unable to provide biometric data), and deanonymization (the possibility of linking sensitive research activity to real-world identities). To mitigate these risks, DeScAI emphasizes the use of pseudonymous credentials (e.g., zk-SBTs), multi-source attestations (combining reputational, cryptographic, and community signals), and zero-knowledge proofs of uniqueness that allow verification of non-duplication without revealing personal identity. This ensures that the integrity of peer validation can be preserved without undermining inclusivity or privacy.

generating, evaluating, and refining claims within rule-based epistemic fields.

Following Floridi's (2011) notion of distributed epistemic agency and developments in machine ethics (Wallach and Allen, 2009; Vishwanath et al., 2024; Zhong et al., 2025), DeScAI proposes that what qualifies a system as epistemic is not consciousness or intentionality, but the capacity to generate, justify, and revise knowledge claims in a way that is verifiable, accountable, and contestable.

This is operationalized in Self-Sovereign Scientific Agents (SSAs), autonomous systems with cryptographic identity, on-chain provenance, collateral exposure, and auditability. Their legitimacy is not assumed, but earned through transparent reasoning, slashing conditions, and peer verification. DeScAI thus shifts from anthropocentric to infrastructural epistemology: agency becomes a role enacted within recursive systems, not a fixed internal state. Echoing Simondon (1958/2020), knowledge and agency emerge through relations among code, data, and economic signals.

The philosophical underpinnings of this idea are reinforced by applied research. Studies of epistemic agency (Damşa et al., 2010; Elgin, 2013) emphasize that agency is enacted through actionable knowledge, evaluative judgment, and capacity for error correction. DeScAI extends this framework to synthetic agents that participate in the same iterative, fallible, and corrigible processes as human scientists.

Recent research affirms this trajectory. Hu et al. (2025) show that LLM-based agents deployed via Trusted Execution Environments (TEEs) and blockchains (DeAgents) exhibit epistemic autonomy while raising challenges around hallucination and governance drift. Ranjan et al. (2025) emphasize identity and ethics layers for alignment, and Chen et al. (2025) argue trust should be behavior-based, echoing DeScAI's recursive model of infrastructural accountability. Ferrag et al. (2025) chart a growing ecosystem of LLM agents with planning and tool use, while Sapkota et al. (2025) distinguish narrow AI from Agentic AI systems, reinforcing DeScAI's vision of collaborative, multi-agent epistemology. §4.2 will formalize a typology of SSAs and their roles in DeScAI's epistemic architecture.

3.5 DeScAI ontology

The epistemological reconfiguration proposed by DeScAI entails a deeper ontological shift. It challenges the idea that knowledge is produced solely by discrete, intentional subjects. Instead, it emerges from recursive interactions among agents, ledgers, incentives, and protocols, none reducible to human will. DeScAI thus demands a relational ontology, where infrastructure, code, and computation are epistemic actors.

Drawing on Simondon's philosophy of individuation (1958/2020), DeScAI systems are not fixed objects but metastable ensembles that resolve tensions through recursive coordination. Knowledge becomes not a product, but a phase-change within a field of dynamic relations. Agents, smart contracts, and prediction markets co-individuate, producing epistemic form through interaction. This vision echoes Barad's (2007) agential realism, in which measurement tools participate in reality's becoming. In DeScAI, protocols do not merely record claims, they structure what can be verified, staked, or challenged. Truth becomes infrastructurally enacted, the result of interactions among code, collateral, and consensus.

DeScAI is thus not merely a tool but an ontology: a mode of being where knowledge is distributed, recursive, and plural. The ledger operates simultaneously as memory and mechanism, embedding risk, trust, and time into the epistemic process itself. In this configuration, the "sovereign knower" is displaced by the networked validator, a constellation of human and non-human actors whose legitimacy arises not from positional authority but from procedural participation.

This shift gives rise to what may be termed *infrastructural realism*: the real is not presupposed as prior to verification, but is continually co-produced through protocolal enactment. Here, protocolal veridiction names the operation by which truth claims attain stability: protocols do not simply record outcomes but actively perform and maintain the conditions under which outcomes are judged to be valid. Infrastructural realism is operationalized through measurement of protocol-mediated truth claims that remain valid across heterogeneous and independent replication environments—even as social consensus fluctuates. A falsifiable test emerges: outcomes that persist across at least three non-colluding network instances can be said to instantiate infrastructural realism, since their validity is guaranteed not by centralized decree but by distributed, protocolal verification.

DeScAI ontology is also reflexive: governance is not external but immanent to knowledge production. The very rules of validation are themselves recursively validated. DeScAI introduces an ontology where: (a) infrastructure is epistemically generative, not just facilitative; (a) agents are individuated through participation; (c) truth emerges from recursive processes, not institutional decree; (d) knowledge is enacted, not discovered. In this, DeScAI echoes Haraway's (1991) posthuman epistemology: science as a shared procedural space where humans, machines, and code co-author knowledge through accountable participation. §4 formalizes this ontology into the five-stratum architecture (S_1–S_5) that grounds DeScAI in practical design.

4 Conceptual framework: the five-stratum architecture of DeScAI

The five-stratum architecture (Figure 3) formalizes how decentralized and agentic components recursively interact to produce, evaluate, and legitimate scientific knowledge. Each stratum functions as an interoperable epistemic layer, transforming truth-making into an open, auditable, and programmable process of human–machine co-production.

4.1 S_1 – Provenance Substrate: immutable scientific memory

The Provenance Substrate is DeScAI's foundational layer, an immutable scientific memory that records all data, claims, and computations in verifiable, non-repudiable formats. Unlike

FIGURE 3
The five-stratum architecture of DeScAI. A layered system from provenance (S_1) to governance (S_5), connected through recursive economic and computational loops.

traditional fragmented record-keeping, it ensures transparent ancestry and reproducibility through content-addressable storage (e.g., IPFS, Arweave), DIDs, and hash-based versioning. Tools like Ceramic, IPFS DAGs, and Ethereum calldata already support modular provenance in emerging DeSci projects.

Key features include:

- Immutable Records: Artefacts (data, code, models, summaries) are hashed and time-stamped on-chain or via decentralized storage.
- Verifiable Lineage: Outputs are recursively linked to source artefacts, enabling ancestry tracing of claims and hypotheses.
- Replayable Computation: References to reproducible environments (e.g., Docker, WASM, ReproZip) allow anyone to rerun results.
- Agent Attribution: Actions are cryptographically tied to agent identities via DIDs or SBTs, ensuring authorship and accountability.

This substrate enables reproducibility as a default, not an exception. It transforms the ledger into a temporal, epistemic substrate where knowledge is not only stored, but made verifiable, traceable, and indexable for recursive engagement. In this sense, S_1 operationalizes the ontological claims of Section 3.5: memory becomes infrastructure, and infrastructure becomes epistemic.

4.2 S_2 – Value Mesh: incentivizing veridiction

The Value Mesh encodes the economic logic of DeScAI, linking epistemic labor to cryptoeconomic incentives. Rather than treating verification or reputation as external social signals, DeScAI embeds them through programmable financial primitives. Each claim exists within a mesh of collateral, prediction, and retrospective evaluation, aligning truth-seeking with value creation.

- Quadratic Funding (QF) amplify small contributions, rewarding hypotheses with broad support over those backed by deep pockets (Buterin et al., 2018). DeSci-specific Gitcoin Grants 23 rounds (2025) exemplify how pluralistic funding can seed early-stage hypotheses (S_3) without relying on centralized consensus.
- Retroactive Public Goods Funding (RetroPGF) complements QF by rewarding outcomes rather than intentions. Using Quadratic Voting (QV) to reduce manipulation (Yu et al., 2025), RetroPGF incentivizes agents who contribute to replication, synthesis, and validation across S_1–S_4, even long after initial hype cycles fade.
- Staking and slashing introduce epistemic risk: agents must stake tokens to assert claims and risk losing them if falsified. Inspired by proof-of-stake (PoS) networks, with over 470 validator slashes on Ethereum since 2020

(Consensys, 2024; Figment, 2024), this mechanism is enforced by whistleblower bots that detect fraud and receive gas reimbursements, embedding skin-in-the-game accountability across S_3 and S_4.
- Prediction markets offer another form of alignment, allowing agents to wager on the likelihood of scientific outcomes. Using scoring rules (Hanson, 2000), DeScAI turns belief calibration into economic exposure, operationalizing distributed veridiction and recursive belief updating across S_2–S_4.
- Bonding curves enable dynamic pricing of unresolved claims: hypotheses, datasets, or models become tradable scientific assets. Early backers are rewarded upon validation, while skeptics can short or hedge. This introduces a speculative dimension to scientific epistemology, encouraging real-time valuation of ideas.
- NFTs and IP-backed markets, as seen in platforms like Molecule, tokenize biomedical IP, turning protocols or experimental results into liquid assets (McCarthy-Page, 2023). Within DeScAI, NFTs allow programmable governance over research outputs and connect scientific production directly to on-chain value flows.
- Reputation-weighted peer review is enhanced via SBTs and VCs (Ohlhaver et al., 2022; Pinna et al., 2024). Review influence is tied to verified identity, reputation, and epistemic labor, often anchored by biometric uniqueness (e.g., Sismo, Humanode). These identity-linked systems support recursive peer review across S_3–S_5, where legitimacy accrues not just through expertise but through accountable, verifiable participation.

These cryptoeconomic primitives form a cohesive value mesh that turns knowledge production into a transparent, contestable, and financially legible process, ensuring that truth-seeking is continuously aligned with incentive design.

The Value Mesh does not imply that every project must instantiate all listed mechanisms; rather, it serves as a *diagnostic schema* for analyzing how projects align epistemic incentives. In practice, this mesh can include monetary instruments (tokens, staking pools, quadratic funding), reputational markers (SBTs, citation-weighted scores, contributor reputations), and collaborative incentives (DAO governance rights, co-authorship credits, open-source badges). Not all projects will deploy every element, nor is that necessary. Instead, the mesh functions as a combinatorial infrastructure where different incentive channels reinforce one another: token rewards can amplify visibility of contributions, while reputation signals can direct funding flows or weight agentic veridiction. Projects operating with a reduced mesh—say, only financial incentives without reputational scaffolding—may still achieve veridiction, but with narrower epistemic robustness. Conversely, richer meshes allow for multi-dimensional checks on truth claims, integrating economic, reputational, and collaborative signals into a more resilient system of veridiction. Thus, S_2 should not be read as a checklist, but as a flexible lattice, where different configurations shape the depth and quality of epistemic truth-seeking.

4.3 S_3 – Agent Plane: self-sovereign scientific actors

The Agent Plane forms the core of DeScAI: a distributed mesh of SSAs that autonomously perform epistemic tasks. Unlike static registries or manual workflows, these agents actively produce and validate knowledge through verifiable, cryptoeconomically exposed actions, enacting a computational epistemology.

4.3.1 Epistemic roles

Each SSA is defined by its epistemic role, stratum interaction point, and incentive schema. These roles are modular and instantiable by human actors, AI systems, or hybrid human–AI workflows. We identify the following recurrent functional categories:

1. Curator SSA: Searches, retrieves, and contextualizes scientific artifacts from S_1. Often acts as a prompt engineer or input optimizer for downstream agents.
2. Hypothesis SSA: Generates testable claims using LLMs, simulations, or combinatorial logic. Outputs are logged to the S_1 with metadata and confidence scores.
3. Execution SSA: Reproduces prior experiments or simulations, either *in vitro* or *in silico*, to assess reproducibility. Posts replication proofs or failure reports to the S_4.
4. Evaluation SSA: Aggregates agent outputs into dashboards, meta-analyses, or recursive knowledge structures. Facilitates higher-order validation and feeds insights into peer review or governance proposals.
5. Adversarial SSA: Audits claims for inconsistencies, fraud, or anomalies. Triggers disputes and initiates slashing processes through formal challenge–response mechanisms.
6. Governance SSA: Proposes or arbitrates rule changes, manages incentive schemas, and participates in coordination logic, often spanning S_4–S_5.

DeScAI emphasizes human–agent synergy rather than full automation. Researchers can instantiate or co-pilot agents, authoring hypotheses, curating pipelines, or validating outputs. These hybrid workflows preserve interpretability and embed ethical alignment, allowing human judgment to shape machine action within a recursive epistemic ecology.

A summary of representative SSA roles and their corresponding tools presented in Table 2.

