



Signals at Scale

REPORT

Value in the emerging AI-mediated information ecosystem

Where value will be created,
where it will be captured,
and what needs to be built.

David Caswell and Shuwei Fang

September 2026

DOI: <https://doi.org/10.5281/zenodo.22894513>

Published under a Creative Commons Attribution 4.0 licence (CC BY 4.0).

Contents

Executive summary	1
Anticipating an AI-mediated information ecosystem	4
Introduction	4
The AI-mediated information ecosystem	5
Previous work	7
The Signals at Scale Summits	9
Objective	9
Design and recruitment	10
Setting up the discussions	12
Situational awareness	12
Early signals	13
The tabletop exercise: News products for billionaires	15
First principles requirements	21
Jobs-to-be-done	31
Investment categories	43
Summary discussions	55
The vibe	60
Value in the AI-mediated information ecosystem	63
Overview	63
Value creation	63
Value capture and moats	67
Detailed categories	69
Infrastructure	75
Conclusion	79
What did we learn about the AI information ecosystem?	79
What did we learn about the process?	80
What could we have done better?	81
Next Steps: The Signals at Scale Studio	82
Acknowledgements	84

Executive summary

This report synthesizes four Signals at Scale summits involving approximately 140 participants from technology, investment, governance, academia, media, and adjacent fields. Its starting premise is that AI will fundamentally reorganize how societal information is observed, produced, represented, verified, distributed, experienced, acted upon, and rewarded. The emergence of AI requires nothing less than a fundamental redesign of society's knowledge infrastructure and of the information flows through which people and institutions perceive reality, form judgments, coordinate, and exercise power.

The report anticipates a transition from an ecosystem organized around fixed content artifacts — articles, videos, podcasts and websites — to one organized around continuous, machine-readable flows of "liquid" information. AI agents will likely gather information at extraordinary scale, manage it as referenceable units, attach provenance and rights, verify and enrich it, combine it with personal and collective context, and assemble moment-specific experiences for humans or for other machines. Distribution may become continuous negotiation among agents over relevance, quality, access, price, permitted use, and action. Information flows may become bidirectional, with signals flowing back from consumption to producers and infrastructure providers.

This restructuring will transform the economics of information. As generic production and synthesis become abundant, value migrates away from finished artifacts and toward scarce observation, proprietary or tacit context, integrity and assurance, personal relevance, decision support, delegated action, and the systems coordinating exchange. The largest new value is likely to arise where information intersects with a person's accumulated context, their immediate intent, and capacity to act. The most powerful positions may therefore belong not to those producing the most information, but to those closest to the consumer's context, intent and transactions.

What the summits produced

The report's central output is a value map of the emerging ecosystem, derived from the first-principles approach the summits ran on. It starts from the most basic needs of the stakeholders in an information ecosystem rather than from the sectors, industries and familiar entities we take for granted. Given the uncertainty of the subject matter this value map is inevitably somewhat convoluted.

The map is based on value-creating functions that describe what the ecosystem does: originating information that does not otherwise exist; representing it so machines can use and retain it; establishing its integrity and permitted use; combining it with personal, community and domain context; interpreting it for a particular purpose; converting it into decisions, actions and experiences; and coordinating the feedback, payment and accountability loops around all of it. These functions are not sequential. They run continuously and in parallel, supported by informational signals flowing back from the point of use — demand, usage, impact, payment, reputation and corrections — and by coordinating roles spanning identity, authorisation, provenance, rights, privacy, security, interoperability, competition and governance.

These functions resolve into fourteen distinct zones of the value map, each with its own economics and in many cases its own market failure, and seventy investment categories distributed across them. Against this, the report identifies likely positions where value is captured and moats that make those positions defensible, together with the layers of infrastructure on which the whole ecosystem depends.

The distribution across those zones is itself a finding. The zones we found hardest to reduce to a small number of categories — ecosystem governance, primary observation, shared public knowledge — are largely those where value is hardest for any single participant to capture. Value creation is distributed across the whole map, but value capture concentrates in a few places close to the consumer's context, intent and transactions. That divergence is the report's central economic finding, and the reason the infrastructure layers are themselves market failures to some degree. It is sharpest at the foundations, where value is hardest for any single participant to hold.

The stakes

The societal stakes are correspondingly immense. Hyper-personalization could give ordinary people informational capabilities once available only to governments, large corporations, and the very wealthy: continuous expert analysis, prediction, representation, and decision support. Yet the same architecture could enable unprecedented surveillance, manipulation, dependency, and concentration. An ecosystem may be highly personalized at the surface while profoundly centralized underneath. Non-portable personal context could create "AI serfdom"; extreme personalization could fragment shared reality into "bubbles of one"; and control of

integrity, identity, or agent authorization could become a new gatekeeping power of extraordinary consequence.

Governance is therefore part of the core product, technical, and market architecture of this new ecosystem. Persistent tensions must be negotiated between personalization and shared knowledge, between useful context and privacy, between frictionless access and sustainable production, between open standards and defensible businesses, and between individual agency and collective welfare. Priority foundations include standards for machine-usable representation and memory provenance; verification, correction, and evidence systems; consumer-controlled context portability; human and agent identity and authorization; attribution, pricing, negotiation, and payment mechanisms; competition safeguards; and public-service infrastructure for audits, emergency channels, shared knowledge, and universal access.

These foundations will not be built by private capital alone. The report points toward a division of labor: philanthropic and research capital should explore uncertain problems and develop open standards and evaluation systems; public capital should protect rights, create demand, and guarantee essential access; and private capital should scale applications, transactions, certification, assurance, and workflow services. Sequencing matters: representation and integrity infrastructure should precede sustainable reward markets.

The report is an initial map of highly uncertain terrain, not a validated blueprint. No mature body of knowledge yet describes a fully AI-mediated information ecosystem, and its eventual form may be stranger, larger, and less legible than current categories suggest. But the direction and urgency are clear. Choices made in the next two to three years may determine which information reaches whom, in what form, under whose authority, and on what economic terms for generations. The required response is not incremental adaptation but coordinated experimentation and institution-building to make the emerging ecosystem truth-seeking, self-correcting, pluralistic, and aligned with human agency and flourishing.

This report was written by humans using AI. The heart of the report – the description of the summit and the reporting of the five sets of discussions – was written entirely by hand, as was the opening section on anticipating the AI-mediated information ecosystem and the conclusion section. The structure of the value map section was produced iteratively – section by section – using Claude Opus 4.8, Claude Fable and ChatGPT 5.6 Sol, based on the handwritten context and instructed by approximately 5,000 words of prompting. The section was then substantially edited and rewritten by hand.

Anticipating an AI-mediated information ecosystem

Introduction

This report reviews outcomes from four summits held between December 2025 and May 2026 in the United States and Europe, focused on exploring the landscape of potential investment opportunities within the rapidly emerging AI-mediated information ecosystem.


The summits themselves, and subsequently this report, assumed some familiarity with the intersection between AI and societal information – information that contributes to the functioning of a healthy and productive society. The central hypothesis behind the summits was that Artificial Intelligence generally, and agentic systems specifically, will fundamentally restructure the ways in which societal information is produced, processed, distributed and consumed in human societies. The summits were explicitly not concerned with the many incremental opportunities available from AI in optimizing of the productivity of the existing, familiar information ecosystem.

The report should be read as an exploratory landscape map rather than a rigorously validated forecast or investment blueprint, or as a systematic and representative account of stakeholder needs. It synthesizes structured discussions among a selected group of expert participants and therefore reflects the assumptions, perspectives and limitations of that process. Its categories are provisional, overlapping and have not been tested against market data, external research or direct evidence from consumers. The report is intended merely to help clarify an as-yet extremely uncertain AI-mediated information landscape and not to provide certainty about any investment opportunity.

These summits were explicitly not political. While some political views inevitably appeared in some conversations, the great majority of the discussions did not assume any particular ideology or motivation other than a desire for a healthy, well-functioning information ecosystem that contributes to human flourishing. This bias toward ‘societally beneficial’ outcomes was an explicit objective of the summits, and while the discussions sometimes referenced possible harms or negative outcomes,

they usually did so in the context of a realistic appraisal of the requirements for beneficial outcomes.

These two themes in the orientation of the summits – the assumption of fundamental transformation and the bias toward positive societal outcomes – were explicitly expressed in the chart below – a chart that was repeatedly referred to throughout each summit:



The summits attempted to focus solely on the new information ecosystem that might emerge as a result of continually improving capabilities and continually increasing adoption of AI. The many other societal, economic and political disruptions that AI is likely to produce were intentionally set aside as much as possible. This is not because these issues are unimportant, or unrelated to the information ecosystem, but because some constraint toward a relatively narrow topic was seen as necessary for progress.

The AI-mediated information ecosystem

The prospect of a digital information ecosystem mediated largely by AI has become widely accepted as increasingly plausible since the launch of ChatGPT in November 2022. The general interpretation of this ecosystem is of AI decision-making at every stage of the information value chain. Early forms of AI deployed in recommender systems and in some personalized curation systems were restricted to choosing the

ordering of traditional content artifacts presented to consumers. In contrast, even the simplest interpretation of a fully AI-mediated information ecosystem anticipates the integration of AI into information gathering and production, into its verification, contextualization and processing, into its distribution and accessibility, and into its consumption in deeply interactive and personalized experiences. More sophisticated interpretations of an AI-mediated information ecosystem anticipate entirely new forms of information and new uses of information arising from the dramatically increased scale, scope and pace of information flows facilitated by AI. Many interpretations anticipate an information ecosystem in which most of the participants are not human but are instead forms of AI agents acting autonomously and intelligently on behalf of their human principals.

The concept of an AI-mediated information ecosystem is closely connected to the concept of an “agentic internet” – an internet in which most of the ‘traffic’ is interaction between AI agents and in which human consumption is relegated to a periphery where experiences are assembled moment-by-moment to serve a particular need or circumstance. This is an internet largely without websites or applications and is focused far less on directly accessing ‘destinations’ and far more on assembling or interpreting the informational ‘raw materials’ needed to achieve a goal. While the agentic internet is not yet apparent to most consumers, it is rapidly emerging and its growth has been accelerated by the emergence of multiple agentic interaction protocols such as MCP, A2A, ACP and AP2.

An AI-mediated information ecosystem is also closely connected to the concept of ‘liquid content’ or ‘liquid information’ – semantic information that is fully separated from the ‘content artifact’ or media container in which it was initially communicated. In an ecosystem of ‘liquid information’ the form of the originating ‘content’ is irrelevant, and both informational and stylistic meaning can be extracted by AI and changed, separated, combined, reused, and repurposed in essentially any way that can be imagined.

It should be clear that the AI-mediated information ecosystem is unlikely to be merely a ‘scaled-up’ or ‘better’ version of our existing information ecosystem, but instead will probably exhibit entirely new architectures, entirely new processes and entirely new uses. Furthermore, given the central role of information in the lives and societies of human beings, as well as our experience of the emergence of the internet, the AI-mediated information ecosystem is likely to have substantial and

unanticipated secondary and tertiary effects. While we can observe a wide range of ‘early signals’ of coming change in the early application of AI to societal information flows, the shape and behavior of an end-stage information ecosystem mediated largely by AI will likely remain opaque for some time. The Signals at Scale summits were an attempt to bring some small measure of additional early clarity to our collective understanding of how it might develop.

Previous work

The organizers of the four Signals at Scale summits – Shuwei Fang and David Caswell – have collaborated on a series of earlier projects focused on the AI-mediated information ecosystem, each of which produced a detailed report similar to this one.

The AI in Journalism Challenge project, conducted from June to November 2023 and funded and operated by the Open Society Foundations, applied an accelerator approach to the rapid application of AI within small, entrepreneurial newsrooms, modeled on the Y Combinator program. This project essentially explored the limits of an incremental application of AI to familiar workflows and products delivering societal information, and the report of its findings is available at: <https://www.opensocietyfoundations.org/publications/open-society-s-applied-ai-in-journalism-challenge>

The AI in Journalism Futures 2024 project was a large-scale scenario development project aimed at understanding the structural possibilities for a fundamentally different information ecosystem mediated by AI. Also funded and operated by the Open Society Foundations, this project fused a broad quantitative approach that collected scenarios from around 1,000 people with a highly structured qualitative scenario development workshop of 60 carefully selected participants. It developed five core scenarios and the report describing its findings can be found here: <https://www.opensocietyfoundations.org/publications/ai-in-journalism-futures-2024>

The AI in Journalism Futures 2025 project was a fully agentic repetition of the AIJF 2024 project. This project was funded by The Tinius Trust and utilized OpenAI’s early ‘Agent Mode’ with the GPT 5.1 model shortly after its introduction in mid-2025. The goal of this project was to provide an apples-to-apples comparison between human and agentic versions of the same large-scale knowledge production project. The project used 1,000 comprehensive AI personas in place of the human contributors, as

well as versions of 'digital twins' to represent experts gathered for the workshop. The final analysis and report was produced entirely agentically and is available here: <https://aijf2025.tinius.com/no>

This body of work, in addition to other work conducted separately by the organizers, heavily informed the design and execution of the Signals at Scale summits, as well as the analysis of their output and the production of this report.

The Signals at Scale Summits

Objective

The core objective of the Signals at Scale summits was to identify useful areas of actionable investment within the emerging AI-mediated information ecosystem. The summits were essentially a reaction to investment trends that had emerged since the launch of ChatGPT in late 2022, across media, journalism and the traditional knowledge-producing sectors, including the first wave of AI-native players. These trends emphasized incremental change, the application of AI to existing tasks, workflows and products, and the default incorporation of many legacy assumptions. Many public discussions about the future of societal information in the AI era remained more focused on interesting conversation than on forging actionable progress and most tended to emphasize potential harms rather than opportunities.

Our approach was to begin with the assumption of fundamental transformation and with the objective of pursuing opportunities for positive societal outcomes, and then to build a structured workshop design that partially constrained the discussions in the direction of those objectives. The organizer's previous experience leading a large scenario development project (the 'AI in Journalism Futures' project) had produced overwhelming feedback about the efficacy of a highly structured approach oriented toward specific outcomes. The design of the Signals at Scale summits incorporated those lessons.

The specific outcomes sought were around establishing a shared understanding of how information might flow in an AI-mediated world – the physics and economics of the emerging ecosystem – as a basis for understanding where resources might be productively invested. This included first understanding what stakeholders in the information ecosystem – primarily consumers – might require from a restructured information ecosystem of radically expanded scale and scope, then understanding what key activities might need to occur in the ecosystem for those requirements to be met ('jobs-to-be-done') and finally understanding where investment might need to be applied to enable those activities.

We ideally wanted to identify where value might naturally accrue in the new ecosystem, where it might be economically defensible and where market failures may occur – a sort of 'value map' of the emerging ecosystem – and to identify a small set

of categories of investable projects for this new environment. This report is the final expression of these goals.

Design and recruitment

The summits were designed to get to the heart of the subject matter quickly, to remain tightly focused on the outcomes and to help forge a community that might continue the discussion following the summit. The design achieved this using eight separate steps:

Pre-work: Prior to the summit each participant was sent some preparatory reading, and a request to identify an ‘early signal’ indicating potential transformative change from AI in their own information environment.

Level set presentation: Each summit began with a one-hour ‘Situational Awareness’ presentation aimed at ensuring that all participants were fully aware of the current state of AI and of its influence on societal information. The presentation covered the latest data on AI benchmarks, adoption, investment, economic effects, capabilities (especially agentic capabilities), broad-spectrum expert opinion, and responses by legacy information-producing institutions. The presentation also covered the economic paradigm shifts in the information ecology caused by AI, the application of an economic framework for businesses subjected to scaling technology, an analysis of when and how existing information business models might break and a “Bull Case for Knowledge” centered on a potential dramatic expansion of the overall market for information enabled by AI. Finally, this presentation set up the summit and its orientation toward transformational and positive societal outcomes and presented an analysis of the early signals submitted by participants.

Tabletop exercise: The purpose of the tabletop exercise was to quickly challenge the baseline assumptions of the participants, to help them adopt a more imaginative, first-person perspective of the dynamics of the competitive environment as set out in the level-set presentation, and to help the group set norms around discussing extreme scenarios. The exercise challenged the participants to first imagine information products and services in a situation essentially free of resource constraints, and then to examine the defensible competitive advantages of those information products and services.

First principles discussion: The opening discussion was intended to build a base for the summit’s outcomes by deeply reconsidering why and how humans value

information, and the core requirements for various stakeholders in the information ecosystem.

Jobs-to-be-done discussion: The second discussion was intended to build directly on the first discussion by exploring how requirements might be translated into high-level functions or capabilities across the emerging ecosystem.

Investment categories discussion: The third and last operational discussion was intended to continue building on the earlier discussions by gathering high-level functions or capabilities into coherent and related categories of potential products, services and infrastructure.

