dTwins: Decentralized Digital Twins for science, engineering and decision-making at scale

0. Summary

- Addressing our most pressing civilizational challenges, i.e., performing effective collective decision-making at all scales from individual and business to society, requires effective and efficient systems for collective sensemaking — here understood as **building shared mental models**.
- The scientific method, aided by probability theory and statistics, is our best tool to build reliable, widely-applicable shared models. However, our **current building blocks for collective systematic modeling and inference are unable to cope** with the scale, complexity and urgency of the problems they should be tackling.
- **dTwins (decentralized digital twins) are a principled solution** to these widespread, foundational challenges. Based on timeless principles of statistics, distributed computing, cybernetics and collective intelligence, they provide a standardized, rigorous yet flexible pathway for scientific and technological communities to **decentralize and reinvent the**

primitives of systematic sensemaking: the acts of scientific model creation, data collection, inference, prediction, and generalization.

- We outline a proposed "tech stack" and generic network architecture for dTwins across any domains, which generalizes Digital Gaia's design and WIP implementation for the ecological and climate impact space.
- Drawing on Digital Gaia's experience, we outline a case study of how the dTwin solution will enable exponentially faster development of the state of the art in regenerative agriculture and its application to generate real-world impact. We draw parallels to other fields and point to the possibility of a unified, cross-domain dTwin network.
- We propose an **ecosystem R&D consortium** around the development of dTwins as an open technology for accelerating collective sensemaking and public good provision at global scale, just as open source software accelerated all of information technology. This consortium will focus on developing the core dTwins technology and applying it to complement and amplify existing DeSci/open science efforts such as Hypercerts.

1. The problem

Part 1: Civilization needs shared mental models

> The future appears alien to us. It differs from the past most notably in that the earth itself is the relevant unit with which to frame and measure that future. Discriminating issues that shape the future are all fundamentally global. We belong to one inescapable network of mutuality: mutuality of ecosystems; mutuality of freer movement of information, ideas, people, capital, goods and services; and mutuality of peace and security. We are tied, indeed, in a single fabric of destiny on Planet Earth. *(Peter Senge)*

As humanity strives to overcome an arguable polycrisis and polities attempt to chart a coherent future, every policy-maker at every level (and indeed, private decision-makers dealing with the far-reaching implications of their local decisions) must come to terms with Senge's "inescapable network of mutuality". Indeed, arguably **the most critical challenges to the commons today are both global and highly contextual**. Issues like climate change or pandemics are obviously relevant to everyone, yet policies must be set locally, and the relevant data points, constraints and objectives form an enormously complex web.

The activity of progressively and adaptively making sense of this web is a process that systems thinkers have dubbed *building mental models.* Mental models are the “water of systems change”: in the words of Donella Meadows,

> … from shared social agreements about the nature of reality, come system goals and information flows, feedbacks, stocks, flows, and everything else about systems.

From the deep interconnectivity of this web, we can conclude that **every instance of contextual understanding and decision-making needs to integrate worldviews and mental models** across the entire spectrum of generality and abstraction, from broad “paradigmatic” assumptions all the way down to context-specific, local/situated (and often embodied/tacit) knowledge is needed. Further, **this integration must be two-way**; just as local knowledge is informed by top-down constraints, global paradigms must and do derive their legitimacy from bottom-up validation against the sharp edge of reality and practice.

However, the crucial task of global integrative sensemaking has been delegated throughout the 20th century to a host of centralized, bureaucratic policy and science institutions. These institutions - whether national or global, governmental or self-authorizing - are simply unable to internalize all this complexity; their conclusions and reports may be dressed in the mantle of authority and the jargon of science and data, but it’s becoming increasingly clear that this appearance is neither a guarantee of good science nor of timely, effective policy-making.

For example, the United Nations’ climate science center has, in the past, produced reports with egregious technical errors that, rather than being easy to paper over due to the ticking time-bomb nature of climate change, have seriously impacted its credibility, damaging the underlying policy cause for years. More recently, the UN’s World Health Organization has experienced significant pushback to its messaging and policy measures to mitigate the COVID-19 pandemic. Whatever the ultimate scientific value of these criticisms, the fact remains that “global science” is quickly losing its efficacy as the hallmark of sane global policy, posing serious danger to the latter’s legitimacy — and this despite the great need for policy that takes effective action in these arenas. **Absent a replacement shared sensemaking framework as a safety net, the very social fabric that’s needed to make policy relevant is threatened.** As Gordon Brander puts it:

> The cost of forking realities has dropped below the Coasean floor, and there's little incentive to merge realities. We fractally fragment understandings, then algorithmically amplify the confusion to maximize engagement. The most effective coordination mechanisms left seem to be memes and conspiracy theories.

There's a painful need for better solutions. However, tackling this will not be a simple matter of replacing institutions, or even of building incentive systems. Rather, **the very primitives of collective sensemaking must be upgraded for the exponential age**.