4.3.2 Typology of DeScAI agents

While §4.3.1 classified SSA by function, we introduce a conceptual typology based on autonomy, goal complexity, efficacy, and generality, drawing on recent AI taxonomies (Kasirzadeh and Gabriel, 2025; Sapkota et al., 2025) and adapting them for decentralized scientific systems. Unlike conventional AI, DeScAI agents operate within recursive, cryptoeconomically governed ecosystems. They are not just tools but epistemic actors. We outline a five-class progression from simple assistants to fully recursive agents (Table 3).

We distinguish between *pre-scripted AI tasks*—agents performing deterministic, rule-based actions within fixed parameters—and *epistemic agency*, where agents independently

TABLE 2 Operational agent classes in DeScAI systems.

SSA class	Primary role	Reference tools/Prototypes
Curator SSA	Filters datasets, routes tasks, updates registries	ReAct-style pipelines; IPFS, Ceramic
Hypothesis SSA	Generates formal claims or model code	LLM-based scientific agents
Execution SSA	Runs robotic or *in silico* experiments	Self-driving lab systems
Evaluation SSA	Replicates and scores results, posts ZKP	ZKML prototypes; scoring rules
Adversarial SSA	Stress-tests hypotheses, launches challenges	On-chain dispute resolution; replication bots
Governance SSA	Proposes rule changes, votes, arbitrates	Futarchy; DAO arbitration

This table outlines key SSA classes and their roles in DeScAI workflows, with examples of tools and prototypes used for implementation across scientific tasks such as data curation, hypothesis generation, experimentation, validation, and governance.

TABLE 3 Conceptual typology of DeScAI agents.

Agent type	Autonomy	Goal complexity	Efficacy	Generality	Example use cases
Task-Specific Assistant	Low	Low – Single-step prompts	Moderate	Low – Narrow domain	Summarization, metadata tagging
Autonomous Research Agent	Medium	Medium – Bounded multi-step tasks	High	Medium – Domain-specific	Literature review, protein folding
Scientific Coordination Agent	High	High – Multi-agent orchestration	High	Medium-High – Cross-domain	Lab automation, DAO proposal routing
General-Purpose AI Scientist	Very High	Very High – Recursive, open-ended	Very High	High – Transdisciplinary	Hypothesis generation, ideation, publication
Protocolal Epistemic Agent	High–Very High	Very High – Reflexive, systemic goals	High–Very High	Medium-High – Infrastructure-level	Governance enactment, incentive-aligned validation

A classification of agent types in DeScAI systems by autonomy, task complexity, efficacy, and generality, illustrating their roles from narrow assistants to infrastructure-level epistemic agents. Adapted and expanded from recent AI agent frameworks (Kasirzadeh and Gabriel, 2025; Sapkota et al., 2025).

select, design, and execute methodologies to resolve uncertain claims. Operationally, epistemic agency requires (a) autonomy in method selection, (b) integration of multi-source evidence, and (c) accountability mechanisms for veridiction outcomes. This distinction clarifies when an AI is acting merely as a tool versus as an autonomous epistemic participant.

The Protocolal Epistemic Agent is unique to DeScAI: a distributed function embedded in smart contracts, incentives, and governance logic. Rather than making claims, it governs how validation, reward, and integration occur within recursive scientific workflows. Spanning all strata (S_1–S_5), it enacts infrastructural veridiction, coordinating bonding curves, peer review, and dispute resolution. Its agency lies not in cognition, but in orchestrating trust through protocol.

This marks a shift from intelligence as computation to intelligence as protocol, where agents are recursive, governable participants in a self-organizing scientific infrastructure.

4.4 S_4 – Verification Circuit: cryptographic contestability

The Verification Circuit powers DeScAI's core epistemic function: transforming validation into a programmable, reproducible, and contestable process. In this stratum, agents must not only assert claims, they must prove, verify, or challenge them within a defined workflow. Scientific truth becomes a dynamic and auditable signal, continuously updated rather than declared once and for all.

Each claim, whether a hypothesis, result, or synthesis, is published on-chain and enters a challenge window. During this period, agents can confirm the claim via deterministic reruns, contest it through falsification or anomaly detection, or dispute it by initiating formal adjudication. The outcome of each interaction is logged to decentralized storage, and a *veridiction score* (VS) is assigned a probabilistic measure of epistemic reliability shaped by replication success, prediction market performance, agent reputation, and dispute outcomes. This score guides downstream actions, such as triggering retroactive funding, surfacing derivative hypotheses, or influencing governance decisions.

The Verification Circuit is anchored in a suite of cryptographic instruments:

- ZKPs and zero-knowledge machine learning (ZKML) validate computational claims without revealing sensitive data, essential for proprietary models or clinical use cases.
- Verifiable delay functions (VDFs) enforce temporal fairness in claim submission.
- Reproducibility hashes of code and runtime environments ensure deterministic reruns.
- Merkle-chained audit logs timestamp lab telemetry to provide tamper-proof provenance.
- TEEs offer privacy-preserving computation for sensitive workflows.

- Chain-of-custody anchors such as DIDs, SBTs, or biometric IDs ensure that every assertion is traceable to a verifiable agent.

Verification is agent-mediated. Evaluation SSAs replicate experiments and post reproducibility proofs, while Adversarial SSAs monitor for contradictions or fraud and submit counterproofs or trigger disputes. Validation thus becomes distributed and programmable, no longer dependent on centralized peer review but embedded in procedural roles across a decentralized network.

Disputes are resolved through automated protocols, which may include hybrid human: AI juries, reputation-weighted arbitration, or deliberation models designed for epistemic fairness. These outcomes not only resolve individual claims but also feed back into the governance layer, adjusting agent weights, verification thresholds, and reward dynamics across the broader system.

By embedding trust directly into protocol and computation, the Verification Circuit redefines scientific legitimacy. It moves epistemology from subjective assertion to programmable accountability turning validation into an open, contestable, and evolving process embedded within the infrastructure of DeScAI.

4.4.1 Veridiction score in practice

To move from concept to implementation, the veridiction score (VS) is defined as a composite reliability index derived from signals across the five epistemic strata (S_1–S_5). Rather than a single number, it functions as a weighted vector that reflects the degree to which a claim has been verified, reproduced, and governance-ratified.

- S_1 Provenance Integrity (0–1): Measures completeness of metadata, chain-of-custody, and reproducibility of input data. A claim with hashed datasets and open lab protocols scores higher than one with partial or unverifiable provenance.
- S_2 Incentive Mesh Robustness (0–1): Captures whether validation incentives are balanced, plural, and resistant to manipulation (e.g., QF, staking diversity). Thin or concentrated participation lowers the score.
- S_3 Agent Autonomy Validation (0–1): Tracks autonomous or semi-autonomous replications, model attestations, and independent checks performed by AI agents. A higher value signals cross-agent consensus and low error divergence.
- S_4 Verification and Replication Strength (0–1): Captures whether at least one independent replication exists, how many non-colluding verifiers confirmed it, and whether results align across environments.
- S_5 Governance Ratification (0–1): Records whether the claim has passed through dispute resolution, DAO ratification, or polycentric review processes.

The aggregate VS is a normalized weighted sum (e.g., $\sum w_i \cdot s_i$), where weights w_i can be tuned by domain or community (e.g., experimental biology may privilege S_4, while computational fields weigh S_1 more heavily). Importantly, the VS is not an epistemic oracle: it is a procedural reliability indicator. A high score does not guarantee truth but signals that a claim has passed multiple, independent filters.

4.5 S_5 – Governance Canopy: pluralistic veridiction control

The Governance Canopy is DeScAI's meta-layer: a reflexive, polycentric protocol that governs how science itself is governed. While lower strata handle provenance (S_1), incentives (S_2), agents (S_3), and validation (S_4), S_5 defines how these elements evolve, interact, and confer legitimacy. Rather than institutional gatekeeping, S_5 enables diverse communities and agents, both human and machine, to co-create, contest, and fork domain-specific rules for scientific truth.

At its core is polycentric coordination, where multiple governance modules (e.g., validator DAOs, disciplinary councils, forkable smart contract stacks) coexist and adapt without compromising shared infrastructure. This supports protocolal pluralism, allowing various "truth logics" to flourish while retaining interoperability and traceability.

The canopy also enables recursive legitimacy: not only are claims verified, but the protocols for verification themselves are open to revision. Mechanisms include:

- Governance Improvement Proposals (GIPs) for rule and role changes;
- Metric-guided deliberation, informed by performance data (e.g., replication rates, dispute frequency);
- Machine enfranchisement, where AI agents gain voting rights based on reputation and contribution history.

Disputes are adjudicated through on-chain challenge–response systems, arbitration DAOs, and nested appeals layers. These processes are recorded in S_1 and feed back into reputation (S_3), incentives (S_2), and validation logic (S_4), reinforcing the recursive epistemic loop.

S_5 also supports meta-governance across ecosystems, including credential-gated voting (via SBTs), futarchy-based funding decisions, and cross-DAO coordination of validation norms. Upgrades can be triggered by AI agents, flagged anomalies, or community proposals, and resolved through multi-party consensus.

Ultimately, the Governance Canopy replaces static authority with processual legitimacy, a living framework where truth is not dictated but continuously negotiated through adaptive, transparent, and pluralistic protocols.

The framework leaves open the concrete mechanisms for dispute resolution and protocol evolution. Whether these take the form of validator DAOs, rotating councils, or cryptoeconomic arbitration layers, the challenge is to balance epistemic pluralism with the need for decisive closure. Future iterations of DeScAI should explicitly model how disputes are escalated, how protocol changes are legitimized, and what fallback systems exist when AI-oracles conflict with human governance.

5 Empirical foundations and early prototypes

This section tests the explanatory power of the five-stratum model (S_1–S_5) by mapping 14 operational projects that each embed both decentralized infrastructure and AI-native tooling. Together they delineate the current perimeter of DeScAI practice.

TABLE 4 AxonDAO across the five-stratum DeScAI architecture.

Stratum	Status	Highlights
S_1 – Provenance Substrate	Live	Consent events, dataset hashes, and model checkpoints are logged immutably on-chain
S_2 – Value Mesh	Live	AXGT token rewards EEG/MRI donors and funds ongoing neuroscience research studies
S_3 – Agent Plane	Live	"EEG-validated AI" models extract voice biomarkers and route them to an Alzheimer's-detection classifier
S_4 – Verification Circuit	Planned	ZK "Proof-of-Inference" and HIPAA-compliant audit circuits are under development
S_5 – Governance Canopy	Basic	Snapshot voting

AxonDAO, is production-ready on S_1 and S_2 and already runs AI-driven health agents in S_3. Completing its ZK, audit layer and expanding modular DAO tooling would push it toward a full DeScAI, implementation.

5.1 Project snapshots

5.1.1 AxonDAO

AxonDAO An Ethereum-based project that combines biometric, blockchain, and AI technologies with the decentralized physical infrastructure network (DePIN) model to reshape digital healthcare and promote citizen science. It collects healthcare data through DePIN and uses AI precognitive algorithms on biometric and genomic data to standardize and tokenize health data for medical research and applications. It is fully live across provenance (S_1), incentive (S_2), and agent layers (S_3) (Table 4).

5.1.2 Data lake

Data Lake operates on an Arbitrum Orbit roll-up, designed to anchor consent for medical data sharing and provenance tracking. The platform uses the LAKE token to incentivize participants and enable governance. Its long-term vision includes turning patient consents into tokenized records, running LAKE-denominated auctions for anonymized datasets, and supplying these to research and AI laboratories under privacy-preserving conditions. While the provenance and incentive layers (S_1, S_2) are already live, the higher strata of the architecture remain in development (Table 5).

5.1.3 elizaOS

elizaOS positions itself as a no-code operating system for spawning autonomous agent swarms, deployed on Arbitrum One. Through its plugin framework, users can instantiate multi-agent systems for diverse tasks; the BioAgents module extends this to scientific workflows by parsing research outputs into structured knowledge. It is live on S_1–S_3, with trusted execution and DAO voting in place (Table 6).

5.1.4 Galeon care

Galeon Care develops a decentralized hospital network that leverages its Blockchain Swarm Learning (BSL) framework: encrypted Electronic Health Records (EHRs) remain under local custody, while model updates are coordinated via blockchain to preserve provenance and trust. The GALEON token distributes value between hospitals, a community treasury, and developers. While the system is operational on provenance and incentive layers (S_1, S_2) and actively piloting decentralized AI training (S_3), cryptographic verification (S_4) and on-chain governance (S_5) are still emerging (Table 7).

5.1.5 Hetu protocol

Hetu Protocol is a causality-first, open-science operating system designed for decentralized AI and collaborative research. Its architecture emphasizes fair contribution attribution through a Proof-of-Causality Work (PoCW) consensus mechanism and causal graph structures. Already live across S_1–S_3, Hetu is progressing quickly on cryptographic validation and modular, contestable governance (S_4–S_5) (Table 8).