Summary discussion: The summary discussion was intended to enable participants to stand back and assess the totality of the summit without constraint, and to surface 'loose ends', reactions, criticism, lessons learned, summaries, reflections, process advice, etc.

Signal channel: Following the conclusion of each summit a channel for all participants was set up to enable further conversation, sharing and contact.

All summits were conducted under Chatham House rules, with participants encouraged to use the protection it provided, and the necessity of sustained, active engagement from all participants was emphasized throughout.

The summits were primarily shaped by the people who participated in them. Attendance at each summit was by invitation, and the identification and recruitment of participants was generally conducted based on three criteria. First, we sought participants who had previously engaged deeply with AI and its transformative implications for the information ecosystem and who were fully aware of the capabilities available from frontier models and agentic systems. Secondly, we sought participants from across a broad range of disciplines. This included some representation from existing information-producing institutions but most representation came from adjacent fields, such as technology, investment, governance, academia, etc. Third, we intentionally sought to avoid participants who were deeply vested in either incremental improvements to the existing ecosystem, or in the potential harms from the emerging AI-mediated ecosystem (acknowledging that while these are important considerations although they were not the focus of these summits). Each summit was essentially a portfolio of diverse but well-prepared

people, primed to engage in discussions that were substantially different from conversations occurring in other venues.

Setting up the discussions

Situational awareness

The workshop opened with a one-hour presentation intended to establish a shared 'baseline' awareness among participants of AI and its implications for the information ecosystem. The presentation traced the trajectory of AI capabilities over the past few years, from GPT-2 in 2019 to the most advanced reasoning and agentic systems that perform at or beyond human expert level on demanding benchmarks. Expert consensus on the possible imminent emergence of Artificial General Intelligence (AGI) and Artificial Superintelligence (ASI) was taken seriously, including recent statements by leadership at frontier labs anticipating dramatic acceleration within the next one to three years.

The presentation reviewed the economic and physical commitments of extraordinary scale currently being applied to extending AI capabilities — including an estimated 7 trillion US dollars of global data center investment through 2030, of which around 5.2 trillion is attributed to AI workloads (McKinsey, 2025). It reviewed adoption surveys showing a third of US consumers already interacting with AI “almost constantly”, and content audits that show around a quarter of published websites and content is already AI-generated. The emerging 'agentic internet' was examined, including tools, orchestration frameworks, and interaction protocols that already enable agentic systems to act, transact, and consume information on consumers' behalf.

A central proposition of the presentation was that the common current strategies of applying AI incrementally to existing tasks, workflows and products are likely to be largely irrelevant in an AI-mediated information ecosystem. It also challenged the assumption that today's ecosystem (journalism, etc.) represents the best possible form of societal information, and reframed 'information overload' as a failure of legacy information artifacts and of our inability to delegate consumption, not a problem of too much information itself¹.

The presentation identified four emergent paradigm shifts that define the new ecosystem: a shift from assumptions of scarcity to assumptions of abundance; a shift

¹ <https://radicallyinformed.substack.com/p/an-ai-enlightenment-the-consumer>

from human audiences to machine (agentic) consumers of information²; a shift from immutable information artifacts to fluid, 'liquid' experiences of information; and a shift from an economy of capturing attention to one of serving intention³ – what people actually know, want, and need. The economic implications of these shifts for today's media economics were reviewed: production ceases to be a defensible moat, the audience becomes an AI intermediary's customer, the market 'barbells' into luxury and commodity ends⁴ while the middle hollows out, and demand signals move upstream into rich, conversational expressions of user intent.

The presentation closed with a 'bull case for knowledge'⁵. Value from information, it argued, will not be destroyed but will migrate from bundled artifacts into separate functions that remain defensible – verification, provenance, context, interpretation, trust and community – while collapsing production costs will make previously uneconomic knowledge viable and unlock latent demand, including entirely new categories of machine audiences. Whether the value created is captured depends on conditions for each stakeholder, such as competition at the point of consumption, demand-signal visibility for producers and shared epistemics for society. The session ended with an admonition that framed the summit: be suspicious of solutions that require the least amount of change.

Early signals

Prior to each summit all participants were asked to submit an example of an 'early signal' that they had observed in the emerging AI-mediated information ecosystem and that might indicate fundamental change. Across all four summits we received 96 usable early signals. The early signals for each summit were presented to participants in the opening presentation, along with affinity mapped clusters. For this report we have re-done that affinity mapping (using Claude Cowork, with substantial editing, and emphasizing the end state that the signals pointed toward) across all 96 early signals, resulting in 12 discrete clusters. These clusters are as follows:

2 <https://www.economist.com/by-invitation/2026/04/16/welcome-to-the-world-of-machine-audiences>

3 <https://shorensteincenter.org/resource/from-attention-merchants-to-intention-architects-the-invisible-infrastructure-reshaping-human-curiosity/>

4 <https://radicallyinformed.substack.com/p/beyond-the-artifact-the-brutal-economics>

5 <https://reutersinstitute.politics.ox.ac.uk/news/information-ecosystem-being-redrawn-ai-might-be-good-news>

1. **The Agent as Interface:** Consumption shifts from visiting destinations to delegating to an agentic layer that reads, filters, and compresses the internet into synthesized, decision-ready answers and actions.
2. **Hyper-Personalization:** Information is dynamically reshaped per person – context, tone, viewpoint balance, cognitive load, time available, “what does this mean for me”, etc. – at a depth of intimacy that is closer to a family member than a service provider.
3. **Content Atomization & Liquid Information:** Information decouples from fixed formats and becomes fully fluid. Information of record breaks into recombining atoms flowing from and through AI-native supply chains, where any input becomes any experience at negligible marginal cost.
4. **Queryable Knowledge Structures:** The unit of value shifts from the story to something like the interrogable model: archives, corpora, civic models, expert reasoning, and structured datasets become grounded, queryable knowledge products – with narrative just one rendering of them.
5. **Disposable, Personal Software:** Vibe-coding and easily configurable agents let non-technical people build their own feeds, filters, briefings, and context-specific apps – rapidly dissolving dependence on one-size-fits-all information products.
6. **Autonomous Newsgathering:** Agent swarms continuously monitor domains, investigate sources, synthesize and evaluate evidence, detect what qualifies as news, and execute multi-step workflows – well beyond what would be possible to coordinate with humans.
7. **Democratized Expertise & Inclusion:** Expert-grade guidance – medical, agronomic, professional – reaches people and low-resource communities previously excluded, while AI brings marginalized or peripheral voices into larger and more central communities.
8. **New Trust, Verification & Provenance Infrastructure:** As synthetic content floods feeds, verification – including AI-supported Open-Source Intelligence (OSINT), factchecking and fraud detection – democratizes, and trust migrates toward cryptographic provenance and machine-readable authoritative sources.
9. **New Information Economics:** The business of information restructures: new content-as-API businesses develop for machine buyers, and bundled media businesses splinter into utilities, services, and embedded data products.
10. **Reconfigured Human Roles & Institutional Lag:** Humans shift toward judgment and context management while AI absorbs infrastructure and process work;

pedagogy, corporate labor models, journalistic practices, and civil society lag behind, and agent accountability remains unresolved.

11. **Collective Sensemaking via Synthetic Relationships:** Emotionally powerful AI personas, simulated memories, participatory fandom, and AI-aggregated collective voice reshape the social and emotional layer of the information ecosystem.
12. **AI as Cognitive Infrastructure:** Agentic AI becomes ambient infrastructure for private reasoning and sensemaking – superhuman at language operations and deeply integrating personal, domain, community, and universal contexts.

Very generally, these possibilities, or possibilities similar to them, were assumed by most of the participants in most of the discussions in all four summits.

The tabletop exercise: News products for billionaires

Overview

The Signals at Scale discussions opened with an exercise intended to encourage participants to explore what information products and services might become if resources were no longer a significant constraining factor. The premise was that AI would give average consumers purchasing power equivalent to that of the billionaire class, and the intent was to give participants the experience of designing a product or service under assumptions of abundance rather than assumptions of scarcity. This prepared the participants for the following discussions by serving as a ‘warm-up’ exercise that encouraged the reimagination of offerings, the reframing of businesses and the reconsideration of assumptions in circumstances where the cost side of information production, processing and communication might approach zero.

The exercise was a heavily adapted and significantly shortened version of a RAND-style tabletop exercise (TTX). This approach begins with a detailed premise (‘the conceit’) describing a complex situation requiring a strategic response. The participants do not play ‘roles’ in the exercise, but instead act as themselves, bringing their own skills, experiences and opinions. Following a substantial discussion about the initial situation and the possible responses to it, an intervention is made with a new development requiring a further strategic response. The discussions are not facilitated in any way, and the participants engage with and respond to each other just as they would in a normal business setting.

For the Signals at Scale version the initial 'conceit' was that a wealthy individual who had been dramatically successful providing financial and business information via proprietary terminals ('the Principal') was now setting up a new enterprise aimed at providing general information products and services to very-high-net-worth individuals, backed by \$8 billion USD in investment. The participants were members of a strategic advisory group assembled to provide guidance in directing that investment toward developing a new class of exceptionally valuable information products and services that justified an exceptionally high subscription fee. Midway through the exercise an intervention was made describing a new product release by a tech platform that offered the same essential capabilities for around one thousandth of the subscription fee of the principal's product. AI was not explicitly featured in the conceit or instructions or required in the discussion.

This tabletop exercise opened the discussions at each summit and lasted approximately 2 hours. The summit organizers merely observed and took notes, intervening only to answer questions about the process or to encourage the efficiency of the conversation.

Results

The following descriptions are drawn from all four summits, based on an 'affinity mapping' of the detailed notes recorded by the organizers. Participants in all summits found this exercise to be difficult and uncomfortable and generally concluded that the discussions had failed to identify a coherent product strategy. Nonetheless, most participants felt that the exercise was successful in its objective of quickly expanding the 'Overton window' of the group regarding information products, services and investments under conditions of radical abundance.

A portion of the discussion was spent on information-adjacent or even non-informational offerings that were considered to have increased value under conditions of radical information abundance. These might be described as 'limbic brain' or 'social information' products and services, and focused on factors like power, influence, relative status, exclusivity or managing the perceptions of others ("looking smart" or "being cool"). They included information about others, access to others, inclusion in exclusive groupings and unlocking participation in 'taste making' and other forms of social leadership. Some of this focus was directed at, essentially, self-improvement via increased cultural and intellectual capital ("improving cultural competitiveness") whereas some was directed at naked social manoeuvring –

possibly under conditions of “cultural inflation” as AI empowered more and more people. Comparisons were made with Medici courtiers and with the World Economic Forum (“an always-on Davos”). While this framing was clearly influenced by the conceit, it was also infused with a keen sense of the power dynamics inherent in information, and of how those dynamics might evolve to accommodate informational superpowers.

A portion of the discussion was spent on products centered on the provision of high-quality proprietary or non-public information. Variants of this included: the commissioning of new information specifically matched to the needs and wants of individual customers or small groups of customers; the offering of bounties for especially valuable non-public information; services that provided human-reported ‘facts-on-demand’ (“DoorDash for facts”); and services that provided proprietary information available only from human engagement with the real world (“sitting outside a congressman’s office”). A similar theme emerged in response to distrust of legacy media’s coverage of events: offering information about “what is really going on”.

A portion of the discussion focused on, essentially, the customer’s “ownership” of their own information sources and consumption. This was sometimes described in terms of loyalty to the customer, sometimes as “cognitive or informational sovereignty”, sometimes in terms of privacy or security and sometimes even as “freedom from fear”. Much of this discussion seemed to be motivated by valuing an escape from “surveillance” of information consumption, even to the point of becoming “digitally invisible”, and an awareness of the increased risks of “digital serfdom” as AI enabled more sophisticated and personalized uses of digital surveillance. An interesting discussion emerged at several of the summits around an inherent contradiction on this topic that might be developed into a potential product – a combination of full digital privacy for the customer while also providing powerful monitoring of the activities of others – effectively “privacy for me, but not for you”.

A portion of the discussion considered products and services that were essentially tools for strategic decision-making based on best-available information under conditions of uncertainty. These tools were generally intended either to decrease the cost of decision-making (in time, in stress) or increase the personal decision-making capacity of the customer, by providing “what I need to know” at the point of decision – a function likened to a royal or presidential aide, or to “Jarvis” from the *Iron Man*

universe. Another variant was in improving the quality of decision-making by providing “an unfair advantage”. Some of these were quite specific, recommending a “next best action”, or providing assessments of the likely impacts of actions – an offering described by one participant as “a personal Palantir”.

A related avenue of discussion centered on offering a predictive capability – an ability to partially see into the future – via the analysis of large amounts of high-quality information to statistically predict future events of direct relevance to customers. This ranged from business and political predictions to social predictions (“where to go and who to meet”) and even predictions about “when to go into hiding”.

Also expanding on the decision-making support concept, a cluster of discussions focused on the delegation of the customer’s decision-making authority to agents (human or AI) empowered to act on the customer’s behalf and to act in their name, and of the role of high-quality information in enabling that. This seemed to relate primarily to more mundane or routine decision-making rather than to high-stakes decisions, although “voting” was a repeated example. This discussion often examined the challenge in initially establishing customer trust in delegated decision-making and action-taking, and trust in the information upon which decisions and actions were taken.

A portion of the discussion focused on the ongoing creation and maintenance of information about the customer themselves (“deep context” or even “radical context”), for use by either human or AI services. The analogy here was a butler or a valet who built up a deep awareness of his employer’s pattern of life and then used that knowledge to provide high-quality service. This was seen as useful across different time scales, from minute-to-minute activity to activity over hours, days, weeks and years. This “awareness of my time” and the ability to bring a “long view context” into current activity was seen as especially valuable. Privacy and extreme discretion were obviously seen as crucially important here, and an ability to usefully and safely combine private personal context with other public contexts was deemed essential. A variant of this discussion focused on the maintenance of a comprehensive digital twin of the customer – a “second me”, complete with “radical context” gleaned from email, texts, family chats, health records, emotions, even always-on cameras and microphones – and then using that twin to identify weaknesses. In this discussion value might come not from information, but from identifying *gaps* in information, including inadequate education, missing knowledge,

cognitive blind spots, deficiencies arising from default “learning frameworks”, etc. Repairing those gaps, once identified, then becomes easy.

A related cluster of discussions centered around provision of, essentially, “life services” – the service-side use of the ‘deep context’ delivered by the human or AI butler or valet. In these potential products a customer’s personal or family context is combined with broader or “universal” context from society-wide information to create highly-personalized information services that improve day-to-day life: coordinating households, scheduling events, providing recommendations, managing health maintenance, complaining, contributing, representing and monitoring, etc.

This discussion sometimes focused on more deeply personal objectives – phrases like “optimizing happiness” were used. These more personal services were of two kinds: Some clustered around activities close to those of therapists or life coaches – a “wise, adult advisor” addressing quasi-philosophical questions like “how should I live?” and “how can I become a better person”. Others were ruthlessly pragmatic, focused on “performance” and a better life through optimization – an objective expressed as “personal ROI”. Participants in one summit described these contrasting services as “life fixation” and “productivity fixation”.

The “life fixation” discussion sometimes veered into a set of products and services that assumed increased value of analog and “real world” experiences as AI “inflated away” the value of informational experiences and of the cultural, social and intellectual outcomes for individuals from those experiences. These products and services – which appeared in all summits – were essentially “anti-informational”, and more connected to terms like “fun” and “joy” and “life”.

Trust as an information product or service also featured in each of the summits, usually as a ‘trust premium’ over background trust in other information sources, but these discussions seemed surprisingly unengaging. Some discussions focused on information products and services for which the need for complete trust was absolute, and in which customers were “buying certainty”. These, however, seemed like relatively niche business use cases and there was little sense that a trust premium alone would be valued sufficiently to justify an exceptional subscription cost. An interesting discussion in one summit explored a potential system of “built-in trust”, in which trust emerged from the system and in which explicit incentives for trustworthiness were not needed.

Despite the summit orientation toward positive societal outcomes from an AI-mediated information ecosystem, the tabletop exercise at several summits included discussion around 'dark' products: the intentional creation of scarcity by engineering constraints on information; intentionally 'poisoning the informational well' in order to increase the value of a 'clean' information source; the sale of deeply personal context from 'freemium' customers to marketers or political campaigns; the use of deep personalization for influence operations; and even blackmail. These brief discussions were either 'tongue in cheek' or intended as warnings of the potential actions of unrestricted actors in an AI-mediated ecosystem.