Part 2: The building blocks for shared mental models are outdated

> Every good regulator must contain a model of the system. *(Roger C. Conant and W. Ross Ashby)*

> If you want to make sense of a complex world, you've got to have an internal system that is equally complex. *(Karl Weick)*

Science is the process of collectively developing justified, generally-applicable beliefs about the world. Technologists, designers and practitioners (including those in creative or strategic activities) use these beliefs to solve problems for themselves or others, and are often the first to find the rough edges between beliefs and reality. Therefore, **the scientific sensemaking machine** — the engine for producing high-confidence, highly-applicable shared mental models — **is a fundamental enabler of public good provision**, in addition to being a hugely valuable public good in and of itself.

Hence, **it matters how the scientific "sausage" is made.** Specifically, the speed, adaptability and reliability of knowledge production rely on the construction of quantitative models and their testing against real-world data, using the tools of probability and statistics. But traditionally, scientific praxis has relegated this core epistemic infrastructure to the domain of convention and authority. Thus **the age of bureaucratic science was also the age of cargo-cult science.**

The practical matters of how science is made (scientific education, publishing and grant-making) have led to a state of affairs where **academic research proceeds as a sequence of one-off**

bureaucratic exercises, siloed efforts of data collection, modeling, and analysis designed to produce static scientific publications (the infamous PDF paper). The results of these exercises more often than not have limited utility to the broader community that supports science (whether directly or through public funds), even to those outside academia proper who are well qualified to engage with scientific research. Worse yet, papers are poorly vetted throughout their lifetime (from peer review to journal acceptance to impact through citations), and tend to fall short of even basic standards of quality control (as memorably shown by McElreath's "Science as Amateur Software Development".)

Once thought to be minor shortcomings that would be overcome by science's natural error-correction process, it's clear now that these basic problems have grave consequences. Since Ioannidis's landmark 2005 finding that most published research findings are false, science as a whole has been immersed in an ongoing replication crisis, with no resolution in sight. To be clear: the problem is not simply that a lot of money is going to waste on bogus studies (though it is); rather, **every single decision-maker who ventures into "following the science" must take into account the very significant possibility that they are building on a foundation of quicksand**. Some particularly consequential examples:

- In medicine, clinical trials and their analyses are so systematically distorted that senior researchers now "assume that health research is fraudulent until proven otherwise". Biomedical research has likewise devolved into a high-stakes forensics exercise.
- Throughout the 2000s, the very foundations of climate science were imperiled, as leading analyses of global temperatures were shown to be critically flawed: the "hockey stick controversy". The above frictions eventually caused a "reality fork" between academics, which was eventually resolved to some degree, but its imprint has had significant implications in global policy and popular debate to this day.
- Finally, in macroeconomics, the mainstream academic theory has become thoroughly divorced from reality — deeply wedded to provably false assumptions, impossible to rigorously justify, and unable to provide any forecasts of actual economic trends. No surprise that both the private sector and policy-makers have almost completely abandoned it, in favor of heuristic and cherry-picked approaches — which are of course more useful on a case-by-case basis, but have led the field to effectively cease development.

The state of the sensemaking process in engineering and business is in a similarly bad shape, though for somewhat different reasons. These domains are largely unaffected by the typical distortions of the publish-or-perish nature of academic research. To the extent that market forces are able to push corporations into making high-quality products and business decisions, incentives are far more aligned towards the production of analyses that conform with reality. However, corporate knowledge management tends to overly rely on trade secrecy, intellectual property and data privacy; meanwhile, organizational biases keep most individual organizations from taking the lead on the adoption of more constructive practices. As a result, **drastic limitations on the dissemination of knowledge hamper sensemaking progress and the adoption of valuable innovations**. Some examples:

- Despite the increasing availability of advanced diagnostic and treatment technologies and vast masses of biomedical research, personalized and precision medicine are still mostly fiction: between incompatible ontologies and data management constraints, no one — neither patients nor doctors — can ever have a complete, holistic picture of health data, let alone the means to estimate treatment outcomes. Both individual and public health decisions end up being made on the basis of gross averages and hunches.
- Precision and sustainable agriculture are likewise hindered: farmers wanting to optimize their land’s yield need to rely on heavily biased industry representatives, expensive (and still biased) consultant agronomists, or do their own research, which is slow, costly and of course, impossible to generalize. The consequence is the continued global dominance of unsustainable industrial monocultures, which are bad for consumers, perpetuate the need for agricultural subsidies, and undermine our planet’s very ability to sustain humanity.