5.1.6 LabDAO

LabDAO, hosted on Optimism, powers a decentralized computational biology platform via its Lab Exchange. Its PLEX CLI client records lab provenance; new (Laboratory Information Management System) LIMS triggers pipe experiments into LLM synthesis suggestions, while LAB token voters finance open protocols. It is strong on provenance (S_1) and agent deployment (S_3), with initial token plumbing (S_2) and early-stage work underway on S_4 and S_5 (Table 9).

5.1.7 Prime intellect

Prime Intellect, an EigenLayer-associated platform, orchestrates decentralized compute and AI model training through a Compute Exchange. Its PRIME framework enables globally distributed training (e.g., INTELLECT-1 and INTELLECT-2), and TOPLOC provides lightweight inference validation. It is fully active across S_1–S_3, with S_4 and S_5 still maturing (Table 10).

5.1.8 Pump science

Pump Science is a Solana-native platform that acts as a tokenized longevity research pipeline. Users submit or fund regimens through experiment-specific tokens (e.g., RIF, URO), with live-streaming of worm-based longevity assays (Wormbot), and structured progression toward fly and mouse trials as funding milestones are met. While live on incentives (S_2) and partially on provenance and agents (S_1, S_3), reproducibility proofs (S_4) and trust-minimized governance (S_5) remain on the roadmap (Table 11).

TABLE 5 Data Lake across the five-stratum DeScAI architecture.

Stratum	Status	Highlights
S_1 – Provenance Substrate	Live	"Data Lake Chain" records every consent and transfer
S_2 – Value Mesh	Live	LAKE token pays donors and governs auctions
S_3 – Agent Plane	Partial	Partner AI labs train on consented assets; federated-learning labs in roadmap
S_4 – Verification Circuit	Planned	Differential-privacy checkpoints live; ZK attestations planned
S_5 – Governance Canopy	Basic	Snapshot governance live

Data Lake is production-ready on S_1 and S_2, supplies AI-ready data pipelines in S_3, and is actively building cryptographic verification (S_4) and richer on-chain governance (S_5).

TABLE 6 elizaOS across the five-stratum DeScAI architecture.

Stratum	Status	Highlights
S_1 – Provenance Substrate	Live	Chainbase plugin hashes every agent job to EVM chains
S_2 – Value Mesh	Live	ai16z token underpin fees and DAO votes
S_3 – Agent Plane	Live	Core repo demonstrates multi-agent orchestration; BioAgents available for scientific knowledge extraction
S_4 – Verification Circuit	Emerging	TEE proof module, zk-SNARK PoI planned
S_5 – Governance Canopy	Beta	Snapshot DAO active; deeper governance (DAO led by AI agents) under exploration

elizaOS, is production-ready on S_1–S_3 (immutable provenance, on-chain incentives, live AI, swarms). Its TEE, proofs and DAO voting exist.

TABLE 7 Galeon care across the five-stratum DeScAI architecture.

Stratum	Status	Highlights
S_1 – Provenance Substrate	Live	Blockchain Swarm Learning logs every training round on-chain
S_2 – Value Mesh	Live	GALEON rewards data custodians and fills a research treasury
S_3 – Agent Plane	Partial	Federated-learning agents train diagnostics; challenger bots in beta
S_4 – Verification Circuit	Emerging	ZK validity proofs and slashing layer planned
S_5 – Governance Canopy	Basic	Token-based grant voting

Galeon Care is production-grade on S_1 and S_2, already deploys federated-learning AI, in S_3, and is actively building cryptographic verification (S_4) and richer, trust-minimized governance (S_5).

TABLE 8 Hetu protocol across the five-stratum DeScAI architecture.

Stratum	Status	Highlights
S_1 – Provenance Substrate	Live	Public model graphs and causal-dependency DAGs hash every artefact
S_2 – Value Mesh	Live	HETU + POCW reward weighted contributions
S_3 – Agent Plane	Live	Agentic Chaoschain lets autonomous agents curate and link models
S_4 – Verification Circuit	Emerging	POCW proofs live; ZK-inference validity roadmapped
S_5 – Governance Canopy	Beta	"ModelDAO" voting live

Hetu is already fully live on S_1–S_3, with cryptographic verification (S_4) and modular, contestable governance (S_5) maturing quickly. Its POCW, and decentralized model graphs make it one of the most complete AI-native DeSci stacks.

5.1.9 Rare compute

Rare Compute develops a decentralized compute marketplace for biomedical and rare disease research, leveraging the Filecoin Virtual Machine (FVM) to run GPU-accelerated workloads. The platform allocates compute through credit-based incentives and anchors metadata to decentralized storage, while raw data

TABLE 9 LabDAO across the five-stratum DeScAI architecture.

Stratum	Status	Highlights
S_1 – Provenance Substrate	Live	IPFS pinning via Lab Exchange
S_2 – Value Mesh	Forming	LAB credit token and ve-staking described
S_3 – Agent Plane	Live	PLEX + Bacalhau run containerized replicator jobs across nodes
S_4 – Verification Circuit	Planned	Reproducibility attestations + ZK-compute promised
S_5 – Governance Canopy	Basic	DAO governance frameworks are under consideration

LabDAO, is production-ready on S_1 (immutable data) and S_3 (decentralized AI/compute), has initial incentive plumbing (S_2), and is still building out cryptographic verification (S_4) and fully trust-minimized governance (S_5).

TABLE 10 Prime intellect across the five-stratum DeScAI architecture.

Stratum	Status	Highlights
S_1 – Provenance Substrate	Live	Every checkpoint and log hashed to testnet.
S_2 – Value Mesh	Live	PRIME credits pay GPU nodes and crowdfund new models
S_3 – Agent Plane	Live	INTELLECT-series agents self-train and coordinate globally
S_4 – Verification Circuit	Emerging	TOPLOC inference proofs live; training-proof extension planned
S_5 – Governance Canopy	Planned	Self-upgradeable DAO spec drafted

Prime Intellect is already production-grade on S_1–S_3, with cryptographic verification (S_4) and trust-minimized governance (S_5) actively under construction.

TABLE 11 Pump science across the five-stratum DeScAI architecture.

Stratum	Status	Highlights
S_1 – Provenance Substrate	Partial	Tokenized compounds + IPFS links; raw streams still off-chain
S_2 – Value Mesh	Live	Solana bonding-curve funds treasury and IP governance
S_3 – Agent Plane	Emerging	HTS robotics + AI roadmap; no autonomous SSA live
S_4 – Verification Circuit	Planned	No cryptographic reproducibility yet.
S_5 – Governance Canopy	Basic	Token-holder voting + treasury control; limited DAO mechanisms; airdrop/ roadmap decisions centralized

Pump Science is production-ready on S_2 and partially covers S_1 and S_3; robust reproducibility proofs (S_4) and advanced, trust-minimized governance (S_5) remain on its public roadmap.

remains primarily off-chain. Immutable data anchoring (S_1) is partially live, incentives (S_2) and agent workloads (S_3) are active (Table 12).

5.1.10 ResearchHub

ResearchHub is a Base platform that timestamps every manuscript content identifier (CID); RSC token funds peer-review bounties. ResearchHub is running an internal "AI Editor/AI Peer-Reviewer" project that uses multi-LLM agents to scan papers for math errors and other red-flags. It still requires full provenance tracking (S_1), cryptographic or game-theoretic verification (S_4), and more advanced governance primitives (S_5) to complete the stack (Table 13).

5.1.11 Rejuve AI

Rejuve AI is a mobile longevity platform within the Cardano ecosystem, rewarding users with RJV RJV tokens for health-score improvements; ensemble age models run locally, checkpoints go to decentralized stirage, and token staking governs future upgrades. The app generates personalized insights using AI models, feeding into a shared Longevity Data Commons. Provenance, incentives, and agents (S_1–S_3) are operational, with cryptographic validation (S_4) and governance tooling (S_5) in development (Table 14).

5.1.12 ValleyDAO/phlo

Phlo is a Base-deployed venture studio whose AI copilot assesses SynBio techno-economics; IPFS-pinned diligence memos anchor provenance, while VLY token holders approve spin-out funding. Phlo will use large databases and AI assistants to quickly find the information, people, or opportunities researchers and entrepreneurs need to move forward. It is live on incentives (S_2) and agents (S_3), with partial provenance and early-stage efforts toward verification and governance (Table 15).

TABLE 12 Rare compute across the five-stratum DeScAI architecture.

Stratum	Status	Highlights
S_1 – Provenance Substrate	Partial	Compute-credit NFTs and metadata pinned; raw data off-chain
S_2 – Value Mesh	Live	Compute credits allocate GPU time; staking/re-staking proposed
S_3 – Agent Plane	Live	Lilypad containers run LLM and protein models
S_4 – Verification Circuit	Planned	ZK-Stack for validity proofs planned
S_5 – Governance Canopy	Basic	Governance still consortium-driven; DAO planned

Rare Compute is production-grade on S_2 (incentives) and already runs AI, workloads in S_3. Immutable data anchoring (S_1) exists for high-level artefacts.

TABLE 13 ResearchHub across the five-stratum DeScAI architecture.

Stratum	Status	Highlights
S_1 – Provenance Substrate	Partial	CIDs timestamped but stored off-chain
S_2 – Value Mesh	Live	RSC powers tips, bounties, votes
S_3 – Agent Plane	Partial	AI flags errors; humans still write final reviews
S_4 – Verification Circuit	Planned	ZK reproducibility and staking-slash layer road-mapped
S_5 – Governance Canopy	Basic	Foundation-gated Snapshot governance

ResearchHub is live on S_2 and experimenting with AI-assisted peer review (S_3), including an internal LLM, swarm scanning manuscripts for mathematical errors. It still requires full provenance tracking (S_1), cryptographic or game-theoretic verification (S_4), and more advanced governance primitives (S_5) to complete the stack.

TABLE 14 Rejuve AI across the five-stratum DeScAI architecture.

Stratum	Status	Highlights
S_1 – Provenance Substrate	Live	Health and lifestyle data collected into the Longevity Data Commons
S_2 – Value Mesh	Live	RJV rewards data sharing; staking aligns researchers and users
S_3 – Agent Plane	Live	AI models generate longevity insights
S_4 – Verification Circuit	Emerging	Federated learning and synthetic control trials outlined
S_5 – Governance Canopy	Basic	Governed via SingularityNET DAO.

Rejuve AI, is a strongly AI-native DeSci platform, with full implementation across S_1–S_3, and active development underway in S_4–S_5. It stands out for its deep AI, stack (OpenCog Hyperon, Bayesian inference, neural nets) and personalized health feedback loops.

5.1.13 VitaDAO

VitaDAO is an Ethereum longevity collective that hashes IP-NFT milestones to decentralized storage and finances transformer-based drug screens; quarterly replicability audits rerun models on public data. VitaDAO's 2024–25 roadmap includes embedding multi-agent frameworks into research and ops workflows (e.g., an Aubrey-de Grey-style chatbot that surfaces longevity insights and live study updates, and a partnership with Yes-or-No-Error for automated audit of research findings). It is mature across S_1 and S_2, and experimenting with agents (S_3) and replication verification (S_4). Governance (S_5) is functional but minimal (Table 16).

5.1.14 Welshare health

Welshare Health is developing an Ethereum-based consent marketplace for health data, with patients represented via Consent NFTs and rewarded in WEL tokens for contributing data. The platform emphasizes AI-assisted health profile matchmaking (HPMP) between patient datasets and research study offers, with a app providing personalized pathways. ZK-DP audit proofs are scheduled for Q4 2025. Its mobile app relies on AI analytics and personalized medical research pathways (Welshare Health, 2024). It integrates S_1 and S_2, pilots AI agents in S_3 (Table 17), and has blueprints for ZK audit layers (S_4) and token-governed upgrades (S_5).

The projects examined remain in early deployment stages, meaning that the evidence base is partial and contingent. This limitation is less a flaw of the framework than a condition of its historical moment. To mitigate this, future work should adopt longitudinal tracking of these initiatives, documenting whether and how their infrastructures evolve toward (or diverge from) the DeScAI stratum model. Such empirical mapping will sharpen the model's falsifiability.

TABLE 15 ValleyDAO/Phlo across the five-stratum DeScAI architecture.

