An Active Inference Ontology for Decentralized Science: from Situated Sensemaking to the Epistemic Commons

Authors (ORCID):

Daniel Friedman [1,2] (0000-0001-6232-9096)

Shaun Applegate-Swanson [1,2] (0000-0001-9474-6038)

Arhan Choudhury [1] (0000-0001-6897-2947)

RJ Cordes [1,2] (0000-0002-9913-7159)

Shady El Damaty [3] (0000-0002-2318-4477)

Avel Guénin—Carlut [1,4] (0000-0001-8239-7264)

V. Bleu Knight [1,2] (0000-0002-9894-1989)

Ivan Metelkin [1] (0000-0003-2899-6690)

Siddhant Shrivastava [1,2,5] (0000-0002-9688-4730)

Amit Kumar Singh [1] (0000-0002-8870-9676)

Jakub Smékal [1,3] (0000-0003-4989-4968)

Caleb Tuttle [3] (0000-0002-8469-2221)

Alexander Vyatkin [1] (0000-0003-1306-4620)

Author Affiliations:
1. Active Inference Lab (activeinference.org/)
2. Complexity Weekend (complexityweekend.com/)
3. OpSci (opsci.io/)
4. Kairos Research (kairos-research.org/)
5. Kernel (kernel.community/)

Email address for correspondence & collaboration: ActiveInference@gmail.com

Abstract

In this work, we examine science from the vantage points of blockchain technology and its connection to decentralized science (DeSci). We consider science as a collective process using Active Inference, an integrative framework that models the cognitive processes of perception, planning, and action selection in terms of Bayesian probabilities and updating. We present the Active Entity Ontology for Science (AEOS, available at coda.io/@active-inference-institute/active-entity-ontology-for-science-aeos) as a composable and versionable system for modeling various science systems, using the Active Inference entity partitioning. Further steps for developing and utilizing AEOS in the context of scientific ecosystems are provided.

Table of Contents

Driving Questions

Centralized Science (CeSci): Operations, Community, and Practice

- What are successful, effective, and productive scientific practices for communities and individuals?
- What frameworks and tools can help form and support communities of practicing researchers to promote long-term collaboration and impact?
- How are preferences and expectations aligned and communicated in ecosystems of scientific collaboration?
- How do we avoid perverse incentives in systems of scientific inquiry?
- How do we ensure effective and transparent resource allocation and funding as research becomes more interdisciplinary, interorganizational, and international?
- How can scientific careers be started and nurtured, increasing the accessibility and vitality of science as a community and body of knowledge?
- What are fair and effective means of elevating voices to scientific thought leadership roles?
- What are good decision trees to identify strategic priorities for scientific funding, intellectual capital allocation, and focused development?

Decentralized Science (DeSci): Coordination Engineering for Knowledge Building

- How can decentralized teams and organizations coordinate research using effective distributed mechanisms?
- What scientific outputs are within the scope of Decentralized Science (DeSci)?
- How will DeSci produce traditional research products such as papers, as well as other outcomes, such as products, services, platforms, tools, and protocols?
- What artifacts and design patterns will stimulate the development of integrated DeSci systems?
- How can Distributed Autonomous Organizations (DAOs) scaffold and catalyze research, grants, education, and community development?
- What methods for value capture and financialization improve the process of scientific discovery, and which methods result in poor outcomes that are at odds with an open science community?
- What improves the reliability of outcomes and accountability of funds in emergent collaborations?
- What are the relations between the trajectories of scientific governance, and social/ecological/political change and collapse?
- How can scientific careers become more accessible to researchers who are non-PhD holders, and those with other commitments such as family, participation in industry, etc?
- How can we use blockchain technology and signature economies to cultivate cultures of fair work attribution, through records of contribution to knowledge artifacts and proof-based protocols?
- How can we maintain a favorable signal-to-noise ratio for DeSci work output?
- Can the Active Inference framework and ontology be used to understand complex sociotechnical systems such as epistemic communities?

Introduction

Science is a cumulative, collective endeavor that augments a single human's sensemaking ability with tools, practices, and processes we use together to observe, measure, and understand the world around us. Even before the invention of modern scientific instruments such as microscopes and particle accelerators, humans relied on their senses to observe the environment and on their brains to make inferences about these observations. This is still true today. Any study of how knowledge is accumulated must consider how the brain works when making predictions based on observations and incomplete data. Additionally, as the information required for scientific research extends far beyond the cognitive capacity of a single unaided human, tools for scientific organization, knowledge management, and communication become increasingly necessary to consider as well.

Science and the Epistemic Commons

A recent shift in science is the move towards greater recognition of the complex nature of systems such as brains or societies [1,2]. For example, cognitive science has largely moved away from the ambition to discover universal laws of behavior to focus more on the causal structure of the brain and its patterns of activity in different scenarios [3]. Utilitarian scientific models today strive to explain natural phenomena by piecewise integration of generative models expected to match the causal structure of the target system at the relevant scale of behavior [4,5]. The development and validation of mechanistic models is based on multiple cognitive strategies, for example involving computational modeling, empirical observation, statistical inference, and domain-specific theoretical considerations. Despite this, the social structure of science has not yet shifted to embrace integrative practices. While the scientific community has long-standing practices which they optimize for, many of these parameters have lost their validity, efficacy, or legitimacy [6,7]. This results in suboptimal integration of complex transdisciplinary knowledge by favoring the partitioning of research along the lines of domain-specific methodologies, organizational processes, thematics, and vocabulary.

Indeed, the modern era has seen the emergence of an unprecedented hegemony of administrative institutions over all other forms of human social organization. Political anthropologist James C. Scott explains the emergence of top-down structures for complex problem solving as a default reflex to compress complexity. Scott claims that the capacity of bureaucratic states and administrators to understand and manipulate the world relies on a systematic reduction of the world onto standardized, legible tokens (teachable facts), such as accounting, and measurement units or legal identities based on a permanent

4

patronym [8]. The need for this "view from above" has placed pressure on research institutions to adapt both their objectives and operating activities to meet the needs of administrators. This "technoscientific" logic has left a complex and pervasive legacy in the organization of contemporary research, including in the way we understand and practice scientific research. In worst-case settings (which are unfortunately common), research quality is evaluated on superficial measures with no principled consideration of deeper relevance, leading to a raging competition between researchers on the number of words written, sheer number of publications of any kind, or the ability to use technical language [9–12], without considering the positive impact of the work. Entire research programs can be funded because of insider networks, or correlation to a hype cycle, while more thoughtful and critical (and ultimately constructive) research may easily go unnoticed. In other words, the key metrics for tracking global scientific progress have become increasingly divorced from their ability to replicate or have meaningful impact [13,14], and more correlated with social and cultural signaling within the scientific community. Further, when the professional scientific community is only a small fraction of society, there is the possibility of detachment or isolation from broader goals.

Given the incentive misalignment and power imbalances between administrators and practitioners, the scarcity of resources, competition over resources both between researchers and research institutions, and the inflexible nature of the bureaucracies which connect and govern extant research institutions, there are few mechanisms available for improving traditional research institutions aside from increasing funding. Luckily, human prosociality has taken many forms throughout our evolutionary history from which we can draw inspiration. We could develop tools and protocols which treat the outcomes of research as contributions to an epistemic commons, a web of informational or knowledge-oriented (epistemic) systems which would allow individuals with disparate motivations, skill sets, and beliefs to collaborate at the intersection of their interests and tackle specific questions or systems with the most contextually relevant tools, in the best interest of cumulative, collective scientific knowledge. In other words, we argue that relaxing the constraints of legibility associated with bureaucracy and its "view from above" could help develop the "view from the ground" (i.e., pragmatic understanding by practitioners) as well as the "view from within" (i.e., scientific knowledge, in the classical sense of an objective description of system organization and activity).

Human prosociality has contributed to the persistence of our species. The emergent structure of social organization is variable but can develop towards hierarchical, top-down control which are transactional and mutually beneficial [15,16]. However, hierarchical control can lead to the sequestration of wealth and power over time, and a centralized motivation that does not always benefit all members of the community. Furthermore, hierarchical structures develop cognitive biases over time that limit the ability to parse

5

multiple, conflicting streams of information that do not fit into standard decision making templates. Decentralized structures may be more capable of acting on divergent information streams, however decentralization alone is not a mechanism that provides intelligent coordination for effective collective action [17,18]. Decentralized Autonomous Organizations (DAOs) are an emerging form of coordination that enables new kinds of cyber-physical communities mediated by immutable, cryptography-protected code [19,20]. DAOs allow for programmable digital social structures, embedded mechanisms for consensus, intermediation of trust-requiring actions, and independent operation of workstreams with multiple leaders instead of a top-down hierarchy. Coordinating such decentralized communities of practice, or more generally governing "complex commons", is novel, difficult, and high-stakes work, often without the scaffolds and norms that traditional offline governance can benefit from [21]. For example, systems that are decentralized in principle may still be functionally centralized in terms of power or asset distribution [22].

We propose that epistemic drives and norms can grow organically from communities of practice, defined by open questions, methodological synergies, personal affinities, diverse contributions, and serendipitous encounters. We model the similarities and differences between Decentralized Science (DeSci) and Centralized Science (CeSci) through the Active Inference framework, understood here as a framework addressing how living systems create an understanding of their environment [23]. The Active Inference formalism affords an integrated account of the multiscale dynamics shaping research, from the constitution of epistemic affordances shaping scientific practices [24] to the coevolution of epistemic beliefs and communities [25], as a process of uncertainty reduction entailed by the descent of a gradient of variational free energy. As applied, this model implies that cognitive understanding derives from adaptive control in agentive engagement rather than explicit representation [26–28], and that the activity of epistemic communities is shaped by the sociocultural constraints defining their structure rather than individual states of mind [29]. Based on these considerations, we articulate a conceptual framing and Active Inference entity ontology model for Decentralized Science.

Before we dive deep into head-on contrastive definitions of CeSci and DeSci, a primer on decentralized systems is crucial. We would like the first definition not to be simply an anti-definition of its counterpart, as we aim to incentivize effective research and set up collaborative relations between emergent DeSci and incumbent CeSci, not simply to ignite controversy. Therefore, even though here we will use the terminology of decentralization, considering nuance in the definition [22,30–32] as well as alternative terms, 'Decentralized' is a good first step. One interesting perspective to apply here is that of "Self Certifying Systems" [33]. These self-certifying systems are those in which an entity validates its own existence in the system, and the incentive mechanism of the system is

6

intrinsic to it. For instance, a dataset in a database system such as InterPlanetary File System (IPFS) [34,35] validates its own legitimacy by hashing the contents using a cryptographic hash function. In these dynamics, the root of trust lies within the network entity itself, and not necessarily the legitimation from sources outside of the system.

Centralized Science (CeSci) and Decentralized Science (DeSci)

Decentralized Science (DeSci) is a relatively new term, introduced by Web3 collectives to describe the use of recent digital tools for funding, training, planning, coordination, execution, dissemination and archival of scientific activities and assets by digitally connected communities. DeSci is an emerging area with multiple, potentially even incompatible, forming senses, so all explorations here should be considered preliminary and partial. The term DeSci suggests, by contrast, the existence of Centralized Science (CeSci), which would stand for the continued status quo organization of science as a highly bureaucratic activity managed by select academic and private institutions. Some motifs or patterns that are found in CeSci and DeSci can be found in Table 1 and Figures 6-8.

DeSci can be defined as the use of Web3 technologies to introduce markets through the deployment of open source finance tools. In this view, DeSci represents the introduction of commodity markets for scientific assets and services where such markets were previously impossible (e.g. financialization through tokenization of intellectual property, scientific platform governance, peer review or curation services, or access privileges to data or infrastructure). In another perspective, DeSci is seen as a set of mechanisms for bottom-up individual sensemaking. From this perspective, DeSci is the capability of individual agents or communities of practice to make sense of the world autonomously, by defining their own questions, language and methodology. These disparate views represent the distinction (but not preference or judgment) between DeSci as a set of emerging Web3 tools and DeSci as a research ecology, respectively.

Tool use, including markets and organizational technologies like DAOs, are fundamental to the ecological view of science [36]. However, proponents of DeSci do not conceive of this distinction as a purely technological one, for example simply describing the fact that DeSci organizes with the Internet while CeSci utilizes paper as well as digital media is an insufficient distinction. For example DAOs don't replace off-chain or offline human organization, they augment it with on-chain guarantees and some automatic processes for human organization (i.e. voting is still a human process, but vote security may be automatic). The broader perspective on DeSci is that it aims to reform the organization of scientific activity itself, which is expected and preferred to translate into an increased ability for Science to fulfill its mandate and align with social values [37]. We fully

7

agree that this expectation is warranted, at least in principle, as the socio-epistemic dynamics we call Science clearly derive, at some level, to the concrete organization producing it. We will therefore attempt to characterize the distinction between CeSci and DeSci in terms of their respective socio-epistemic organizations.

Both CeSci and DeSci aim at knowledge discovery and curation, but they differ in their means to this end in terms of their incentives, structure, and norms, or more generally, in terms of their research ecology. In CeSci, the language, methods for discovery and strategic priorities are imposed from the top down by a core group of decision makers acting out the mandate of government policy positions. Core groups of decision makers also receive external input and constraints, for example in the case of peer review. In DeSci, methods and norms emerge from the bottom-up interactions of a web of loosely connected communities of practices with diverse coordination and prioritization mechanisms. Therefore, CeSci and DeSci diverge significantly in terms of their sensemaking; one imposes specific meaning, understanding and other cognitive constraints to scientific activity so as to keep it legible to the "view from above", the other cultivates the "view from the ground" by allowing those same cognitive objects to unfold spontaneously throughout the course of scientific activity.

Though there is fundamental overlap, and perhaps more similarity than not between CeSci and DeSci at this incipient stage, those terms stand for a fundamental divergence in the social and epistemic dynamics shaping the organization of research. Here we explore some of those dynamics by considering key differences between DeSci and CeSci.

CeSci: centralizing institutions against situated sensemaking

CeSci is characterized by the presence of centralizing agents (such as government institutions, philanthropic foundations, private businesses, or universities) having the power and willingness to establish and enforce rules, goals, ontologies, and other informal constraints within the wider scientific ecosystem. Its purpose can be understood as pragmatic rather than purely epistemic, in that CeSci institutions are funded and motivated, internally and externally, to build epistemic value (i.e. knowledge) that can assist other institutional actors and stakeholders understand and manipulate their reality in terms of their objectives, for better or worse. An example of this is the 19th Century development of statistics, as it was explicitly intended to help the developing modern States to perceive and control their population through the collection and analysis of demographic data, and effectively shaped what it meant to "think like a State" [29]. Due to extrinsic mandates, CeSci orients the sensemaking activity of its participant researchers toward specific technical and pragmatic purposes.

Although the goal of scientific research has largely extended beyond war and administration, the modern logic of centralization and technoscience is largely continued in contemporary research. The managerial role over the epistemic commons was granted to late modern academic institutions, and complemented with the role of supporting technological development through research on the property of chemicals, mechanical systems, and the like. Techno-scientific institutions have a core mandate of assisting outside institutions in their pragmatic endeavors. Their overarching goals are to accumulate epistemic value in that direction. Influence over research institutions by external organizations can be direct or indirect. In the first case, it may mean cadres of administrators exercising direct control over what research objectives should be, what funding agencies can be or should be approached, and recruiting based on compliance with external mandates. Indirect influence can be exerted, for example, by external organizations offering funding based on use of specific language or alignment with pragmatic research outcomes.