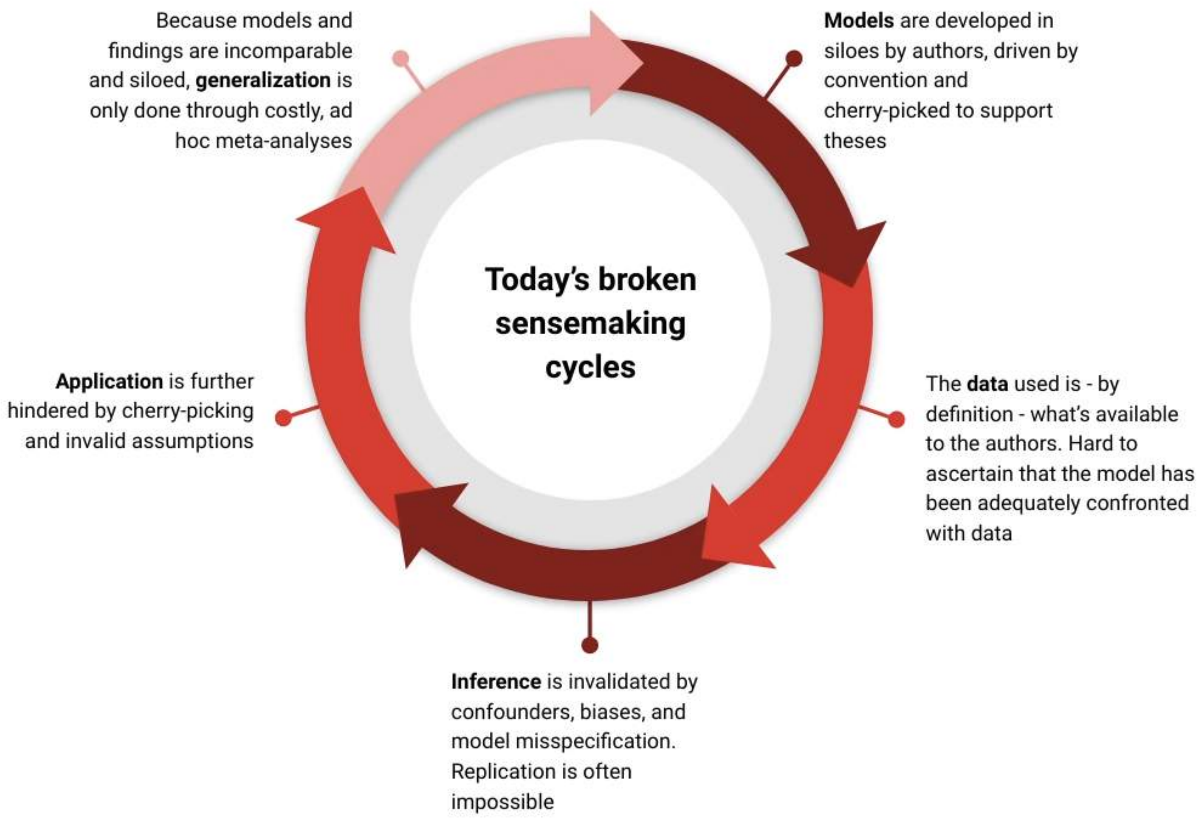

This is not to say that the methods have stagnated, quite the contrary. The state of the art of statistics (and its new subfields like machine learning) has benefitted enormously from the last few decades of computational and data-gathering advances: powerful new tools like probabilistic programming languages, Pearl's do-calculus of causality and counterfactuals, and of course artificial neural networks are all available for free in multiple, highly efficient open-source implementations in all major programming languages, in both local and cloud-based environments (with Python-based Google Colab being a favorite). And by the same measure, the wide availability of these tools (and of vast amounts of cheap compute, bandwidth and storage) means that statistical models and analyses can be shared, compared, consolidated, collaborated on, and componentized, providing the breeding ground for a truly exponential evolution of knowledge.

What would a future look like where all fields of science and engineering had fully embraced these new advances and turned them into primitives for a fully adaptive, global flow of sensemaking?

2. In a future with dTwins… (A case study)

> Alice the Agronomist has been asked by Fatima the Farmer to develop a plan for improving the yield and sustainability of Fatima's farm in Jordan. She starts by loading the farm's dTwin, syncing it from the live network to her local environment.
>
> By default, she sees a public view of the farm, including satellite imagery, weather, and other open data. She can also see that the dTwin is using a generic farm model, selected for Fatima when she opted her farm into the network. The generic model includes a soil health module that's calibrated to a regional average based on the last national soil survey, as well as modules with generic growth curves for all the crops that Fatima has been planting.
>
> Fatima gives Alice full permissions to her dTwin. Now she can see all data from the drone flyby service that Fatima hired last year, as well as harvest bookkeeping records and photos.
>
> Using the dTwin's built-in inference engine, Alice can quickly spot an interesting phenomenon: one of Fatima's wheat fields is 50% more productive than the generic model predicts, given her growing conditions and agricultural practices! The engine had already taken this into account when the harvest data was loaded, and adjusted farm-specific parameters such as soil quality estimates.
>
> Maybe that wheat field has something special? Alice decides to investigate. She takes some quick soil samples for the whole farm and

adds them to the dTwin. The engine immediately notes that Fatima’s soil is a bit more sandy than the regional average, but that doesn’t explain the productivity for that one wheat field…

Fatima mentions that her family has been using seeds from previous harvests on that field for generations, since it grows with little water. This gives Alice an idea. She runs a genomic analysis of the wheat and adds it to the dTwin; the engine notices an infrequent combination of seed, soil and irrigation characteristics that matches an obscure model developed by an experimental biologist in Portugal. The engine automatically overrides the generic model with the more specific one, yielding a growth curve prediction that matches Fatima’s wheat perfectly.