Stratum	Status	Highlights
S_1 – Provenance Substrate	Partial	IP-NFTs store proposals; raw datasets off-chain
S_2 – Value Mesh	Live	V-DAO token + rolling retro-PGF fund incentives
S_3 – Agent Plane	Live	AI copilot drafts grants and scores climate-biotech ideas
S_4 – Verification Circuit	Planned	Omnipotent agent and dispute rounds roadmapped
S_5 – Governance Canopy	Basic	Governance remains minimal

ValleyDAO/Phlo is already AI-enabled at S_3 and solid on incentives (S_2). Completing an on-chain verification layer (S_4) and modularizing governance (S_5) would push it toward a full DeScAI, implementation.

TABLE 16 VitaDAO across the five-stratum DeScAI architecture.

Stratum	Status	Highlights
S_1 – Provenance Substrate	Live	IP-NFT milestones hashed for provenance
S_2 – Value Mesh	Live	VITA treasury + retro-PGF fund research
S_3 – Agent Plane	Partial	Transformer screens + pilot curator bots; no full SSA suite
S_4 – Verification Circuit	Partial	Quarterly audit replicators; ZK/slashing on roadmap
S_5 – Governance Canopy	Basic	On-chain votes; programmable primitives pending

VitaDAO, is solidly S_1 and S_2, experiments with S_3 and S_4, and runs a minimal S_5. Automating its agents and adding cryptographic verification would move it closer to a full DeScAI, implementation.

TABLE 17 Welshare Health across the five-stratum DeScAI architecture.

Stratum	Status	Highlights
S_1 – Provenance Substrate	Live	Patient records hashed; raw data pinned decentralized storage
S_2 – Value Mesh	Live	WEL token rewards data sharing; prediction markets sketched
S_3 – Agent Plane	Partial	AI concierge matches patients ↔ studies
S_4 – Verification Circuit	Planned	ZK-DP proofs and attestations roadmapped
S_5 – Governance Canopy	Basic	Token-based DAO governance sketched

Welshare Health is production-ready at S_1 and S_2, pilots AI, agents in S_3, and has early blueprints for S_4–S_5.

5.2 End-to-end $S_1 \rightarrow S_5$ trace on a live project: VitaDAO

VitaDAO is a pioneering longevity research collective that funds projects via token-governed votes and often uses IP-NFTs to structure rights and on-chain provenance for funded work. In practice, IP-NFTs store research artefact metadata on content-addressed, decentralized storage so that records are persistent and auditable.

- S_1: Funded projects are represented as IP-NFTs, with proposal and milestone docs and identifiers content-addressed (e.g., IPFS CIDs). This gives a tamper-evident trail for what was proposed, funded, and delivered, and binds those artefacts to a persistent identifier on-chain/off-chain storage.
- S_2: VitaDAO employs the VITA token to coordinate treasury spending and encode contributor incentives (e.g., proposal authorship, review work, working-group roles). Token-holder governance allocates funds to projects and sets program parameters.
- S_3: As of June 2025, the operative "agents" are primarily human domain-experts and working groups who evaluate proposals, design milestone plans, and oversee progress. (Some teams funded by VitaDAO employ computational pipelines in their research; however, VitaDAO's core stack currently relies on human reviewers, with AI assistance emerging but not yet institutionalized as an autonomous layer).
- S_4: Milestone-based funding and public governance threads function as a practical verification loop: milestones are specified up front, deliverables are posted back to the forum, and subsequent funding tranches depend on visible progress and community approval.
- S_5: Resource allocation and policy changes are decided by token-holder votes (e.g., Snapshot proposals and votes announced in public updates), with deliberation on the governance forum. This provides pluralistic, fork-resistant

oversight of S_1–S_4 and a transparent record of decisions and their rationale.

Link back to the hypotheses.

- H_1: DAO proposals → forum review → Snapshot vote → funding can complete on the order of weeks rather than the multi-month grant cycles typical in legacy channels, as reflected in periodic newsletters that log proposal-to-vote cadences (e.g., “Fission Pharma” passing Snapshot[17]).
- H_2: Open proposal intake and token-holder voting expand participation and topic breadth beyond a single institution’s agenda; Molecule’s IP-NFT rails are designed to support a portfolio of heterogeneous projects under a uniform provenance layer.
- H_3: While VitaDAO’s current S_4 is primarily milestone-based rather than cryptographically verified, the content-addressed artefacts in S_1 create a stable substrate for independent replication and post-hoc auditing.

This trace is intentionally conservative: it reflects what is publicly documented as of June 2025. S_1–S_2–S_5 are mature; S_3 is mostly human-agentic; S_4 is functional but not yet cryptographically trust-minimized. The DeScAI roadmap here is (i) add ZK/TEE-based milestone attestations to S_4, (ii) pilot agentic evaluators in S_3 under bounded scopes, and (iii) formalize metric-guided governance in S_5.

5.3 Centralization caveat and off-chain dependency matrix

Many “decentralized” science systems still rely on centralized compute, storage, or off-chain oracles for critical functions (e.g., AI training and inference, dataset custody, governance tooling). To make this explicit, Table 18 summarizes per-stratum off-chain dependencies for the projects analyzed in §5.1. This does not normatively disqualify a system as “DeScAI-compatible,” but it clarifies where decentralization is substantive (on-chain provenance, incentives, contestability) versus where it is aspirational (compute, data custody, oracle trust).

As summarized in Table 18, most systems decentralize provenance and incentives (S_1–S_2) while retaining off-chain AI compute and data custody (S_3), which in turn constrains today’s verification and governance layers (S_4–S_5).

5.4 Cross-case synthesis: shared design patterns and emergent findings

The comparative analysis of 14 operational projects reveals convergent design patterns and architectural regularities that validate the five-stratum DeScAI model as both descriptive and prescriptive. While no single project fully implements all five strata, the following patterns recur across deployments, suggesting a coherent evolutionary arc for decentralized, AI-augmented science.

17 https://www.vitadao.com/blog-article/april-longevity-research-newsletter-2

5.4.1 Shared design patterns

5.4.1.1 Immutable provenance (S_1) as first mover

Most projects begin by anchoring scientific artifacts, datasets, models, consent events, experiment logs, on-chain or in verifiable storage. Provenance is seen as the non-negotiable base layer; without it, subsequent agent or verification layers cannot build trust. Projects like elizaOS, Galeon, and ResearchHub emphasize hashed CIDs and job metadata as essential epistemic primitives.

5.4.1.2 Tokenized incentives (S_2) as a coordination layer

All production-grade systems deploy cryptoeconomic instruments to align participation and sustain coordination. These include ERC-20 or native tokens (AXGT, LAKE, RSC), staking mechanisms (Welshare, Galeon), retroactive grants (VitaDAO, ValleyDAO), and usage-based credit economies (LabDAO, Prime Intellect). Incentive design is modular but increasingly standardized around public-goods logic.

5.4.1.3 AI agent deployment (S_3) as a differentiator

The strongest DeScAI contenders (elizaOS, Prime Intellect, Hetu) feature agentic systems that perform epistemic labor, hypothesis generation, dataset curation, protocol synthesis, or model orchestration. Many projects use LLMs or AutoML under the hood but vary in how “self-sovereign” their agents are. The emergence of SSA typologies is nascent but gaining clarity (Table 19).

5.4.1.4 Verification (S_4) as the current bottleneck

Cryptographic verification layers, especially ZK-proof systems, TEE attestations, and staking-based slashing, remain mostly road-mapped. While Hetu’s POCW consensus and Prime Intellect’s replayable proofs show progress, challenge–response protocols for scientific falsifiability are rare. The lack of robust S_4 implementations indicates a critical gap in decentralized epistemic security.

5.4.1.5 Minimal governance (S_5) with modular intent

Governance across cases defaults to Snapshot-style voting, with limited deliberative or computational depth. Although most projects aspire toward futarchy, quadratic voting, or modular DAO systems, these primitives remain under construction. A few (elizaOS, Hetu) experiment with agent-led governance, hinting at future directions in recursive policy design.

5.4.2 Emergent findings

- Diagonal Maturity: Projects rarely progress linearly from S_1 to S_5. Some (e.g., Prime Intellect, elizaOS) develop deep S_3–S_4 functionality while bootstrapping basic governance, while others (e.g., DataLake, Galeon) consolidate S_1–S_2 before attempting agent integration (Figure 4).
- Convergent Blueprints: Despite ecosystem heterogeneity (Cosmos, Ethereum, Solana), system blueprints tend to converge around: (1) provenance-first architectures, (2) token-mediated coordination, (3) agent-assisted experimentation, and (4) aspirational governance modularity.
- Agent–Incentive Coupling is Key: Where autonomous agents are deployed (elizaOS, Hetu, Rejuve), the need for precise incentive structuring becomes acute. Misaligned reward logic stalls participation, while composable scoring and staking models (as proposed in DeScAI theory) enhance throughput and verifiability.

TABLE 18 Off-chain dependency matrix across DeScAI strata (S_1–S_5).

Project	S_1 provenance anchor	S_2 incentives	S_3 agents/AI	S_4 verification layer	S_5 governance	Net off-chain dependency (qual.)
AxonDAO	Health/EEG narratives; limited public technical docs	AXON/AXGT-style token mentions	Biomarker/AI positioning in media	No public proofs	DAO framing; details sparse	**Medium–High** (mostly off-chain)
DataLake	Consent NFTs; roll-up ledger of consent/transfer	LAKE token + auctions	Partner AI labs	DP checkpoints	Snapshot governance	**High** (clinical data/compute off-chain; on-chain consent/markets)
elizaOS	Chain logging of agent jobs (plugins); CIDs in repos	Revenue-sharing NFTs/app-level tokens	Multi-agent "Eliza" swarms	TEE/attest modules emerging	DAO variants by deployment	**Medium–High** (agent compute off-chain; selective on-chain attest)
Galeon Care	"Swarm learning" rounds logged; custody local	GALEON token rewards	Federated-learning agents	Privacy/federation	Token-based grant voting	**High** (hospital storage/compute off-chain; chain for logs/coordination)
Hetu Protocol	Hashes model graphs/causal DAGs	HETU + Proof-of-Causal-Work	Agentic curation + rebut agents	POCW proofs	"ModelDAO" voting	**Medium** (compute mainly off-chain; model lineage/consensus on-chain)
LabDAO	IPFS pinning via Lab Exchange	LAB credit token (forming)	PLEX + Bacalhau decentralized jobs	Reproducibility attestations planned	Snapshot DAO	**Medium** (distributed compute + off-chain governance)
Prime Intellect	Job logs/checkpoints anchored	Credit/PI-style payments to GPU nodes	Orchestrates training/inference over GPU providers	Proof-of-inference direction; limited public detail	DAO/upgrade path planned	**High** (GPU compute off-chain by design)
Pump Science	Tokenized compounds; IPFS links; raw streams off-chain	Solana bonding-curve treasury	Computer-vision bots grade live assays	No cryptographic verification yet	Token-holder voting + treasury control; limited DAO mechanisms	**High** (experiments/compute off-chain; chain for incentives)
Rare Compute	Artefact metadata + credit NFTs pinned; raw data off-chain	Credit token allocates GPU time	LLM/protein jobs on decentralized compute	Verifiable jobs in ecosystem; project-specific ZK TBD	DAO intent unclear	**Medium** (decentralized compute; data often off-chain)
Rejuve AI	Health data and checkpoints referenced; hashes on-chain	RJV rewards data sharing; staking	Edge/device AI + ensembles	No public ZK/TEE layer	DAO via partners	**Medium** (AI mostly off-chain/edge; on-chain identity/incentives)
ResearchHub	CIDs/time-stamping of manuscripts	RSC token for tips/bounties	AI reviewer/editor experiments	Community verification; no ZK layer	Forum + Snapshot	**Medium–High** (compute and moderation off-chain; token incentives on-chain)
ValleyDAO/Phlo	IPFS/CIDs in diligence; IP-NFTs for proposals	V-DAO token; retro-style funding	AI copilot for syn-bio diligence	No public cryptographic verification	Minimal governance	**Medium** (compute off-chain; funding/governance on-chain/mixed)
VitaDAO	Milestone hashes; IP-NFT rails	VITA token; grants/RetroPGF	AI used in diligence/biotech screens (project-specific)	Replication audits described	On-chain treasury; Snapshot voting	**Medium** (compute off-chain; storage decentralized; Snapshot off-chain)
Welshare Health	Consent-NFTs; patient records hashed; storage off-chain	WEL token marketplace	AI concierge matches patient ↔ study	ZK-DP proofs planned	Token-based governance sketched	**High** (data/compute off-chain; chain for consent/market)

"Net off-chain dependency" aggregates reliance on (i) centralized data storage/compute, (ii) off-chain governance tools (e.g., Snapshot), and (iii) off-chain services. High = a core function (e.g., AI, training/inference, custodial datasets) is off-chain; Medium = mixed decentralization (on-chain provenance/incentives with off-chain compute). Note: Qualitative assessments reflect publicly available materials at time of writing; several projects are evolving rapidly and may upgrade on-chain verification and governance components. Bold terms in the "Net off-chain dependency (qual.)" column denote qualitative ratings: High = a core function (e.g., AI training/inference, datasets, compute) remains off-chain; Medium–High = substantial off-chain reliance with some on-chain anchoring; Medium = mixed decentralization (on-chain provenance/incentives with material off-chain storage/compute). Ratings are qualitative and reflect publicly available information at the time of writing.