On a less concerning note, several summits also featured discussion about products and services that acted like 'court jesters' whose value was in keeping customers grounded in an environment of information abundance and extreme personalization. These products and services would challenge customers, perhaps in heretical ways, and introduce serendipity and randomness into their information space. This was sometimes described as an information experience "without guardrails" in the manner of a frank opinion from a trusted friend, and as a way for consumers to remain grounded, humble and 'real'.

Conclusions

These tabletop exercises were intended to 'loosen up' the thinking and to stretch the imagination of the participants regarding radically abundant and hyper-personalized information within an AI-mediated environment. The idea was, obviously, that each of us will become the high-net-worth customers that were the focus of the exercise, and the recipients of informational products and services that are currently only available to the very rich. This assumption was apparent in the many comparisons made with various informational servants of wealthy and powerful people, such as a royal aide, a presidential briefer or Jarvis from Iron Man. The discussions were, perhaps inevitably, sometimes infused with political overtones but the end objective was always extreme service. As one participant put it, the information ecosystem had progressed from "you get what you see" in the print era to "you get what you are recommended" in the social media era to "you get what you want" in the AI era. "You get what you want" is as good a summary of the billionaire news product as any.

The exercise did not produce a coherent premium product strategy, but its failure was analytically productive. Once the cost of producing, processing and communicating ordinary information was assumed to approach zero - and once a

platform could replicate a product at a thousandth of its price – the discussion moved away from conventional information products and toward the remaining forms of scarcity.

Those scarcities included exclusive access to primary or non-public information; deep knowledge of the customer; trusted representation of the customer's interests; superior prediction, decision support and execution; human relationships and social access; privacy and cognitive sovereignty; and embodied, social or joyful experiences that cannot be reduced to abundant digital information. The discussion therefore suggests that the value of future information services may lie less in delivering information than in improving a person's position, decisions, agency, relationships or life.

The exercise also revealed the dual-use nature of these capabilities. The same deep context and personalization that can protect and empower a consumer can be used for surveillance, manipulation, engineered scarcity and influence. Trust did not appear to command a sufficient premium as a stand-alone product; it was more compelling when embedded in a system that delivered consequential outcomes. The durable opportunity is therefore not "better content," but a trusted, context-rich and action-capable relationship with the user.

First principles requirements

Overview

The first structured discussion at all summits focused on the foundational, or 'first principles' requirements for each of five categories of stakeholders in the information ecosystem: information consumers, information producers, builders of information products, investors in the ecosystem, and society generally.

Participants were asked to focus on the most elementary needs and wants for the stakeholders discussed within an AI-mediated information ecosystem. They were explicitly requested to disregard solutions or ways of serving those needs and wants as much as possible, and to attempt to be detached and clinical in their discussions. Some examples of possible topics within each category were provided in order to 'level set' the discussion, but these were clearly described as such and did not dominate the conversations.

The arrangement of these discussion groups varied slightly between summits, due to specific timing requirements and lessons learned from earlier summits. In most discussions the consumer's requirements were prioritized, with two separate groups discussing this category. In some summits the 'builder' and 'investor' discussions were combined. The selection of participants into each discussion group was done randomly in the first summit and then using Claude and participant biographies in subsequent summits.

Results

Information consumers

The First Principles discussions at all summits produced some variant of Maslow's hierarchy, adapted for information needs and wants. This sometimes followed a relatively abstract ladder (truth, understanding, status and then self-actualization) but more often centered on more immediately pragmatic social needs. These usually began with safety and security at the base, including dependable and reliable access to information, and not being exploited, manipulated or trapped, etc. The middle of the hierarchy was typically where 'comfort' needs like trust, provenance and aspects of personalization appeared. The needs in the upper portion of the hierarchy were often social, variously described as social belonging, cultural relevance, recognition, validation and esteem, whereas the apex was where personal growth, agency and 'freedom' appeared.

Entertainment, in both active and passive forms, was frequently recognized as a fundamental requirement, often in opposition to "aspirational" uses of information. This was sometimes expressed as "fun", "joy", "delight" or "playfulness", and sometimes as comfort or as a way to consume "surprise" in a safe way.

The accuracy or factuality of information was, as expected, seen by many as a fundamental requirement, but there were also several conversations that questioned how important accuracy really was to many consumers, and whether it should sometimes be subservient to criteria like identity or utility. One discussion structured this position by categorizing trust as objective, intersubjective and subjective: "true for everyone"; "true for the tribe"; and "true for me". Much of the discussion on accuracy went beyond big words like "truth" and "trust", which were seen as somewhat philosophical, and focused instead on the need for sources of authority, for assurances of basic verification, for provenance (often using dietary metaphors) and for transparency about the influence of intermediaries like editors and platforms.

Direct consumer control over information consumption, and the *feeling* of control over information consumption, were often proposed as core requirements. This was often in the context of a 'trade-off' with effort, and the unwillingness of consumers to explicitly customize news products via settings was raised several times. This desired state was described in different conversations as "effortlessly being in the know", "control without responsibility" and control accompanied by "cognitive ease".

A related requirement that accompanied many of the discussions was for safety or "harmlessness" in information consumption – the ability to consume information without feeling exploited. This was sometimes expressed in terms of "control over the algorithm", but also often in terms of privacy or personal agency and even as consumer choice provided by portability within a competitive marketplace. The value of this requirement even featured as a potential service – payment for the avoidance of tracking and influence.

Sense-making as a fundamental requirement for information consumers was rarely discussed explicitly, even though the use of AI chatbots as sense-making tools was a frequent topic. The access to essentially 'all' available information via AI chatbots and the ability to personalize and contextualize the consumption of information via AI chatbots seemed to be taken as a given. This was usually discussed in terms of convenience and personalization, without the need to depend on or interact with the interpretations of others, and comparisons with Waymo driverless taxis were common.

Extreme personalization, especially in terms of the personal relevance of information, was overwhelmingly seen as a fundamental requirement for consumers. This was often talked about as much more than just convenience, for example as "anticipation of information needs", as information that "fits my day" or "fits with my tribe" or as addressing "what am I missing" gaps in understanding. These discussions were emphatically not about personalization as recommendations of artifacts based on clicks but were instead about the experience of information itself. This often extended into something that one participant referred to as an "aspirational algorithm" – a form of personalization that proactively took a "whole person" perspective and considered the long-term goals and wellbeing of the consumer.

A related discussion topic that partly contradicted and partly complemented the discussion on personalization was the social or societal component of consumer needs. This was a separate subject but in the consumer context it appeared as

something like the individual's need for a shared reality with their community, including as a basis for belonging, identity and shared governance ("democracy"). In one discussion this was described in terms of the intersection of two books from the 2010s – "The Filter Bubble" and "Bowling Alone", whereas in another it was discussed in more philosophical "individual vs society" terms.

Each of the summits featured some form of a discussion about AI agents as consumers of information, usually in terms of their human principals ("the humans of agents") but sometimes in terms of their own 'needs'. Agents were acknowledged as very different 'kinds' of consumers, but there was very little specific detail to these discussions. The information needs of AI agents seemed to be quite uncertain or even alien to the summit participants.

Information producers

The discussions around fundamental requirements of information producers took a broad definition of 'information producer', including legacy media, creators and influencers, governments, universities, companies, NGOs, individuals and autonomous AI agents. The discussion became surprisingly philosophical at most of the summits, with "incentive" or "motivation" seen as a primary requirement. In these conversations the originating reason ("objective function") behind the activities of information producers – whether money, power, impact or something else – was seen as the starting point.

A second fundamental requirement for information producers was stability, expressed in one discussion as "a need to exist". More specifically this requirement was often around clarity of business models, often closely connected to an ability to protect the value created by information production in some way but also connected with stable and transparent conditions created by ecosystem participants like platforms and regulators. Several discussions brought up the need for stable "accountability structures" – a clear assignment of who was responsible for what – and this view was seen as increasingly important in an agentic ecosystem in which private context would be combined with public information to create personalized experiences.

Several discussions brought up "relevance in society" as a fundamental need, using terms like "attention", "recognition" and "individual value" (interpreted by one participant as "career visibility"). This need was sometimes contested as a personal or socially-motivated desire of some individuals instead of as a fundamental need

and sometimes intersected with discussions of “influence”; however it was also positioned as an important motivator for both individuals and institutions to participate in information production.

Other fundamental requirements discussed were more immediately pragmatic. Most fundamental was the need for information producers to have continued access to supplies of primary information – the “raw materials” of information products and services. This was variously described as access to open data, the continued integrity of open data, access to situations, decision-makers or processes that generate societally relevant information, the unrestricted ability to interact with communities, etc. In several discussions the potential for a restriction of such access as AI enabled more systematic, personalized and effective use of it was identified as a risk.

A significant requirement that was intensely discussed at all summits was the need for access to the demand signal from the end-use of information, so that information producers could direct their resources efficiently and value their contribution accurately. This was proposed as direct access to users and user communities, however the intense privacy risks associated with personal intent and personal context accessible to AI were very much appreciated. Discussion generally turned to some form of feedback loop, perhaps in the form of an aggregated or otherwise anonymized demand signal, as well as to the many challenges associated with such a feedback loop – for example the units of information upon which feedback was provided.

Related to the requirement for access to the end-use demand signal was a proposed requirement for a measure of control of the end-use itself. This was often positioned in terms of the value of reputation and demonstrated trustworthiness in the provision of accurate information, and the associated need to prevent a producer’s information product from being misused, misread or intentionally distorted – with resulting consequences for the reputation value of the producer.

Another pragmatic requirement was the talent, skills and capabilities necessary to implement what was likely to become a far more technical form of information production. Behind this was an assumption that future information production was generally unlikely to be like past information production, and that there was perhaps a mismatch between the talent and skills of current information producers and future needs. These discussions were clearly tinged with the uncertainty of how the application of AI would actually play out within organizations, institutions and teams.

Finally, the need for information producers to protect the value of the information that they produce, and to capture enough of that value to justify their continued existence and participation in the ecosystem, was seen as a key requirement. This was not only discussed in terms of legacy publishers seeking compensation for the unrestricted scraping of their content, but also in terms of the feasibility of building for new AI-native information production processes. A “monetization mechanism” for information in an ecosystem in which its semantic meaning was separable from the information artifact that carried it was seen as essential. Specific proposals for what that mechanism might be were, however, absent.

Builders of information products and services

The first principles discussions at all summits identified variants of predictability, or “line-of-sight”, as a key need for builders of information products – whether in start-ups or within existing organizations. This was especially framed in terms of uncertainty about whether products, functions and capabilities might become functional within the major AI platforms, and anecdotes about projects delayed because of this kind of uncertain future relevance, or abandoned when Claude or ChatGPT became competent in the function, were common. To a lesser degree this need for visibility was also discussed in relation to regulation, security needs, protection of intellectual property, privacy and ethical norms, however it was the risk of capabilities being subsumed by the AI platforms that dominated these conversations.

A portion of these discussions narrowed in on stable, predictable marketplaces for information as being at the heart of builders’ need for visibility, because they offered at least the possibility of growth and acceleration. This was seen as applicable in both public and private spheres, with Wikipedia and ‘crypto’ being offered as examples of each. The idea of ‘free marketplaces’ for information was explicitly discussed several times, but more often from the perspective of marketplaces free from monopoly and capture by AI platforms rather than merely free from regulation or from restrictions on expression.

Standards of various kinds emerged as frequently cited needs for builders of information products. This included standards for enabling products to access data and to deliver data, standards for attribution and therefore for monetization, and even standards for access to pooled data available for collective use (“a data commons”). These discussions were sometimes accompanied by discussions about enforcement

of standards at AI scale and speed, including via an “agentic police force” that audited compliance and conducted investigations into misuse.

A special category of standards often discussed was standards around interoperability or portability of data of various kinds – largely as a way of avoiding builders being “locked out” from access to consumers (or agents) who were “locked in”, or being similarly “locked out” from access to the flows or “pipes” that carried data, facts, context, or other information.

There were also discussions around distribution and the difficulty of identifying and reaching potential customers in a heavily intermediated agentic internet, and around the uncertainty in business models centered on complex “intention” (or “intention plus context”) in contrast to the ‘attention’ focus of the existing information ecosystem.

Investors

Discussion of the first principles needs of investors in the AI-mediated information ecosystem centered on the necessity of a ‘big picture’ strategic vision of the contours and components of the emerging ecosystem – a ‘point-of-view’ that carried enough conviction and enough potential for impact to warrant investment. Discussion of this requirement also featured the need for talent that could formulate and articulate such a vision, as well as for more future-facing conversations that might help to develop it. It was mentioned several times that developing such a credible vision would likely require participation from a much more diverse community of participants than assembled for the Signals at Scale summits.

A second fundamental need of investors in the emerging ecosystem was identified as capital that was both patient and risk-tolerant. This led to conversation about the different kinds of capital, each with different incentive structures, including investment centered on access to innovation (often exited as ‘acquihires’), investment centered on societal impact (such as Effective Altruism investments) and investment centered on optionality and speed of adaptation. Conversely, the value of time for an investment to explore a problem, as opposed to merely capital and resources, was also discussed and appreciated.

A special objective of investment that was intensely discussed at one summit was investment that sought to create an entirely new marketplace, instead of merely succeeding within an existing marketplace. This was positioned as a key difference

between “California VC” and “New York VC”, with the latter seen as more about succeeding within the status quo. This discussion pointed out the very specific incentive structures that often lay behind investment in entirely new marketplaces.

Like both information producers and builders of information products, investors in the AI-mediated information ecosystem valued “stability, confidence and clarity”, including the protection of intellectual property and defensibility of profits, regulatory clarity, etc. But unlike those similar discussions the investment-focused conversations also explored the need for a much better understanding of risk in the AI economy generally, and the need for new ways to value or assess projects – including through “outsourced quality assessments” that insulated investors from the need to make fine-grained assessments themselves (termed “proof of value” by one participant). At the highest level a basic understanding of what might constitute “public goods” and what might constitute “private goods” in the emerging AI-mediated information ecosystem was considered a requirement.

A significant requirement, valued by investors, that was perceived as lacking in the existing information ecosystem was a body of continuing applied research on AI-mediated information. This included data and analysis around the application of AI to information (especially the “human-to-AI interface”), data from the foundation model companies around demand signals, knowledge of the ecosystem “upstream of AI” (including training data, chip availability, data centers, energy access, etc), detailed information about model control, ownership and governance, and foundations for due diligence of technology, people and relationships.

At a more practical level the investment discussions identified the development and maintenance of situational awareness amid an accelerating rate of change as a key requirement, especially in relation to identification of and access to “deal flow”. The identification of talent, both for investment and as investees, was an accompanying need.

Finally, as in the builders’ discussions, several of the investment discussions brought up “protection from monopoly” as a need, but with a significant difference. Like builders, investors were seen as requiring protection against seeing the offerings of their investments being incorporated into the capabilities of general-purpose AI models or agentic systems, but they were also seen as requiring moats or competitive advantages (“unfair advantage”) of their investments. This sometimes came up as a tangible contradiction, for example in discussions around standards –

which can both protect against dominance by the AI platforms, but which can also undermine the moats of start-ups.

Society

Despite the goal for the Signals at Scale summits of concentrating on the positive potential of AI in the information ecosystem, many of the discussions for societal needs within the AI-mediated information ecosystem focused on the prevention of harms, and on the need for various mechanisms that might mitigate the risk of societal harms.

Protection against monopoly behavior by AI platforms, and especially against the lock-in of consumer context by platforms, was a major theme of these discussions. The requirements that were considered included antitrust legislation (and its enforcement), regulated interoperability and the establishment of decentralized or public service alternatives to platforms. The trade-off between protection from monopoly on one hand and the need for AI platforms to recoup their investment on the other was discussed, and potential solutions such as the sunseting of moats after recoupment were proposed. This class of requirements was often described in terms of pluralism of information sources, or even as “a new pluralism” in recognition of the differences with historical media pluralism.

A societal need for “accountability infrastructure” was discussed in different ways, including through transparency requirements, confidential reporting requirements to regulators, standardized metrics, sampling programs, guardrail enforcement, etc. These were sometimes accompanied by suggestions for dedicated “monitoring infrastructure” or “observatories”, as well as by concerns about potential abuse or the emergence of new ‘gatekeeping’ structures. Infrastructure and processes that protect consumers from violations of the privacy of accumulated context from ongoing interactions with AI were seen as especially required.

Positive needs for society in an AI-mediated information ecosystem included many commonly discussed prospective benefits such as truly universal access to information, new means for societal accountability via ‘watchdog’ functions at unprecedented scale and new relevance of information in the lived lives of consumers (guidance in voting was raised several times as an example).

Accompanying these positive requirements, however, was a clear need for incentives for the production of new civic and societally useful information – a variant of a similar

discussion in the information producer conversations. This version of that requirement was more focused on information as a public good, and also more focused on potential roles for governments in funding and sustaining civic information products or data sources.