Alice connects the dTwin to her local optimization tool, and quickly finds the patches where soil conditions are best suited for the heritage wheat. With other constraints in mind, she and Fatima quickly come up with a plan to expand that crop to several other patches, replacing underperforming or expensive seeds. As the next season progresses, they will be able to track whether their experiment works in real time, just by following the dTwin’s updated inferences.

Fatima’s dTwin automatically shares all these updated parameters with other dTwins in the network. The detailed data like the wheat genomics are configured to be private, but Fatima opts to make the provenance public. And in a few weeks, other farmers reach out: their own dTwins have alerted them that there is a new wheat varietal that is perfect for their sandy, low-water conditions, and all of a sudden, they want to buy seeds from Fatima!

> Throughout this process, Delia the Data Scientist has been keeping the detailed model specifications well-groomed, and keeping close watch on the network's automated diagnostics. Anyone can inspect them and make sure that inferences are valid, leaving no room for cherry-picking or p-hacking.
>
> And Anna the Agroforestry Expert can reuse the models and extend the network to cover agroforestry farms as well. Next is Roberta the Reforester, and so on…

Stories like the above are what Digital Gaia is aiming to catalyze by building a dTwin-based solution for regenerative agriculture, called Fangorn. A diagrammatic view of this solution is illustrated below.

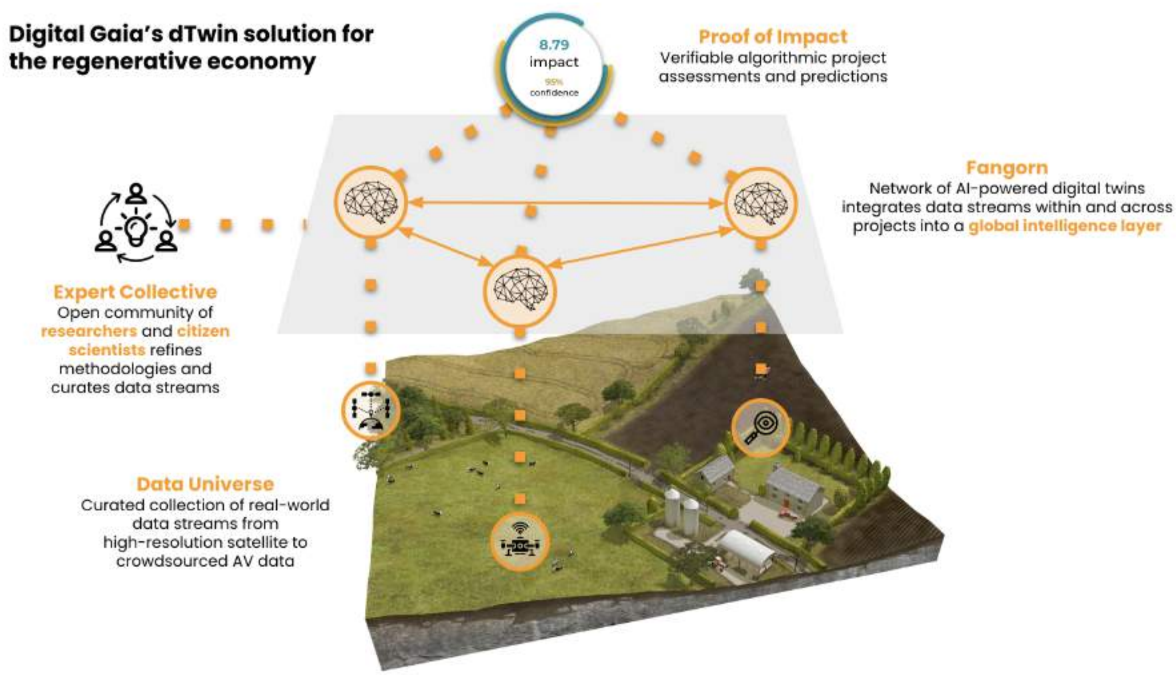

A **fully analogous story could be written for a medical practitioner, fitness instructor, nutritionist, urban planner, educator, conservationist, economic policy-maker…** All these domains struggle with analogous challenges and could benefit from the advantages of a dTwin network (while likely interfacing with their own dTwin networks through different connectors and applications).

Indeed, the generality of the dTwin architecture suggests that ultimately there might be a single, cross-domain dTwin network underlying all domains. Through our applied R&D at Digital Gaia, we have already found several opportunities for such cross-domain integrations, between agroecology and fields as diverse as climate, development economics, urban planning, food science, and microbiology. Senge's "inescapable network of mutuality" is becoming visible right before our eyes.