Accordingly, the production of knowledge artifacts, such as journal articles, is currently enacted according to a hierarchical structure of funding agencies, scientists, institutions, and publishers. There are in-groups and out-groups, as well as unspoken rules to obtain funding for and publish results of scientific literature. Moreover, there are agendas driving all agencies that fund scientific research, which are not necessarily transparent. Many basic research questions are framed in a way that adheres to these agendas. For example, researchers trying to characterize a mechanism by which a brain protein folds will need to emphasize its supposed role in curing Alzheimer's disease, when in fact any actionable outcome may be unclear or distant from the research [38]. This situation provides a strong incentive for researchers to leverage the information advantage they have over administrators to grab their attention by making incredibly grandiose, possibly dishonest claims on the meaning of their research. Administrators simply lack time or incentive to check on such claims, and active researchers lack the power to modify the incentive structure and would risk their job if they refused to follow them.

One aspect of the difficulty in communicating between administrators and active researchers is related to a discrepancy between the relevant time scales and perspectives for their respective activities. Bureaucratic funding agencies (private or governmental) have specific, long-standing research directives that cannot be modified easily, both because of long-term commitments on organization structure or mandate and because of the limited time officials have to understand, evaluate and integrate feedback. For example, an academic administrator will tend to lack domain expertise in the details of protein folding, and to lack motivation to discuss how category theory or information geometry could lead to major scientific breakthroughs. Administrators, in virtue of their function, are motivated to discuss how the research they manage contributes to the

competitiveness and prestige of the institution they answer to, and why it is therefore worthy of more funding. They may encourage their employees working on protein folding to claim in bad faith that their research will solve cancer, to publish a lot of papers into highly competitive (although not necessarily domain-relevant) journals, or to collaborate with companies for patents. If, for example, experts on protein folding want to express the benefits of future conceptual advancements, the importance may not be conveyed among fellow researchers and administrators due to a lack of common language, and therefore an inability to make decisions in a framework of mutual understanding.

Because of this fundamental asymmetry in power, information, and incentives, the centralizing hierarchical structure of science interferes with the rapid adaptation and evolution of scientific ideas. The need for researchers to make their ideas and objectives legible from the perspective of administrators conflicts with their ability to develop a context-relevant ontology and to focus on actually understanding the target system, or aiming to understand phenomena beyond the scope of expected research narratives. Centralizing institutions need to build largely artificial semantic and incentive silos, both as a way to compete against rival institutions and to measure and control the activity of researchers. Centralizing institutions organize science around outcomes which are legible and manageable, therefore hampering the kind of situated and transdisciplinary sensemaking which enables and grows our understanding of the world, and our ability to ground this understanding within a common language.

DeSci: governing epistemic commons

In contrast to the centralizing institutional view presented above, research can be organized as a decentralized federation of communities of practice, each striving to address specific issues in a contextually-relevant way. Such communities, by definition, lack a direct dependence on any centralizing institutions, or at the very least such institutions lack the ability to step in and micromanage their activity. Of course, the mere fact that those communities do not need to adapt their language or objectives to the expectations of administrative cadres does not make their activity scientific in nature. However, decentralization offers specific opportunities to scientific communities by allowing them to focus on developing pragmatic solutions to concrete problems and/or purely advancing the epistemic understanding of the natural world, free of centralizing institutional bias. As decentralized communities of practice lack the coercive power which defines centralizing institutions, their continued existence relies entirely on the voluntary participation of members, and therefore necessitates an institutional structure facilitating and incentivizing participation. Consequently, and also due to the lack of semantic silos (see below), decentralized scientific communities will have a strong incentive to cooperate with each other, to make relevant issues discoverable to wider sections of the population,

10

to allow participation purely out of voluntary interest, and to develop unifying languages easing their integration.

The opportunity for knowledge appropriation (via legal enclosure through intellectual property institutions as well as artificial semantic siloing) relates the development of scientific knowledge to a broader, well-studied category of collective action problems: the governance of common goods [39]. Common goods (commons) are defined as resources that are rivalrous or non-rivalrous yet non-excludable (one cannot be forbidden from accessing them). The potential for misuse or pollution of the commons consequently necessitates governance at the scale of communities of stakeholders, and regulation needs to follow certain conditions that were initially formalized in economic sociology [11,12] and recently generalized in evolutionary theory [40–42].

Roughly, the implementation of governance systems for commons entails:
1. the unambiguous definition of a problem to solve;
2. the construction of a community of stakeholders willing to help solve this problem;
3. the institution of a decision system allowing participation on an equal footing by all members of the community;
4. agreement on a set of rules (which are interactive and mutable with validity checks over time), which fairly reward active cooperation within the community;
5. the iteration of (1)-(4) at higher scales where the community as a whole is confronted with a situation requiring governance.

Although the definition of knowledge as a commons (a public good) entails a discussion of how exactly knowledge can be appropriated, the relation of decentralized science to common goods management is rather straightforward. Communities of practice organize around specific problems or questions, define rules for participation, and coordinate with each other to define and solve meta-problems - such as the development of standard tools or a common language or unifying framework which facilitate progress within each community.

Thus far, the distinction we have made between Centralized and Decentralized Science is largely organizational. For example, early European academia emerged from the clerical legitimation of scholarly "guilds," professional associations similar to contemporary trade unions or cartels, which effectively restricted who could practice a certain craft or access a specific kind of knowledge within a certain area (generally a city). Thus, patterns of sensemaking were immune from any outside oversight, and were imposed by the community on each and every one of its individual members. Centralizing institutions robbed community members of important aspects of their intellectual autonomy, and clearly enacted autonomous agency (see Thinking like a State [29] for a

discussion of institutional autonomy). In contrast, decentralization is facilitated through processes that facilitate broad access to the epistemic commons, as described above in the example of the printing press. When we highlight the opposition between CeSci and DeSci, we are prefiguring a discussion on how specific kinds of organization can facilitate differences in the integration of information about the world over multiple dynamical scales, and what it means for the scientific community.

In the contemporary age, digital technologies provide clear opportunities for proponents of DeSci. Indeed, they facilitate the widespread diffusion of information with a relatively low entry bar, and therefore the voluntary cooperation of multiple communities of practice engaged in situated sensemaking. This encourages the decentralized adoption of standardized methodologies and languages, transparency over code and data management, and willingness to produce information of clear value to outsiders. Indeed, because semantic silos cannot be enforced in order to increase one's influence, the digital dynamic tends to encourage collaboration in the interest of the broader system because it is the only way by which an actor can fulfill their objectives (see DeSci Token Engineering below). We see the success of such a decentralized yet unifying project in Wikipedia, which is beyond any doubt the most important example of a largely standardized, accessible and transparent knowledge system, and is based entirely on voluntary and decentralized cooperation. At the moment, there is however no appropriate, analogous system to use in the organization of scientific activity, as an alternative (or complement) to the currently dominant role of centralizing institutions.

Implementing Decentralized Science in the Web3 era

Contemporary digital tools have the potential to establish a distributed, transparent, scientific community with open participation, where every participant has root level access to all knowledge and the rules for its production. Online communities have seized the opportunity to develop and utilize such digital tools for funding, training, planning, coordinating, executing, disseminating, and archiving scientific assets and activities, therefore establishing the blueprint of a Web3 DeSci. There is no defining characteristic of the whole set of Web3 DeSci tools. Nonetheless, some features characterize a larger subset of these tools and are worth highlighting.

First, Web3 tools allow for more interactions (including economic transactions) between economically- and geographically-diverse groups. Second, Web3 tools can increase the transparency of products, activities, and interactions that are themselves digital, or that can be represented digitally. Third, Web3 tools allow for the introduction of novel markets that, in some pockets, might increase decentralization but in others, might increase centralization. Instead of using general terms to elaborate on each of these features, we hope to demonstrate their reality and significance as we overview more specific features of and developments toward a Web3 DeSci implementation.

Web3 for DeSci

Before heading into some details about the implementation of DeSci (which may just be one iteration or genre of the infinite game of Science), the reader should be equipped with familiarity around some basic Web3 terms. Note: the following terms are just current technological means to building a desired framework, and are presented here as introductory, not as comprehensive reviews or final answers.

- Web3
- Blockchain
- Smart contract
- Token (fungible and nonfungible)

Web3

First, "Web3" is an imprecise term referring to an ethos that values digital self-sovereignty and, more important for our purposes, to recent developments in peer-to-peer protocols that service this ethos. Part of the Web3 ethos is to value open source software and the personal control of one's data [43]. These values have influenced

and continue to influence the design and adoption of these protocols. Indeed, the term "Web 3.0" was introduced by Gavin Wood [44], a key figure in the development of the Ethereum and Polkadot protocols, who proposed Web3 as a "post-Snowden" Web, a decentralized Web similar to the Web1 of the early internet and in contrast to Web2 where a small number of centralized agents control people's data, access to websites, and the web applications that many use every day [45]. Web3 protocols are diverse, but the ones at the heart of most functional systems are so-called "blockchain" or distributed ledger technology (DLT) protocols that support complex smart contracts.

Blockchain

Recent developments in peer-to-peer (i.e. decentralized, distributed) protocols, particularly blockchain protocols such as ones that facilitate Distributed Autonomous Organizations (DAOs [20,21,46]), are the biggest enablers of a Web3 DeSci implementation. Without going into detail, a blockchain protocol enables a network of self-interested, competing computer nodes to agree on the state and immutable history of a shared ledger. Anyone with an internet connection and enough funds can append the ledger or verify its integrity, though it is still possible to implement access control with smart contracts. Blockchains are a natural fit for building currencies, which historically emerged as an extension to bookkeeping by institutional actors [47], and currency building in fact constitutes their most widely known application today. However, the principles of data integrity and access apply just as well to running any record or application that must be tamperproof. Blockchains can be used to establish ownership of digital assets (e.g. datasets or intellectual property [48]), mediate between pseudonymous actors (e.g. during peer review), establish decentralized identities used for reputation (e.g. when applying for grants), prove an actor has done something (e.g. stored a dataset or run a compute service), and more. The key here for DeSci is that blockchains enable actors to keep track of scientific activity in permissionless systems, and to build systems of exchange and accounting around any kind of digital information – which includes much of the products and tools of modern science.

Smart Contract

Blockchains can establish digital records via programs called smart contracts. A smart contract (in the context of blockchains) is a computer program whose code is publicly available but cannot be changed by anyone, is executed on computers across a network, and can be invoked (i.e., started) by anyone [49,50]. Smart contracts run on blockchains. When the blockchain's record is updated by a smart contract, the update is verified by the distributed nodes on the blockchain network. Despite the presence of "contract" in the term, smart contract does not refer to a legal contract, though it is still

14

possible to implement certain legal agreements in smart contracts, such as the triggering of payments in the case of predetermined on-chain activity. Not all blockchain protocols support the complex smart contracts that enable the features listed in the previous paragraph.

Token (fungible and non-fungible)

One kind of smart contract is called a token. A token is a digital asset. It can be bought, sold, or transferred on the blockchain. The implementation of a token in code is little more than a record of which addresses (i.e., users, loosely speaking) own how many units of the token. When units of a token are exchanged, the record is updated, not unlike what happens in a bank's database when two of its customers use the bank to transfer money to each other. Tokens come in various forms, broadly separable into fungible and non-fungible (though see [51] for a more complete discussion). A token of the latter form is often referred to as an NFT (Non-Fungible Token). The salient difference between these types of token is in the name; there are many interchangeable units of any given fungible token, and there is only one unit of any given non-fungible token. Because it is easy to use tokens to record ownership, they have been used by DAOs to designate voting rights (similar to shares of stock in a corporation) or provide proof of participation [52,53]. Tokens have many other use cases and have even enabled the emergence of a discipline known as token engineering [36], which studies the use of tokens to incentivize certain behaviors within a market.

DeSci Organization: DAOs, Blockchain, and Markets

The blockchain technology affords certain advantages over alternative systems in terms of transparency, decentralization, and tamperless regulation. However, its development was quite explicitly motivated by the implementation of market mechanisms at the exclusion of other forms of social organization, and therefore focused heavily on the implementation of monetary transactions. The issue with that position is straightforward. The very existence of multiplicative investments in a monetary economy (as entailed, e.g., by the legal status of companies in capitalist countries) entails the emergence of an extremely polarized power law distribution in wealth, and can (even in the absence of strategic interaction or information asymmetry) lead to the control of a large proportion of the economy by a few individuals [26]. In reality, strategic interaction and information asymmetry are ubiquitous, and the power entailed by the wealth of a few aristocrats gives them the capability to build a population of dependents and influence the rules of the game to their advantage. A complete disinvestment of public powers or otherwise democratic institutions from economic redistribution and regulation would therefore expectedly lead to

15

the development of a feudal-like economy, where a few wealthy individuals use their economic leverage to yield near absolute political power over the many.

The fundamental vulnerability of marketization-oriented Web3 technologies is deeply aggravated by preexisting asymmetries in wealth and power. Any token that can be traded freely (while not pegged to a stable asset) is vulnerable to price instability entailed by speculative investment, which would make it essentially worthless as a currency (i.e., a means of holding and exchanging value). On the other hand, any token that is not coupled in some way to the mainstream economy cannot be used to pay for basic goods such as food or shelter, and therefore lacks any economic value. Historical currencies were instituted by States through 3 main leverages: by creating demand through taxation, by backing their value to some external value holder (generally rare resources such as precious metals or recognition of debt by institutional actors), and by forcefully crushing the non-monetary economy and the communal institutions regulating it [47]. None of these leverages are accessible to decentralizing collectives, at least not without significant State / capital backing and an open betrayal of their explicit mandate [1]. Although the problem of monetary value is specific to currencies, the very demarch of marketization entails the ability to create economic value - or, in other words, to find a buyer for what you're trying to sell. For example, if no one is willing to pay you (in fiat currency, crypto, or some other value) the equivalent of a living wage for cleaning data, you are not going to make a living cleaning data.

Despite the aforementioned shortcomings, marketization can still play a substantial role in scientific decentralization. Markets offer an opportunity for large scale economic integration between peers, and therefore for the basis for an "exodus" from the grasp of centralizing institutions [55]. For example, tokens could provide an explicit and unfalsifiable account of traditionally less-valorised work, such as data management. However, blockchain-powered marketization cannot provide a sufficient solution for a successful decentralization of scientific activity. The blockchain only ensures decentralization in the limited sense that it makes breaking the rules that were written onto the blockchain extremely difficult, and therefore helps enforce the libertarian notion of liberty, defined negatively by the absence of explicit coercion [56]. Building DeSci, on the other hand, entails a broader shift toward positive liberty to do science in a decentralized way, i.e. the empowerment of individuals and collectives to organize scientific activity around decentralizing values and institutions. This process entails that we pay attention to the actual dynamics of decentralizing systems, and ensure that individuals and collectives can effectively enact their formal freedom of collaborative sensemaking.

---

[1] see [54] for a general discussion of value generation and regulation in cryptocurrencies.