Finally the concept of “sovereignty” was discussed at different levels of application, along with the requirement to protect it. The “cognitive sovereignty” of the individual was seen as most in need of protection, including from risks of loss of agency through undue persuasion as well as from the erosion of capability through the “cognitive offloading” of information consumption and use to AI. The “cognitive sovereignty” of the nation-state was also discussed, with requirements largely centered on security (including of continued access to AI and thus to AI-mediated information), on protection of identity, culture and language and even on the risks of the gradual dissolution of societies as their citizens enter AI-mediated informational “bubbles of one”. An interesting conversation also emerged around the “sovereignty” of groups or communities distinct from their individuals or enclosing nation-states, with potential requirements around the self-formation and self-maintenance of communities as purely social (rather than AI-mediated) phenomena.

Conclusions

The first-principles discussions describe a multi-sided ecosystem whose participants have related but frequently conflicting requirements.

Consumers seek safety, relevance, cognitive ease, agency, privacy, enjoyment and belonging. They want highly personalized information without the burden of constant configuration, and control without excessive responsibility. They also need some connection to a shared reality, creating an unresolved tension between extreme personalization and social cohesion.

Information producers need motivation, access to primary inputs, visibility into demand, protection against misuse and a credible mechanism for capturing enough of the value they create to justify their participation in the ecosystem.

Builders and investors need predictable markets, interoperability, standards, defensibility and protection from platform enclosure.

Society needs pluralism, public-interest information, accountability, universal access and protection of individual, community and national cognitive sovereignty.

The central conclusion is that the ecosystem cannot be designed around a single objective such as accuracy, consumer utility or commercial efficiency. It must negotiate persistent tensions between personalization and shared knowledge, privacy and useful context, open standards and private moats, frictionless access and sustainable production, and individual agency and collective welfare. These tensions and negotiations are not peripheral governance questions; they are part of the core product and market architecture.

Jobs-to-be-done

Overview

The second structured discussion at all summits focused on the specific functions, or 'jobs-to-be-done', for each of five parts of the value chain of the information ecosystem: production; processing; distribution; consumption experience; and coordination and governance. These were positioned as discussions around value creation within the emerging ecosystem.

Participants were asked to focus on the most elementary functions that must be available in the AI-mediated information ecosystem in order to satisfy the requirements discussed in the first session. Participants were asked not to focus on specific products or services, but instead on the larger scale functions that might be needed to fulfill stakeholder needs and wants. As with the first discussion, some examples of possible topics within each part of the value chain were provided in order to 'level set' the discussion, but again these were clearly described as such and did not dominate the conversations.

In some summits, due to timing or lessons learned, the distribution and consumer experience discussions were combined. Again the selection of participants into each discussion group was done randomly in the first summit and then using Claude and participant biographies in subsequent summits.

Results

Production

A significant cluster of discussion centered on the use of AI to scale the gathering and collection of societally relevant information by many orders of magnitude. The emergence of this ‘aggregation’ function, and its immense scale relative to the scale of current information gathering, was assumed to be a given. This AI information gathering, applied to all digitally accessible sources, was seen as providing the raw material for “hyper-personalization” of information experiences. It would, it was generally assumed, eventually include essentially every token produced or exposed in public in any form – whether in text, speech, imagery, video, and data – including from governments, institutions, public figures, social media, databases and even microphones, cameras and other sensors, perhaps eventually mounted on drones, humanoid robots or other forms of AI embodiment.

In more specific terms these discussions debated the manner in which this ‘universe’ of information would be collected and made accessible to AI. Options that were considered include centralized platforms, distributed information gathering facilitated by standards, and a “brokerage layer” in which specialist information might be bought from domain ‘gatherers’, subjected to quality vetting and then aggregated and resold for general agentic use. There were several proposals for a public service component of information gathering, and an interesting discussion around the need for “trusted execution” in the gathering of information.

In essentially all these conversations it was acknowledged that certain kinds of information would likely remain impervious to AI gathering – information obtained via relationships, information obtained in private settings, etc – but these conversations also generally anticipated (or included examples of) AI interviewing sources, commissioning freelance reporters and engaging in other human-like information gathering activity.

A variant of these debates centered specifically on the difficulty of capturing (or even recognizing) ‘deep context’ – the rich and nuanced historical background of individuals, communities, institutions, stories or situations, often involving subtext, social undercurrents or “tribal knowledge”. In one discussion it was pointed out that if humans could not easily articulate such ‘deep context’ then it would likely be difficult to capture it in any conceivable schema or ontology.

The discussions of information gathering quickly included questions about the units in which information might be gathered, and the need for some structure of that information in order to accommodate association of provenance, corrections, standardized distribution, trading and persistence in a historical record. The advantages and disadvantages of a common schema or ontology for societal information collected at AI scale were debated, with general agreement that this could introduce a significant bottleneck, bias or 'gatekeeping' at the very top of the AI information funnel.

Part of this discussion illustrated the challenge of building bidirectional linking of information between the point of collection and the point of its use in generating an experience for a consumer. This was seen as not merely important functionality for monetization, but also for efficient use of information collection resources (the "demand signal"), for conditional use of proprietary information and for enablement of correction in the event of error. As in the ontology discussion this linkage was also seen as potentially dangerous from a security and privacy perspective, because it necessarily enabled the tracking of which consumer was consuming which information.

An important component of the discussion on gathering information was the need for infrastructure, and not merely the basic infrastructure of information capture and aggregation. The need for infrastructure to enable the uncorrupted tracking of provenance was discussed ("provenance by design"), including in terms of existing techniques such as watermarking (e.g. SynthID) or standards for the cryptographic tracking of provenance (e.g. C2PA) but also in terms of hypothetical future techniques. Infrastructure for working with and managing information stored in AI-scale units (potentially a "Fact Management System", or FMS) was another repeated topic.

Infrastructure that enabled the protection of the value of collected information by "claiming your work at the point of creation" was discussed in particular detail, including in the context of new units. The potential for some method of applying intellectual property protection and assignment "at the fact level" was discussed, possibly providing protection for a relatively short time following the collection of rare but valuable facts. Infrastructure useful for the conversion of rights into revenue was also discussed but generally defaulted to "micropayments" in some form. Most discussion around monetization infrastructure included some concept of tracking of

use, and relatively little discussion focused on untracked access to information flows via subscriptions.

Processing

The discussions of processing jobs-to-be-done all followed a similar template that first required the extraction of meaning from raw information, and then undertook steps to make that meaning trustworthy, available and valuable.

These discussions often began with a topic that partly overlapped with parts of the discussions on information production: the need for a process of “normalization” of information into “atomic units” that provided structure and reference points. This was variously described as “unlocking information”, “de-verticalizing information” and “separating signal from noise”, and a formal system for identifying relevance and “signal” was seen as unavoidable.

The processes that would do this were also explored in some detail, including consideration of compliance infrastructure for legal, regulatory, privacy and ethical compliance, and of enrichment infrastructure that provided cleaning processes, first-line quality checks, and the assignment of reference IDs and other metadata – possibly including metadata about cost or expense of collection and perhaps pricing. Other enrichment topics included the attachment of potential restrictions on the use of the information ‘atom’, a ‘certification’ of the level of processing or confidence (potentially with associated indemnity) and the associated acceptance of responsibility.

A substantial part of the discussions on processing jobs-to-be-done focused on verification processes. This included provenance or “chain of custody” verification and the assignment of reputation scores from provenance, as well as elementary cross-checks and more substantial automated fact-checks. The central challenge of enabling verification at AI scale (i.e. well beyond a scale that would be feasible for human fact-checkers) came up frequently, with suggestions around statistical methods (i.e. sampling), just-in-time fact-checking prior to use, or exotic forms of “truth tech” appearing in the discussions. An important thread in these discussions was the need to maintain independence between information collection processes and information verification processes – quite possibly extending to different systems, different organizations, and even different sectors of the information industry as AI enabled separation and specialization became feasible.

A third substantial topic in these discussions was the need to optimize collected information for the production of experiences, and especially for the production of hyper-personalized experiences. Two of these discussions got into considerable detail on this job-to-be-done, focussing in particular on the challenge of integrating information with both public and private context – i.e. the background informational context (whether universal or at the level of communities, situations or stories) and an individual's personal context as accumulated from prior interactions, etc. There was debate about whether this integration was a processing task or a task that could be undertaken in the creation of a specific experience, but the need to provide information optimized for this function was generally agreed.

Lastly a series of discussions focused on jobs that were variants of quite common functions in any large-scale information processing operation. One discussed the need to optimize for information storage and retrieval at vast scale, a non-trivial challenge that was made more complex by the need to index or organize informational units in ways that enabled efficient access at useful speed. It was pointed out by one participant that traditional information producers often have archives that are, relatively speaking, quite small compared with the potential for datastores of AI-produced information, and that the need to provide and maintain context for the assembly of information experiences might require sophisticated explicit indexing or new forms of vectorization. Another set of discussions narrowed in on the ordinary mechanics of information governance, including controlling the access to information for different uses and functions.

Distribution

Discussion of jobs-to-be-done in the distribution of information usually began with discovery, including questions about what discovery meant in an AI-mediated information ecosystem. The starting point was active discovery via interaction with chatbots, or even directly with search, but quickly included passive discovery via personalization of various kinds – including personally-directed passive experiences like bulletins from ChatGPT Pulse or Google's Information Agents. These information-oriented approaches were seen as far more relevant for an AI-mediated environment than the kinds of content-oriented recommender systems that dominated the social media era, and could potentially extend or evolve into "predictive informational experiences" based on a consumer's accumulated context and intent.

'Negotiation' emerged as a significant job-to-be-done in distribution, including negotiation between consumer agents and producer agents around potential value of information to an experience being assembled for a consumer and the cost of that information, and negotiation within agentic processes around the balance of experiences (for example the balance between truthful, factual information and entertainment-oriented information). One discussion imagined an 'agent' operating in a manner similar to that of a talent 'agent', mediating between provider and consumer, and one participant remarked that in an agentic ecosystem driven by consumer intent, "distribution IS negotiation".

A more familiar function of the distribution layer that was raised was the filtering of information; however, this was discussed much more in terms of active 'protection' instead of in terms of passive restriction, and these discussions also veered close to the concept of 'negotiation'. In this interpretation the primary 'job-to-be-done' was the protection of the scarce attention of the consumer from those who sought to attract, waste or otherwise commandeer it. Closely related to this was the task of protecting the 'intent' of the consumer, ensuring that the consumer got what she or he desired instead of what others in the ecosystem (marketers, influence-seekers) desired for them. As in the discovery discussions, the function of 'filtering' seemed to be closely related to the task of predicting the needs of the consumer from the context of their prior intent and current situation.

A set of jobs-to-be-done in the distribution layer involved the visibility that this layer would have over aggregated consumption by many consumers. Some of these tasks were straightforward – the potential assembly of aggregated, anonymized demand signals, the identification of new opportunities for information production, or the potential measurement of aggregated impact – but some were related to solving collective action problems in information gathering. An example provided in one discussion was a 'Groupon-style' approach to investigative journalism, in which expensive information gathering could be triggered upon reaching some threshold number of commitments or intent signals from aggregated consumers.

As with the information production discussions, the discussions on distribution identified a set of relatively obvious jobs necessitated by the rest of the ecosystem and enabled by the potential visibility from distribution into both production and consumption activity. This included the facilitation of monetization transactions, including identifying consumers, consumption and payment events, producing the

pricing and analytics to enable payment processing, distinguishing between different tiers of consumption, etc., as well as perhaps the application of informational 'nutrition labels', and formatting information for different forms of consumption (individual humans, enterprises, agentic consumption, etc). These tasks suggest that the complexity of information distribution may be such that it may become a specialized brokerage or 'middleman' function, separate from AI platforms.

Lastly, a set of distribution jobs were discussed that seemed to require an 'overriding' or 'bypassing' of personalization in order to provide broad distribution of specific information. An obvious example of this would be emergency 'broadcasts', but it could also involve certain public service messages around elections or other critical themes. This 'overriding' or 'bypassing' in distribution could also work in the other direction, enforcing protection against broadly targeted influence operations, cybersecurity attacks, 'context poisoning' or other forms of manipulation.

Experience

Unsurprisingly, the jobs-to-be-done discussed for the consumer experience of information were closely integrated with the jobs described in the distribution layer, and even with the processing and production layers. The consumption experience is the originator of demand signals (intent) and the consumer of attribution and other metadata from the production layer. It both produces and consumes context, and it is the controller of 'personalization' and of access to the chain of information necessary to produce it.

The job of personalization was seen in all discussions as central to AI-mediated consumption experience, and the definition of 'personalization' was generally assumed to be broad enough to include personalized behavior as well as personalized information. One participant described the experience of AI-mediated consumption as a transformation from "how you need to interact with it" to "how it needs to interact with you". This extreme flexibility of the consumption experience, and its reliance on diverse forms of 'intent', was referred to in one discussion as "situational consumption", and applied as much to enterprise or agentic consumption as to individual human consumption.

Somewhat in tension with this "situational consumption" is the provision of consumer 'agency' – which was frequently identified as a core function of consumption experiences. This was seen in both simple terms – transparency, the ability to opt in or out of certain functionality – and in more nuanced terms, including the value to

consumers of feeling in control at all times (“the perception of agency”). The job of providing consumer agency was further complicated by a distinction between immediate agency – “I get what I want when I want it” and more diffuse agency oriented toward the longer-term objectives of the consumer. Several discussions touched on providing a means to enable shared narratives or shared understanding within communities through the provision of information about what others in the community were consuming – possibly by voluntarily ceding control of part of the consumption experience to others.

A set of jobs-to-be-done around the “experience of trust” was discussed several times, including via the communication of verification processes, the communication of reputation from attribution sources, the accumulation of “trustworthiness” in the experience over time, and the provision of access to provenance chains or decision-making criteria. A closely-related job centered on privacy and the protection of the consumer’s intent and context, which would likely become increasingly intimate as the benefits of such intimacy became available in the experience (a phenomenon that has become increasingly common over the last several years in chatbots and character bots, and which was assumed likely to extend to informational consumption). The possibility that in the context of AI-mediated information consumption the idea of ‘trust’ might become much harder and more safety-related than it is in static media consumption was raised, which might make the job of trust maintenance both more necessary and more difficult.

One discussion took place about whether advertising was a job-to-be-done here, and if so then what form it might take, however this mainly served to draw attention to how little advertising (the use of AI-mediated consumption experiences to influence consumers for commercial or political objectives) was discussed. Several participants explicitly mentioned the distasteful nature of discussing intentional influence or persuasion in an AI context. A loosely related discussion was around the potentially unavoidable need to collect or measure feedback from information consumption, assessing whether the information provided value or had an impact.

Coordination and governance

The discussions on coordination and governance jobs-to-be-done in an AI-mediated information ecosystem typically began with an attempt to list the various kinds of coordination that might apply – lists that quickly became long and complex. Many of these kinds of coordination overlapped or intersected, and almost all were

accompanied by identification of specific early or emerging examples. Some forms that were discussed included:

- Coordination through technical standards and other forms of technical cooperation, especially around interoperability (e.g. MCP and other agentic protocols).
- Economic coordination, ranging from the coordination provided by open and transparent markets enabled by adequate anti-monopoly enforcement on one hand, to constraints or prescriptions on the definitions of value, units of trade, conditions of trade and pricing on the other.
- Regulation as coordination, directly applying political decisions to AI-mediated flows of societal information (e.g. The EU AI Act).
- Platform dominance as coordination (e.g. Google's dominance of search and the resulting ecosystem around 'search engine optimization').
- Coordination via legal action and legal precedent (e.g. New York Times lawsuit against OpenAI).
- Third-party 'voluntary' coordination, potentially with various degrees of 'voluntariness' (e.g. NewsGuard's provision of editorially assessed 'trust lists').
- Coordination via international cooperation (e.g. A CERN or ICANN for AI-mediated information flows, or an information variant of emerging AI Safety infrastructure).
- Benchmarking and transparency as coordination (e.g. the benchmarking of language models and the radical transparency of AI providers like Anthropic).
- Coordination via funding or purchasing choices, for example by national governments.
- Cultural coordination and cultural sovereignty initiatives (e.g. Danish national foundation models, and the role of the National Library of Norway's NCC corpus).
- Coordination via public ownership or public service (e.g. the partial coordination of the UK's media landscape by the presence of the BBC).
- Coordination through provision of tools (e.g. PyTorch or CUDA).
- Coordination via restrictions on hardware for inference (an "OPEC for GPUs"), with implications far beyond the information ecosystem.
- Coordination around the creation and protection of intellectual property, including traditional copyright and perhaps new kinds of protection. (e.g. the SPUR coalition).