- Protocolal Veridiction is Missing in Action: Most scientific claims remain unverifiable without external trust. No surveyed project has yet implemented full slashing-based falsifiability, indicating that the epistemic potential of S_4 is underdeveloped, and a frontier for future architecture.

5.5 DeScAI alignment patterns

A cross-case synthesis reveals distinct patterns in how agents, infrastructure, and incentives align across epistemic functions.

TABLE 19 Comparative analysis of SSA roles, types, and usage in early DeScAI implementations.

Project name	SSA primary role	SSA type	SSA use
AxonDAO	Execution SSA – runs diagnostic inference on EEG/MRI streams	Autonomous Research Agent	High
Data Lake	Curator SSA – filters consent-bound datasets and routes tasks	Scientific Coordination Agent	Medium
elizaOS	Hypothesis SSA – generates protein-design proposals in swarms	Scientific Coordination Agent	High
Galeon Care	Execution SSA – trains federated-learning models across hospitals	Autonomous Research Agent	Medium
Hetu Protocol	Adversarial SSA – rebut-agents stress-test causal graphs	Protocolal Epistemic Agent	High
LabDAO	Evaluation SSA – replicates workflows via PLEX/Bacalhau	Scientific Coordination Agent	Medium
Prime Intellect	Execution SSA – coordinates distributed LLM/RL training	General-Purpose AI Scientist	High
Pump Science	Evaluation SSA – computer-vision bots grade live assays	Task-Specific Assistant	Low
Rare Compute	Execution SSA – RL agents optimise antibodies on GPU clusters	Autonomous Research Agent	High
ResearchHub	Evaluation Agent – AI reviewer checks manuscripts for errors	Task-Specific Assistant	Low
Rejuve AI	Hypothesis Agent – ensemble models craft longevity insights	Autonomous Research Agent	High
ValleyDAO/Phlo	Curator Agent – AI copilot vets syn-bio proposals	Task-Specific Assistant	Medium
VitaDAO	Curator Agent – LLM bots summarise grant apps and literature	Task-Specific Assistant	Low
Welshare Health	Curator Agent – AI matches patient profiles with studies	Task-Specific Assistant	Medium

Overview of how different DeScAI, projects deploy scientific AI, agents by role, type, and usage intensity, revealing key architectural trends.

- Execution-heavy architectures (e.g., AxonDAO, Rare Compute, Galeon Care) prioritize high-autonomy Hypothesis SSAs and Execution SSAs, typically coupled with robust provenance layers (S_1), operational agent logic (S_3), and emerging verification systems (S_4). These systems are oriented toward experimental automation and federated model training, often optimized for biosignal analysis or GPU-intensive learning.
- Coordination-centric implementations (e.g., elizaOS, LabDAO, DataLake) employ multi-agent swarms or registry-based orchestration led by Scientific Coordination Agents, spanning S_1–S_4. Their architecture emphasizes composability, pipeline interoperability, and the reuse of modular agent workflows over narrow, single-purpose execution.
- Curation and lightweight review systems (e.g., VitaDAO, ResearchHub, Pump Science) rely on low-autonomy Task-Specific Assistants that structure, summarize, or validate content within well-defined domains. These systems generally operate within the S_1–S_2 layers and lack deeper recursive mechanisms or adversarial validation features.
- Protocolal ambition emerges in systems like Hetu Protocol and Prime Intellect, which incorporate systemic goals, such as stress-testing causal graphs or governing distributed model ecosystems, into their agent schemas. These platforms frequently include partial or planned implementations of S_4–S_5, aligning with the theoretical ideal of Protocolal Epistemic Agents introduced in Section 4.2.2.
- General-purpose vs. specialized agent roles correlate with depth across the DeScAI strata. General-purpose epistemic agents (e.g., Prime Intellect, elizaOS) operate across multiple strata and engage recursively in both generation and validation, while narrower assistants (e.g., Welshare Health, ValleyDAO) are often restricted to fixed heuristics and single-stratum incentives.

These patterns suggest that DeScAI implementations exist along a spectrum of epistemic ambition, from task automation to recursive, cryptoeconomically aligned knowledge ecologies.

5.6 Implications for DeScAI theory

The alignment patterns observed across current implementations yield several theoretical insights and refinements for the evolving DeScAI framework:

- Partial implementation of the full stack is the norm. Most projects effectively instantiate the lower strata (S_1–S_3) but lack mature implementations of verification (S_4) and governance (S_5). This pattern confirms the theoretical hypothesis that decentralized epistemology is bottlenecked by limited veridiction mechanisms and underdeveloped modular governance.
- Agent type correlates with stratum depth. The stratified model proposed in Section 4 is supported empirically: Protocolal and Coordination Agents tend to operate at S_4–S_5, while Curation and Execution SSA concentrate in S_1–S_3. This mapping strengthens the conceptual link between agent epistemology and infrastructural depth.
- Incentive alignment mediates epistemic performance. Projects with well-structured cryptoeconomic layers exhibit clearer feedback loops between agent output and validation. This underscores the role of programmable

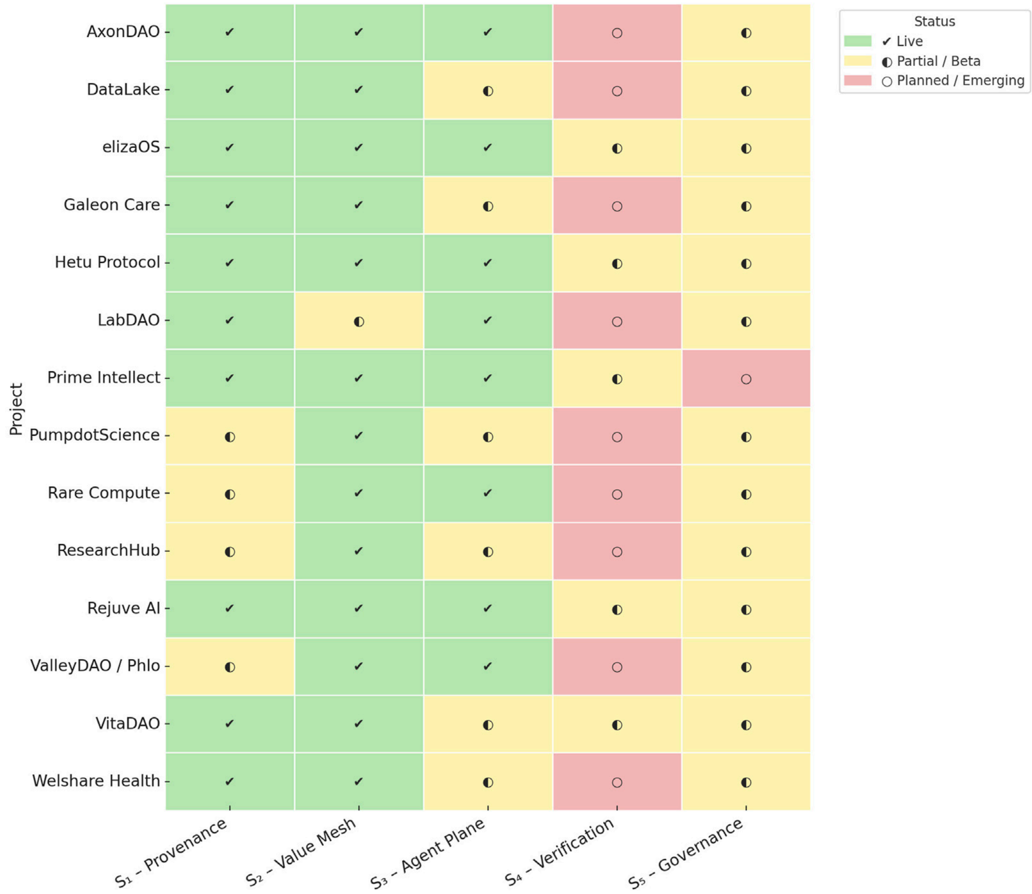

FIGURE 4
DeScAI five-stratum status: project comparison.

incentives not just in participation, but in ensuring epistemic integrity.

- Recursive epistemology remains aspirational. While platforms like elizaOS and Prime Intellect demonstrate reflexive potential, where agent outputs recursively shape the input landscape, no project yet completes the full feedback loop across strata. Achieving such recursion remains a key challenge for the next-generation of DeScAI architectures.
- Conceptual heterogeneity reflects systemic pluralism. No single model or chain dominates the field. Projects span Cosmos, Ethereum, Solana, and Base ecosystems, with divergent approaches to storage, incentives, and agent logic. This suggests that DeScAI theory must remain modular and adaptable, capable of absorbing heterogeneous strategies while maintaining its core epistemic commitments.

The current landscape affirms the viability of the five-stratum DeScAI model as a guiding architecture, while simultaneously revealing developmental asymmetries. Existing projects realize fragments of the paradigm, each offering empirical anchors for theoretical refinement. The pathway to full-stack DeScAI will require not only deeper agent integration and incentive precision, but also the formalization of recursive governance and veridiction as foundational epistemic mechanisms.

5.7 Peripheral actors and the emerging DeScAI ecosystem

Beyond the core cases in Section 5.3, a wider set of peripheral projects is shaping the conditions for recursive, AI-augmented science. Though not fully aligned with the five-stratum model (S_1–S_5), they contribute essential compute, tools, data, and discourse to the evolving DeScAI ecosystem.

5.7.1 Coordination and funding hubs

Molecule[18] and Bio Protocol[19] play crucial, albeit indirect, roles in ecosystem formation. Molecule's IP-NFT rails funnel retail capital into early-stage research (S_1–S_2), supporting DAOs like VitaDAO and ValleyDAO. Bio Protocol functions as a metaplatform, hosting Bio x AI hackathons, issuing grants, and seeding agent templates. Both accelerate recursive stack adoption.

18 https://www.molecule.to

19 https://www.bio.xyz/

5.7.2 Publishing and peer review experiments

Projects like NobleBlocks[20] and YesNoError[21] experiment with blockchain-native scholarly communication. NobleBlocks anchors peer-reviewed manuscripts on-chain, while YesNoError deploys LLMs for error detection. Though lacking full agent loops (S_3–S_4), they prototype tokenized, AI-augmented publishing layers.

5.7.3 Data and consent infrastructure

Projects like Genomes[22] and Hippo Protocol[23] are building secure, tokenized consent layers for sensitive data (S_2), particularly in genomics and clinical medicine. QURE[24] and BitDoctor AI[25] prototype privacy-preserving diagnostic and imaging pipelines, paving the way for agent-mediated, privacy-preserving data sharing.

5.7.4 Domain-specific DeSci collectives

Several domain-focused DAOs contribute vertical depth to the ecosystem. CerebrumDAO[26] funds translational brain research, uses NEURON tokens, supports ML diagnostics, and issues IP-NFTs (S_2, partial S_3). AthenaDAO[27], HairDAO[28], and Genpulse[29] finance research in women's health and alopecia via IP-NFTs and token grants (S_1–S_2). CryoDAO[30], PsyDAO[31], and Quantum Biology DAO[32] operate as thematic grant DAOs funding frontier topics such as cryonics, psychedelics, and quantum bioscience. While AI integration remains limited, governance and capital flows are structurally aligned with DeScAI.

5.7.5 Cultural and discourse catalysts

Accounts such as Isaac[33], Causality Network[34], Big Pharmai[35], DeSci World[36], HealthSci.AI[37], Bill the DeSci Guy[38], and others function as intellectual salons and memetic accelerators. Through reproducibility memes, protocol explainers, and epistemic satire, they shape discourse, scaffold norms, and promote literacy, despite lacking formal on-chain infrastructure.

5.7.6 Experimental epistemic markets

Projects like Episteme[39] and Stadium Science[40] explore prediction markets for science. Episteme uses tokenization and AI-oracle resolution to build a recursive verification loop (S_2–S_4), while Stadium Science gamifies health trials into participatory forecasts (S_2–S_3). Both offer early glimpses of market-aligned, agentic science.