16

DAOs, whether enforced by blockchain technology, or more classical means, provide the direct structure for their own activity. DAOs are organizations defined by an explicit mandate and specific rules for participation and decision, and are usually open to anyone willing to accept them. In the extent that these rules are reasonably fair, and that individual participants get a say in the decision system and overall organization, DAOs constitute a natural fit for the management of commons.

Of course, none of this excludes the value problem discussed above: science mandates that scientists be paid a living wage while they focus on their work, or otherwise to have enough time and resources to be able to conduct scientific work. In consequence, a radical shift to DeSci would likely necessitate broader social reform, such as the institution of a right to work in the research domain of one's choice, or perhaps a solution such as Universal Basic Income. Some cryptocurrencies, like the Ğ1 currency [57], try to build such change form the ground up by distributing a predefined rate of monetary creation equally between members, and explicitly grounding the networks to social webs of trust and exchange. But implementing such changes, especially at a global scale and within a reasonable time frame, would be entirely beyond the ability of contemporary DeSci communities. In the remainder of the document, we will therefore focus on steps that DeSci can practically take right now, many of which are related to tokenization and marketization mechanisms.

DeSci Vision and Open Questions

In our society today, a relatively small number of individuals have full-time careers in scientific research.hose with family commitments, a significant role in industry, or an aversion to academic social systems are largely absent from the process of creating and curating accessible knowledge artifacts and truth-seeking using the scientific method.

Can we support and develop communities that train researchers from many diverse backgrounds and provide opportunities for scientific collaboration, funding, and publication outside of the narrow constraints of CeSci career trajectories? How can we motivate and reward productive scientific engagement in this system, where there are less rigid barriers for who is allowed to conduct research? Can we benefit from neurodiversity and large-scale harmonization of different perspectives to reach epistemic consensus states beyond the reach of the narrow confines of what constitutes a "traditional scientific mind," or collection of such minds, in our society today? Decentralized communities of practice focused on research are already self-organizing, and will begin to answer these questions over the coming decade.

DeSci Token Engineering

Blockchains and tokens can be used to create incentive mechanisms to promote desired behavior and punish undesired behavior, and these incentive mechanisms can be used in DeSci to encourage valuable scientific activity. The practice of developing incentive mechanisms that utilize tokens is referred to as "token engineering" [20]. There are many ways to use tokens to reward behavior. In Filecoin, for example, a storage node is rewarded tokens for storing data only if it submits successful cryptographic proofs demonstrating it has stored a unique copy of the data over a predetermined period of time; the node's collateral is reduced if it fails such proofs. Not only can DeSci utilize systems that use token engineering, but DeSci can create token-engineered ecosystems. Token-engineered systems might be created for research-specific use cases, for example [58]. Token engineering can also be used to reward researchers and establish a value flow for a scientific ecosystem, a topic which we turn to now.

In the case of scientific journals, recent adoption of blockchain technology in the field of data economy can serve as an inspiration for designing a decentralized knowledge ecosystem. Blockchain offers the tools to observe, design, and control complex system behavior through the means of token engineering [59]. Were we to model entire communities as a collection of interacting agents which transact, perform work, and make economic decisions (e.g., investment, staking), this simulation could provide insight into the effectiveness of the overarching incentive structure, the resilience of the system to various risks, and the capacity of the organization for interoperability. We leave this as an opportunity for future study.

Token system engineering might ask questions like:
- How are stakeholders rewarded for meaningful participation?
- What is the value flow? Where does value start/enter? Where does value go/exit?
- What is the value distribution, internal and external to the system of interest?
- How does the system shape patterns of attention?

One example of token engineering used for value allocation is based on the Ocean Data Marketplace [60], in which a community-funded treasury disburses funds in the form of fungible tokens to research teams based on criteria defined by the community. At this point, the research team has full autonomy and is responsible for efficiently distributing the funds across different researchers. Researchers can send tokens to the decentralized knowledge market to gain access to datasets, papers, or any other scientific work published by other researchers, who will receive these tokens as compensation for their contribution to the DeSci ecosystem. Every transaction made in this decentralized knowledge market is facilitated through smart contracts, which collect transaction fees that are fed back to the community treasury.

18

The discussion on science token engineering leads to a natural question: what are the limits of decentralized science? Alternative value flows designed by the practice of token engineering assume new laws are enforced by digital contracts deployed on a transparent blockchain. As such, their ability to enforce particular actions outside the digital world become very limited. In its current form, DeSci is constrained by the affordances available to digital entities; hence the scope of scientific outputs is also constrained to research that has digital elements.

DeSci Research

Currently, most active DeSci projects primarily focus on building the infrastructure necessary to allow the dynamics described above, while only a fraction use their resources on research [2]. Examples of research-focused DAOs include VitaDAO and Molecule which, at the time of writing, have funded active projects in longevity research. Other active groups include the Governance Research Institute, Governauts, and Metagov DAO, researching governance, smart contracts, and other aspects of coordination in the Web3 space, all of which are closely related to the development of the DeSci infrastructure. As such, there is no clear distinction between infrastructure research and development. Furthermore, many DeSci DAOs have a clear vision of the type of research they want to conduct, but are currently in the process of getting funding or curating an initial community of active contributors.

There are many open questions relating to creating and sustaining impactful communities of practice in a Web3 DeSci setting. Technological changes to the process of science in DeSci may be exciting and impactful, but cannot circumvent the fundamental challenge of carrying out high-quality research. Even with hypothetically-successful systems in place for proposing projects, knowledge management, financial support, and decentralized publishing, individual researchers must still be motivated, committed, and engaged to produce useful work. Creating and sustaining impactful DeSci communities is challenged by the scale of information available for review or meta-analysis projects, the size of the design space of possible experiments, the multiple individual ways of working, and the difficulty in producing quality research. Without long-term mentorship, training, and support for scientists at all career stages, it is difficult to imagine that merely increasing the extent or type of system decentralization would result in a successful scientific ecosystem.

---

[2] A non-exhaustive list of the DAOs and groups mentioned is available here, building on [61], and will be updated.

DeSci Publishing and Review

DeSci seeks to use Web3 tools to improve the transparency and incentives within academic publishing and to improve the accessibility of published products. Efforts in DeSci around publishing are, in most cases, responding to the occasionally-misaligned incentives within current academic publishing, such as the lack of incentive for reviewers to provide unbiased and quality reviews, the incentive for journals to publish revenue-generating papers (even at the expense of rejecting epistemically-profitable, but less monetarily-profitable, research), the disincentive to share negative results, etc. [62]. With respect to transparency and incentives, Web3 DeSci projects are experimenting with using tokens (i.e. cryptocurrencies), markets, reputation metrics, and forum-like platforms. Increasing accessibility involves using decentralized storage networks to store published products (e.g. papers, reviews, datasets, code). This section hints at the direction of a DeSci approach to publishing and peer review with a discussion of current DeSci projects dealing with academic publishing and how blockchain-based technologies could potentially fit into the picture.

Since 2017, a number of authors have proposed blockchain-based systems for peer review [36,63–68]. Various works such as [36,67], [68] and [66] propose the use of blockchains to make the peer review process transparent and incentivize good behavior through reputation points recorded on the blockchain, though they differ in details, with [36,67] and [66] focusing on the entire pipeline even including publication. Below, we discuss several of these systems.

Avital [64] focuses on the motivation for a blockchain-based peer review market, where authors hire reviewers, and all participants interact through smart contracts on blockchains. Ants-Review builds on this idea of a blockchain-based peer review market [63] and outlines an implementation that uses Ethereum [69] and the InterPlanetary File System (IPFS) [34,35]. In Ants-Review [63], an issuer (i.e. author) creates an AntReview for a paper. The AntReview specifies the paper's issuers, file address (on IPFS), requirements, deadline, and approver. The approver is responsible for approving and paying reviewers (using the funds of the AntReview) for their reviews. A reviewer submits a review by uploading it to IPFS and sending the review's IPFS address to the AntReview for the approver's assessment. Anyone can add funds to an AntReview. These funds can be distributed only by the approver and can be sent only to reviewers. If the funds aren't spent before the review's deadline, those who contributed funds are refunded. Note that, in AntsReview, all parties are pseudonymous, associated only with their Ethereum addresses, and that the use of IPFS and Ethereum means everything is open access.

With Reddit-like features for curating, discussing, and incentivizing research, ResearchHub [66] is an even further departure from the current peer review and

20

publishing system. Its users can upload, comment on, summarize, upvote, and downvote papers and posts. Users are rewarded RSC, an Ethereum-based token, for these activities. Papers and posts are aggregated within "Hubs" which function like journals or subreddits. Allowing ideas to circulate in the open, ResearchHub shuns the traditional peer review model of waiting until a paper is mostly or entirely finished before having a small number of reviewers comment on it. ResearchHub relies on the community and the incentives created by the RSC token to create a system in which the better a post or paper is, the more it is seen and its author(s) rewarded. ResearchHub also plans to create a feature allowing researchers to raise funds from others in RSC. The platform differs from other projects in the DeSci space in that it does not currently rely much on peer-to-peer, Web3 protocols, such as Ethereum and IPFS.

At this point, to better grasp the opportunities that Web3 protocols have for publishing and reviewing in science, we might benefit from a brief look at the protocols themselves. Affordances for increasing transparency and accessibility are primarily in decentralized storage networks, such as IPFS and Arweave [70]. IPFS, the InterPlanetary File System, addresses files with content addressing based on cryptographic hashes. In practice, this allows two things. First, a file can be stored on any IPFS node in the world and can still be accessed by anyone who has the file's address. Second, an address can never correspond to more than one file. These properties mean that any file uploaded to IPFS is available to anyone with an internet connection (this increases transparency and accessibility), and that a file cannot be changed without its address also changing (this increases transparency). Arweave increases transparency and accessibility by providing permanent and open storage. It does this with a blockchain-inspired data structure called a blockweave and a proof for the network's consensus mechanism called Proof of Access in which, to receive a reward, a miner must prove it is storing some randomly determined previous block of data. There are other storage protocols used to store papers and related data, but these examples sufficiently illustrate the common benefits of these systems.

DeSci Funding and Finances

A more traditional method of funding a DeSci research project is through grant applications. In this case, the process starts with a research investigator submitting a proposal for a particular research endeavor, which is reviewed by the individuals or funding agencies who allocate the awards based on their own criteria determining merit. In the case of DeSci and DAO-based organizations, funding can be provided to the DAO within their internally-defined structures and further distributed within the DAO in order to fulfill the initiative of the grant. Alternatively, stakeholders or investors may provide funding due to their initially genuine research interest or other incentives, and ultimately become

detached from the original research motive, thus allowing the DAO to distribute resources to fulfill emergent agendas.

One alternative approach for funding DeSci is the creation and maintenance of markets for projects (bounties) or the outcomes of projects such as data or intellectual property. When an organization/individual is interested in conducting research based on interests, a systematically-managed decentralized exchange for task performance can play a pivotal role. The marketplace could have research proposals by organizations/individuals and researchers as well. A base level requirement for any market-supported project would be to produce some monetary or other form of value investors agree on. Key to the synergetic researcher-investor dynamic would be finding an alignment across multiple aspects of the project such as focus, constraints, cost, timelines. The machine learning platform Kaggle implements methods similar to this, where only researchers that perform best according to the project-defined metrics are rewarded.

Once a project has begun, it can benefit from mutual researcher-investor direction and redirection. Regulation of such direction is imperative; DeSci could build on top of existing blockchain infrastructure with a few modifications. For instance, a structure known as a DAICO (DAO+ICO [71]) proposes to lock all the proceeds of an ICO (Initial Coin Offering) into a decentralized autonomous organization (DAO) smart contract and put the governance over that DAO in the hands of the investors. This locking needs to be consensual from both ends and needs to be recalibrated as projects change. A roadmap with the specific, measurable, achievable, relevant and time-bound targets can be agreed upon; the investor gets the required output and the researchers get a well-defined source of funding. Other than direct funding, people interested in the project can involve themselves passively, forming a niche supporting or keeping a watchful eye on the project. Various options for this type of implicit contribution can help build a community, with social outlets such as newsletters and citation value networks. These communities can contribute collectively in terms of contributions to reduce the reliance on initial major investors and even foster the project in the case that investors pull out, hence decentralizing the investor influence. One particular project in this regard, Science Fund [72] (Figure 1), allows donors to contribute to science funding pools opened for a specific topic with an amount of their choosing. In return donors receive a NFT receipt (a "Science Funding Token", or SFT).

Figure 1. Dynamic research funding model of Science Fund [72].

These mechanics allow for the following affordances:
- **Curated Funding Pools**: Each donation is added to a chosen funding pool. Researchers can apply through an open protocol with their track record from early published groundwork. As new results emerge from the funded work, the NFT's metadata is updated.
- **Impact Tracing**: Donors have a much more significant participation and visibility into the projects they helped fund. Their SFTs are a personal key to unlock a continuously-evolving experience data trail about their donation's impact.
- **Future Participation:** The NFT receipts become the key to future interactions with the scientists they supported, with evolving forms of participatory governance in the field(s) they contributed to.
- **Funding allocation:** When researchers apply to one of the topic-specific funding pools, they are not coerced to reveal all their invaluable insight upfront. They simply need to share their most influential scientific work (any already-published scientific artifact they created, such as a publication, preprint, or even just a cool dataset) to a given pool. Selected scientists' works then form the basis of a growing evidence trail that is matched to the donor's minted NFT. The NFT receipts become the key to future interactions among scientists and donors, and open a universe of endless new possibilities of mutual exchange and future interactions.

DeSci Data and Knowledge Markets

Web3 tools enable markets for data storage and compute services that could be used by DeSci advocates to develop a "knowledge market" or decentralized data commons [73]. Agents in this potential market exchange data and other services such as data analysis and model training. Note that, because of their digital nature, such markets welcome opportunities to integrate human and machine agents, allowing for a significant increase in automation. While the base layer here is a data marketplace, it is plausible that future developments could emerge layered on top of this, due to the incentives and informational integrity provided by the market. This section covers the base layer of a data marketplace and how it encourages certain beneficial developments. In particular, data marketplaces have the potential to facilitate data archival, interoperable data standards, data discoverability, and increased automation, and they provide an easier way for industry to build on and fund work in academia.

The base layer of a DeSci knowledge market is first a data market. Multiple blockchain-based data marketplaces have emerged [74–76] since the Ethereum mainnet launch in 2015 [69]. Blockchains provide a number of benefits over traditional data markets, principally the ability to create datatokens, a construct pioneered by Ocean [74]. A data token is, just like any other blockchain token, a digital asset. However, unlike regular tokens, datatokens can be used to access data services. Consuming a data service involves either, 1) being sent the associated dataset, or 2) having the data provider run a compute job on the dataset. This datatoken construct is significant because it allows data services to be purchased on decentralized (blockchain-based) markets. Anyone can publish a dataset as a type of token, and anyone with sufficient funds can purchase one such token to consume the data service. The datatoken construct is even more valuable within the context of another mechanism developed by Ocean, namely, one that incentivizes the curation of quality data: Users can lock-up funds to signal support for a particular token (called staking) and earn a portion of the transaction fees associated with purchases of the token. Rational actors will seek out and stake on high quality datasets to receive a higher portion of the total fees, thereby curating quality data.