3. The dTwins Mission

> If you want to teach people a new way of thinking, don't bother trying to teach them.
> Instead, give them a tool, the use of which will lead to new ways of thinking. *(Buckminster Fuller)*

In Section 1, we argued that the business-as-usual of scientific and business sensemaking can't keep up with the interconnected, dynamic character of our present collective challenges. This is not by itself a novel diagnosis: there have been many calls and efforts towards rethinking this system, both from within and from outside. Notably, the decentralized science (DeSci), open science and citizen science movements have made great strides in setting up parallel institutions, structures and tools in the scientific realm. The goal of dTwins is not to supplant these efforts, but to complement and amplify them by intervening along a dimension of practice and tooling that has so far been relatively invisible and stagnant: the dimension of **effective inference**, or simply, of how well the knowledge community's cycle of knowledge integration and evolution adheres to information-theoretical principles (and therefore performs effectively in converging towards a useful representation of reality). The **concrete transformations** we aim to enable are summarized below:

- From **siloed, monolithic models** to **libraries of semantically composable model components**
- From **availability bias as default** to **automatic model-data confrontations**

- From **ad hoc inference** to **rigorous, replicable inference** using automated tools and guardrails
- From **"trying all the models until we find one that is applicable"** to **streamlined, efficient and justifiable application**
- From **costly, ad hoc generalization through meta-analysis** to **automatic generalization via parameter propagation** through the network

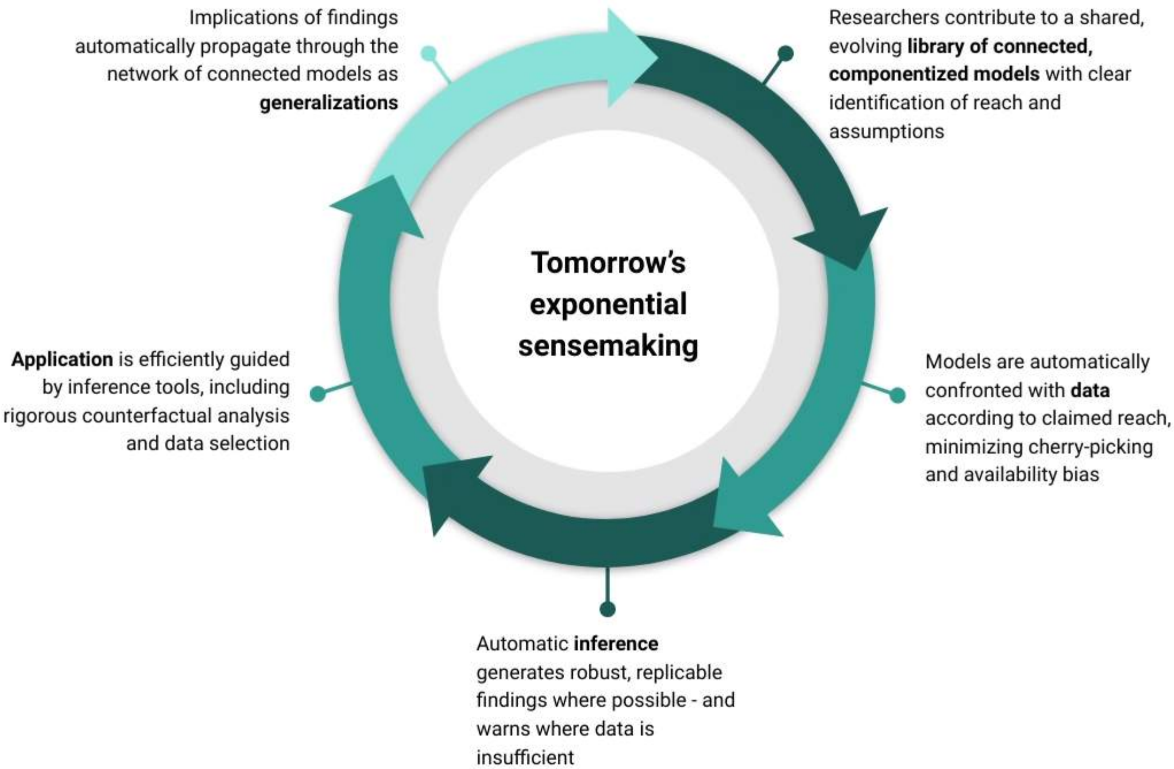

Effective inference is one of two key maturity dimensions where solutions in the current landscape are limited. The second is decentralization of the knowledge process and knowledge community. **Simultaneous intervention across both of these dimensions** has the uniquely potential to generate dramatically greater results, by **enabling the creation of new ecosystems of augmented collective intelligence** that are able to deal with the overriding ground-level heterogeneity, contextuality and ambiguity characterizing the experiences, needs, and interactions of real people like Fatima the Farmer. These deeply granular and adaptive institutions will be able to **see, hear and integrate** what she and others have to say about their **situated, embodied, enactive relation** to their physical and socioeconomic environments.