- Coordination on the deliberate imposition of friction on the flow of certain kinds of information through the AI-mediated ecosystem (e.g. as with CSAM content).
- Coordination of knowledge pools or grounding data, including protection against the “poisoning of knowledge pools” and possibly against “knowledge inflation” by acting as a “Federal Reserve of grounding data” (e.g. the new arXiv endorsement system and “one-strike penalty”).
- Coordination at the level of distinguishing between individual humans and AI (e.g. WorldID’s “proof of personhood”).
- The potential coordination of human information work by AI, including in the production and processing of information (e.g. Early AI use of TaskRabbit contractors).

An interesting discussion examined more high-level interpretations of coordination of the AI-mediated information ecosystem that might operate much like the scientific method or like the pluralistic journalism industry as a whole – the potential for coordination through feedback loops. These forms of coordination would be dependent on methods for conducting repeatable evaluation of information, and on independent systems of evaluation that applied those methods. These feedback-based mechanisms were repeatedly described as extremely challenging to develop, but also as having the advantages of operating “beyond the influence of” specific groups or organizations and thus bypassing many of the political challenges that might be present in other forms of information coordination. These discussions also included several references to the “good faith” requirements for the effective functioning of these kinds of feedback-based systems in the pre-AI information ecosystem.

Several broad discussions occurred around the entire rationale for coordination at all – essentially about the fundamental job-to-be-done. Some participants saw this job as essentially negotiating between the needs and wants of individual consumers and the needs and wants of society (including via the restraint of technology and information providers), and many considered an uncoordinated AI-mediated information ecosystem to be inevitably “a mess”. At the heart of these discussions seemed to be the question of when ‘coordination’ became “coercion, control and censorship”. Uncertainty around the “measure of value” that coordination was pursuing was common, with one participant remarking that “value can’t just be compliance to protocol”. Similar uncertainty appeared around the question of what

constituted “abuse”. This was clearly a very “values-laden” discussion that centered at least partly on control of potential new sources of power in an AI-mediated information ecosystem.

All of the jobs-to-be-done discussions on coordination attempted to move beyond these more abstract and hypothetical challenges and return to more tangible and practical tasks. There was broad agreement on the benefits and desirability of interoperability in general and of interoperability of personal context specifically. The desirability for personal context to be “portable” between suppliers of intelligence was unanimous, and most participants also expressed a fundamental need for privacy-protecting handling of personal context. One possible solution discussed at two summits was potential coordination or negotiation between small on-device models and larger off-device models around private context, which might be especially valuable in relation to intimate context like emotional or neural information. Other areas of broad agreement included the desirability of open data and of the need to “protect open data from closing as soon as it becomes useful using AI”. The enablement of extreme personalization and the resulting increased relevance of information to consumers was also seen as a universal good, despite some acknowledgment of the benefits of shared societal knowledge. The potential for coordination to be a substantial accelerant of societal benefits from AI-mediated information flows was a theme in these discussions at all summits.

One particularly nuanced coordination job discussed at one summit was the need to coordinate (or “manage”) the impersonation of individuals by AI. This was discussed in terms of ‘management’ rather than ‘prohibition’ because it quickly became clear that the lines between ‘delegation’ to agents and ‘impersonation’ by agents were surprisingly blurred. For example, an agent with “explicit authority” to obtain or communicate information on behalf of its human principal might be said to impersonate its principal – an uncertainty that is exacerbated when the agent’s authority becomes implicit. This was seen as a larger issue for the AI economy than just for information, but could have special significance for information that might include “new kinds of speech” via agents.

Several discussions focused on the mechanics of possible regulation of information flows. Focusing on regulation of sources and outputs, rather than on use and method, was generally seen as more practical. The impracticality of applying human-scale methods to AI-scale processes was discussed, with the failure of the labeling

requirements of the EU AI Act used repeatedly as an example. The importance of clearly defined thresholds was discussed, not only as politically desirable but also as a technical necessity. Some expressions of the difficulty in applying substantial regulation to products and services that were already in daily use by hundreds of millions of people were countered by examples of rapid regulatory course-changes in China and Russia, and also during the COVID pandemic.

Finally, a fascinating set of discussions emerged around the enforcement of coordination in an agentic internet. Proposals for an “agentic police force” that might patrol the agentic internet seeking out noncompliance in other agents on information flow restrictions quickly transitioned into concerns about hyper-authoritarian dystopias (an “agentic Stasi”). The prospect of self-dealing by agents was raised – i.e. agents biasing information flowing through them for their own purposes – and questions about “how to punish an agent” were raised. In one discussion this escalated into a broader discussion about potential authoritarian abuse of the coordination layer of the AI-mediated information ecosystem, with one participant questioning whether “there was a difference between AI sovereignty and AI nationalism”.

Conclusions

The jobs-to-be-done discussion was essentially about replacing the legacy “produce, package and distribute” model of societal information flows with a modular, liquid, machine-oriented and bidirectional model, operating at vastly greater scale.

The forward path begins with observing or collecting information, representing it in usable units, attaching provenance and rights, verifying and enriching it, retrieving and combining it with relevant context, and assembling a personalized experience. It then extends beyond consumption into recommendations, negotiation, delegated decisions and action.

The reverse path is equally important. Consumption generates intent, demand, usage, impact, correction and payment signals that must flow back toward processors and producers. In this architecture, distribution becomes negotiation: agents may negotiate over access, price, permitted use, evidentiary quality and the balance between factual, entertaining and socially useful information. The experience layer becomes the principal control surface because it holds the consumer relationship, observes intent and determines which information is used and valued.

Coordination and governance are therefore operational functions rather than external constraints. Interoperability, identity, authorization, privacy, corrections, market rules, emergency overrides and agent accountability must work at machine speed and ecosystem scale. One of the clearest implications is that collection, verification, distribution and coordination may become distinct specialisms rather than functions bundled inside a single information institution.

Investment categories

Overview

The third structured discussion at all summits focused on the broad categories for potential investment. These categories were chosen by the organizers for each summit, informed by the earlier discussions, but gravitated toward the same four categories: the representation layer (the units in which societal information was captured, traded, accessed and used); the integrity layer (the processes by which information is made trustworthy and valuable); the interpretation and discovery layer (the ways in which information is merchandised and consumed); and the reward layer (the ways in which information producers and processors are incentivized to add value to the ecosystem). These were positioned as discussions around value capture within the emerging ecosystem, and participants were informed that they should subdivide their assigned category in any way that might be useful.

Participants were asked to focus particularly on the economic dynamics within the category – the control points, network effects, switching costs, other moats, market failure, etc. They were also requested to pay special attention to infrastructural aspects of their category, including such infrastructure as standards, common technology platforms, regulation, etc. including infrastructure that might address potential market failure. Participants were also asked to imagine what a ‘category winner’ in the manner of Google’s dominance of search might look like in their category.

In one summit, the interpretation & discovery category was split, and again the selection of participants into each discussion group was done randomly in the first summit and then using Claude and participant biographies in subsequent summits.

Results

The Representation Layer

Discussions on value within an AI-mediated information ecosystem were built on an assumption of a tradable unit of information as the store or communicator of that value. There was near-unanimous agreement that for most of the ecosystem that unit would not be a stand-alone discrete artifact like a text article, a video or an audio recording, because such artifacts were 'designed' primarily for human production, human processing and human consumption at human speeds and at human scale. It was pointed out that even in the pre-AI distributed information ecosystem of social media the valuable structures have not been artifacts or small curations of artifacts, but larger structures like indexes and vast searchable aggregations, often of information generated by consumers.

In trying to pinpoint the potential sources of specific value in the representation of information within an AI-mediated information ecosystem, participants systematically examined several options. The most basic was the "clear act of collection or origination" and whether that had value independently from its representation – essentially seeing value as attached to an action rather than to a static representation. On representation the starting point was seen in most discussions as 'facts', and the discussion focused on whether and under what conditions facts had value, and how the value of facts might be priced. This led to recency as a source of value for facts – in effect a scarcity imposed by delayed propagation in the ecosystem – and to the prospect of intellectual property protection being applied to some classes of facts. Some participants asked "whether intellectual property might cease to exist" whereas others suggested that the object of protection of intellectual property constraints might change from artifacts to usage rights for some kinds of represented facts, similar to trade secrets.

Representation of information – ontologies, schemas, knowledge graphs, catalogs, vector representations, causal models, etc. – was most commonly seen as standards, and especially standards as public infrastructure, rather than as systems or centralized records. This standards perspective was seen as necessary for adoption and for a pluralistic information environment inhabited by diverse private and public participants, but was also seen as an inhibitor of the development of competitive advantage ("moats") around representation and therefore of value capture in the representation layer. It was pointed out in one discussion that 'public good'

information that was already difficult to monetize may become even more difficult to monetize when represented in atomized or structured form.

A clear theme that emerged at all of the summits was that the value created in representation was likely to be not in the representation itself, but in the surrounding activities – the selection of data for representation, the gathering or creation of that data, the categorization of it, the structuring or “atomizing” of it into the representation scheme, maintenance and quality control of information within that representation scheme, making it available for use from the scheme and delivering it to market (MCP servers, LLMs, agentic tools, etc) via some kind of merchandising approach, etc.

One discussion explored the relationship between the representation of information and the demand for information by aggregated consumers at the experience layer, considering that the aggregation of demand at different levels of granularity might be part of the representation layer – essentially representing demand instead of representing supply, or valuing questions rather than valuing answers. In this interpretation value might be created primarily by aggregating demand (including potentially from a diversity of sources separate from the AI platforms) and only secondarily by satisfying that demand. This could work in various ways, including the “Groupon-like” approach raised in a separate discussion in which the creation of new information (for example an investigation) is driven by registered demand from many consumers.

The representation scheme itself was seen as a hard problem with no obvious solution, although also likely to become stable and ‘solved’ eventually (“What is it in 100 years?”). In an attempt to pin it down several conversations considered existing forms of structured information, such as clinical trial information, scientific publishing generally and even Google Street View, and concluded that a key distinction lay in whether the “world” being represented was considered to be closed and self-consistent (like scientific or spatial information) or open and inherently inconsistent. The latter interpretation might orient value toward reputation and provenance built on a foundation of “fact claims”, whereas the former might orient toward completeness and internal connections built on a foundation of “facts”, with value emerging essentially from the network effects of their interconnection.

The Integrity Layer

The central theme in discussions about investment in the integrity layer was, unsurprisingly, trust and the value of trust. The context for these discussions was an acceptance that the ‘trust-generating’ processes in the existing information ecosystem were broken, producing a “market failure” in which public demand for “trustable” information was not being met.

In these discussions two ‘grand themes’ emerged around trust-related value – one around essentially a ‘social’ basis of trust (“trust in the reputations of people and brands”, whether in isolation or in aggregate) and one around a more ‘science-like’ basis of trust (focused on direct evidence and logical inference).

On the ‘social basis’ theme one proposed store of value was ‘trustworthiness’ scores for individuals and agents, based on reputations created automatically by comparing previous fact claims with settled facts. This was extended in one proposal to enabling ‘trustworthy entities’ to contribute to the trustworthiness scores of others – an idea that seemed to have significant ethical challenges, but which is also the basis of generally well-trusted information production systems in the current ecosystem, such as Wikipedia. A potentially expanded role for ‘bridging algorithms’ that allocated trustworthiness scores by assessing the diversity of people or agents supporting or rejecting a fact claim – an approach taken by ‘community notes’ on X and Facebook – was proposed. Several discussions asked whether it might be possible to build network effects around a socially interactive interpretation of trust, for example by assuming the trustworthiness of entities (people, agents, institutions) that were trusted by others. One proposal suggested adding a small financial cost to agents communicating publicly, in order to raise the cost and thus the quality, of AI-mediated speech at scale while also generating value.

On the ‘science-like’ theme there was a general assumption that evidence in different forms could become much more accessible, observable and valuable to information consumers in an AI-mediated information ecosystem. In one discussion it was observed that much of current journalism, for example, intentionally isolated the consumer from the underlying evidence (the documents, the transcripts, the data, the validity of the process, etc) because of the scarcity of distribution bandwidth and of human attention, and that most of that evidence could now be made available to a consumer or, more likely, their agent – unlocking new value from it. A related discussion explored the analysis of information, and seemed to conclude that AI

could make interpretations of facts more trusted and thus valuable by incorporating personal context into interpretation, by clarifying (for example by enhancing sensemaking or by separating opinion from facts/logic) as well as by enabling easier comparisons between interpretations.

Each discussion around value from 'trust' invariably extended into a conversation about 'truth'. This seemed uncomfortable in an investment-oriented discussion, with many participants seeing more philosophical 'what is truth?' questions as difficult and impractical. There was general agreement that 'accuracy' – the fidelity of the report to the information source – was a more immediately valuable objective than 'truth', and that both 'provenance' and 'authenticity' were not valued the same as 'accuracy' (or as 'trustworthiness'). Several participants expressed surprise that there wasn't more academic study of 'applied truth', and proposed phrases like "truth engineering" and "truth tech" were offered as new fields that might produce new value.

A more practical discussion, in several different forms, examined accumulated value from the maintenance of the integrity of information in its representation and especially in its communication with an agentic internet. The role of cryptographic and distributed ledger/blockchain approaches was a well-discussed topic, especially as a 'sealed envelope' for guaranteeing the value of information passed from agent to agent toward a consumption experience. The origin, or 'sender', of such information objects was seen as especially important with the certification or at least standardization of provenance seen as a basis of value by many. It was assumed that this could extend beyond C2PA-style provenance-tracking into more semantically-oriented forms.

Another potential source of new value that came up in two discussions was the potential of decentralized 'information marketplaces' to promote integrity, or at least, in the words of one participant, to "provide an integrity floor" within an agentic internet. This elicited comments about the potential for gaming such marketplaces, and it was pointed out that the existing human-only information marketplace had essentially failed in relation to trust. A response was that new potential in an agentic internet to access revealed preferences of consumers in the consumption experience as indicators of trust, perhaps even to the point of using some kind of 'trust unit' as the currency of the marketplace itself.

A portion of each discussion of sources of value in the integrity layer was spent on options for verification or validation of claims in the AI-mediated ecosystem. Some

participants saw the potential new value in the separation of information verification services from information production services, but others expressed doubts that a market for such services would develop. A variant of this involved a 'trusted third party' information verifier or validator, but this was seen as difficult to separate from just a new attempt at gatekeeping or bias, with the example of NewsGuard in the US and its investigation by the FCC⁶ used to illustrate how this might fail. Conversations developed about the need to separate validation and verification from inevitably political human influence and bias, and proposals for agentic forms of verification were common – some kind of extension of existing AI-enabled automatic fact-checking as a service or even as foundational infrastructure. A distinction between validation as fact-checking and validation as chain-of-custody emerged several times, emphasizing reference to evidence or security of communication formats respectively. The impracticality of doing vast numbers of "mini-validations" was pointed out, which extended to a linkage between representation and validation, and the potential of using the 'fit' of facts within a database of other facts for validation was mentioned – a possibility that might generate considerable network effects.

Some participants questioned whether verification, or even integrity, had much consumer value, pointing out that while it might be essential for some information (e.g. weather) it might be less so for others (e.g. 'horse-race' politics or celebrity news). One participant hypothesized that the value of integrity may just be "brand value" in another guise, while others said integrity as the foundation of all value from information. The perspective of integrity as security – protection from manipulation and exploitation, etc – was revisited and was seen as increasingly valuable as the risks of, and sophistication of, manipulation and exploitation increased in an agentic information ecosystem. Integrity was most often seen as a key differentiator and competitive advantage, and in some fields (scientific, medical, financial, etc) as the entirety of value.

All of the integrity layer discussions produced less-discussed but otherwise interesting possible value sources that might develop in the AI-mediated information ecosystem. Several discussed a "registration system" for agents for certain communication tasks, with identification via blockchain and "certification organizations" to administer it and charge for access to it (for example in ways similar to credit card fees) and 'agentic audits' to ensure compliance. The need for truth-

6 <https://apnews.com/article/trump-administration-newsguard-media-monitoring-ftc-investigation-291c71ebc2d00020326a56807ffccf5f>

seeking “motivation” of some kind to be “built into” information agents was expressed, thus aligning value in the ecosystem with societally beneficial outcomes. Several participants insisted that integrity systems must be fully oriented toward agents first, rather than toward humans, because the dominance of machines in the information ecosystem was inevitable and thus value would necessarily accumulate around agentic processes rather than human processes. The necessity of, and difficulty of, a “corrections process” for correcting erroneous information was reiterated several times, but its value to consumers was questioned. A key discussion occurred around the dual needs for verification and validation in the assembly of information experiences for consumers, as well as in the individual units of information that passed through the ecosystem to be used in the assembly of experiences. This suggested that much of the value in the integrity layer might develop nearer to the consumption experience, rather than nearer to the sources of information or in the piping in between.