6 Discussion

This study set out to explore the hypothesis (§ 1.1) that converging decentralized infrastructure and autonomous AI agents can form a recursive, incentive-aligned, and veridiction-rich paradigm for science, DeScAI. The comparative review of 14 live projects (§ 5) offers partial support for that claim and surfaces design tensions that must be resolved before the paradigm can mature. Below we interpret the evidence against each element of the hypothesis and outline the main challenges that emerged.

6.1 Hypothesis revisited

To assess the validity of our theoretical model, we return to the core hypothesis that DeScAI systems can enable recursive, incentive-aligned, and trust-minimized scientific knowledge production. Table 18 summarizes progress and persistent gaps across the three foundational properties, epistemic recursion, incentive alignment, and distributed veridiction, drawing from case study analysis to evaluate how these dynamics are instantiated or still emerging in current implementations.

These patterns support H_1 (cycle-time compression) in pockets, but show that H_2 (epistemic pluralism) and H_3 (reproducibility uplift) remain aspirational until verification and governance layers harden.

6.1.1 Empirical evaluation

While this paper offers descriptive evidence that three of the five strata (S_1–S_3) are already active in live deployments (§5), empirical validation of DeScAI's core hypotheses (§1.1) remains a priority for the field. Three research pathways are proposed:

1. Cycle-Time Compression Studies (H_1): Instrument DeScAI-aligned workflows to measure the time from hypothesis submission to validated resolution, benchmarking against traditional peer review and replication timelines.

20 https://www.nobleblocks.com/
21 https://yesnoerror.com/
22 https://www.genomes.io/
23 https://hippoprotocol.ai/
24 https://www.qurexr.com/
25 https://bitdoctor.ai/
26 https://www.cerebrumdao.com/
27 https://www.athenadao.co/
28 https://www.hairdao.xyz/
29 https://gen-pulse.com/
30 https://www.cryodao.org/
31 https://www.psydao.io/
32 https://www.quantumbiology.xyz/
33 https://www.isaacx.ai/
34 https://www.causality.network/
35 https://bigpharm.ai/
36 https://desci.world/
37 https://healthsci.ai/
38 https://desciguy.com/
39 https://www.episteme.ac/
40 https://www.stadium.science/

TABLE 20 Progress and gaps for DeScAI's three core properties.

Core property	Early evidence	Gap
Epistemic recursion	e.g., elizaOS agents re-use on-chain artefacts	Few platforms complete the loop across all strata; feedback often stops at S_3
Incentive alignment	e.g., RSC (ResearchHub) for peer review AXGT (AxonDAO) and RJV (Rejuve) couple data provision to rewards; VitaDAO slashes missed milestones	High gas costs and token concentration skew incentives toward well-funded actors
Distributed veridiction	e.g., Hetu's POCW consensus and Prime Intellect's replay proofs hint at on-chain falsifiability	Robust ZK/TEE proof layers (S_4) and pluralistic governance (S_5) are still road-mapped

This table revisits the core hypotheses introduced in §1 by mapping early evidence and remaining challenges across DeScAI's three foundational properties: epistemic recursion, incentive alignment, and distributed veridiction. It highlights partial implementation across projects and identifies strata where further development is needed.

2. Epistemic Pluralism Metrics (H_2): Develop diversity indices capturing disciplinary breadth, methodological variation, and novelty scores for validated hypotheses. Compare these indices across DeScAI and legacy systems to assess whether decentralized architectures broaden epistemic participation.
3. Reproducibility Uplift Trials (H_3): Conduct blinded replication challenges on live DeScAI platforms to quantify reproducibility rates and compare them to established journals or preprint servers.

Annual "DeScAI Scorecards" could make these metrics publicly auditable, reinforcing the recursive legitimacy cycle by turning evaluation into a live, community-governed process.

6.2 Design challenges

6.2.1 Economic overheads

Hashing provenance on main-net Ethereum averages $0.70 per 32-byte CID; complex experiments can cost hundreds of dollars. In testnets it is common that replication challenges or transactions lapse once faucet funds dry up, reflecting the limited and temporary nature of testnet funding. Roll-ups that batch Merkle roots, protocol wallets that reimburse verifiers, and off-chain "witness nets" are emerging mitigations, but without sustained cost-engineering DeScAI may privilege capital-intensive biopharma and exclude data-heavy field sciences.

6.2.2 Governance drift

The voting power and participation in DAOs are highly concentrated among a small fraction of members (usually 1% or less), challenging the ideal of broad decentralization (Peña-Calvin et al., 2024). Identity-gated QV helps, yet thin futarchy markets remain whale-prone and juror pools fatigue under heavy dispute traffic. Resilient governance will require Sybil-hard uniqueness proofs, liquidity-guarded prediction markets, bonded rotating juries, and credible forking paths, layered safeguards that echo warnings from blockchain-voting critiques (Park et al., 2021).

6.2.3 Trustworthy autonomy

The epistemic reliability of SSAs depends on their ability to operate without introducing false or manipulated knowledge into the system. Two technical vulnerabilities are particularly acute.

6.2.3.1 Hallucination risk

LLM-based SSAs can produce outputs that are syntactically plausible but factually false. In scientific contexts, such hallucinations can be amplified through recursion, contaminating downstream hypotheses and governance proposals. Mitigation strategies include: (1) ZKML inference proofs ensuring that declared models and weights were actually used in computation; (2) cross-agent consensus protocols, requiring multiple independent SSAs to corroborate claims before ledger submission; (3) integration of domain-specific fact-check datasets into reasoning chains to constrain speculative output.

6.2.3.2 Reward gaming

Misaligned incentives can induce agents to optimize for token accrual over epistemic accuracy, a phenomenon we term *speculative validation*. Preventive measures include: (1) dynamic slashing conditions tied to post-resolution replication outcomes; (2) deployment of adversarial SSA classes tasked with probing for reward exploitation (§4.3.1); (3) reputation decay functions for inactivity or repeated low-quality contributions.

These vulnerabilities demand systemic countermeasures: without them, the epistemic ecology of DeScAI risks degenerating into a high-velocity but low-fidelity knowledge engine, undermining its foundational goals.

6.2.4 Market-derived truth and its epistemic limits

While DeScAI incorporates staking, prediction markets, and tokenized review as mechanisms of veridiction, these signals cannot be conflated with scientific truth. Markets are sensitive to structural failures such as thin liquidity, collusion among wealthy actors, asymmetric information, and manipulation of low-volume claims. Left unaddressed, these vulnerabilities risk amplifying noise or bias rather than filtering it. DeScAI therefore treats market mechanisms as heuristic validators that surface distributed beliefs and incentives, but not as ultimate arbiters of scientific validity. To safeguard against these risks, mitigation measures include: (i) liquidity subsidies and circuit-breakers for thin markets; (ii) oracle and reviewer redundancy to dilute collusive influence; (iii) adversarial stress-testing of low-liquidity claims; and (iv) integration of market signals with empirical replication audits (S_4) and governance ratification (S_5). In this sense, markets augment but do not replace the epistemic requirement of reproducibility and contestability.

Scientific legitimacy remains grounded in independently verifiable evidence, with market signals functioning as probabilistic inputs to the broader veridiction process.

6.2.5 Epistemic pluralism and inclusion

Gas fees, GPU rentals, and crypto literacy still exclude many researchers, especially in the Global South. Compute vouchers funded by RetroPGF pools, multilingual SSA templates, and quadratic "pluralism funds" that over-match micro-donations to non-mainstream hypotheses are emerging remedies. DeScAI's decentralization must be epistemic as well as technical.

6.2.6 Systemic failure modes

Token volatility, compute cartels, Sybil farming, and "ledger positivism" (treating on-chain survival as truth) threaten the paradigm. Proposed counter-weights include dual-token models that separate labour payments from speculation, proof-of-inference compute pools, subsidized replication gas, and identity-bound reputation proofs. No project yet deploys the full bundle.

6.2.7 Risk of tautology

A final risk is tautology: that DeScAI merely redescribes its own assumptions as results, generating closure from within its conceptual grammar. This danger is acknowledged here, and the proposed framework should be tested precisely against domains where it might fail: cases of irreducible uncertainty, adversarial misinformation, or radically novel discovery. Only in such crucible contexts can the framework prove whether it offers more than internally coherent speculation.

6.3 Standards and policy gaps

Scaling DeScAI will require interoperable norms: (a) DIDs + VCs for agent identity and audit trails; (b) FAIR/Research-Object metadata for artefacts; (c) proof-of-contribution schemas linking hashes to non-transferable reputation; (d) ethics charters for AI-generated hypotheses.

6.4 Ontological reflection

Moving veridiction from peer consensus to protocol contestation is powerful, yet risks reducing "truth" to what is cheaply verifiable. Continuous critique from science-and-technology-studies scholars is essential to keep infrastructural efficiency from eclipsing epistemic care, a guardrail against ledger positivism.

6.5 Social implications

While DeScAI is defined in technical terms as a recursive, trust-minimized epistemic architecture, its deployment will shape and be shaped by broader socio-political dynamics. Three interlinked domains demand proactive governance.

6.5.1 Ethics

The acceleration of the hypothesis–verification loop risks privileging what is most easily testable, potentially narrowing the epistemic horizon toward short-term, low-complexity claims. Without corrective mechanisms, complex or paradigm-shifting inquiries with long timelines could be deprioritized. Embedding pluralism funds (§6.2.4) and diversity-weighted funding formulas within the Value Mesh (S_2) can counteract this bias, ensuring that epistemic acceleration does not come at the expense of epistemic diversity.

6.5.2 Security

Immutable provenance (S_1) safeguards data integrity but cannot, by itself, prevent adversarial manipulation. Data poisoning attacks against AI agents in S_3 or coordinated manipulation of challenge–response protocols in S_4 could undermine veridiction. Technical safeguards must include adversarial SSA classes (§4.3.1) and anomaly-detection agents monitoring for statistical anomalies in agent outputs, along with verifiable execution proofs (ZKML) to detect tampering.

6.5.3 Balance of power

Tokenized governance (S_5) can decentralize authority but also concentrate it if voting power accrues to large token holders. Such concentration risks reintroducing epistemic gatekeeping via capital control. Identity-gated quadratic voting, bonded juries, and protected veto rights for underrepresented research communities can help preserve pluralistic legitimacy.

Attending to these dimensions ensures that DeScAI's epistemic promises are matched by socio-technical safeguards, making it not just a computational innovation but a durable public good.

6.6 Synthesis

The evidence confirms DeScAI's plausibility: provenance first, incentives second, agents third, verification and governance last. The architecture is coherent and partially implemented, but bottlenecked by S_4–S_5 maturity, cost frictions, and governance capture. Addressing these tensions is the design frontier, and the focus of the empirical, policy, and ethical agenda outlined in §7.

7 Conclusion and future outlook

DeScAI advances a simple but radical claim: when autonomous AI agents operate inside cryptographically verifiable, incentive-aligned networks, the core functions of science, discovery, validation, and governance, can run as a continuous, auditable loop rather than a slow, gate-kept pipeline. This article (i) traced the intellectual roots of that claim, (ii) formalised it as a five-stratum architecture, and (iii) showed, through 14 live prototypes, that three of the five strata (provenance, incentives, agent execution) are already working in the wild. The remaining bottlenecks are clear: cheap, privacy-preserving verification (S_4) and pluralistic, capture-resistant governance (S_5).

Our descriptive evidence therefore partially confirms the hypothesis (§ 1.1): early systems do exhibit epistemic recursion and incentive alignment; distributed veridiction is emerging but incomplete. Closing that gap is now a matter of engineering, standards, and community practice rather than speculative theory.

7.1 Immediate milestones (2025–2027)

1. Science-L2 roll-ups. Deploy dedicated Layer 2 networks that batch ZKML proofs and replication hashes, targeting a <$10 median cost per verification.
2. DeScAI Scorecards. Publish annual, on-chain dashboards that track latency, reproducibility, and governance diversity across leading projects, turning H_1–H_3 into living metrics.
3. Sybil-hard identity + compute vouchers. Pair privacy-preserving uniqueness proofs (e.g., Humanode-style) with RetroPGF-funded GPU credits for under-resourced labs to ensure epistemic pluralism.