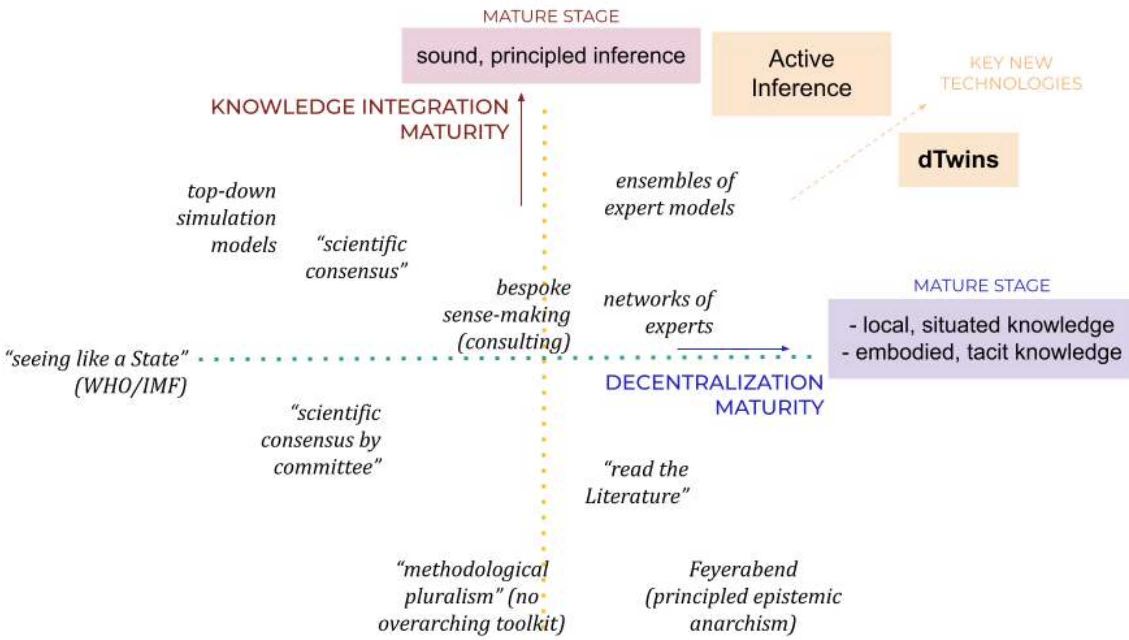

Thus:

👬 **The mission of dTwins is hence to *integrate*, through sociotechnical system design based on first principles of information theory, statistics and computer science, the knowledge present in *decentralized* networks of experts and practitioners.**

Learning from other superminds' evolution

The emerging field of collective intelligence has established the concept of a supermind: a cohesive agent emerging from the collective learning and behavior of a human community, augmented by the right shared beliefs, structures, processes and technologies (including notably software tools). This concept has given rise to a coherent design philosophy aimed at tackling wicked problems through the application of supermind design principles. The dTwins vision explicitly adopts this philosophy, which has also been increasingly a key part of the Active Inference research program for modeling and designing intelligent agents and ecosystems.

In particular, a lot of the above closely parallels (and is closely inspired by) the history of open source software. In particular, the patient and rigorous development of an open source stack, largely bottom up from the Linux microkernel and programming language compilers, that eventually stood up to (and then quickly surpassed) the best proprietary systems built by global behemoths like IBM and Microsoft, and within little more than a decade, led to the present-day status quo where application developers' work consists mostly of stitching together calls to open source libraries and APIs (using open source development tooling and deploying using an open source DevOps stack). As Robert Wright prophetically wrote in 1999:

> The vast, fast collaboration allowed by information technologies slowly turned the multinational technical community into an almost unified consciousness.

As this software development supermind continues to accelerate, the gap between it and the practices of science and engineering become ever more shocking. (Modular libraries with clearly defined semantics and documentation, organized in neatly navigable dependency graphs; standard issue tracking enabled by a culture of quality control and accountability; version control and dependency management that allow bug fixes and improvements to be incorporated in all deployments within hours — these are just some features of the developer experience that the software world now takes for granted, yet are entirely unavailable for the majority of science and engineering fields.)

dTwins are an explicit attempt to bridge these gaps, not just to make certain areas of decision-making more computational efficient, but to weave software and real-world decisions together — as indeed they have always been. For the logical conclusion is a global mind that asymptotically integrates these autonomous subsystems into a coherent whole that respects both their freedom and Peter Senge's "single fabric of destiny". As Thomas Malone, one of the founders of collective intelligence, writes:

> The global mind isn't something that either exists or doesn't. Instead, like collective intelligence and superminds, it is a perspective—a way of looking at the world. If you choose to look at the world this way, then the global mind exists by definition. And as more and more intelligent individuals become more and more closely connected, this perspective will become more and more useful.

4. Notes Towards the dTwin Protocol and Network Architecture

The dTwin system will follow a singular architectural pattern, driven by first principles of statistics, cybernetics, and systems design. We will not provide an extensive technical description here (if only because many of the details must evolve as specific requirements across use cases become clearer). We focus instead on some key concepts.