While this base layer is only a curated data market, it indirectly incentivizes the development of additional components relevant to a DeSci knowledge market, such as an increase in archival of scientific data. With a marketplace, data collectors have the opportunity to be financially rewarded and otherwise recognized for simply storing and selling their data. With the transparency of blockchain-based markets, it would be easy to keep track of how many times a researcher's or lab's datasets have been requested, allowing for the emergence of reputational metrics associated with data. It has been observed that currently, "there is little incentive to invest time in archiving or repackaging data sets" because "investing time in a project beyond its usefulness for publication is

counterproductive, given the high expectations for producing research publications" [77]. Filling this incentive gap, a data marketplace could provide another avenue of being rewarded and recognized for the valuable contribution of collecting quality data, and could therefore lead to increased data archival and availability.

Interoperable data standards are incentivized by a data market. The benefits of interoperable data standards can be seen in, for example, the neuroscience community, with its adoption of BIDS, "a standard for organizing and describing MRI datasets" [78]. This standard allows neuroscientists to more effectively collaborate within and across labs. It has reduced the amount of time they spend reorganizing datasets, leaving them more time for productive scientific work such as data analysis. In a data marketplace, where data suppliers are in competition and where the data users prefer certain formats over others, the data suppliers can be relied upon to provide data in those preferred formats, given sufficient competition. Even if a dataset is originally published in an esoteric data structure, someone can profit by purchasing it, organizing it into a known format, and reselling it. In a market where data purchasers may buy data from varied sources to improve their models, they are likely to seek out datasets that are compatible with each other. That is, they will seek out data structured in interoperable formats. The result is a data marketplace in which suppliers are actively organizing data to suit the needs of purchasers, converging on interoperable data formats.

Another potential emergent benefit of the data marketplace is increased data discoverability. Currently, finding scientific data is cumbersome [79]. Researchers often rely on social connections and fragmented data repositories. A data marketplace helps to address this issue through the competition between data suppliers who can gain from increasing the discoverability of their data. The more discoverable a dataset is, the more likely it is to be purchased. Data suppliers can increase discoverability by, for example, adding relevant tags and other metadata. Tags can help data search engines sort datasets. Metadata might even include a cryptographic signature from the researcher(s) who collected the data, attesting to the source and quality of the data. Additional organization and discoverability of datasets might come in the form of a knowledge graph, in which datasets can be related to each other, to researchers, to papers, to code that has been used to analyze it, etc. Note again that the data marketplace can be read and operated on by both humans and machines; much of the sorting of datasets in the marketplace might be done by automatons created by data suppliers. In the end, the suppliers within a data marketplace have a strong incentive to make their datasets easy to find, so they will likely adopt whatever methods best accomplish this.

Moreover, a data marketplace where data are curated, highly available, structured in known and interoperable formats, meaningfully-tagged, and discoverable, provides

25

good conditions for an increase in automated scientific discovery and, more specifically, better conditions for meeting the Nobel Turing Challenge introduced by Hiroaki Kitano [80] and further refined by The Alan Turing Institute [81]. This challenge involves creating "AI systems capable of making Nobel-quality scientific discoveries highly autonomously at a level comparable, and possibly superior, to the best human scientists by 2050" [81]. Smaller challenges within this grand challenge include ensuring data are shared, linked, and machine-readable. The data marketplace described above could help address these and possibly other challenges, such as simply financially incentivizing researchers to make their valuable datasets available.

Finally, a data marketplace affords profitable interactions between academic labs and organizations in industry. Labs can sell data to companies that need data and have AI expertise but that might not have the expertise or equipment to collect such data. This not only allows companies to more easily find and use scientific data, but by buying such data, companies would also be funding the scientists who collected them. This opens up a market in which data with both epistemic and pragmatic value can be easily priced and discovered. It also makes it easier for companies and the economy at large to communicate their practical data needs to researchers; this could enhance the current system in which a study's perceived pragmatic value is a function of how well its value is communicated through a grant proposal. A data marketplace could lead to a greater flow of information among academic and non-academic (e.g. citizen, government, non-profit, profit-oriented) organizations, with respect to the value of various kinds of data, allowing easier and potentially better funding of data with pragmatic value (though these benefits of funding may not extend to data which has only epistemic value). A well-designed data marketplace enables fair price discovery of digital knowledge assets and encourages productive applications of meticulously-collected, high-quality scientific datasets outside of basic research, boosting research applicability and downstream innovation.

Teamwork Modeling with Systems Engineering

To decrease the complexity and uncertainty that individuals face within large-scale DeSci projects, it is useful to introduce the systems engineering approach [82,83]. This process of sensemaking for individuals and teams can be thought of as the use of a map, enabled with modern features like personalization and dynamic rendering [84]. For better communication and team attention management, it is important to have consensus on the team objectives, project lifecycle management, practices in use, and metrics for tracking the performance and relevant behavior of all team members. Targets and goals of the projects are set by the needs of external stakeholders and individuals fulfilling specific project roles. Governance for decisions, assignments, and resources are based on unequivocal agreement according to the team members' affiliation and authority to

26

manage each other's labor. We consider such a group of individuals as an organization or organizational unit in the case of specific tasks to resolve.

Akin to the use of Systems Engineering in the building of complex infrastructure (e.g. bridges, planes), we propose that people involved in DeSci projects document the roles involved in these team dynamics, as well as their concerns and preferences. This allows DeSci activity to be structured according to an established lifecycle with needed practices [85,86]. Here, we are not concerned with the domain-specific practices of a project lifecycle, but with general practices that are relevant for DeSci development, such as social modeling, economics, governance, and development in a broad sense. During any DeSci project, we can associate a human or other entities as actors who are fulfilling the roles that need to be performed. These entities fulfilling project roles engage in practices that entail completing tasks or cases. All the relevant objects and open cases must be resolved during the project for it to be a success (for instance, knowledge artifact production or fundraising). These cases need to be documented with reference to the roles, practices, and actions required for successfully achieving the target goal. This case/issue tracker functionality is an attention management tool for the people involved in the project. This process allows for productive, auditable, and coherent work on projects.

Like CeSci, DeSci faces challenges and failure modes across multiple domains, some of which were briefly described above. At this incipient and formative stage, what may provide utility for DeSci would be a useful unifying framework to allow for system conceptualization, design, and implementation. In the following sections, we describe a possible starting point for the integration of DeSci and CeSci systems, through the combined application of Active Inference and Systems Engineering [83].

27

BOLTS of the CeSci to DeSci Integration

The difficulties of successful implementation of novel systems, even where they are simple or viable to be implemented in parallel rather than as replacements to extant systems, lend themselves to underestimation. Successful, ubiquitous systems often feel obvious in terms of their value proposition and their reasons for adoption. For example, the product "bar code" would appear to be ubiquitous and obvious in terms of its benefits, but despite its conception in the mid-20th Century, by 2004 only around 80 to 90% of the United States' top 500 companies had adopted them [87]. In terms of its roughly 50-year climb to this broad adoption, it took a quarter of a Century just to find adoption within a single market segment after a small collection of large institutions began courting potential standards [87,88]. The difficulties of implementation, adoption, and integration of novel systems relevant to our purposes can be summarized through a qualitative use of the Active Inference Conflict Model, which helps to frame conflict across business, operations, legal, technical, and social (BOLTS) surfaces, especially where it relates to information differentials [89]. Below, the bar code is used as an example that characterizes how the BOLTS surfaces apply to complex patterns of adoption and operation:

- **Business**. Integration with and replacement of extant receipt and payment systems, few of which would have been compatible with any bar code standard, represented a significant difficulty. Additionally there are the financial and non-financial costs of barcode adoption that an organization would have to consider.
- **Operations**. Integration with and replacement of extant inventory systems, few of which would have been compatible with any bar code standard, represented a significant difficulty.
- **Legal**. For some products, technical standards and instantiations of product IDs would have to map to regulatory compliance systems, contracts, and be available for use in advanced records systems.
- **Technical**. While the idea of the bar code was conceived in the mid-20th Century, it took decades to confirm widely adopted standards. Adoption at any given point in the climb to broad adoption meant the risk of implementing a pattern that might become obsolete.
- **Social**. People didn't necessarily trust machines, and were concerned with the potential for pricing errors at the time-of-scan. Companies were also concerned with how bar-codes would affect the shelf-appeal and artistic design of products.

This "bar code" example helps to illustrate how the implementation, adoption, and integration of new systems is roughly proportionate with the abstract and real integration conflicts with extant systems. While modern, shared technological infrastructure has

allowed for more expedient implementation of new systems, conflicts across BOLTS processes still exist and may be greatly exacerbated when systems are intended to be used by institutions and organizations, which have their own formally-defined BOLTS protocols, in addition to individuals where BOLTS protocols are likely to be encoded in norms, habit, and narrative [89].

While DeSci may be implemented in parallel to CeSci, the overwhelming majority of research funding is channeled through large organizations, foundations, and especially government agencies. Each of these and their respective channels are beholden to a complex, complicated, and interconnected web of BOLTS standards, norms, and controls, which are a struggle to navigate. While it is tempting to make do with what funds might be available, a bridge between CeSci and DeSci is not just about access to funds at scale but also about access to tangible resources, influence on institutional funding and agendas, ability to collaborate, and avoiding scarcity in funding. In addition, solving conflicts relevant to a DeSci-CeSci bridge may lead to greater accountability, reliability, and impact of DeSci information products, and therefore improve the ability for potential investors to trust in the new marketplace of ideas. Areas of tension and interoperation between DeSci and CeSci are summarized through the use of BOLTS below:

Business

While government agencies and non-government organizations are not operating under the return on investment model found in for-profit businesses, they still consider the dollar cost of impact, opportunity costs, and likelihood of impact. Agencies have to consider how to structure "portfolios" of funded projects and initiatives and do so based on their current mission, the missions of their stakeholders, their sources of funding, and the track-records of the organizations and people they fund. In order to make DeSci compatible with or even superior to CeSci projects and organizations in this context, DeSci systems would have to enable a high level of situational awareness for funders in terms of both progress and potential for impact, while offering advantages in terms of cost and comparability. This may not be immediately achieved – like most new services and systems, DeSci may start off with inefficient processes, and will become cheaper and more efficient in producing informational products of higher quality over time.

Operations

Accountability and use of funds is taken very seriously in many of the CeSci funding channels, especially where funding flows through multiple hierarchies. Even where funding flows through less steps, such as in NGOs, demonstrating success in order to raise funds constitutes its own stream of work requiring a substantial allocation of attention and funding. Documentation and records management in operations serves to fulfill the "evidential function", pro forma evidence of action and use [90]. These records should be

29

systematic and need to meet relevant standards for reliability, integrity, compliance, and comprehensiveness [91]. It is here that DeSci offers challenges and opportunities relative to traditional systems [20]. Where CeSci organizations rely heavily on manual processes in order to meet documentation requirements, DeSci is predisposed toward generating records by merit of reliance on formalized, computational protocols. However, these advantages will only be realized if there is a directed effort to ensure automatically generated records are adapted to map to the myriad CeSci requirements.

Legal

Many of the business-, technical-, and operations-related documentation requirements extend to legal use. It is in this domain that records serve their "warranty function", clarifying intents, deliverables, requirements, and other agreements [90]. Project documentation and both personal and organizational identity become key obstacles to DeSci's compatibility with CeSci, especially where teams are emergent or may change in structure over time. Extensive work is required in this domain in order to overcome legal compatibility challenges, given that no amount of advantages of efficiency, cost, quality, or impact can allow DeSci to circumvent problems related to CeSci funders being able to exercise their options related to legal recourse.

Technical

Technical aspects of CeSci funding blend with operations and legal aspects, as would be the case with most bureaucratic systems. The synthesis between these two sections (operations and legal) is where the potential bridge between DeSci and CeSci is currently most precarious, given that an inability to meet legal requirements will eliminate this bridge entirely. However, it is also where the bridge has the potential to outperform traditionally-funded organizations, as meeting operating and legal requirements computationally would, as previously suggested, offer a number of advantages.

Social

The proverbial "pipeline" of government and institutional funding to research and development is monolithic and interniccine, and at times quite obscure [92]. Those involved in the channeling of funds have standard operating procedures encoded both in terms of business, operations, legal, and technical surfaces, as well as in social norms and in narratives. The preference for DeSci to offer more accountability or other efficiencies for research is compatible with the preferences and requirements of those who currently occupy positions which channel funds in CeSci. Case studies, post mortems, and high levels of visibility in both successes and failures, as well as test-runs in incubators or regulatory sandboxes, could assist in communicating the benefits of DeSci to potential stakeholders. Further, finding ways to allow DeSci systems to be flexible enough to map directly to CeSci standards for accountability could also facilitate DeSci integration. For

30

example, even though DeSci allows for and can encourage purely horizontal teams, one can start with teams which can: assign a higher authority and offer them power of direct oversight, allow them to measure and assess compliance and performance, and be responsible for reporting and conveying that information to funders (even where this could be automated) [93] to help to introduce CeSci standards to DeSci without requiring excessive accommodation.

BOLTS and Reliability for Science Systems

It is important to provide reliability in all of these BOLTS dimensions, as new technologies and systems can easily fail if there are repeated bad experiences or unexpected incompatibilities. Even if the technical and legal aspects of DeSci funding out-perform CeSci funding, failure of DeSci to provide a smooth integration from the social and business perspectives could delay or prevent wide adoption. For example, traditional funding agencies often come in at middle-to-late stages of research projects, which is why so much preliminary work is required for a NSF/NIH grant. It is often said grants are not "for ideas", but rather they are for seriously-developed avenues of specific tractable research. In DeSci, some more speculative vehicles might exist (e.g. a "Learner's Fund"), but we cannot expect these new vehicles to back-propagate the risk-tolerance of DeSci into legally- and organizationally-bound CeSci funding agencies.

A priority of both CeSci and DeSci is to ensure reliability of outcomes and high levels of accountability, in complex scientific commons. If a team has been funded, disintegration or failure to produce work doesn't just mean a negative outcome for the team, it means a negative outcome for the funder, which has the potential for network impacts on trust within the entire market. This means that teams within the DeSci space would not only need to vastly outperform CeSci in terms of cost, time-to-impact, and efficiency, but would also need to effectively be perceived as high-reliability organizations in order for the market to survive. Given that DeSci teams are likely to be emergent, they may not be able to rely on intimate trust between one another, and instead will have to rely heavily on shared Ontology, Narrative, Formal documents, and Tools [83]. The team working environment should be designed to create the best possible conditions for the cognitive security of the individuals on teams, such that they stay on task, maintain consistency in terms of goals and intents, and understand their impact on the team and the future of the market as a whole. This type of system design also means high cognitive security on the part of the funder, which means high visibility and high quality narrative information management tools at their disposal [38], the ability to do retroactive analysis of track-records of individuals and collections of individuals, the ability to rapidly compare projects, and the ability to negotiate cost as a basis to reduce risk based on current information. Additionally, transparency ensures that when things do go wrong, it reduces

31

the negative sentiment of the funder on the network, and might be mitigated by insurance options (e.g. community stake in research) or in-house recourse options (e.g. sanctions and penalties on negligent and bad-faith actors)

While we detailed the BOLTS conflicts related to funding processes as DeSci practices are integrated into existing CeSci practices, this framework can be applied to several other aspects of this integration such as the publication of knowledge artifacts, scientific review of knowledge artifacts, communicating research results at conferences, analyzing data, training and mentoring new researchers, and developing scientific software. CeSci and other fields already have successful aspects which should be preserved as DeSci is integrated, such as specific mentorship, individual, and peer development programs. As a framework that furthers the holistic design of transdisciplinary systems for science, we turn to Active Inference for additional insight into DeSci which will be outlined in the following section.