Finally, most discussions on the integrity layer as a source of value included conversation around governance of integrity-generating functions. One prospect that was raised multiple times was some kind of fiduciary duty to society attached to use of the infrastructure of the AI-mediated information ecosystem by agents. Another was the possible role for AI-equipped humans to act as governors of integrity in an agentic ecosystem, possibly at some significant scale as part of public or civic service. Several discussions emphasized the risks of full agentic access to “root control” of the integrity layer, and the prospect that this control might become a “sacred responsibility” of some kind, existing above the needs of the commercial marketplace, was suggested.

Interpretation and discovery (The experience layer)

Value created in the experience of information by consumers in an agentic AI-mediated information ecosystem was a rich discussion. A starting point in each of the summits was the importance to consumers of privacy-protecting and interoperable personal context – privacy protecting because of the intimate nature of both individual expressions of intent and of accumulated historical context, and interoperable because of the extremely powerful ‘lock-in’ effect of non-transferable personal context, potentially leading to essentially a form of AI ‘serfdom’. An assumption for the investment discussions was therefore that there would likely be no switching costs between AI platforms and providers, for example because of consumer ownership of their own personal context in a ‘databox’-like system, or

because of consumer access to agentic negotiators able to easily move context from place to place in search of the best value for the consumer.

In all discussions some variant of 'agentic representation' for the consumer was assumed to exist. These representations were seen as likely to be 'owned' or 'controlled' by the consumer, with a duty to serve the consumer and to advance the consumer's interests, however the prospect of a two-tier system emerging in which some people 'owned' their own agentic representations and others did not was also considered. The prospect of not owning one's own agentic representation was seen as so dangerous as to be referred to by one participant as "soul buy-out". These representations were also seen as presenting "many surfaces for trading information", including both buying and selling, and were at the heart of a possibly dramatic expansion in the demand for information.

Two distinct kinds of agentic representations appeared repeatedly in these conversations. One kind was an agentic representation as a separate character, with a separate identity, acting as a servant to the consumer. This kind was variously likened to fictional agentic characters like Jarvis, C-3PO from *Star Wars*, or 'computer' from *Star Trek*, or to advanced versions of Alexa or Siri with more internal agency, or to butlers or concierges. The second kind was essentially a 'digital twin' of the consumer – a constantly-maintained 'model' of the consumer based on their full historical context and on the explicit and implicit needs and wants of the consumer (a "life machine"). This was also described as "your life as a prompt".

It was not clear from these discussions whether the differences between these two kinds of agentic representation were trivial or profound, and the distinction seemed to center on the agency of the representation – the 'servant' model having less top-line agency than the "life machine" model. This might have consequences for the nature of new demand for information from each kind, in the sense that the servant model may be more likely to consume familiar information (although perhaps at greater scale and scope) because its needs would ultimately originate in the explicit wants and needs of its human principal, whereas the life machine model might behave in unexpected ways as it sought to optimize the wellbeing of its 'twin' based on its own assessment of its own wants and needs. This distinction was referred to by one participant as a distinction "between task mode and experience mode".

A separate thread of discussion centered on the bidirectional connections between information consumption and information production, and the opportunities for

value-adding exchanges in both directions. This included the communication of provenance, trust signals, usage constraints, etc. forward from production to consumption, and the communication of demand signals, usage information, etc. backward from consumption to production – value creating activity that is already emerging in the form of attribution services or citation frameworks for AI products. Value from developing, maintaining and communicating ecosystem-wide analytics were also discussed, potentially even for use in payments.

This also included discussion of the challenge of ‘merging’ different kinds of context for use in generating highly personalized, relevant experiences – especially merging private personal context with broad or ‘universal’ context from public sources, but also including access to other pools of context that might be traded commercially, such as community context, or niche domain context. Opportunities from the provision of context and from the blending of multiple contexts with consumer intent were mentioned.

A different source of value creation emerged in some conversations around providing consumers with new and potentially valuable cognitive experiences from AI-mediated experiences of information. This might include experiences of control or mastery over information, experiences of validation through information (providing “the right to believe”), experiences of harmony or security through understanding or experiences that were attuned to the learning frameworks of individual consumers. This was a more speculative conversation, and intersected heavily with discussions around trust, but it illustrated that sources of value in a radically different agentic information ecosystem may be unexpected.

The Reward Layer

There was a general consensus at all summits that information consumers would likely claim most of the new value created by an AI-mediated information ecosystem, that inference providers and marketplaces would likely claim substantial value, and that information producers would find it more difficult to differentiate their offerings than previously. Some alternative opinions also emerged, however. One prospect was that inference costs might rise significantly when inference stopped being subsidized by investors in hyper-scalers, possibly to the point at which use of powerful AI in routine information gathering and processing became infeasible. Another was the prospect of a substantially degraded or polluted information ecosystem, flooded with AI-generated noise, and the associated prospect of significant value remaining in

manual processes. These alternatives were exceptions, however, and so the question of how to incentivize and reward information producers in a much less differentiated marketplace was seen as urgent.

Most discussions began with questions about where the marketplace for information would be, and about what would be traded in that marketplace. A core question about market structure was whether information providers would have direct relationships with information consumers, or whether their product would be aggregated by others positioned between production and consumption – the metaphor of farmers having their product aggregated by grocery stores before reaching consumers was used several times. A core question around product was whether there would be commercial demand for verified, truthful information, and if so then what form that information might take (i.e. what scope, what granularity, what format, what recency, what certification, etc).

Options for market structure that were discussed included one-to-one relationships (with aggregation occurring in the consumption experience across all of the consumer's relationships), a centralized marketplace in the form of an Amazon or a financial exchange, a decentralized marketplace in the form of employment or dating markets, an auction in the form of the programmatic advertising marketplaces, a business-to-business marketplace in which information producers contract with AI platforms (i.e. an extension of today's deals between publishers and LLM companies), and information producers acquired by or otherwise employed by AI platforms (perhaps even in the 'DoorDash for facts' model).

Options for a reward mechanism centered disproportionately on micropayments, which were seen as being freshly relevant for use in agenticly mediated information flows despite their repeated failure in the pre-AI digital information ecosystem. This new potential was repeatedly identified as originating in the decision-making capability of AI, and the resulting possibility of agents using micropayments to create consumer experiences without any requirement for added cognitive friction for the consumer. Investment opportunities were seen in the required micropayment infrastructure, in new kinds of payment barriers and paywalls that operated on "microtransactions" and on aggregated "super subscriptions" across many information producers, facilitated by micropayments. Open questions regarding these opportunities remained around how to set prices for information, how to charge for access to domain context, how existing agentic payment protocols may or may

not work for micropayments and, most significantly, how a healthy liquid marketplace for societal information might be created and sustained – especially in an ecosystem in which so much information would be freely available, and no-cost ‘good enough’ experiences might be widely available.

The potential for public service organizations to act as “market makers” to anchor such a marketplace was suggested – variously described as a CERN of AI information (backstopping science ‘markets’), as a BBC-like entity (providing a ‘floor’ for provision of societal information) or as something like the Chicago Board of Trade (making a market in informational commodities). Other forms of this included more fundamental approaches that provided “public service compute”, issued either as “public service tokens” to citizens or, more likely, some kind of compute credits for access to consumer-friendly public service information agents. Several participants remarked that this approach might quickly become, essentially, “information soup kitchens” or “information food stamps”, leading to significant societal divides around unrestricted access to high-quality information.

An alternative to this, proposed in one discussion, was creating the ability for all participants in the marketplace – including consumers and agents – to earn “public service credits” under the coordination of a public service information entity. Human ‘creators’ and AI agents would earn credits when they provided some kind of public service information value – perhaps by providing health information or voting information. Those public service credits could then be spent on access to other information or compute – for example financially useful information or celebrity or lifestyle information.

Regarding the ‘product’ or ‘tradable unit’ side of the marketplace discussions, a prospect that was repeatedly raised was that this might be merely “access to an accumulating archive of information fragments” – an option seen as a kind of emerging default. This was seen as “anathema” to traditional information providers, and also additionally threatening because only true “primary sources” would presumably have value – with “secondary sources” such as many traditional articles seen as merely “token processing”. Other options were similar to those raised in the representation discussions, however the ‘meta-problem’ of price-setting, or “value sorting” was raised repeatedly without an accompanying satisfying solution. The discussed proposals – process certification, provenance, speed or recency, etc. – were recognized as largely value-carriers from the existing information ecosystem.

The one structurally different proposal that was offered in several separate discussions was that of trading consumer intent data (the “demand signal”) for provided information.

These opportunities for reward mechanisms for consumers were discussed repeatedly, usually around selling access to both accumulated personal context and to ongoing minute-by-minute intent. This was seen as requiring ownership of personal data by consumers, and these conversations returned to secure “data box” information vaults, perhaps built using something like the Solid⁷ framework. In several discussions it was pointed out that markets for personal context and intent might not be financial but might instead involve trading access to data and compute. Investment opportunities in facilitating, aggregating and perhaps brokering personal information for use in the wider information ecosystem were proposed. The raw value of the ‘demand signal’ was seen as relatively certain, with many indications from OpenAI and Anthropic indicating that they were already extracting significant value from that source. In terms of its value for information producers, however, several participants questioned whether this might just become a low-margin ‘DoorDash for facts’ relationship when implemented at the limit, rather than a healthy marketplace of equals.

Throughout the discussions on reward systems and marketplaces for information the conversation repeatedly returned to pricing in some form, and the designation of the value of information. As mentioned earlier, this was seen as a hard problem that was also deeply connected to trust and social utility. Potential adjacent systems that might offer lessons were often mentioned, including various certification or aggregated rating systems like the Wikipedia Source List or TrustPilot, and systems that captured “commitment signals” like prediction markets.

Conclusions

The four original layers – representation, integrity, experience and reward – remain useful, but the discussions suggest that value will be distributed unevenly across them. The market failure takes a similar form in each: the functions that generate the most value for the ecosystem as a whole are generally those whose value is hardest for any single participant to capture, and they are therefore the functions least likely to be built by commercial means alone.

⁷ <https://solidproject.org/>

Representation is foundational but may largely behave like infrastructure. Common schemas, protocols and reference systems create broad value, yet standardization can limit private capture. Commercial opportunities are more likely to arise in collecting, structuring, maintaining, enriching and delivering represented information than in owning a universal representation standard.

Integrity is likely to be essential but variably monetizable. It may support strong commercial models in high-consequence domains such as finance, health, science, law and security, while remaining difficult to sell separately in lower-stakes consumer information. Its greatest value may emerge close to the point of use, where evidence, provenance and verification affect a decision, action or liability.

The experience layer (Interpretation and discovery) appears to be the largest site of both new value and potential private capture. Personal context, intent, memory, agentic representation and control of the interface can generate powerful learning effects and switching costs. This makes consumer ownership of and portability of context especially important.

The reward layer is the least resolved and perhaps the most consequential. The summits identified attribution, pricing, demand aggregation, marketplaces and agentic payments as necessary, but did not establish what the tradable unit should be or how prices should be set. Without workable reward systems, the ecosystem may create substantial consumer and platform value while failing to sustain the production of socially useful primary information.

Summary discussions

Overview

The summary discussions were set up to enable participants to step outside the structure of the first three discussion sessions and engage in an unrestricted, open-ended conversation about the entire scope of the summits and about the process itself. These discussions were held among all participants and were lightly moderated by the organizers.

Each summary discussion began with a free-ranging review of what the group had collectively learned over the course of the summit, including ‘hot takes’, expansions of various points or concepts, corrections and criticisms and some “connecting of dots”

between different conversations. Generally speaking, these conversations produced relatively few new ideas beyond those contributed earlier, but considerable deliberation and analysis – including around the considerations and trade-offs associated with various proposals.

A portion of each summary discussion was explicitly reserved for discussion of the process itself, including about how well the lessons learned were captured by the note-takers. The question of what was missed was explicitly raised, including the ideas and concepts missed but also the process steps missed. This typically led to a discussion about how well the process worked, and suggestions for improvement. Each summary discussion concluded with a brief conversation about the shaping of deliverables – essentially this report.

Results

None of the summary discussions assessed that the discussions had reached any kind of conclusion, much less any actionable conclusions, however all of them pointed to engagement with investment opportunities in the emerging AI-mediated information ecosystem as a “maturity process” that had been materially advanced by the summits. Many participants remarked on the difficulty of maintaining situational awareness during accelerating improvements in AI and agentic capabilities. Many remarked on the difficulty of being both transformative and positive, especially given the magnitude of anticipated transformation – a cognitive “negativity bias” in the words of one participant. The lack of a common vocabulary for many of the central concepts was mentioned repeatedly, as was the difficulty in “discussing these topics for the first time” – an experience of many participants (“it was hard to jump right in”). Some participants remarked that norms around “difficult conversations” had not yet formed. It was generally assumed by all participants that we were in the very early stages of a very long transition period – the “flip phone era of AI” – and that the end state might have remained unclear because “it is stranger than we can imagine”. The usefulness of terms like “human flourishing” as objectives for an end state were questioned as being too vague to be actionable or clarifying.

Many relevant concepts and subjects were missed in the summit discussions. Some of these under-discussed or undiscussed topics related to expressions or opportunities at the frontier of AI and technology, including: AI-enabled prediction markets; the prospect of deep and intimate relationships with agents by many consumers; the potential emergent effects of networks of agents (“a trillion agents

with strong interests”); and the potential influence of open weight (and true open-source) models on the information ecosystem. Many more underdiscussed and undiscussed topics, however, were extensions of concerns from the existing digital ecosystem: AI literacy and education for an AI-mediated ecosystem; shared informational goods and narratives and the risks of a “bubble-of-one”; biases and other constraints “baked into” models; the need to preserve archives and libraries that may quickly become obsolete and inaccessible; access to AI in low connectivity environments; the potential resurgence of analog forms of informational media as consumers actively avoid AI; the need for critical information services for managing emergencies, etc. Participants also generally felt that the discussions missed connections with how other fields (law, intelligence, enterprise knowledge management, etc.) were approaching similar challenges, and failed to explore the extent to which the information ecosystem was already changing in extraordinary ways under the influence of AI.

More broadly, the summary discussions revealed a lack of baseline knowledge about knowledge – what is societal information? Why do people consume it? What is ‘understanding’? – and about ‘truth’ – how do objective, subjective and inter-subjective truths work in specific ways in actual practice within flows of societal information, and what is behind human incentives toward truth? Participants in all summary discussions also remarked that more attention should be given to what might not change, no matter how weird the AI-mediated ecosystem became – deeply cognitive and deeply social needs were cited.

A portion of each summary discussion focused on the persistence of assumptions from the legacy information ecosystem and of legacy information. The concept of discrete artifacts often remained in various ways, as well as the associated “package and distribute” model of information distribution. This “gravitational pull” of assumptions from legacy journalism occurred even in participants with backgrounds in technical innovation and in discussions at the San Francisco summit – held in the heart of the “AI Gulch”. Several participants even proposed that traditional forms of journalism should not be expected to change and should be “committed to” – although this was an isolated view.

All summary discussions included a political or ideological dimension in some places, and it was clear that some participants saw AI and AI-induced transformation of the information ecosystem through a political lens – often in terms of ‘socialism’ and

'capitalism', or in terms of more or less regulation by governments. Most participants anticipated some kind of broad public or political backlash against AI as threats to livelihoods and identity emerged, and several saw the acceleration of AI as a "pandemic moment" for governments.

The summary discussions were also forums in which participants expressed ideas, concepts and questions that perhaps did not easily fit within the quite structured format of the summits. One suggested that AI might begin acting like a "membrane" that expressed and reacted very differently to different things in different contexts – an information ecosystem that was continually and inextricably influenced by the information it carried. One asked whether we were solving for utility or identity in any transformation of the information ecosystem, suggesting that there might be utterly different functions that were only served previously by the same products due to accidental historical constraints. The prospect that the unit economics of information production might have no viable solution was raised, with severe consequences not only for information producers but also for AI platforms and for society. Several participants pointed out that humans placed significant value in fun, joy and "a good life", and that facilitating these things through information was fundamental. Several participants questioned how much of knowledge was embodied experience, and whether multi-modal experience and 'experience-in-the-loop' might remain uniquely human contributions well into the AI era. Interesting discussions of AI-mediated politics occurred, including observations that political representation is driven by decision-making constraints that may diminish or disappear. The concept of 'super-empowerment' of individual information producers was discussed and several participants suggested that the "User Generated Content (UGC) vs professional framing" was antiquated and that new assumptions of universal participation may be more appropriate. The prospect of far more social (i.e. human-to-human) agentic products and services was mentioned, with agentic networks that were primarily focused on sociality, dating, 'employment', friend development and the like.