The dTwin stack

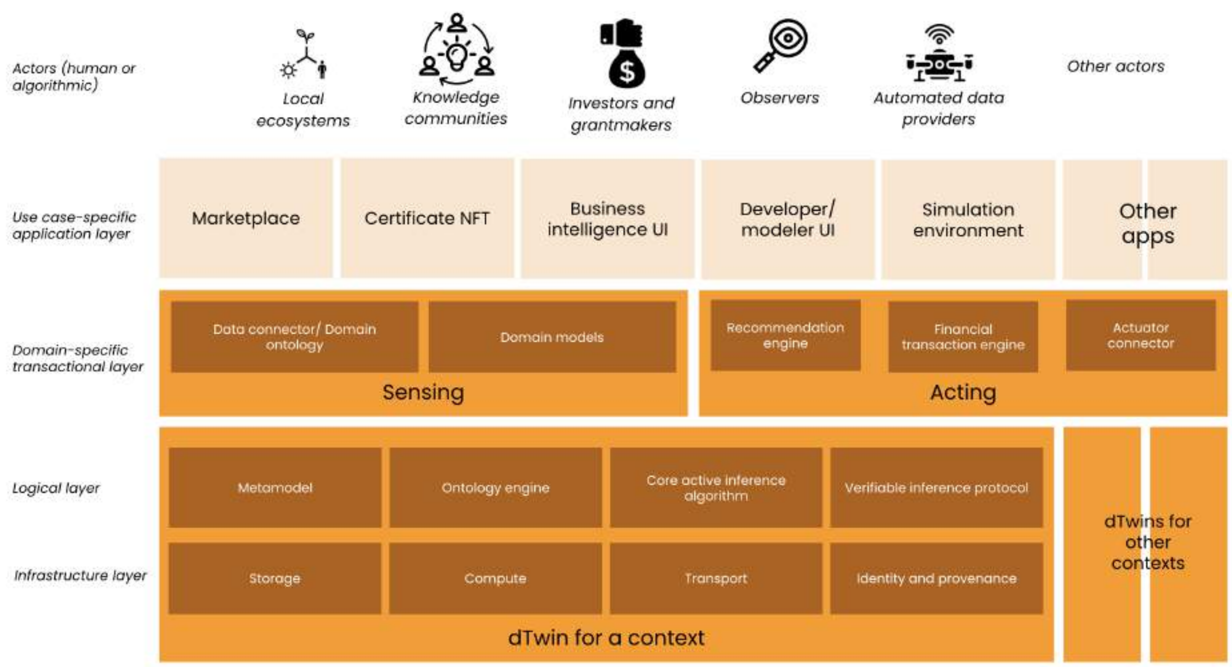

Aggregating and integrating knowledge

We can highlight some basic concepts powering sound and principled inference, that connect basic notions of how knowledge itself works to mathematical foundations:

- **Bayesian inference** is the principled way to reason about beliefs and uncertainty. By casting knowledge and working theories/hypotheses in terms of probability distributions, Bayesian methods enable them to be revised and updated given structural assumptions about the world (”models”) and observational data.
- **Planning as inference:** By casting policy values and preferences also as probability distributions, selecting optimal courses of action can be brought into the domain of Bayesian

belief updating. This neatly extends the deep mathematics of Bayesian inference towards a practical understanding of the knowledge as something that isn't just passively accumulated, but which *arises* out of a dynamic interaction between actors and their environment. It's worth noting that planning-as-inference is also perhaps the core concept of Active Inference, a closely related framework.

- **Hierarchical models** are a standard way to generalize across multiple contexts through *partial pooling*, where context-specific inferences can be generalized ("pooled") across multiple contexts by assuming some structural relationship between different contexts (which can be related as general graphs without cycles, and not just as trees).
- **Variational inference** (sometimes known as Variational Bayes): a technique to produce effective solutions to (analytically or numerically) intractable Bayesian inference problems by approximating them with feasibly solvable approximations.