Active Inference, Systems Engineering, and Science

In this section we justify the use of Active Inference in systems modeling, introduce an entity description for science modeling consistent with Active Inference and Systems Engineering models, and revisit patterns in CeSci and DeSci.

What is Active Inference?

Active Inference is a framework for the integrated modeling of entity perception, cognition, action, and impact in the niche. Active Inference draws upon the formalism of the free energy principle, which provides a single statistical imperative underpinning the assembly of organic and inorganic matter: the minimization of informational free energy [94–96]. This principle applies to systems at every scale. The process of free energy minimization has been positioned as the driver of autopoietic processes in living systems ranging from cells to cities [97,98]. Its core meaning is that it allows a teleonomic, nearly-agentive account of self-organization. For example, a set of interacting systems may be likely to spend a lot of time in an aggregated form, reflected by persistence in a low free energy region of their coupled dynamics corresponding to an aggregated state [95]. The minimisation of free energy at the system's scale entails the minimisation of variational free energy for each of its constituting particles, which can be interpreted as the optimisation of a Bayesian model entailed by the particle's structure (more specifically, the difference between model Complexity and Accuracy) [99]. Because of this correspondence, we can claim that entities aggregate in virtue of reducing their uncertainty about future sensorimotor states [100,101].

The related framework of Active Inference leverages this formalism to model one or multiple interacting entities as involved in continual perception and action angled towards reducing the expected uncertainty about their sensorimotor flow [102,103]. They can do so either by altering internal states so as to infer the causes of sensory states, or by altering external states so as to cause expected states. Both processes participate in a common control architecture based on predictive processing, which can be interpreted as a hierarchy of "algorithms" matching input to output over nested time scales [104]. In this hierarchical model, previously-acquired beliefs about controlling behavior in different contexts are internally encoded, and result in adaptive arbitration between epistemic and pragmatic value choices (i.e. exploration vs exploitation, [105]). Thus, the framework of Active Inference allows us to model the observations, beliefs, and predictions (i.e. expected states) of agents as statistical dependencies, or Bayesian probability distributions.

Much of the focus in the Active Inference literature has focused on organisms, especially in the context of human behavioral neuroscience [106]. However, Active Inference has also been applied to various other informational and cyberphysical settings such as human communication [107], remote teams [83], human conflict [89], trust interactions with robotics [108], and epistemic communities [25]. Here we apply Active Inference and Systems Engineering to the case of modeling centralized and decentralized science.

Active Inference and Science

Evidence of collective intelligence, understood as the capability of human societies to solve problems, surrounds us. In most cases, it is instantiated in cultural artifacts or cognitive gadgets [109] that follow from a process of cumulative cultural evolution dating from the branching of the *Homo* genus [110,111]. We are surrounded by living fossils of this cumulative process, from smartphones to hospitals, from net fishing to general agreements about the recognition of private property. Our capability to understand and navigate the world does not simply follow from information in our minds, but it is (to some extent) entailed by the very structure of the ecological niche and the social constraints that shape our behavior. This means that entities are, to some extent, autonomous or agentic regarding their own cognition [29]. Entities engaged in Science as a cognitive endeavor are no exception.

Like other instances of collective intelligence involving stigmergy [112,113], Science is a cumulative endeavor that relies on previously-acquired beliefs and established communities of practice. Scientists and other cognitive entities perceive stigmergic cues from the niche (including epistemic resources, as well as other social entities), update their cognitive models (learning), and act to modify the niche in the form of research outputs (Figure 2). As scientists follow new leads or choose to ignore them, communities form coalesce and disperse, and scientific theories are created, developed, and forgotten. Frequently, scientific endeavors are undertaken with multiple, potentially competing, goals in mind. For example, scientists may undertake their work with an aim to achieve some or all of the following outcomes: earn a paycheck, improve their reputation, improve the reputation of their company or institution, conduct empirical research without bias, solve a particular problem, or discover something novel, simple, or beautiful.
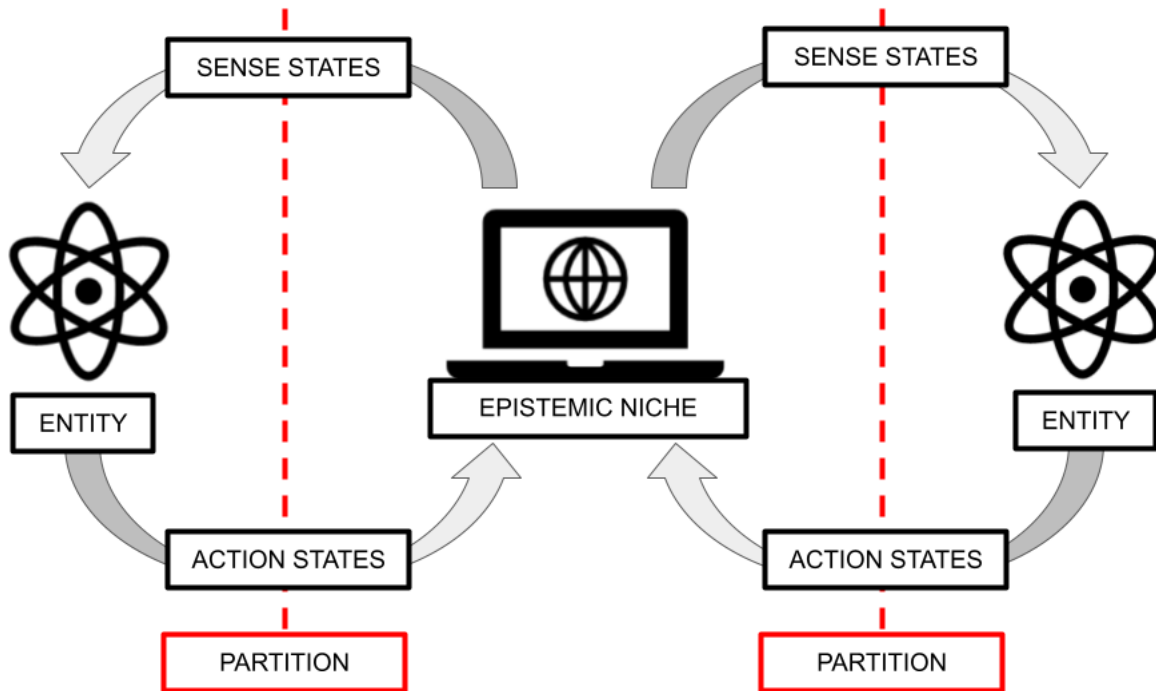
34

Figure 2. Two active entities interacting via their shared epistemic niche. For more detail on the entity perception-cognition-action loop in Active Inference, and the Markov Blanket formalism comprising the partition shown here, see [95].

Active Inference is relevant for model-based understanding of scientific systems because it offers a composable multiscale ontology and integrative framework for entity and system behavior [98,114]. Thus the usage of Active Inference can help us understand the entanglement of multiple processes occurring at nested scales of behavior. In Active Inference, actions across scales are modeled as a systematic attempt by cognitive entities (i.e. scientists and institutions) to reduce uncertainty about aligning their future perceptions with their preferences, so as to ensure their persistence of their defining organization. For example, the development of administrative systems and agriculture was described within Active Inference formalism as a product of City-States' attempt to understand and control their reality [29], analogously to the process of morphogenesis as a process of cells manifesting a body plan which reduces their uncertainty about their cellular niche [100].

Importantly, the Active Inference account is not limited to conservative (homeostatic) behavior, but also to more complex phenomena such as memory, anticipation, abductive logic, and autopoiesis. For example, Active Inference has been used to model the systematic alteration of one's niche [104,115,116], as well as the creation of new tools and their incorporation into one's cognitive identity [100]. Therefore, it very much provides grip on innovation, by contextualizing subjective systems of meaning in relationship to the concrete (biological and sociocultural) organizations which underlie cognition [117].

35

In the context of action selection related to research (e.g. experimental design, decisions about scientific communication), Active Inference provides a division between two kinds of motivation in research: epistemic value (informational gain related to e.g. new research findings or knowledge) and pragmatic value (reward associated with increases in wealth, longevity, or status). Any innovation in institutional or cognitive processes will affect the broader system both epistemically, by defining the social norms which entail scientific activity, its cognitive underpinnings, and the kind of cues and actions scientific activity involves; as well as pragmatically, by reinforcing the specific behaviors which help the relevant (institutional or human) agent to reduce (subjective) uncertainty about future sensorimotor states. When viewed through this lens, the scientific process balances action selection based upon the both epistemic and pragmatic value, and this process plays out across multiple scales of nested organization. For example, at different stages of their day/career, the individual human researcher prioritizes more exploratory/learning activities, while at other times prioritizes more exploitative/productive activities. At the scale of an individual, institution, or even scientific field, the epistemic activities are reflected by basic research, while the pragmatic activities can be considered as translational or application-oriented research, or engineering.

As metascience and sociology of science typically focus on the motivations and systems of incentive on the scale of individuals and institutions, we consider the multiscale nature of Active Inference to provide a relevant angle for future research. This work reflects early steps towards leveraging this framework into a more detailed understanding of how scientific entities interact, and how contemporary changes associated with the challenges and opportunities of DeSci could improve on the status quo.

Active Entity Ontology for Science (AEOS)

Here we describe the attributes and structural basis of a descriptive entity-oriented ontology for Science, consistent with the principles of Active Inference and Systems Engineering. This project has the working title and acronym Active Entity Ontology for Science (AEOS). Previously, among many other use cases, ontologies have been used for schematic mapping of cryptographic systems and knowledge ecosystems [118]. Additionally, approaches for complex systems modeling based upon entities and their affordances have precedence in economics [119] and other areas. We build on this work to create a useful and versionable resource which stands to be developed and extended by researchers and practitioners in the coming years. This work is meant to be a starting point for the construction of active entity-oriented ontologies of science. See coda.io/@active-inference-institute/active-entity-ontology-for-science-aeos for an interactive site that provides an updated version of AEOS.

Table 1 provides example entity interaction motifs for CeSci and DeSci ecosystem models. These motifs are presented as initial drafts on the specification and visualization of subsystems/modules that are important for larger networks of interacting entities.

Area of Concern	Example CeSci Motif	Example DeSci Motif
Funding	Team Publishes Knowledge Artifact (Grant), then Requests Funding From Funding Agency, which Scientifically Assess the Grant. Funding Agency Provides Opportunity for Researcher and Provides Constraints for Researcher and Sends Fungible Asset to the Team.	Team Publishes Knowledge Artifact (Grant), then Requests Funding From a DAO which Scientifically Assesses the Grant The DAO Engages in governance regarding the Grant and Sends Fungible Asset to the Team.
Communication	Academic University Communicates Asynchronously To Lab via Computer System	Human Communicates Synchronously To Human via the Computer System
Cryptographic relationship	Research investigator uses Cryptographic Identity on Computer System to Provides Funding To Team and Provides Updates to Funder like Research Agency (Government)	Research investigator Team transfers Intellectual Property as Non-Fungible Asset to a DAO Cryptographic Identity that is a Smart Contract .
Scientific Review	A Research investigator Sends Fungible Asset to Journal/Publisher to Publish Knowledge Artifact like a Paper. The Journal/Publisher Requests Scientific Review from other Research investigators.	After Publishing a preprint Paper, the Research investigator receives a Fungible Asset as a function of the use of the Knowledge Artifact
Research & Analysis	Research investigators in Lab Team Scientifically Assess Primary Knowledge Artifact then Publishes Knowledge Artifact	Research investigators in DAO Community of Practice Scientifically Assess Primary Knowledge Artifact then Publishes Knowledge Artifact
Publishing	After Scientific Review, a Research investigator Publishes Knowledge Artifact at Journal/Publisher.	A Research investigator Publishes Knowledge Artifact like a Paper on a DAO-produced knowledge market.

Table 1. Some example motifs of Research and Funding in CeSci (Centralized Science) and DeSci (Decentralized Science). See the AEOS site for an interactive version of this Table, where all entities and affordances can be unpacked and explored.

AEOS Entity partitioning

Any AEOS entity consists of 3 kinds of states that are partitioned off from the entity's environment: Internal states (states inside the entity that are hidden from the outside, for example those involved in cognition), Sense states (incoming statistical dependencies), and Active states (outgoing statistical dependencies). Together these three states (Internal, Active, and Sense) constitute the "particular" or autonomous states of the entity [99]. One of the advantages of Active Inference as an integrative framework is this: used instrumentally, it can model various phenomena and ensembles across scales from the quantum [120] to the social and planetary [121,122], including heterogeneous types of entities which enact the scientific ecosystems of today and tomorrow.

Using the Markov blanket partitioning of variables on a Bayesian graph: Entities consist of specified Internal states, reflecting the (observer's model of) the Entity's generative model of perception, cognition, and action. External states are those that are inferred or acted on by Entities. Internal states are partitioned from external states via a Markov blanket, which is further divided into nodes with incoming statistical dependencies (Sense states) and outgoing statistical dependencies (Action states) [95,98,99]. Active states are reached via an internal process of policy selection, which incorporates information about what affordances (capacities for action) exist and the preferences that the entities would like to reduce uncertainty about the realization of (for example a preference for comfortable temperatures or financial success).

Entity models of this type have advantages in that they can be compared across entities and niche surroundings [115], and enable variational Bayesian inference and message passing algorithms [123]. In this case of online DeSci, a type of digital message passage algorithm is occurring on a bipartite graph composed of active and informational entities in the AEOS. This may be represented as a factor graph, which presents with an interesting avenue for future research and application [123–125].

AEOS Entity Classes and Types

The fundamental object or kind of thing in the AEOS is the Active Entity, which in online settings is almost always interacting with Informational Entities [83]. Drawing upon the Active Inference framing of a system, this entity is a system of interest that is partitioned from its niche surroundings based upon its capacity for persistence, thus making it a "thing" [126,127].

In AEOS, there are two Classes of Entity (Figure 3): Active epistemic entities (which select action policies and enact affordances), and Informational entities (which contain information and are acted upon). There are multiple Entity Types within each Class. In the

38

digital/online setting, Active Entities are always intermediated by Informational entities (e.g. People can talk online through modification of a shared epistemic niche reflected by a computer system). Informational entities are always intermediated by Active entities (e.g. two databases can only communicate with each other through a computer system). "Adaptive" and "mere" Active Inference entity notion is via [128].