A topic that came up repeatedly during the summary discussions, as it did in many of the earlier structured discussions, was the need for consumer ownership of their own context ("sovereignty of context"). It was suggested several times that ownership of one's own historical context and 'intent signal' was so important as to perhaps become "a new human right", and that an AI-mediated information ecosystem centered on consumers as owners of what several participants described as "the

source of value” might become far healthier than our current, perhaps more exploitative, information ecosystem.

Lastly, all summary discussions included conversation about the process of the summits. Many participants mentioned that the heavily structured nature of the summit was very useful and helped to maintain focus on specific objectives with specific outcomes. Several participants contrasted this to more open-ended and free-flowing discussions which were “interesting but not satisfying” and suggested that these kinds of heavily structured discussions were particularly useful in exploring such uncertain and “mysterious” subject matter. One participant expressed this perspective as “the method was the message”. Most agreed that the summits produced “glimpses of insight, but not a systematic excavation”, in the words of one participant, and many felt that much more time was needed (one participant suggested a “summer school” and another proposed a second summit). Multiple participants from each summit volunteered that the discussions “began where other discussions about AI and the information ecosystem left off”.

Conclusions

The summits advanced a maturity process rather than producing a settled blueprint. Participants lacked a shared vocabulary, carried persistent assumptions from the artifact-based legacy ecosystem and were attempting to reason about an end state that may be substantially stranger than current product categories imply.

The most consistent conclusion across the summary discussions was the importance of consumer ownership – effectively sovereignty – over personal context and intent. Context was repeatedly treated as a principal source of value, a source of platform power and a potential basis for a healthier ecosystem if controlled by the individual rather than captured by an intermediary.

The summary sessions also exposed important omissions: AI literacy, archives, low-connectivity access, emergency information, intimate human-agent relationships, open models, embodied knowledge, analog alternatives, shared narratives and the emergent behavior of large networks of agents. More fundamentally, the discussions lacked a sufficiently developed account of what information, understanding, truth and knowledge do for individuals and societies.

The structured process was nevertheless valuable. It produced recurring hypotheses and exposed foundational tensions that less disciplined discussions often avoid. Its

outputs should be understood as an initial map of the problem space – not as a forecast, market taxonomy or investment recommendation that has already been validated.

The vibe

All four summits followed the same structured design and so it is not surprising that they shared many commonalities, dynamics and themes. Some general observations from the organizers/authors from across all summits are as follows:

Participants got it. All summits immediately and fully embraced the focus on an entirely AI-mediated information ecosystem and mostly set aside the legacy-oriented assumptions that tend to take up considerable airtime and energy at ‘AI and Journalism’ events. There was a general agreement that AI was already, or would soon become, fully capable of doing most familiar information production tasks, and also that the changes brought to the ecosystem by AI would eventually be radically transformational. It was generally assumed that processes, products and institutions of the legacy information ecosystem would likely not survive in their current form, and that the capabilities of AI – in particular the scale of informational work it could enable, in both production and consumption – would substantially expand the information ecosystem.

A fully agentic internet was assumed. Participants generally seemed keenly aware of the growing practical capabilities of AI agents and of the special potential of their role as intelligent (or “editorial”) nodes in networked communication. This assumption of an agentic internet appeared to introduce significant uncertainty into the discussions in each summit, because it significantly raised the potential for unexpected emergent behaviors in the ecosystem. Several participants fondly recalled the relative simplicity of “AI before agents”.

Tech participants didn’t have special perspectives. There was less of an obvious divide between ‘technical’ and ‘non-technical’ participants than expected, especially in the substantial discussions around fundamental issues. This seemed to be because once the broad capabilities of AI were accepted – as they were – the discussions quickly centered on human and societal experiences of information and on the practicalities of serving the informational needs of humans. Many participants from technical backgrounds had not engaged with those questions deeply and some had

relatively theoretical or untested perspectives on how humans use information in practice.

There were some ideological overtones. Some participants seemed to position AI as a facet of a larger ideology centered on, roughly, rationalism or perhaps naïve empiricism. Although this was referred to as ‘techno-capitalism’ several times, it seemed different from specific concerns about inequality exacerbated by AI, lack of regulation of AI or the potential for corporate exploitation of people using AI, all of which were raised in various summits but usually in more of a problem-solving capacity. On the other hand there were no obvious signs of any triumphalist or even technophilic ideology, with even the most AI-forward participants firmly focused on outcomes and benefits rather than on any belief system.

There was some optimism. The summits were explicitly focused on opportunities for positive societal outcomes from AI-mediated information, and so optimism was actively sought out. Generally speaking most of the optimism emerging from the summits was around the increased empowerment of individuals via better access to relevant information, and more opportunities to use that knowledge to improve their lives. There was also some optimism around new forms of information production, and even about some forms of manually produced information within an AI-mediated ecosystem, but this was less common.

There was some pessimism. Despite the proposed focus on positive societal outcomes, and despite the optimism around empowerment of individuals, the general feeling at each of the summits was vaguely pessimistic – or at best neutral. Some participants seemed wary of the prospect of individuals being informationally super-empowered by AI and some expressed a distrust of human nature when provided with additional powers – the prospect of a world in which “everyone had their own secret police force” came up independently several times. Many participants were deeply unsettled by the extreme uncertainty ahead, and participants who were already deeply engaged with frontier models on ambitious projects seemed to be disproportionately likely to harbour free-floating unease arising from what they were experiencing.

Acceleration was palpable. There was a notable change in the general sense of urgency among participants between the first summit in New York in December 2025 and the last summit in Copenhagen in May 2026. The intervening five months had seen the launch of new models that seemed to deliver step-function improvements

in performance on many tasks (verified by benchmarks) as well as the rapid improvement of agentic toolkits like Claude Code/CoWork, OpenAI Codex and OpenClaw. The increasing difficulty of maintaining situational awareness was a common theme, and some participants mentioned a kind of vertigo associated with awareness of AI capabilities.

This was exhausting work. It was very clear that participants were exerting great cognitive effort in trying to imagine the characteristics of a radically different AI-mediated ecosystem and in responding to new ideas, concepts, possibilities and assumptions as they appeared. At each summit this mental exertion was explicitly mentioned so many times that it became a kind of cliché, and the mounting exhaustion of participants was apparent over the day. Many participants interpreted this as a need for more time, ranging from a frequent desire for an additional day to a suggestion from one participant of a 6-week “summer school”.

In addition to these general observations each summit also had its own particular characteristics, although these were not pronounced. The New York summit felt a bit more influenced by the higher societal functions of the legacy information ecosystem – authoritativeness, systems of record, shared informational environments, etc. The San Francisco summit seemed slightly more willing to entertain the possibility of radically different futures, although in general rather than specific terms, but ironically was also a bit more grounded on key information functions like scientific publishing. The Harvard summit was perhaps a bit more wide-ranging, with more discussion of trade-offs and interactions and more of an ecological perspective emerging than at other summits. The Copenhagen summit had, not surprisingly given the European venue, a little more focus on sovereignty, regulation and standards than some of the others. In general, however, the similarities between discussions were more pronounced than the differences.

Value in the AI-mediated information ecosystem

Overview

This analysis synthesizes the full body of notes from all five sessions (including the tabletop exercises and summary discussions) across all four summits – a total of 20 substantial discussions, each among 30–35 engaged experts. It maps where value is likely to be created in an AI-mediated information ecosystem, where – and by whom – that value is likely to be captured, and what follows for investment. Its central finding is a divergence between creation and capture: the functions on which the ecosystem most depends are, in general, the hardest to monetize, while value capture gravitates toward a small number of positions close to the consumer's context, intent and transactions.

Value creation

The strongest, highest-level proposition from across the 20 discussions is that AI-mediated informational abundance shifts value away from finished informational artifacts and toward systems that, from the bottom (i.e. production) up and in order, do the following:

1. Observe or obtain information that is not otherwise available;
2. Make information machine-usable and persistent;
3. Establish its integrity and permitted use;
4. Combine it with personal, community or domain context;
5. Interpret it for a particular purpose;
6. Convert it into decisions, actions and experiences;
7. Coordinate the feedback, payment and accountability loops around those activities.

Flowing in the opposite direction, from the top (i.e. consumption) down and unordered, is a set of informational signals communicating **demand, usage, impact, payment, reputation assessment** and **corrections**.

Cutting across the entire system are systems for coordinating **identity, authorization, provenance, rights, privacy, security, interoperability, competition** and **governance**.

This is not simply a longer version of the existing publishing value chain. It changes the object being produced from an artifact for human consumption to an adaptive service that continuously connects 'liquid' information with context, intent and action.

This highest-level proposition suggests that in the emerging AI-mediated information ecosystem value will be created in 14 distinct zones:

Primary observation and origination. Value is created by discovering facts, collecting real-world evidence, gaining access to people or situations, generating new data and preserving tacit or local knowledge. Value is captured by owners of exclusive access; specialist data providers; sensor operators; field networks; domain institutions and commissioned human gatherers. The principal tension or market failure is that generic digital information is easily commoditized, while costly civic information may lack sufficient paying demand.

Representation and persistent memory. Value is created by converting raw information into uniquely identifiable, structured, updatable, retrievable and tradable units with metadata and historical continuity. Value is captured by managed repositories, domain knowledge systems, indexing services, enrichment providers and maintenance of proprietary high-quality datasets. The principal tension or market failure is that while great ecosystem value may come from open standards, the commercial value of the standard itself is limited. Also, a universal ontology could embed bias or become a bottleneck.

Integrity and assurance. Value is created by establishing provenance, authenticity, evidentiary support, confidence, permitted use and correction history, and by protecting information against manipulation. Value is captured by verification services, compliance platforms, certification bodies, security providers, epistemic underwriters and high-trust domain brands. The principal tensions or market failures are that consumers may not pay directly for integrity outside high-stakes settings, and that central verifiers risk becoming politicized gatekeepers.

Personal context and intent. Value is created by maintaining long-term knowledge of the individual's circumstances, preferences, relationships, goals, vulnerabilities and present intentions. Value is captured by consumer-controlled agents and vaults; operating-system or model providers; identity and consent platforms and context brokers. The principal tension or market failure is that this is both a major source of utility and also the most powerful source of lock-in, surveillance and influence.

Ownership, portability and privacy are therefore key unresolved issues in this zone of value creation.

Community and domain context. Value is created by capturing the tacit knowledge, norms, history and specialized understanding of the community or domain needed to interpret facts correctly. Value is captured by domain networks, professional communities, cultural institutions, specialist agents, and proprietary context pools. The principal tensions or market failures are that deep context is difficult to formalize and may be lost through over-standardization, and that communities may be exploited without sharing in the resulting value.

Interpretation, prediction and sense-making. Value is created by connecting facts, comparing explanations, identifying gaps, simulating consequences, reducing uncertainty and constructing analytical frameworks. Value is captured by domain-specific agents; analytical platforms; decision-support providers; prediction and scenario systems. The principal tension or market failure is that general-purpose models may absorb generic functions, leaving defensibility dependent on proprietary context, workflow integration, accountability or demonstrably better outcomes.

Negotiation and protection. Value is created by negotiation between consumer and producer agents, quality vetting and filtering as active defense of attention and intent. Value is captured by trusted agents with a fiduciary duty to consumers, by managing the scale and complexity of ecosystem-wide negotiations and by reputation aggregation for use by agents on behalf of consumers. The principal tension or market failure is the possibility that the expense of negotiation and protection may render it unavailable to substantial numbers of consumers.

Experience and cognitive services. Value is created by delivering information in a form suited to the user's situation, learning style, emotional state and longer-term objectives, while protecting their attention and intent. Value is captured by the agent or interface controlling the consumer relationship; providers with accumulated trust and context and specialized experience designers – especially cognitively-oriented designers. The principal tensions or market failures are that the interface may become the dominant chokepoint and may optimize for the provider's interests rather than the consumer's, and also that extreme personalization may erode shared reality.

Decision delegation and action. Value is created by recommending a next action, negotiating, transacting, representing the user and carrying out authorized decisions.

Value is captured by agent providers, transaction platforms, workflow systems and providers accepting responsibility for outcomes. The principal tension or market failure is that greater agency increases liability, security and authorization risks while the boundary between assistance, delegation and impersonation remains unclear.

Demand aggregation and commissioning. Value is created by combining intent from many users to identify unmet needs, commission new information or finance expensive investigations. Value is captured by marketplaces, brokers, Groupon-like commissioning platforms and organizations with privileged visibility into aggregated demand. The principal tensions or market failures are that demand signals are highly valuable but potentially intimate, and that aggregators may reduce producers to low-margin suppliers – leading to a reduced supply of quality information.

Attribution, exchange and reward. Value is created by connecting use with origin, setting prices, processing payment and allocating rewards among contributors. Value is captured by marketplace operators, payment rails like micropayments, attribution services, rights managers and analytics providers. The principal tension or market failure is that information is non-rival, frequently abundant and difficult to price – for example micropayments solve payment friction but not the underlying problem of valuation or demand.

Shared knowledge access and public-service information. Value is created by maintaining common factual resources, civic information, emergency communication, cultural memory and universal access. Value is captured diffusely by society as a whole, primarily through better societal decisions, accountability, resilience and participation rather than through direct financial returns. The principal tension or market failure is that the benefits are diffuse while the costs are concentrated, producing chronic underinvestment and dependence on public, philanthropic or cross-subsidy models.

Ecosystem governance. Value is created by facilitating the healthy functioning of the entire ecosystem, including its marketplaces, by establishing a stable, simple and known set of rules within which the ecosystem can develop. Value is captured through the authentic recognition by participants in the ecosystem of the role of governance in the health of the ecosystem. The principal tension or market failure is that while the benefits of a healthy, well-governed ecosystem are realized by all participants, the costs of governance – both direct and indirect – are incurred only by some.

Socially-oriented, embodied and anti-informational experience. Value is created by creating belonging, status, challenge, serendipity, joy, human contact and real-world experience as digital information becomes abundant. Value is captured by exclusive communities, trusted human networks, cultural institutions, event and experience providers and human-centered services. The principal tension or market failure is that these markets may reproduce inequality and status competition, because their value may derive precisely from exclusion and scarcity.

The largest net new value is likely to be created near the intersection of information with personal context, intent and action on behalf of the consumer. Generic synthesis produces convenience, but it is deep context that can make information personally consequential. Decision support can reduce uncertainty and cognitive load for consumers, and delegated action can turn understanding into measurable outcomes.

Important value will also be created upstream by expanding the supply of primary information and downstream through social and public benefits. These areas may be economically weaker because the party creating the value is not always able to identify or charge the eventual beneficiary.

Value capture and moats

The notes from across all discussions in all summits, and the ecosystem structure described above, point to six likely positions where value is captured. These are the roles in the ecosystem where new value will likely accumulate. They are not just 'moats', or defensible competitive advantage, but the roles for which a combination of moats and structural advantages will likely cause value to accumulate. The consumer, for example, has no 'moat', and yet accumulates much of the new value from the ecosystem. The six roles accumulating value are:

The consumer. Consumers may receive a large share of the value as surplus, from lower search and decision costs, better services and greater personal capability.

Inference and platform providers. Providers that supply intelligence at scale can capture value through compute, aggregation, distribution and platform economics, particularly where they can incorporate adjacent functions into general-purpose products.

Holders of the user relationship. The provider controlling the agent, interface, memory and personal context may occupy the most powerful commercial position. It

sees intent, knows demand, controls discovery and can mediate every downstream transaction.

Marketplace and coordination operators. Organizations that aggregate demand, establish rules, verify identities, process payments or provide compliance can collect transaction rents without producing most of the underlying information.

Owners of genuinely scarce inputs. Primary sources, proprietary data, exclusive access, domain context and trusted real-world networks can retain pricing power where their inputs cannot be reconstructed from abundant public material.

Integrity and liability-bearing providers. In high-stakes domains, providers able to certify, indemnify or accept responsibility for information and decisions can capture a premium that ordinary “trust signals” may not command.

Traditional secondary information producers – for example those moving information from sources to fixed consumption artifacts – are in the most vulnerable position. Unless they control scarce origination, a trusted relationship, proprietary context, valuable rights or a direct demand signal, their output may be treated as an easily substitutable processing layer.

The moats that act to accumulate value are more specific than the roles to which value accrues and are essentially control points within the ecosystem. *Context is the master moat*, and every other capture position is partly downstream of this one because whoever controls context controls the experience, the demand signal, and the negotiating position with producers. Value capture tends to concentrate near the consumer, however other transaction and ‘certification’ positions are the durable fee-takers because of two-sided network effects and standards control. Producers capture value only at the edges of abundance, in primary origination, exclusivity and relationships, recency, possible fact-level usage rights, and reputation. The most important potential control points are therefore not necessarily at the places where the most information is produced.