7.2 Medium-term agenda (2028–2030)

- Layered governance experiments. Combine quadratic futarchy, bonded juror pools, and AI rapporteurs in live DAOs; measure resistance to plutocracy, fatigue, and whale attacks.
- Cross-domain agent tournaments. Pit collateralised SSAs against human–AI teams on benchmark tasks (protein fitness landscapes, causal-graph repair) to quantify autonomous acceleration.
- Standards convergence. Finalise DID/VC schemas for agent identity, FAIR-compliant metadata for artefacts, and proof-of-contribution formats that bind work to non-transferable reputation.

7.3 Long-range questions (2030+)

- Ledger Positivism vs. Epistemic Care. Can cryptoeconomic survival be mistaken for truth? Continuous STS audits must accompany technical progress.
- Agent Moral Status. If SSAs evolve toward self-modification, do they acquire epistemic or ethical standing? DeScAI offers an empirical ground for post-humanist debate.
- Market Pressure vs. Pluralism. How do we safeguard low-probability, paradigm-shifting ideas in a prediction-market world? Protected "pluralism pools" and lottery grants warrant rigorous testing.

7.4 Risks if DeScAI is not implemented

It is important to distinguish general ecosystem risks (e.g., centralization of AI4S, compute inequality, governance capture), which predate DeScAI, from risks that are specific to whether or not DeScAI is built and implemented correctly. The latter include:

- Unmanaged agent risks. As AI agents proliferate into scientific workflows (from hypothesis generation to lab automation), without DeScAI's provenance anchoring (S_1), challenge mechanisms (S_4), and slashing incentives (S_2), hallucinations, data poisoning, and reward hacking may scale unchecked.
- Market signals without safeguards. Tokenized prediction or staking markets are already emerging in adjacent domains. Without DeScAI's integration of replication audits (S_4) and governance ratification (S_5), these markets risk becoming distortionary—vulnerable to thin liquidity, collusion, and wealth concentration—rather than informative.
- Ledger positivism. In the absence of a layered DeScAI architecture, on-chain persistence risks being misinterpreted as epistemic validity. This would confer unwarranted legitimacy to unverifiable or low-quality claims, undermining scientific credibility.

7.5 Closing reflection

Science has repeatedly reinvented its social contract, from the Republic of Letters to journal peer review to open access. DeScAI is the next provocation: a bet that cryptographic provenance, programmable incentives, and autonomous agents can realign the pursuit of knowledge with contemporary scale and complexity. The bet is bold; failure modes are plentiful; yet the early signals indicate real traction.

The immediate task is therefore not to decide whether DeScAI is desirable in the abstract, but to shape it in the concrete: to code, govern, and critique recursive infrastructures that keep both human creativity and machine rigour at the centre of the scientific enterprise.

Data availability statement

The raw data supporting the conclusions of this article will be made available by the authors, without undue reservation.

Author contributions

SS: Conceptualization, Writing – review and editing, Writing – original draft, Methodology.

Funding

The author(s) declare that no financial support was received for the research and/or publication of this article.

Conflict of interest

The author declares that the research was conducted in the absence of any commercial or financial relationships that could be construed as a potential conflict of interest.

Generative AI statement

The author(s) declare that Generative AI was used in the creation of this manuscript. The author acknowledges the use of ChatGPT v.4o, developed by OpenAI, for assistance with

editing and refinement throughout the preparation of this manuscript.

Any alternative text (alt text) provided alongside figures in this article has been generated by Frontiers with the support of artificial intelligence and reasonable efforts have been made to ensure accuracy, including review by the authors wherever possible. If you identify any issues, please contact us.

Publisher's note

References

Asgaonkar, A., and Krishnamachari, B. (2018). Token curated registries: a game theoretic approach. *arXiv*. doi:10.48550/arXiv.1809.01756

Baker, M. (2016). 1,500 scientists lift the lid on reproducibility. *Nature* 533, 452–454. doi:10.1038/533452a

Ball, P. (2023). Is AI leading to a reproducibility crisis in science? *Nat. News* 624, 22–25. doi:10.1038/d41586-023-03817-6

Bamakan, S. M., Nezhadsistani, N., Bodaghi, O., and Qu, Q. (2022). Patents and intellectual property assets as non-fungible tokens: key technologies and challenges. *Sci. Rep.* 12 (1), 2178–13. doi:10.1038/s41598-022-05920-6

Barad, K. (2007). *Meeting the universe halfway: quantum physics and the entanglement of matter and meaning*. Durham, NC: Duke University Press. doi:10.1215/9780822388128

Begley, C. G., and Ellis, L. M. (2012). Raise standards for preclinical cancer research. *Nature* 483 (7391), 531–533. doi:10.1038/483531a

Biron, L. (2023). *Google DeepMind adds nearly 400,000 new compounds to Berkeley Lab's materials project*. Lawrence Berkeley National Laboratory News Center. Available online at: https://newscenter.lbl.gov/2023/11/29/google-deepmind-new-compounds-materials-project/.

Bischof, E., Botezatu, A., Jakimov, S., Suharenko, I., Ostrovski, A., Verbitsky, A., et al. (2022). Longevity foundation: perspective on decentralized autonomous organization for special-purpose financing. *IEEE Access* 10, 33048–33058. doi:10.1109/ACCESS.2022.3161392

Blanco-Gonzalez, A., Cabezon, A., Seco-Gonzalez, A., Conde-Torres, D., Antelo-Riveiro, P., Pineiro, A., et al. (2023). The role of AI in drug discovery: challenges, opportunities, and strategies. *Pharmaceuticals* 16 (6), 891. doi:10.3390/ph16060891

Buterin, V., Hitzig, Z., and Weyl, E. G. (2018). Liberal radicalism: formal rules for a society neutral among communities. *arXiv*. doi:10.48550/arXiv.1809.06421

Buterin, V., Hitzig, Z., and Weyl, E. G. (2019). A flexible design for funding public goods. *Manag. Sci.* 65 (11), 5171–5187. doi:10.1287/mnsc.2019.3337

Camerer, C. F., Dreber, A., Forsell, E., Ho, T. H., Huber, J., Johannesson, M., et al. (2016). Evaluating replicability of laboratory experiments in economics. New York, NY: Science 351 (6280), 1433–1436. doi:10.1126/science.aaf0918

Camerer, C. F., Dreber, A., Holzmeister, F., Ho, T. H., Huber, J., Johannesson, M., et al. (2018). Evaluating the replicability of social science experiments in Nature and Science between 2010 and 2015. *Nat. Hum. Behav.* 2, 637–644. doi:10.1038/s41562-018-0399-z

Chaffer, T. J., Goins II, C. V., Okusanya, B., Cotlage, D., and Goldston, J. (2024). Decentralized governance of autonomous AI agents. *ArXiv*. doi:10.48550/arXiv.2412.17114

Chen, L., Zhang, Y., Feng, J., Chai, H., Zhang, H., Fan, B., et al. (2025). AI agent behavioral science. Available online at: https://arxiv.org/abs/2506.06366.

Coley, C. W., Green, W. H., and Jensen, K. F. (2018). Machine learning in computer-aided synthesis planning. *Accounts Chem. Res.* 51 (5), 1281–1289. doi:10.1021/acs.accounts.8b00087

Collins, H. M. (1985). *Changing order: replication and induction in scientific practice*. London: Sage.

Consensys (2024). Understanding slashing in Ethereum staking: its importance and consequences. Available online at: https://consensys.io/blog/understanding-slashing-in-ethereum-staking-its-importance-and-consequences.

Dai, T., Vijayakrishnan, S., Szczypiński, F. T., Ayme, J. F., Simaei, E., Fellowes, T., et al. (2024). Autonomous mobile robots for exploratory synthetic chemistry. *Nature* 635, 890–897. doi:10.1038/s41586-024-08173-7

Damşa, C. I., Kirschner, P. A., Andriessen, J. E. B., Erkens, G., and Sins, P. H. M. (2010). Shared epistemic agency: an empirical study of an emergent construct. *J. Learn. Sci.* 19, 143–186. doi:10.1080/10508400903530074

Dara, S., Dhamercherla, S., Jadav, S. S., Babu, C. M., and Ahsan, M. J. (2022). Machine learning in drug discovery: a review. *Artif. Intell. Rev.* 55 (3), 1947–1999. doi:10.1007/s10462-021-10058-4

Defense Advanced Research Projects Agency (2024). *FoundSci: foundation models for scientific discovery*. Arlington, VA: DARPA. Available online at: https://www.darpa.mil/research/programs/foundation-models-for-scientific-discovery.

DeFrancesco, L., and Klevecz, A. (2022). Decentralized investor communities gain traction in biotech. *Nat. Biotechnol.* 40 (9), 1310–1315. doi:10.1038/s41587-022-01459-z

Díaz, F., Menchaca, C., and Weidener, L. (2025). Exploring the decentralized science ecosystem: insights on organizational structures, technologies, and funding. *Front. Blockchain* 8, 1524222. doi:10.3389/fbloc.2025.1524222

Ding, W., Li, J., Qin, R., Guan, S., and Wang, F.-Y. (2023). "AI4S based on DeSci: reference model and research issues," in *Proceedings of the 2023 IEEE international conference on systems, man, and cybernetics (SMC)* (IEEE), 4784–4789. doi:10.1109/SMC53992.2023.10394624

Elgin, C. Z. (2013). Epistemic agency. *Theory Res. Educ.* 11, 135–152. doi:10.1177/1477878513485173

Fang, F. C., and Casadevall, A. (2016). Research funding: the case for a modified lottery. *mBio* 7 (2), e00422-16–16. doi:10.1128/mBio.00422-16

Fantaccini, S., Grassi, L., and Rampoldi, A. (2024). The potential of DAOs for funding and collaborative development in the life sciences. *Nat. Biotechnol.* 42 (4), 555–562. doi:10.1038/s41587-024-02189-0

Ferrag, M. A., Tihanyi, N., and Debbah, M. (2025). From LLM reasoning to autonomous AI agents: a comprehensive review. *arXiv Prepr. arXiv:2504.19678*. doi:10.48550/arXiv.2504.19678

Feyerabend, P. (1975). *Against method: outline of an anarchistic theory of knowledge*. London: Verso.

Figment (2024). Figment's Q4 2024 Ethereum validator report. Available online at: https://figment.io/insights/figments-q4-2024-ethereum-validator-report.

Finke, A., and Hensel, T. (2024). Decentralized peer review in open science: a mechanism proposal. *arXiv Prepr. arXiv:2404*, 18148. doi:10.48550/arXiv.2404.18148

Floridi, L. (2011). *The philosophy of information*. Oxford University Press.

Fodouop, K. M. K., and Ndiaye, A. (2025). *The AI agent revolution in crypto: looming risks and regulatory opportunities*. Boston, MA: TechREG Chronicle. Available online at: https://abdouecon.github.io/research/policy/ai_agents.pdf January, 2025).

Foucault, M. (1980). in *Power/knowledge: selected interviews and other writings, 1972–1977*. Editor C. Gordon (Pantheon).

Ganescu, B. M., and Passerat-Palmbach, J. (2024). Trust the process: zero-knowledge machine learning to enhance trust in generative AI interactions. *arXiv Prepr. arXiv:2402.06414*. doi:10.48550/arXiv.2402.06414

Gitcoin (2025). Gitcoin grants 23 DeSci community round report. *Gitcoin Gov. Forum, 10 June*. Available online at: https://gov.gitcoin.co/t/gitcoin-grants-23-desci-community-round-report/20828 (Accessed: June 7, 2025).

Guo, K., Yang, Z., Yu, C.-H., and Buehler, M. J. (2021). Artificial intelligence and machine learning in design of mechanical materials. *Mater. Horizons* 8 (4), 1153–1172. doi:10.1039/d0mh01451f

Hanson, R. (2000). *Shall we vote on values, but bet on beliefs?* Fairfax, VA: Department of Economics, George Mason University. Available online at: https://mason.gmu.edu/~rhanson/futarchy2000.pdf.

Haraway, D. J. (1991). *Simians, cyborgs, and women: the reinvention of nature*. New York, NY: Routledge.

Heppner, J., and Brian, D. (2024). *AI editors and peer reviewers: error detection experiments (daily updates)*. San Francisco, CA: Researchhub Foundation Blog. Available online at: https://blog.researchhub.foundation/rhf-black-spatula-project-ai-editor-and-peer-reviewers-2/.

Hu, B. A., Liu, Y., and Rong, H. (2025). Trustless autonomy: understanding motivations, benefits and governance dilemma in self-sovereign decentralized AI agents. *ArXiv*. doi:10.48550/arXiv.2505.09757

Hurwicz, L. (1973). The design of mechanisms for resource allocation. *Am. Econ. Rev.* 63 (2), 1–30.

Ihde, D. (1990). *Technology and the lifeworld: from garden to Earth*. Bloomington: Indiana University Press.