Active Entity Ontology for Science (AEOS) Classes

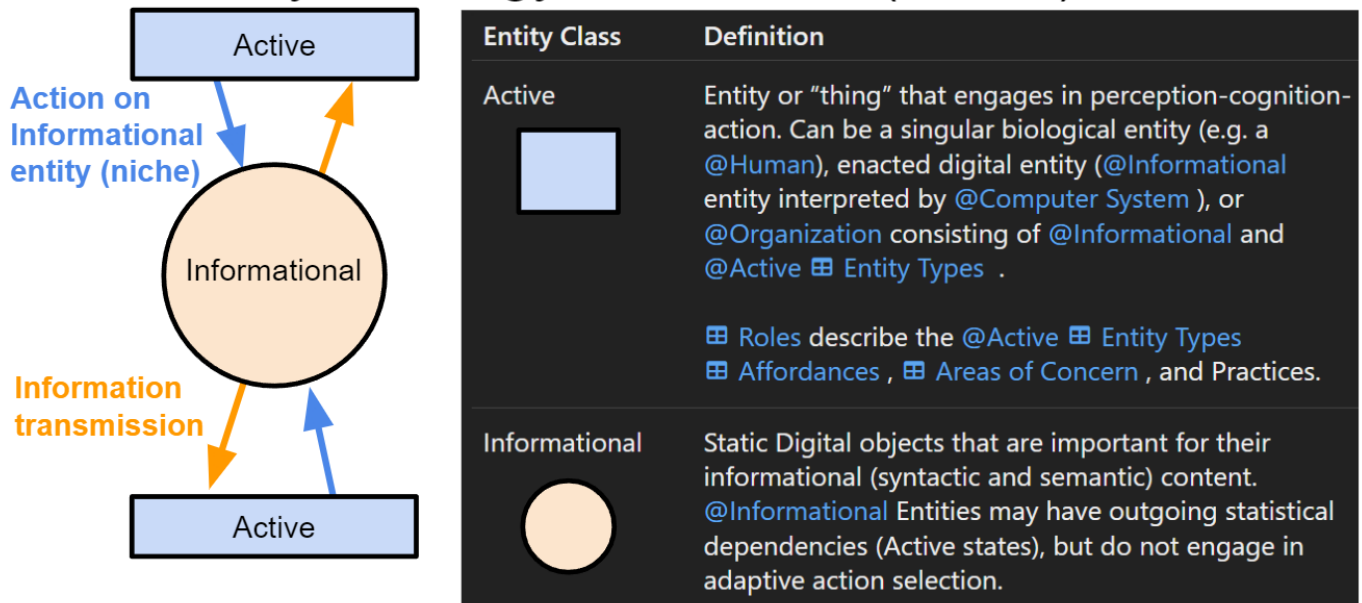| Entity Class	Definition
Active	Entity or "thing" that engages in perception-cognition-action. Can be a singular biological entity (e.g. a @Human), enacted digital entity (@Informational entity interpreted by @Computer System), or @Organization consisting of @Informational and @Active ⊞ Entity Types . ⊞ Roles describe the @Active ⊞ Entity Types ⊞ Affordances , ⊞ Areas of Concern , and Practices.
Informational	Static Digital objects that are important for their informational (syntactic and semantic) content. @Informational Entities may have outgoing statistical dependencies (Active states), but do not engage in adaptive action selection.

Figure 3. Entity Classes in Active Entity Ontology for Science (AEOS). Active Entities (blue boxes) are any kind of organization with agency at any scale, for example humans, teams, organizations, institutions, and DAOs. Informational Entities (orange circle) is any type of computational data or epistemic resource, for example a dataset, grant, article, or fungible/non-fungible asset.

Both classes of entities (Active and Informational) inherit the same type of Active Inference partitioning scheme described above. See the AEOS site for an updated list of the diverse types of entities currently modeled. Some examples of Active entities in AEOS include Human, Computer System, Organization, Team, DAO, and Academic University, Community. Of special note are the Organizational entities, composed of groups of humans, computer systems, or other organizations (e.g. a university made up of departments, made up of labs, made up of people), along with their informational resources. As with any other Active Entity, Organizational entities engage in outgoing actions from a set of discrete or continuous options (affordances), and also can be the recipient/target of action from other entities.

Entities of the Informational class are those which are being modeled for their informational content, and are solely acted upon. Some types of Informational entities may also display agentic qualities when used by Organizational and Computer System entities, for example a software program. The Informational entities are used as epistemic resources in the niche of other Agent-class entities. Some examples of Informational entities in AEOS include Knowledge Artifact, Fungible Asset, Non-Fungible Asset, Dataset, Metadata, Paper, Intellectual Property, Code, Grant, Blockchain, and Project Catechism

AEOS Policy Selection, Areas of Concern, and Roles

Active entities engage in policy selection based upon their internal generative model (Figure 4). The behavioral policies that active entities enact are selected from a set of possible action sequences for that entity: their affordances. In the AEOS, the affordances are determined by the type of entity, its roles, and its context – for example a human entity with administrator role may have the affordance to modify a file in a certain situation. The specific affordances that an entity engages in, in a specific setting, reflect the role (formal/assigned or informal) that the entity has (Figure 5). This role-based approach towards organization draws on Systems Engineering [83], and allows the flexibility to both describe current systems as well as design new systems. For example if a human is providing grant funding to a team, that human entity has the role "Funding provider". With respect to the affordance of providing funding to a research team, this role could also have been fulfilled by a funding agency or a DAO.
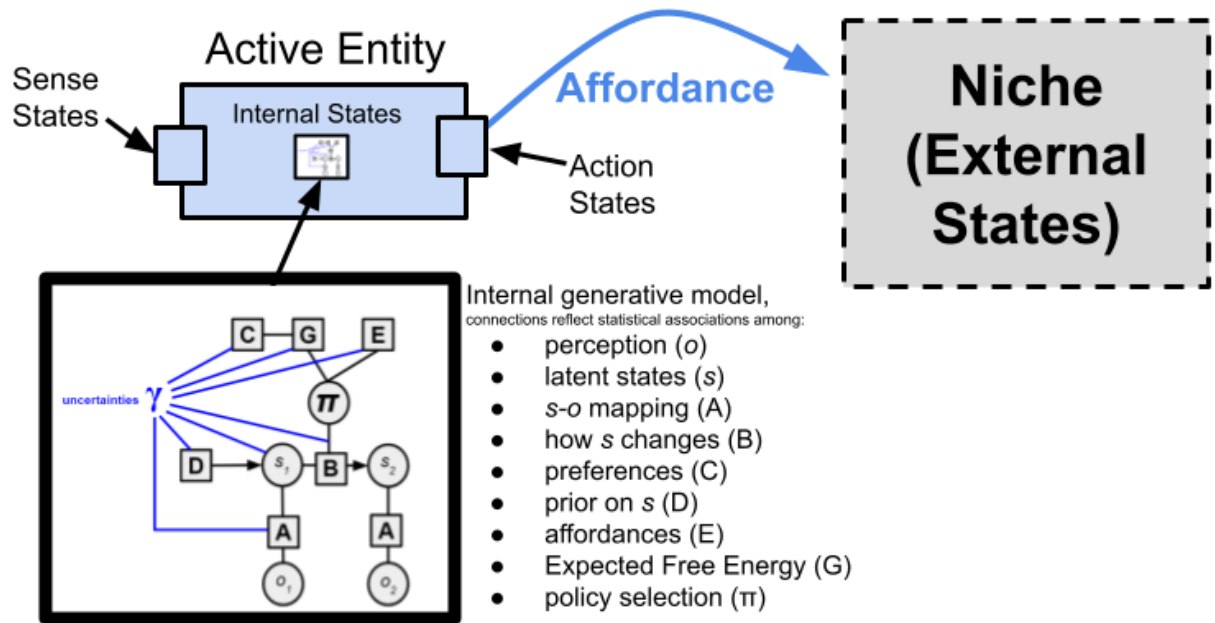
Figure 4. Active entity action selection. Using the Active Inference partitioning [106,127], the Active Entity at top left can be modeled as having incoming Sense States and outgoing Action States which interact in an affordance-based fashion with the niche. The internal generative model of the entity, here modeled as a partially observable Markov Decision Process, describes the process of action selection given an incoming stream of observations.
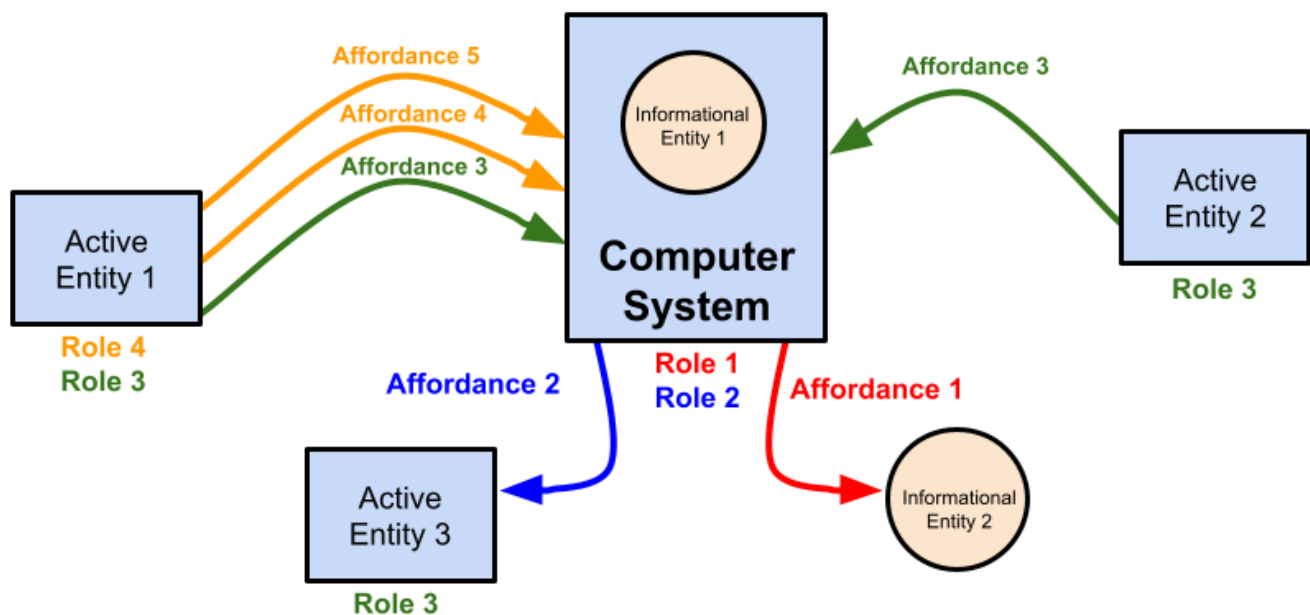

Figure 5. Entity Roles. Each Entity has Affordances (capacities for action) and can have Roles in the organizational context (assigned performances to be done) In the online context, all Active entities are intermediated by computer systems, and implicitly informational entities (shared epistemic niche).

41

Below in Figure 6-8 are several initial graphical examples of applications of the AEOS. The natural language sentences at the bottom of each image, reflect CeSci and DeSci motifs to be found within each Area of Concern (Communication, Scientific Review, Funding). Words in blue reflect AEOS entities and active relationships (affordances) that link the entities. These examples show the capacity of AEOS to model various key motifs and patterns found in both CeSci and DeSci systems. One possible advantage of a system such as AEOS is that scientific motifs can be described in natural language, have an interactive visual representation, and also provide an avenue for formal modeling. We are currently exploring the integration of AEOS models with complex system modeling software such as cadCAD [129] or darcspice [130], in order to enable description, analysis, and simulations of epistemic ecosystems.
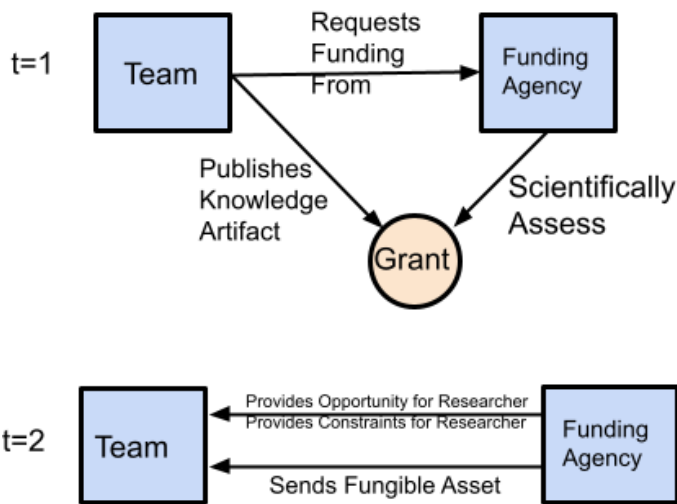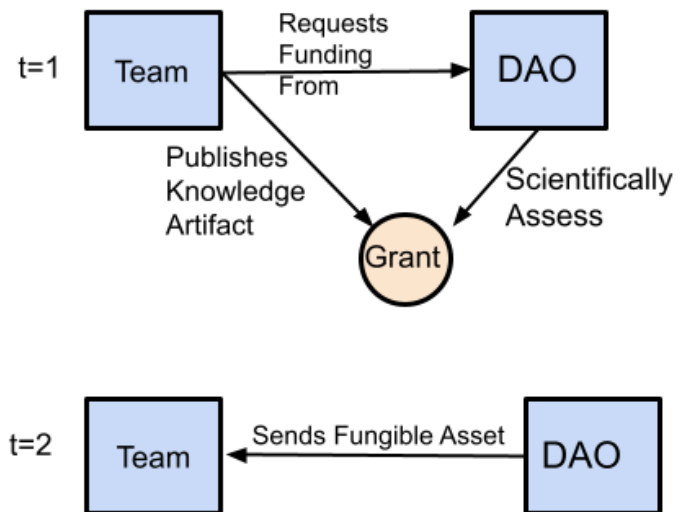
Team Publishes Knowledge Artifact (Grant) , then Requests Funding From Funding Agency , which Scientifically Assess the Grant.

Funding Agency Provides Opportunity for Researcher and Provides Constraints for Researcher and Sends Fungible Asset to the Team .

Team Publishes Knowledge Artifact (Grant) , then Requests Funding From a DAO which Scientifically Assess the Grant

The DAO Engages in governance regarding the Grant and Sends Fungible Asset to the Team.

Figure 6. Example motifs of Funding in CeSci (left) and DeSci (right). Blue boxes are Active Entities in AEOS, orange circles are Informational Entities. In this example, the initial timestep t=1 is noted on the top of each side, and the following timestep t=2 is below.
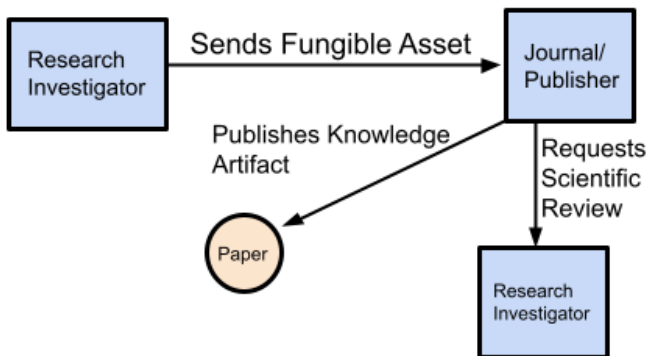
42

Communication

CeSci

Academic University Communicates Asynchronously To Lab via Computer System

DeSci

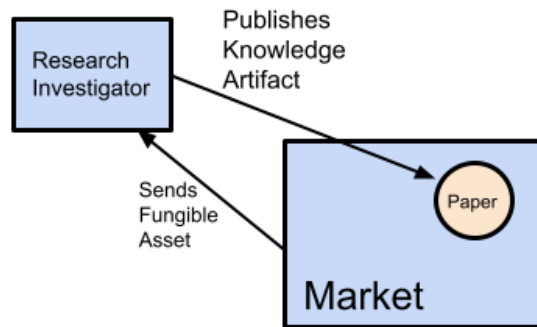

Human Communicates Synchronously To Human via the Computer System

Figure 7. Example motifs of Communication in CeSci (left) and DeSci (right).