Learning as a network

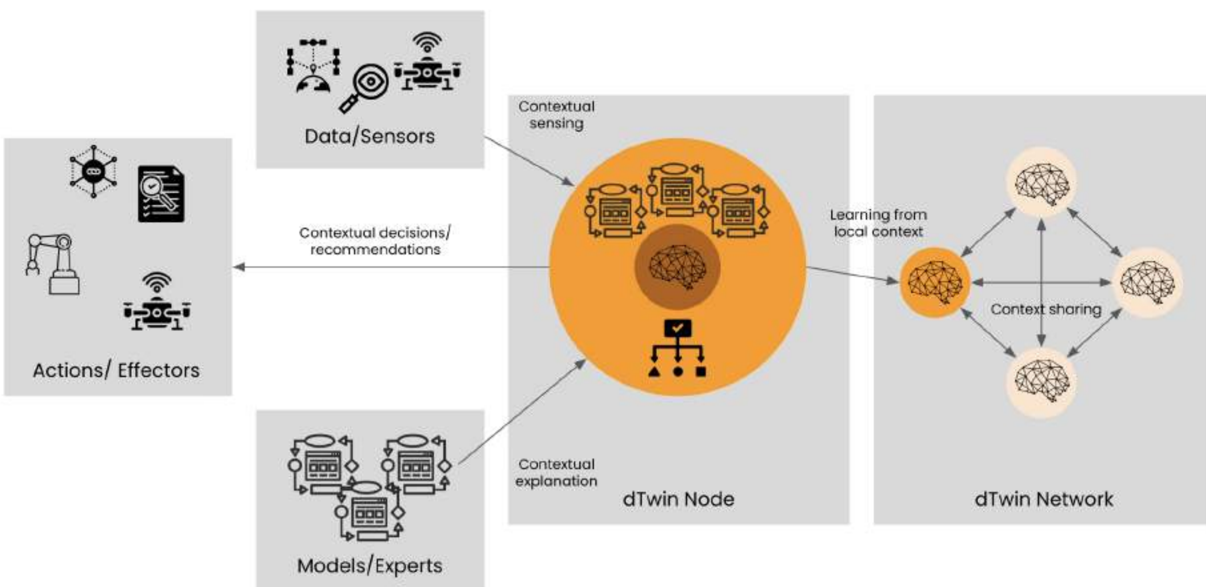

Centralized policy thinking has classically relied on simulations to make sense of the world. Sometimes these are large top-down models with some empirical validity; sometimes they're patchworks of small unrelated models. The concept of dTwins takes a giant step ahead by enabling a collective understanding of complex interdependent phenomena to emerge out of local nodes that stand in for concrete physical counterparts. In other words, dTwin technology is a way in which nodes can perform joint inference (in the sense just discussed). Through planning-as-inference, this also enables optimal decision-making (although this may mean different things *locally* and *globally).*

Recasting inference into a decentralized network setting requires some additional concepts and technologies:

- **Variational message-passing** (VMP) is a common algorithm to solve variational inference problems on Bayesian networks, where variational parameters (roughly, variables in the model) for each node are found by an iterative procedure of belief updating and message-passing between nodes; this procedure requires only local information and converges to optimal parameters for the entire network that balance “bottom-up” requirements (fit to observations) and “top-down” requirements (fit to structural assumptions and expert priors).
- **Verifiable computing** is a set of techniques for distributing computational loads to untrusted partners in such a way that results are still verifiable. Verifiable computing has a long history in citizen science through projects like SETI@home and Folding@home. The proof-of-work and proof-of-stake algorithms that power blockchain technologies are closely related ideas.
- **Data chain of custody**, like verifiable computing, has to do with the problem of distributing workloads to untrusted agents. A chain of custody enables nodes to have an immutable record of the provenance of data, avoiding issues like corruption and adversarial data manipulation.
- **Dynamic network topology** refers to the idea that nodes and edges may be added and removed throughout the lifetime of the network. Long used in basic internet engineering to produce network resiliency, dynamic topologies also have complex effects on convention emergence and coordination.
- **Verifiable inference** is the application of verifiable, decentralized computation techniques to the inference function taking the form “data, model, starting parameters → new parameters”. Each dTwin is able to form its own local chain of verifiable inferences, based on local models and data (specialization); these are broadcast to peers, which independently choose how much to trust and when/where to generalize based on their own local models and data.

5. The dTwin Consortium

> Let us abandon convention and consider, without encumbrance, a more optimized scientific ecosystem. Humans are the most powerful architects in the world. Let us reinvent the innovation space. Let us take on the creativity crisis. *(Roberta Ness)*

At Digital Gaia, we are building our entire network and product around dTwins, and we believe that many other organizations will do likewise, forming the basis for a large, thriving, interdisciplinary ecosystem. To accelerate this vision, we are creating a **dTwin Consortium**.

The consortium is targeting a **total $1 million USD commitment** for 2023, between funds and in-kind contributions in the form of R&D expertise and other resources. These funds will be dedicated to open-source R&D on dTwins as a public good infrastructure. To bootstrap the consortium, Digital Gaia commits to dedicating **$200,000 in R&D resources**.

Governance

In the short-term (until mid-2023), we propose to **incubate the consortium within Digital Gaia**'s legal structure, in the format of a partnership contract between Digital Gaia and other consortium members.

In the medium-term, we aim to **transition the consortium to its own independent foundation**, with open, decentralized governance intentionally designed following the principles and values mentioned above, in a public forum with widespread engagement.

Ethical commitments

We are deeply cognizant of the dangers of intelligence run rampant; one clear dark side of "software eating the world" has been the appropriation of economic value and sociopolitical power by organizations especially skilled at wielding that software. From this perspective, we aim to make ethical values such as **responsibility**, **fairness** and **sovereignty** an explicit part of the dTwin charter from day one, and we commit to building these values into our operating protocols through intentional, adaptive design and great humility. In due time (as binding agreements are hammered out), a Statement on Ethical Values and a Code of Conduct shall be prepared.