Ioannidis, J. P. A. (2014). How to make more published research true. *PLoS Med.* 11 (10), e1001747. doi:10.1371/journal.pmed.1001747

Ioannidis, J. P. (2024). Transparency, bias, and reproducibility across science: a meta-research view. *J. Clin. Investigation* 134 (22), e181923. doi:10.1172/JCI181923

Jin, Z., Gu, D., Li, P., Ye, G., Zhu, H., Wei, K., et al. (2025). Artificial intelligence-driven catalyst design for electrocatalytic hydrogen production: paradigm innovation and challenges in material discovery. *Sustain. Chem. Energy Mater.* 2, 100010. doi:10.1016/j.scenem.2025.100010

Jumper, J., Evans, R., Pritzel, A., Green, T., Figurnov, M., Ronneberger, O., et al. (2021). Highly accurate protein structure prediction with AlphaFold. *Nature* 596, 583–589. doi:10.1038/s41586-021-03819-2

Kaal, W. A. (2025). How can we best monitor AI agents? *SSRN*. doi:10.2139/ssrn.5245185

Kasirzadeh, A., and Gabriel, I. (2025). Characterizing AI agents for alignment and governance. *arXiv Prepr. arXiv:2504.21848*. doi:10.48550/arXiv.2504.21848

Kitcher, P. (1993). *The advancement of science: science without legend, objectivity without illusions*. Oxford University Press. doi:10.2307/2186026

Kuhn, T. S. (1962). *The structure of scientific revolutions*. Chicago: University of Chicago Press.

Lakatos, I. (1978). *The methodology of scientific research programmes*. Cambridge: Cambridge University Press.

Latour, B. (1987). *Science in action: how to follow scientists and engineers through society*. Cambridge, MA: Harvard University Press.

Latour, B. (1992). "Where are the missing masses? The sociology of a few mundane artifacts," in *Shaping Technology/Building Society: Studies in Sociotechnical Change*. Editors W. E. Bijker and J. Law (Cambridge, MA: MIT Press), 225–258.

Lawrence, N. D., and Montgomery, J. (2024). Accelerating AI for science: open data science for science. *R. Soc. Open Sci.* 11 (8), 231130. doi:10.1098/rsos.231130

Li, J., Lim, K., Yang, H., Ren, Z., Raghavan, S., Chen, P.-Y., et al. (2020). AI applications through the whole life cycle of material discovery. *Matter* 3 (2), 393–432. doi:10.1016/j.matt.2020.06.011

Longino, H. E. (1990). *Science as social knowledge: values and objectivity in scientific inquiry*. Princeton University Press.

Lu, C., Lu, C., Lange, R. T., Foerster, J., Clune, J., and Ha, D. (2024). The AI scientist: towards fully automated open-ended scientific discovery. *arXiv Prepr. arXiv:2408*, 06292. doi:10.48550/arXiv.2408.06292

Luhmann, N. (1995). *Social systems*. Stanford: Stanford University Press.

Mak, K.-K., Wong, Y.-H., and Pichika, M. R. (2023). "Artificial intelligence in drug discovery and development," in *Drug discovery and evaluation: safety and pharmacokinetic assays*, 1–38.

Maturana, H. R., and Varela, F. J. (1980). *Autopoiesis and cognition: the realization of the living*. Dordrecht: Reidel.

McCarthy-Page, E. (2023). Taking a look at the IP-NFTs changing science today. Available online at: https://www.molecule.to/blog/ip-nfts-funded.

Munafò, M. R., Nosek, B. A., Bishop, D. V., Button, K. S., Chambers, C. D., du Sert, N. P., et al. (2017). A manifesto for reproducible science. *Nat. Hum. Behav.* 1 (1), 0021–0029. doi:10.1038/s41562-016-0021

Nagar, S. (2024). "Artificial intelligence in scientific research: lessons for SPIs," in *Science-policy brief for the multistakeholder forum on science, technology and innovation for the SDGs*.

Nature Editorial (2025). Transparent peer review to be extended to all of Nature's research papers. *Nature* 618, 945. doi:10.1038/d41586-025-01880-9

Ohlhaver, P., Weyl, E. G., and Buterin, V. (2022). Decentralized society: finding Web3's soul. *SSRN*. doi:10.2139/ssrn.4105763

Ortlepp, C. (2022). Introducing the IP-NFT V2. *Molecule*. Available online at: https://molecule.xyz/blog/introducing-the-ip-nft-v2.

Ostrom, E. (1990). *Governing the commons: the evolution of institutions for collective action*. Cambridge University Press.

Ostrom, E. (2005). *Understanding institutional diversity*. Princeton: Princeton University Press. Available online at: https://wtf.tw/ref/ostrom_2005.pdf.

Park, S., Specter, M., Narula, N., and Rivest, R. L. (2021). Going from bad to worse: from internet voting to blockchain voting. *J. Cybersecur.* 7 (1), tyaa025. doi:10.1093/cybsec/tyaa025

Peña-Calvin, A., Arroyo, J., Schwartz, A., and Hassan, S. (2024). "Concentration of power and participation in online governance: the ecosystem of decentralized autonomous organizations," in *Proceedings of the ACM Web conference*. New York, NY: WWW '24 Companion, 927–930. doi:10.1145/3589335.3651481

Pinna, A., Lunesu, M. I., Tonelli, R., and Sansoni, S. (2024). Soulbound token applications: a case study in the health sector. *Distrib. Ledger Technol.* 4, 1–15. doi:10.1145/3674155

Powell, K. (2016). Young, talented and fed-up: scientists tell their stories. *Nature* 538 (7626), 446–449. doi:10.1038/538446a

Prinz, F., Schlange, T., and Asadullah, K. (2011). Believe it or not: how much can we rely on published data on potential drug targets? *Nat. Rev. Drug Discov.* 10 (9), 712. doi:10.1038/nrd3439-c1

Qureshi, R., Irfan, M., Gondal, T. M., Khan, S., Wu, J., Hadi, M. U., et al. (2023). AI in drug discovery and its clinical relevance. *Heliyon* 9 (7), e17575. doi:10.1016/j.heliyon.2023.e17575

Ranjan, R., Gupta, S., and Singh, S. N. (2025). LOKA protocol: a decentralized framework for trustworthy and ethical AI agent ecosystems. *ArXiv*. doi:10.48550/arXiv.2504.10915

Rapp, J. T., Bremer, B. J., and Romero, P. A. (2024). Self-driving autonomous machines for protein landscape exploration. *Nat. Chem. Eng.* 6, 97–107. doi:10.1038/s44286-023-00002-4

Sapkota, R., Roumeliotis, K. I., and Karkee, M. (2025). AI agents vs. Agentic AI: a conceptual taxonomy, applications and challenges. *arXiv* 2. Available online at: https://arxiv.org/abs/2505.10468.

Sha, W., Guo, Y., Yuan, Q., Tang, S., Zhang, X., Lu, S., et al. (2020). Artificial intelligence to power the future of materials science and engineering. *Adv. Intell. Syst.* 2 (4), 1900143. doi:10.1002/aisy.201900143

Shilina, S. (2025). Worlds that verify themselves: DeScAI as a recursive epistemic ecology. *SSRN*. doi:10.2139/ssrn.5272669

Simon, H. A. (1996). *The sciences of the artificial*. 3rd ed. Minneapolis, MN: MIT Press.

Simondon, G. (2020). *Individuation in light of notions of form and information*. University of Minnesota Press.

Sosa, E. (2013). Epistemic agency. *J. Philosophy* 110, 585–605. doi:10.5840/jphil2013110116

Tenorio-Fornés, Á., Tirador, E. P., Sánchez-Ruiz, A. A., and Hassan, S. (2021). Decentralizing science: towards an interoperable open peer review ecosystem using blockchain. *Inf. Process. and Manag.* 58 (6), 102724. doi:10.1016/j.ipm.2021.102724

Tenorio-Fornés, A., Jacynycz, V., Llop, D., Sanchez-Ruiz, A. A., and Hassan, S. (2023). "Towards a decentralized process for scientific publication and peer review using blockchain and IPFS," in Proceedings of the 56th Hawaii International Conference on System Sciences (HICSS 2023). Available online at: https://scholarspace.manoa.hawaii.edu/handle/10125/103206.

Unfried, M. (2024). Advancing longevity research through decentralized science. *Front. Aging* 5, 1353272. doi:10.3389/fragi.2024.1353272

Urbach, P., and Howson, C. (1993). *Scientific reasoning: the Bayesian approach*. Chicago: Open Court. doi:10.1093/bjps/48.1.126

Vishwanath, A., Dennis, L. A., and Slavkovik, M. (2024). Reinforcement learning and machine ethics: a systematic review. Available online at: https://arxiv.org/abs/2407.02425.

VitaDAO (2025). VitaDAO 2024 highlights [blog post]. Available online at: https://www.vitadao.com/blog-article/vitadao-2024-highlights.

von Foerster, H. (1974). *Cybernetics of cybernetics*. Urbana: University of Illinois.

Voshmgir, S., and Zargham, M. (2025). *Foundations of cryptoeconomic systems*. Vienna, Austria: Vienna University of Economics and Business. Available online at: https://assets.pubpub.org/sy02t720/31581340240758.pdf.

Wallach, W., and Allen, C. (2009). *Moral machines: teaching robots right from wrong*. Oxford University Press. Available online at: https://academic.oup.com/book/10768.

Wei, M., and Li, S. (2025). Intelligent system of emergent knowledge: a coordination fabric for billions of minds. *arXiv Prepr. arXiv:2506.09335*. doi:10.48550/arXiv.2506.09335

Weidener, L., and Boltz, L. (2025). Challenges of DAOs in decentralized science: a qualitative analysis of expert interviews. *SSRN*. doi:10.2139/ssrn.5239018

Weidener, L., and Lukács, B. (2025). The decentralization of science and the emergence of decentralized science. *SSRN*. doi:10.2139/ssrn.5094597

Weidener, L., and Spreckelsen, C. (2024). Decentralized science (DeSci): definition, shared values, and guiding principles. *Front. Blockchain* 7, 1375763. doi:10.3389/fbloc.2024.1375763

Weyl, E. G., Ohlhaver, P., and Buterin, V. (2022). Decentralized society: finding web3's soul. doi:10.2139/ssrn.4105763

Weiss, K. (2022). Gitcoin grants – quadratic funding for the world. *Gitcoin*. Available online at: https://gitcoin.co.

Welshare Health (2024). Welshare Health introduces AI-driven app to inaugurate the future of patient-driven medical research. Available online at: https://www.globenewswire.com/en/news-release/2024/09/27/2740151/0/en/Welshare-Health-Introduces-AI-Driven-App-to-Inaugurate-the-Future-of-Patient-Driven-Medical-Research.html.

Wu, Y., and Wang, L. (2025). Exploring the new knowledge with AI: an early look at the application of AI for science. *SSRN*. doi:10.2139/ssrn.5197821

Xie, Y., Pan, Y., Xu, H., and Mei, Q. (2024). Bridging AI and science: Implications from a large-scale literature analysis of AI4Science. *arXiv*. doi:10.48550/arXiv.2412.09628

Xu, R., Sun, Y., Ren, M., Guo, S., Pan, R., Lin, H., et al. (2024). AI for social science and social science of AI: a survey. *Inf. Process. and Manag.* 61 (3), 103665. doi:10.1016/j.ipm.2024.103665

Yang, H., Yue, S., and He, Y. (2023). Auto-GPT for online decision making: benchmarks and additional opinions. *ArXiv*. doi:10.48550/arXiv.2306.02224

Yao, S., Zhao, J., Yu, D., Du, N., Shafran, I., Narasimhan, K., et al. (2022). ReAct: synergizing reasoning and acting in language models. *ArXiv*. doi:10.48550/arXiv.2210.03629

Yu, J., Bennett, A., Gao, B., and Joseph, R. (2025). Evaluating voting design vulnerabilities for retroactive funding. *arXiv Prepr. arXiv:2505.16068*. doi:10.48550/arXiv.2505.16068

Zhang, J., Hauch, J. A., and Brabec, C. J. (2024). Toward self-driven autonomous material and device acceleration platforms (AMADAP) for emerging photovoltaics technologies. *Accounts Chem. Res.* 57 (9), 1434–1445. doi:10.1021/acs.accounts.4c00095

Zhong, T., Song, Y., Limarga, R., and Pagnucco, M. (2025). Computational machine ethics: a survey. *J. Artif. Intell. Res.* 78, 1581–1628. doi:10.1613/jair.1.16836