Scientific Review

CeSci

A Research investigator Sends Fungible Asset to Journal/Publisher to Publishes Knowledge Artifact like a Paper. The Journal/Publisher Requests Scientific Review from other Research investigator.

DeSci

After Publishing a pre-print Paper listed on a knowledge Market. The Market Sends Fungible Asset to Research investigator as a function of the use of the Knowledge Artifact

Figure 8. Example motifs of Scientific Review in CeSci (left) and DeSci (right).

43

Discussion

This paper evaluated the opportunity to implement decentralized science in Web3 technologies from the perspective of Active Inference.

We first highlighted the distinction between two approaches to scientific research, respectively Centralized Science (CeSci), where research is driven from the top-down perspective of centralizing institutions, and Decentralized Science (DeSci), where it is driven from the bottom up through situated sensemaking by communities of practices. We argued the second stance can facilitate the governance of scientific knowledge (understood as an epistemic common good) by providing both the incentive and the opportunity to build a detailed understanding of various systems and problems in an integrated language.

We consequently reviewed potential avenues to implement DeSci in the era of Web3. We underlied a major limitation to marketization and tokenization, one of the main mechanisms of Web3 technologies, namely that creating real world value for the token renders it vulnerable to speculative investment and hoarding. We discussed the potential roles and benefits of modern affordances, including online teams, blockchain, tokenization, and smart contracts. Finally, we discussed several practical aspects of integrating DeSci concepts and practices into CeSci, such as business, operations, legal, technical, and social (BOLTS) aspects of this transition.

Finally, we discussed the role of the Active Inference framework in describing scientific ecosystems. As a formalism relating the basic activity of dynamical systems to epistemic inference, we argued it is a natural fit to describe science as a socio-institutional system. We outlined an Active Entity Ontology for Science (AEOS), intended to define key Entity Classes, Entity Types, Areas of Concern, Roles, Affordances, and Action-Perception States involved in conducting scientific work, whether CeSci or DeSci, with examples illustrating particular embodiments of this ontology. We hope that further development of research within the scope of AEOS would yield fruitful progress in our ability to design and evaluate possible forms of scientific organization.

Some specific next steps include:
- Implementation of AEOS using complex systems modeling frameworks such as cadCAD [129], which have already been applied to relevant cyberphysical systems in Web3.
- Exploration of Active Inference ontologies and relationships with category theory [120] and other formalisms.

- Creation of Graphical User Interface for DeSci (DeSciCAD), to enable the real-time use of AEOS in science ecosystem design and operation. For example, a user could select entity types, assign them roles, connect the entities with affordance-based edges, and engage in simulation and parameter optimization. The front end of such a design interface would ideally be accessible and graphical, on the backend there could be a composable Active Inference simulation.
- Exploring system designs in AEOS that scaffold and incentivize different kinds of research such as basic/theoretical, translational/applied, and quantitative/qualitative work. One specific application would be to use AEOS in behavioral modeling of incentive gaps, such as getting people to publish data even if they don't publish the paper, or finding ways to accept modular contributions, so that in the case of retraction or deciding not to publish, pieces can still be made available and citable.

We hope that this work has provided a useful step towards the interactive modeling of scientific knowledge as a process of multiscale active inference.

Author contributions using CRediT – Contributor Roles Taxonomy [131]:

- Conceptualization (AGC, AKS, SED, VBK, AC, DAF)
- Funding acquisition – No funding was acquired specifically for this project.
- Investigation (AGC, JS, SED, DAF)
- Project administration (DAF)
- Visualization (VBK, DAF)
- Writing – original draft (AGC, AKS, AC JS, SS, SED, VBK, RJC)
- Writing – review & editing (DAF, SAS, RJC, AC, SED, AGC, VBK, IM, SS, AKS, JS CT, ASV)

Works cited

1. Weaver W. Science and complexity. Am Sci. 1948;36: 536–544. Available: https://www.ncbi.nlm.nih.gov/pubmed/18882675

2. Gershenson C. The Implications of Interactions for Science and Philosophy. arXiv [nlin.AO]. 2013. pp. 781–790. doi:10.1007/s10699-012-9305-8

3. Wright C, Bechtel W. - Mechanisms and psychological explanation. In: Thagard P, editor. Philosophy of Psychology and Cognitive Science. Amsterdam: North-Holland; 2007. pp. 31–79. doi:10.1016/B978-044451540-7/50019-0

4. Chan ME, Arvey RD. Meta-Analysis and the Development of Knowledge. Perspect Psychol Sci. 2012;7: 79–92. doi:10.1177/1745691611429355

5. Lawrence MG, Williams S, Nanz P, Renn O. Characteristics, potentials, and challenges of transdisciplinary research. One Earth. 2022;5: 44–61. doi:10.1016/j.oneear.2021.12.010

6. Head ML, Holman L, Lanfear R, Kahn AT, Jennions MD. The extent and consequences of p-hacking in science. PLoS Biol. 2015;13: e1002106. doi:10.1371/journal.pbio.1002106

7. Rivera H. Fake Peer Review and Inappropriate Authorship Are Real Evils. J Korean Med Sci. 2019;34: e6. doi:10.3346/jkms.2019.34.e6

8. Scott JC. The Trouble with the View from Above. 2010 [cited 30 Jan 2022]. Available: https://www.cato-unbound.org/2010/09/08/james-c-scott/trouble-view-above/

9. Fanelli D. "Positive" results increase down the Hierarchy of the Sciences. PLoS One. 2010;5: e10068. doi:10.1371/journal.pone.0010068

10. Vazire S. Implications of the Credibility Revolution for Productivity, Creativity, and Progress. Perspect Psychol Sci. 2018;13: 411–417. doi:10.1177/1745691617751884

11. Munafò MR, Nosek BA, Bishop DVM, Button KS, Chambers CD, Percie du Sert N, et al. A manifesto for reproducible science. Nature Human Behaviour. 2017;1: 1–9. doi:10.1038/s41562-016-0021

12. Reproducibility and reliability of biomedical research. [cited 26 Feb 2022]. Available:

https://acmedsci.ac.uk/policy/policy-projects/reproducibility-and-reliability-of-biomedical-research

13. Hills TT. The Dark Side of Information Proliferation. Perspect Psychol Sci. 2018; 1745691618803647. doi:10.1177/1745691618803647

14. Barker Bausell R. The Problem with Science: The Reproducibility Crisis and What to do About It. Oxford University Press; 2021. Available: https://play.google.com/store/books/details?id=oHEWEAAAQBAJ

15. Bar-Yam Y. Multiscale variety in complex systems. Complexity. 2004;9: 37–45. doi:10.1002/cplx.20014

16. Bruni LE, Giorgi F. Towards a heterarchical approach to biology and cognition. Prog Biophys Mol Biol. 2015;119: 481–492. doi:10.1016/j.pbiomolbio.2015.07.005

17. Feinerman O, Korman A. Individual versus collective cognition in social insects. J Exp Biol. 2017;220: 73–82. doi:10.1242/jeb.143891

18. Rossi F, Bandyopadhyay S, Wolf M, Pavone M. Review of Multi-Agent Algorithms for Collective Behavior: a Structural Taxonomy. IFAC-PapersOnLine. 2018;51: 112–117. doi:10.1016/j.ifacol.2018.07.097

19. Mohanty D. Ethereum Architecture. In: Mohanty D, editor. Ethereum for Architects and Developers: With Case Studies and Code Samples in Solidity. Berkeley, CA: Apress; 2018. pp. 37–54. doi:10.1007/978-1-4842-4075-5_2

20. Nabben K. Decentralised Autonomous Organisations (DAOs) as Data Trusts: A General-purpose Data Governance Framework for Decentralised Data Ownership, Storage, and Utilisation. 2021. doi:10.2139/ssrn.4009205

21. Schneider N, De Filippi P, Frey S, Tan JZ, Zhang AX. Modular Politics: Toward a Governance Layer for Online Communities. Proc ACM Hum-Comput Interact. 2021;5: 1–26. doi:10.1145/3449090

22. Gencer AE, Basu S, Eyal I, van Renesse R, Sirer EG. Decentralization in Bitcoin and Ethereum Networks. arXiv [cs.CR]. 2018. Available: http://arxiv.org/abs/1801.03998

23. Fields C, Levin M. How Do Living Systems Create Meaning? Philosophies. 2020;5: 36. doi:10.3390/philosophies5040036

24. Veissière SPL, Constant A, Ramstead MJD, Friston KJ, Kirmayer LJ. Thinking through other minds: A variational approach to cognition and culture. Behav Brain Sci. 2019;43: 1–97. doi:10.1017/S0140525X19001213

25. Albarracin M, Demekas D, Ramstead M, Heins C. Epistemic Communities under Active Inference. 2022. Available: https://www.preprints.org/manuscript/202201.0124

26. Bruineberg J, Kiverstein J, Rietveld E. The anticipating brain is not a scientist: the free-energy principle from an ecological-enactive perspective. Synthese. 2018;195: 2417–2444. doi:10.1007/s11229-016-1239-1

27. Constant A, Clark A, Friston KJ. Representation Wars: Enacting an Armistice Through Active

47

Inference. Front Psychol. 2020;11: 598733. doi:10.3389/fpsyg.2020.598733

28. Hipólito I. Cognition Without Neural Representation: Dynamics of a Complex System. Front Psychol. 2021;12: 643276. doi:10.3389/fpsyg.2021.643276

29. Guénin--Carlut A. Thinking like a State - Embodied intelligence in the deep history of our collective minds. 2021 [cited 25 Jan 2022]. doi:10.31219/osf.io/dxnzt

30. Buterin V. The meaning of decentralization. Medium com. 2017;6.

31. Nabben K. Is a DAO a Panopticon? Algorithmic Governance as Creating and Mitigating Vulnerabilities in "Decentralised Autonomous Organisations." 2021. doi:10.2139/ssrn.3907693

32. Kiayias A, Lazos P. SoK: Blockchain Governance. arXiv [cs.CR]. 2022. Available: http://arxiv.org/abs/2201.07188

33. Graber J. Web3 is Self-Certifying. In: Medium [Internet]. 23 Dec 2021 [cited 26 Feb 2022]. Available: https://jaygraber.medium.com/web3-is-self-certifying-9dad77fd8d81

34. Benet J. IPFS - Content Addressed, Versioned, P2P File System. arXiv [cs.NI]. 2014. Available: http://arxiv.org/abs/1407.3561

35. IPFS powers the distributed web. [cited 15 Feb 2022]. Available: https://ipfs.io/

36. Tenorio-Fornés Á, Tirador EP, Sánchez-Ruiz AA, Hassan S. Decentralizing science: Towards an interoperable open peer review ecosystem using blockchain. Inf Process Manag. 2021;58: 102724. doi:10.1016/j.ipm.2021.102724

37. Hamburg S. A Guide to DeSci, the Latest Web3 Movement. In: Future [Internet]. 9 Feb 2022 [cited 25 Feb 2022]. Available: https://future.a16z.com/what-is-decentralized-science-aka-desci/

38. Cordes RJ, Applegate-Swanson S, Friedman DA, Knight VB, Mikhailova A. Narrative Information Management. Zenodo. COGSEC; 2021. doi:10.5281/zenodo.5573287

39. From Theory to Practice. The MIT Press; 2007. Available: http://www.jstor.org/stable/j.ctt5hhdf6

40. Wilson DS, Ostrom E, Cox ME. Generalizing the core design principles for the efficacy of groups. J Econ Behav Organ. 2013;90: S21–S32. doi:10.1016/j.jebo.2012.12.010

41. Paul W.B. Atkins, David Sloan Wilson, Steven C. Hayes. Prosocial. 2019. Available: https://www.newharbinger.com/9781684030248/prosocial/

42. Wilson DS, Philip MM, MacDonald IF, Atkins PWB, Kniffin KM. Core design principles for nurturing organization-level selection. Sci Rep. 2020;10: 13989. doi:10.1038/s41598-020-70632-8

43. Beck J. What is Web3? Here Are Some Ways To Explain It To A Friend. In: ConsenSys [Internet]. 12 Jan 2022 [cited 22 Feb 2022]. Available: https://consensys.net/blog/blockchain-explained/what-is-web3-here-are-some-ways-to-explain -it-to-a-friend/

44. Edelman G. The Father of Web3 Wants You to Trust Less. Wired. 29 Nov 2021. Available: https://www.wired.com/story/web3-gavin-wood-interview/. Accessed 26 Feb 2022.

45. ÐApps: What Web 3.0 Looks Like. [cited 22 Feb 2022]. Available: https://gavwood.com/dappsweb3.html

46. Kim HM, Laskowski M, Nan N. A First Step in the Co-Evolution of Blockchain and Ontologies: Towards Engineering an Ontology of Governance at the Blockchain Protocol Level. arXiv [cs.CY]. 2018. Available: http://arxiv.org/abs/1801.02027

47. Graeber D. Debt : The First 5000 Years. In: Melville House Books [Internet]. 2011 [cited 18 Feb 2022]. Available: https://www.mhpbooks.com/books/debt/

48. clepp. Molecule's biopharma IP-NFTs — A technical description. In: Molecule Blog [Internet]. 6 Aug 2021 [cited 15 Feb 2022]. Available: https://medium.com/molecule-blog/molecules-biopharma-ipnfts-a-technical-description-4dcfc6bf77f8

49. Bartoletti M, Pompianu L. An empirical analysis of smart contracts: platforms, applications, and design patterns. arXiv [cs.CR]. 2017. Available: http://arxiv.org/abs/1703.06322

50. Khan SN, Loukil F, Ghedira-Guegan C, Benkhelifa E, Bani-Hani A. Blockchain smart contracts: Applications, challenges, and future trends. Peer Peer Netw Appl. 2021; 1–25. doi:10.1007/s12083-021-01127-0

51. Shorish J, Stephenson M, Zargham M. A Practical Theory Of Fungibility. WU Vienna University of Economics and Business; 2021 May. Report No.: 8137. Available: https://ideas.repec.org/p/wiw/wus051/8137.html

52. Nandwani A, Gupta M, Thakur N. Proof-of-Participation: Implementation of Proof-of-Stake Through Proof-of-Work. International Conference on Innovative Computing and Communications. Springer Singapore; 2019. pp. 17–24. doi:10.1007/978-981-13-2324-9_3

53. poapxyz. POAP. [cited 1 Mar 2022]. Available: https://www.poap.xyz/

54. Peters G, Panayi E, Chapelle A. Trends in Cryptocurrencies and Blockchain Technologies: A Monetary Theory and Regulation Perspective. 2015. Available: https://papers.ssrn.com/abstract=3084011

55. Carson KA. Exodus. 2021. Available: https://theanarchistlibrary.org/library/kevin-a-carson-exodus?v=1624008858

56. Carter I. Positive and Negative Liberty. Zalta EN, editor. The Stanford Encyclopedia of Philosophy. Metaphysics Research Lab, Stanford University; 2022. Available: https://plato.stanford.edu/archives/spr2022/entries/liberty-positive-negative/

57. Duniter. In: Duniter [Internet]. [cited 1 Mar 2022]. Available: https://duniter.org/

58. Analysis D. GRIDCOIN WHITE PAPER. Available: https://gridcoin.us/assets/docs/whitepaper.pdf

59. McConaghy T. Towards a Practice of Token Engineering. In: Ocean Protocol [Internet]. 1 Mar 2018 [cited 19 Jan 2022]. Available:

49

https://blog.oceanprotocol.com/towards-a-practice-of-token-engineering-b02feeeff7ca

60. Tools for the Web3 Data Economy. In: Ocean Protocol [Internet]. [cited 22 Jan 2022].
     Available: https://oceanprotocol.com/

61. Smart contract research forum. In: Smart Contract Research Forum [Internet]. [cited 15 Feb
     2022]. Available: https://www.smartcontractresearch.org/

62. DeSci Foundation. Why we need to fundamentally rethink scientific publishing. In: Medium
     [Internet]. 1 Dec 2021 [cited 15 Feb 2022]. Available:
     https://desci.medium.com/why-we-need-to-fundamentally-rethink-scientific-publishing-43f2ae3
     9af76

63. Trovò B, Massari N. Ants-Review: A Privacy-Oriented Protocol for Incentivized Open Peer
     Reviews on Ethereum. Euro-Par 2020: Parallel Processing Workshops. Springer International
     Publishing; 2021. pp. 18–29. doi:10.1007/978-3-030-71593-9_2

64. Avital M. Peer Review: Toward a Blockchain-enabled Market-based Ecosystem.
     Communications of the Association for Information Systems. 2018;42: 28.
     doi:10.17705/1CAIS.04228

65. van Rossum J. The blockchain and its potential for science and academic publishing. Inf Serv
     Use. 2018;38: 95–98. doi:10.3233/isu-180003

66. The ResearchCoin Whitepaper. In: ResearchHub [Internet]. [cited 15 Feb 2022]. Available:
     https://www.researchhub.com/paper/819400/the-researchcoin-whitepaper

67. Tenorio-Fornés A, Jacynycz V, Llop-Vila D, Sánchez-Ruiz A, Hassan S. Towards a
     Decentralized Process for Scientific Publication and Peer Review using Blockchain and IPFS.
     2019. doi:10.24251/HICSS.2019.560

68. Jan Z, Third A, Ibanez L-D, Bachler M, Simperl E, Domingue J. ScienceMiles: Digital Currency
     for Researchers. Companion Proceedings of the The Web Conference 2018. Republic and
     Canton of Geneva, CHE: International World Wide Web Conferences Steering Committee;
     2018. pp. 1183–1186. doi:10.1145/3184558.3191556

69. Ethereum. Ethereum Whitepaper. In: ethereum.org [Internet]. [cited 14 Feb 2022]. Available:
     https://ethereum.org/en/whitepaper/

70. Williams S, Berman L, Raybould I, Uemlianin I. Arweave: A protocol for economically
     sustainable information permanence. [cited 15 Feb 2022]. Available:
     https://www.arweave.org/yellow-paper.pdf

71. Explanation of DAICOs. In: Ethereum Research [Internet]. 6 Jan 2018 [cited 30 Jan 2022].
     Available: https://ethresear.ch/t/explanation-of-daicos/465

72. Science Fund. [cited 28 Feb 2022]. Available: https://www.sciencefund.io/

73. El Damaty S. Rich in Data, Poor in Wisdom: Science Needs a Decentralized Data Commons.
     In: OpSci [Internet]. 4 Dec 2021 [cited 27 Feb 2022]. Available:
     https://pulse.opsci.io/rich-in-data-poor-in-wisdom-science-needs-a-decentralized-data-commo
     ns-98c7ffdb56a1

74. Ocean Protocol Foundation with BigchainDB GmbH. Ocean Protocol: Tools for the Web Data Economy. Available: https://oceanprotocol.com/tech-whitepaper.pdf

75. Datum. Datum-WhitePaper.pdf. Available: https://datum.org/assets/Datum-WhitePaper.pdf

76. Dataspace. datapace_whitepaper.pdf. Available: https://www.datapace.io/datapace_whitepaper.pdf

77. The Problem of Data: Data Management and Curation Practices Among University Researchers. In: CLIR [Internet]. 13 Dec 2016 [cited 18 Feb 2022]. Available: https://www.clir.org/pubs/reports/pub154/problem-of-data/

78. Gorgolewski KJ, Auer T, Calhoun VD, Craddock RC, Das S, Duff EP, et al. The brain imaging data structure, a format for organizing and describing outputs of neuroimaging experiments. Sci Data. 2016;3: 160044. doi:10.1038/sdata.2016.44

79. Gregory K, Groth P, Scharnhorst A, Wyatt S. Lost or found? Discovering data needed for research. Harvard Data Science Review. 2020. doi:10.1162/99608f92.e38165eb

80. Kitano H. Nobel Turing Challenge: creating the engine for scientific discovery. NPJ Syst Biol Appl. 2021;7: 29. doi:10.1038/s41540-021-00189-3

81. The Turing AI scientist grand challenge. In: The Alan Turing Institute [Internet]. [cited 16 Feb 2022]. Available: https://www.turing.ac.uk/research/research-projects/turing-ai-scientist-grand-challenge

82. SEBoK. [cited 24 Jan 2022]. Available: https://sebokwiki.org/wiki/Guide_to_the_Systems_Engineering_Body_of_Knowledge_(SEBoK)

83. Vyatkin A, Metelkin I, Mikhailova A, Cordes RJ, Friedman DA. Active Inference & Behavior Engineering for Teams. 2020. doi:10.5281/zenodo.4021163

84. Cordes RJ, Friedman DA, Maron M. Reimagining Maps. 2020. doi:10.5281/zenodo.4170026

85. ISO/IEC/IEEE 15288:2015. In: ISO [Internet]. 2020 [cited 23 Feb 2022]. Available: https://www.iso.org/standard/63711.html

86. Essence 1.2 - SEMAT. [cited 23 Feb 2022]. Available: http://semat.org/essence-1.2

87. Weightman G. The History of the Bar Code. In: smithsonianmag.com [Internet]. 23 Sep 2015 [cited 28 Jan 2022]. Available: https://www.smithsonianmag.com/innovation/history-bar-code-180956704/

88. Silverman L. Barcodes: A Brief History. In: corp.trackabout.com [Internet]. 19 Nov 2015 [cited 28 Jan 2022]. Available: https://corp.trackabout.com/blog/barcodes-brief-history

89. David S, Cordes RJ, Friedman DA. Active Inference in Modeling Conflict. 2021. doi:10.5281/zenodo.5759807

90. Dikopoulou A, Mihiotis A. The contribution of records management to good governance. The TQM Journal. 2012;24: 123–141. doi:10.1108/17542731211215071

91. ISO. ISO 15489 - Information and documentation - Records Management. 2001.

51

92. Cordes RJ. The Pipeline. The Pipeline: EdTech, Knowledge Management, and NATSEC. COGSEC; 2021.

93. Cox RJ, Wallace DA, Wallace D. Archives and the Public Good: Accountability and Records in Modern Society. Greenwood Publishing Group; 2002. Available: https://play.google.com/store/books/details?id=sH5BaIWT7o0C

94. Tumiel J. Spinning up in active inference and the free energy principle. In: jared tumiel [Internet]. 14 Oct 2020 [cited 21 Jan 2022]. Available: https://jaredtumiel.github.io/blog/2020/10/14/spinning-up-in-ai.html

95. Friston K, Heins C, Ueltzhöffer K, Da Costa L, Parr T. Stochastic Chaos and Markov Blankets. Entropy. 2021;23. doi:10.3390/e23091220

96. Parr T, Pezzulo G, Friston KJ. Active Inference: The Free Energy Principle in Mind, Brain, and Behavior. MIT Press; 2022. Available: https://scholar.google.ca/scholar?cluster=1067966391259043561&hl=en&as_sdt=0,5&sciodt=0,5

97. Sims M. How to count biological minds: symbiosis, the free energy principle, and reciprocal multiscale integration. Synthese. 2020. doi:10.1007/s11229-020-02876-w

98. Ramstead MJD, Kirchhoff MD, Constant A, Friston KJ. Multiscale integration: beyond internalism and externalism. Synthese. 2021;198: 41–70. doi:10.1007/s11229-019-02115-x

99. Friston K. A free energy principle for a particular physics. arXiv [q-bio.NC]. 2019. Available: http://arxiv.org/abs/1906.10184

100. Kuchling F, Friston K, Georgiev G, Levin M. Morphogenesis as Bayesian inference: A variational approach to pattern formation and control in complex biological systems. Phys Life Rev. 2019. doi:10.1016/j.plrev.2019.06.001

101. Fields C, Levin M. Scale-Free Biology: Integrating Evolutionary and Developmental Thinking. Bioessays. 2020;42: e1900228. doi:10.1002/bies.201900228

102. Friston K. Active inference and agency. Cognitive neuroscience. 2014. pp. 119–121. doi:10.1080/17588928.2014.905517

103. Pezzulo G, Rigoli F, Friston K. Active Inference, homeostatic regulation and adaptive behavioural control. Prog Neurobiol. 2015;134: 17–35. doi:10.1016/j.pneurobio.2015.09.001

104. Hipólito I, Baltieri M, Friston K, Ramstead MJD. Embodied skillful performance: where the action is. Synthese. 2021;199: 4457–4481. doi:10.1007/s11229-020-02986-5

105. Marković D, Goschke T, Kiebel SJ. Meta-control of the exploration-exploitation dilemma emerges from probabilistic inference over a hierarchy of time scales. Cogn Affect Behav Neurosci. 2021;21: 509–533. doi:10.3758/s13415-020-00837-x

106. Da Costa L, Parr T, Sajid N, Veselic S, Neacsu V, Friston K. Active inference on discrete state-spaces: A synthesis. J Math Psychol. 2020;99: 102447. doi:10.1016/j.jmp.2020.102447

107. Vasil J, Badcock PB, Constant A, Friston K, Ramstead MJD. A World Unto Itself: Human Communication as Active Inference. Front Psychol. 2020;11: 417.

doi:10.3389/fpsyg.2020.00417

108.   Schoeller F, Miller M, Salomon R, Friston K. Trust as Extended Control: Active Inference and User Feedback During Human-Robot Collaboration. ArXiv. 2021 [cited 21 Sep 2021]. Available: https://www.semanticscholar.org/paper/7c39921b6050d7be82e18df2a31973f2c3f864ea

109.   Heyes C. Précis of Cognitive Gadgets: The Cultural Evolution of Thinking. Behav Brain Sci. 2019;42. doi:10.1017/S0140525X18002145

110.   Muthukrishna M, Doebeli M, Chudek M, Henrich J. The Cultural Brain Hypothesis: How culture drives brain expansion, sociality, and life history. PLoS Comput Biol. 2018;14: e1006504. doi:10.1371/journal.pcbi.1006504

111.   Graeber D, Wengrow D. The dawn of everything -A new history of humanity. 2021 [cited 25 Feb 2022]. Available: https://uujec.org/Portals/0/The%20Dawn%20of%20Everything.pdf

112.   Doyle MJ, Marsh L. Stigmergy 3.0: From ants to economies. Cogn Syst Res. 2013;21: 1–6. doi:10.1016/j.cogsys.2012.06.001

113.   Correia L, Sebastião AM, Santana P. On the role of stigmergy in cognition. Progress in Artificial Intelligence. 2017;6: 79–86. doi:10.1007/s13748-016-0107-z

114.   Ramstead MJD, Constant A, Badcock PB, Friston KJ. Variational ecology and the physics of sentient systems. Phys Life Rev. 2019. doi:10.1016/j.plrev.2018.12.002

115.   Constant A, Ramstead MJD, Veissière SPL, Campbell JO, Friston KJ. A variational approach to niche construction. J R Soc Interface. 2018;15: 20170685. doi:10.1098/rsif.2017.0685

116.   Friedman D, Tschantz A, Ramstead MJD, Friston K, Constant A. Active inferants: The basis for an active inference framework for ant colony behavior. Front Behav Neurosci. 2021;15: 126. doi:10.3389/fnbeh.2021.647732

117.   Guénin--Carlut A. Cognitive agency in sociocultural evolution. 2022. doi:10.31219/osf.io/x7yr4

118.   Coelho FC, Brandão A. Decentralising scientific publishing: can the blockchain improve science communication? Mem Inst Oswaldo Cruz. 2019;114: e190257. doi:10.1590/0074-02760190257

119.   Brian Arthur W. Economics in Nouns and Verbs. arXiv [econ.GN]. 2021. Available: http://arxiv.org/abs/2104.01868

120.   Fields C, Friston K, Glazebrook JF, Levin M. A free energy principle for generic quantum systems. arXiv [quant-ph]. 2021. Available: http://arxiv.org/abs/2112.15242

121.   Boik JC. Science-Driven Societal Transformation, Part I: Worldview. Sustain Sci Pract Policy. 2020;12: 6881. doi:10.3390/su12176881

122.   Rubin S, Parr T, Da Costa L, Friston K. Future climates: Markov blankets and active inference in the biosphere. J R Soc Interface. 2020;17: 20200503. doi:10.1098/rsif.2020.0503

123.   Friston KJ, Parr T, de Vries B. The graphical brain: Belief propagation and active inference. Netw Neurosci. 2017;1: 381–414. doi:10.1162/NETN_a_00018

124.   van de Laar TW, de Vries B. Simulating Active Inference Processes by Message Passing. Frontiers in Robotics and AI. 2019;6: 20. doi:10.3389/frobt.2019.00020

125.   ForneyLab.jl: Julia package for automatically generating Bayesian inference algorithms through message passing on Forney-style factor graphs. Github; Available: https://github.com/biaslab/ForneyLab.jl

126.   Friston K, Douglass DS, Iennaco de Vasconcelos ML, Sillett S, Sganzerla L, Tickles D, et al. Transcript of: Karl Friston, 1st Applied Active Inference Symposium, Active Inference Lab, June 21, 2021. 2021. doi:10.5281/zenodo.5797041

127.   Friston K, Da Costa L, Sajid N, Heins C, Ueltzhöffer K, Pavliotis GA, et al. The free energy principle made simpler but not too simple. arXiv [cond-mat.stat-mech]. 2022. Available: http://arxiv.org/abs/2201.06387

128.   Kirchhoff M, Parr T, Palacios E, Friston K, Kiverstein J. The Markov blankets of life: autonomy, active inference and the free energy principle. J R Soc Interface. 2018;15: 20170792. doi:10.1098/rsif.2017.0792

129.   cadCAD – A Python package for designing, testing and validating complex systems through simulation. [cited 15 Feb 2022]. Available: https://cadcad.org/

130.   darcspice: EVM agent-based token simulator 🐟🌀. Github; Available: https://github.com/opscientia/darcspice

131.   CRediT - Contributor Roles Taxonomy. In: CASRAI [Internet]. 14 Sep 2019 [cited 28 Feb 2022]. Available: https://casrai.org/credit/