The major moats are:

The personal context store: whoever holds the individual’s history, preferences and relationships.

The consumer agent: whoever represents the user and determines which providers can reach them.

The intent and demand signal: whoever observes what individuals and aggregated populations want or need.

The provenance and rights graph: whoever can connect information to origin, permitted use, correction and compensation.

The distribution and negotiation layer: whoever selects, prices and assembles information into an experience.

Identity and authorization infrastructure: whoever determines whether a human or agent may speak, transact or act for another party.

Verification and compliance rails: whoever defines sufficient evidence or conformity for consequential uses.

Public-service grounding: whoever maintains the information floor upon which private services depend.

These positions can create legitimate efficiencies but can also become infrastructure monopolies – including even the public-service grounding floor. Investment analysis should therefore examine not only the value of a capability, but whether it creates an avoidable chokepoint and whether its governance is compatible with its importance to the ecosystem.

Detailed categories

The following 70 categories consolidate every investable domain identified across the five sessions in all four summits. They come directly from the raw discussion notes rather than from the value map, however they have been mapped into the 14 zones of the value map. These categories are not complete, their categorization is not exact and their boundaries are porous, but they provide very thorough coverage of the investment space at considerable granularity.

These categories fit differently with different kinds of capital. Most are a good fit for private capital, but several favor blended capital (private and philanthropic), whereas the public-service concepts are proposed precisely because private capital is an inappropriate fit.

The categories are:

Primary observation and origination

1. **AI-native public-source collection:** Systems that continuously identify, collect and organize public information across text, speech, imagery, video and structured data.
2. **Human-in-the-loop field gathering:** Systems and networks that enable agents to commission humans for interviews, observations, document retrieval and other real-world reporting.
3. **Specialist primary-data providers:** High-quality domain data in areas such as health, science, finance, climate, law, public administration and local communities.
4. **Sensor and embodied observation networks:** Responsible collection from physical sensors, remote systems and other machine-observable environments.
5. **Public-record modernization:** Digitization, structuring, maintenance and machine accessibility of government and regulated records and proceedings.
6. **Open-data protection and expansion:** Preventing valuable public data from becoming closed, degraded or inaccessible once AI increases its commercial utility.
7. **Public-interest commissioning funds:** Financing investigations, local coverage, civic data and other information with high social value but inadequate commercial demand.
8. **AI-native producer tools:** Systems, skills and processes that help information producers transition from artifact production to structured, service-oriented production.

Representation and persistent memory

9. **Archives, libraries and cultural memory:** Preservation, migration, indexing and accessibility of historical materials that may otherwise become technically obsolete.
10. **Atomic information and fact-management systems:** Systems for maintaining granular, referenceable, updatable and correctable informational units at AI scale.
11. **Domain ontologies and knowledge graphs:** Structured representations optimized for bounded fields where vocabulary, evidence and relationships can be defined with useful precision.
12. **Causal and temporal representations:** Tools that represent changes, sequences, dependencies and competing causal explanations rather than static facts alone.

13. **Metadata and enrichment services:** Assignment of identifiers, provenance, confidence, collection cost, rights, restrictions and other machine-readable metadata.
14. **Large-scale indexing, storage and retrieval:** Infrastructure capable of maintaining historical context, versions and fast access across very large information stores.

Integrity and assurance

15. **Provenance by design infrastructure:** Capture of origin and chain-of-custody information at creation rather than retrospectively.
16. **Authenticity and tamper-evidence:** Cryptographic or other mechanisms for establishing whether information and communications have been altered.
17. **Automated verification and evidence retrieval:** Fact-checking at machine scale, source comparison, claim decomposition, sampling and just-in-time verification.
18. **Corrections, versioning and propagation:** Systems that record corrections and transmit them through repositories, models, agents and downstream experiences.
19. **Reputation and performance histories:** Evidence-based records of the accuracy and behavior of sources, institutions and agents over time.
20. **Truth tech:** Emerging techniques for structural orientation toward truth, such as consistency systems, bridging-algorithm systems, Bayesian crowdsourcing, etc.

Personal context and intent

21. **Personal context vaults or “data boxes”:** Secure stores controlled by the individual, with granular permissioning and the ability to serve multiple agents or providers.
22. **Context portability and migration services:** Tools and protocols allowing personal memory, preferences and agent history to move between providers.
23. **Privacy-preserving context integration:** On-device models, confidential computing and selective-disclosure systems that combine private and public context without unnecessary exposure.

Community and domain context

24. **Community and domain context pools:** Governed resources that preserve specialist, local, cultural or institutional context for authorized uses.
25. **Domain-specific information agents:** Specialized agents combining proprietary knowledge, workflows, evidence standards and accountability in bounded fields.

Interpretation, prediction and sense-making

- 26. **Sense-making and comparative interpretation:** Tools that separate evidence from inference, compare explanations and help users understand disagreement.
- 27. **Prediction and scenario analysis:** Systems that estimate likely developments and consequences relevant to individuals, enterprises or public institutions.

Negotiation and protection

- 28. **Delegated action and negotiation agents:** Authorized systems that obtain information, bargain, transact, schedule, communicate or act for their principals.
- 29. **Attention and intent protection:** Filters and agent services designed to resist manipulation and protect the user from unwanted commercial or political influence.
- 30. **Privacy and cognitive security:** Privacy preservation, anti-surveillance, “digital invisibility”, protection of intent against manipulation and personal context poisoning.

Experience and cognitive services

- 31. **Serendipity and challenge mechanisms:** Systems that intentionally introduce diversity, dissent, randomness or “jester” functions into personalized environments.
- 32. **Consumer-aligned or fiduciary agents:** Agents contractually and technically obligated to protect the consumer’s interests, attention, context and intent.
- 33. **Applied agentic services:** Consumer agents that act as either servant/concierge agents or as digital-twin “life machines” representing their principals.
- 34. **Long-term memory:** User-controlled systems that maintain longitudinal context and identify gaps, risks or opportunities.
- 35. **Experience innovation:** Exploration of deep experiences, including “aspirational algorithms” oriented to long-term wellbeing and personal groundedness.

Decision delegation and action

- 36. **Assurance, indemnity and insurance:** Products that accept or price responsibility for the quality of information or agent-supported decisions.
- 37. **Agent identity, registration and audit:** Registration systems for agents, certification organizations and fee structures, agentic audit services.
- 38. **Agent authorization and delegation controls:** Management of who an agent represents, what it may do and when its behavior becomes unauthorized.

- 39. **Decision-support and next-best-action services:** Contextual and intent driven recommendations tied to explicit objectives, constraints and evidence.
- 40. **Life services and applied decision support:** Applied decision services for households, health, scheduling, careers, cultural life, social life, etc.

Demand aggregation and commissioning

- 41. **Demand aggregation and bounty systems:** Combining intent signals to commission information whose cost would exceed the willingness of any single user to pay.
- 42. **Pricing and “proof of value” systems:** Measurement of information quality, impact, commitment and willingness to pay (Dynamic pricing).
- 43. **Producer demand intelligence:** Privacy-preserving signals showing information producers where unmet demand exists and what contribution was useful.

Attribution, exchange and reward

- 44. **Aggregated subscriptions and bundles:** Agent-managed access to multiple information providers without separate consumer decisions for every transaction.
- 45. **Rights declaration and automated licensing:** Machine-readable conditions governing access, reuse, modification, attribution and payment.
- 46. **Agentic payments and micropayment rails:** Low-friction settlement initiated by agents across many suppliers.
- 47. **Information marketplaces and brokerages:** Discovery, negotiation, quality control and exchange of primary information, context or access.
- 48. **Consumer-controlled context exchanges:** Mechanisms through which individuals grant access to context or intent in return for money, services or compute.
- 49. **Machine-speed compliance infrastructure:** Automated enforcement of privacy, rights, use restrictions and sector-specific rules across agentic flows.
- 50. **Attribution and usage accounting:** Connecting information use to contributors and recording how it affected an assembled experience or action.

Shared knowledge access and public-service information

- 51. **Public-service market making:** Public or nonprofit institutions that create a demand floor for socially useful information or anchor an otherwise illiquid market.
- 52. **Public-service compute or information credits:** Entitlements that provide access to high-quality agents and information services irrespective of income.

- 53. **Shared reality and community experiences:** Products that preserve common reference points or allow communities to understand what others are encountering.
- 54. **Accessibility, language and low-connectivity services:** AI-mediated information for underserved languages, disabilities and environments without reliable high-bandwidth access.
- 55. **Civic and public-service information agents:** Publicly accountable alternatives providing essential information, accountability and universal service.
- 56. **Cultural and language sovereignty infrastructure:** National models, corpora and institutions preserving the ability of communities and nations to represent themselves.
- 57. **Emergency and public-interest override channels:** Secure ways to bypass ordinary personalization for critical warnings, elections or other essential public communication.
- 58. **Civic information incentives:** Funding and sustaining civic information products and data sources as public goods, including watchdog capacity at scale.

Ecosystem governance

- 59. **Standards conformance laboratories:** Shared environments for testing whether products and agents implement standards consistently and safely.
- 60. **Competition and portability enforcement:** Regulatory and technical capacity to prevent context lock-in, discriminatory access and infrastructure enclosure.
- 61. **International coordination institutions:** Shared rules, monitoring, incident response and standards for cross-border AI-mediated information flows.
- 62. **Regulatory and competition capacity:** Anti-monopoly enforcement, interoperability and context-portability mandates, moat-sunsetting after recoupment, etc.
- 63. **Investment intelligence and due-diligence infrastructure:** Independent technical and market assessment for investors operating in a fast-changing ecosystem.
- 64. **Independent audits and conformance testing:** Technical testing of agents, information pipelines and standards compliance.
- 65. **Information observatories:** Continuous measurement of ecosystem behavior, concentration, model outputs, information quality, harms and access.

- 66. **Information-security and context-poisoning defenses:** Protection against malicious inputs, manipulation of personal context and attacks on shared knowledge resources.
- 67. **Interoperability standards and reference implementations:** Open protocols for data, context, attribution, identity, agent communication and rights.
- 68. **Applied research on AI-mediated information:** Continuing research into interfaces, demand, trust, behavior, market structures and societal effects.

Socially-oriented, embodied and anti-informational experience

- 69. **Analog and anti-informational offerings:** Real-world, embodied experiences whose relative value rises as AI “inflates away” informational experiences.
- 70. **‘Cognitively organic’ certification services:** Rubrics, standards, evaluation processes and branding infrastructure for AI-free, or even ‘digital-free’ experiences or lifestyles.

Infrastructure

Infrastructure surfaced in every session, and the same stack loosely recurred regardless of the discussion’s starting point. The notes from across all four summits imply eight layers of infrastructure:

Physical infrastructure, such as compute, chips, energy, connectivity and devices. This is indispensable but not unique to the information ecosystem. Its principal information-specific implications are unequal access, concentration and the possibility that inference costs materially shape which services are viable.

Information collection infrastructure, such as connectors into systems, sensors, public-data interfaces, human commissioning systems and trusted execution at the point of observation. Drones and robots may eventually become part of this layer, and the previous one.

Representation infrastructure, such as unique identifiers, schemas, metadata, historical repositories, indexing systems, versioning systems and systems capable of handling granular information units at AI scale.

Integrity infrastructure, such as systems for establishing and maintaining provenance, authentication, verification, corrections, reputation, security and evidence chains.

Context and interoperability infrastructure, such as personal context vaults (in either software or hardware), authorization systems, consent processes, context portability and secure communication between local and remote models.

Identity and registration infrastructure for both humans and agents, enabling ownership, reputation, payment processing and many other vital functions across the ecosystem.

Economic infrastructure, such as attribution systems, usage accounting, rights declaration and propagation, price discovery, negotiation and payment.

Accountability and public-service infrastructure, such as observatories, audits, regulatory interfaces, competition enforcement, emergency and override channels, shared grounding resources and universal-service provision.

Some form of this infrastructure stack emerged independently at the discussion, summit and series levels. Discussions on infrastructure were often associated with related discussions about the interaction between infrastructure and products or processes, and about the implications of infrastructure at each level of the stack. Some general observations that appeared repeatedly and that are related to investment-related implications are as follows:

Infrastructure will create more value than it can capture. Many of the most important capabilities – standards, open data, corrections, interoperability and public grounding resources – produce positive externalities across the ecosystem. A purely commercial investor may underfund them because competitors can share in the benefit. Philanthropic and government capital therefore have a central market-forming role, not merely a harm-mitigation role.

Infrastructure should be neutral and separated from competitive services. A promising model is an open or regulated core combined with competing commercial implementations and managed services. For example, an attribution standard might be open while analytics, compliance, hosting and settlement services built around it are commercial. This reduces the danger that a single provider converts a necessary standard into a proprietary chokepoint.

Universal representation should not be the initial objective. The discussions repeatedly identified the risk that a single ontology becomes a source of bias or constraint. Investment should favor modular, extensible and domain-specific

representations that can interoperate, rather than assuming that a universal “fact format” can be designed in advance.

Infrastructure must be agent-native. Human-scale labels, reviews and enforcement mechanisms will not be sufficient for machine-scale information flows. Identity, permissioning, evidence transmission, auditing, correction and compliance need machine-readable rules and automated execution, with meaningful human oversight at defined thresholds.

Governance is an integral part of the infrastructure. Infrastructure that manages identity, personal context, provenance or verification will exercise significant power. Its governance, appeal mechanisms, transparency and resistance to capture should be designed alongside its technical architecture. These cannot be easily postponed until after adoption, because the scale and complexity of its use will likely prohibit after-the-fact shoehorning of governance.

Public service can help create well-functioning markets. Governments and public-service institutions can accelerate the ecosystem by purchasing structured public data, verification, accessible information services and interoperable systems. Procurement can establish demand and quality standards without requiring government to build every service directly.

Infrastructure investment will likely require blended capital. A practical sequence would be as follows: Philanthropic and research capital should explore uncertain problems, develops public-interest prototypes and organizes standards; Government protects rights, funds public goods, creates procurement demand and requires interoperability where necessary; and private capital scales products and services that can generate revenue from transactions, compliance, assurance, workflow integration or improved outcomes.

The largest infrastructure risk is enclosure. The ecosystem could become nominally personalized but structurally centralized if one provider controls context, identity, distribution and payment. Infrastructure investment should therefore be tested against a simple question: does it increase the ability of participants to enter, leave and interoperate, or does it create another dependency that becomes difficult to escape?

Four more general insights about infrastructure follow from these implications.

First, most of this stack is heavily standards-dependent, and therefore probably a public good. The layers on which everything else depends are precisely the layers private capital is least able to capture, creating a chronic underinvestment risk at the ecosystem's foundations.

Second, sequencing matters. The reward layer cannot function without representation and integrity infrastructure beneath it, and so early catalytic investment by philanthropic and government capital in representation, standards, provenance and verification is plausibly the highest-leverage intervention available – it unlocks the private markets that follow.

Third, the exceptions are fee-taking control points – registration and certification, payment rails, brokerage – which are the investable infrastructure positions for private capital, defensible through network effects and functioning under any context-ownership regime.

Fourth, infrastructure is where power concentrates: control of provenance, identity or the integrity layer's "root" is a gatekeeping position, and the summits' own language ("sacred responsibility", "fiduciary duty", "agentic Stasi") makes clear that governance must be designed into these systems from the start – pointing toward neutral, consortium or public stewardship models rather than conventional ownership for the most sensitive layers.

The practical implication is a division of labor between categories of investors: private capital should own the flows (transactions, brokerage, certification) and the applications built on the stack; philanthropic capital should build the standards, evaluation systems and fund the research the stack depends on; and public capital should anchor the markets and guarantee the floors – universal access, emergency channels, production of low-demand/high-importance information, and the market-making functions – that no private actor is incentivized to provide.

Conclusion

What did we learn about the AI information ecosystem?

A top-level lesson from these summits is that there is no existing body of knowledge about what an information ecosystem fully mediated by AI might look like – no communities, no organizations, no academic theories, no books, no industry group, no cohort of start-ups. Over the course of five months, in sustained engagement with around 140 experts from the frontier of AI applied to societal information, we did not uncover any significant experience or expertise with a complete, well-formed, well-evidenced vision for what the AI-mediated information ecosystem might look like. We are truly in uncharted territory, and these summits felt like genuine exploration.

On the ecosystem itself, the same core lessons surfaced repeatedly across all four summits. The ecosystem is unlikely to be organized primarily around the production and distribution of content. Its central economic activity may instead be the continuous conversion of observation into structured information, of information into context-sensitive understanding of how consumer intent should be served, and of understanding into decisions and action – including the creation of experiences. The ecosystem will likely transition from one centered on a relatively static ‘stock’ of discrete artifacts to one centered on the continuous ‘flow’ of information from source to use. The scale of the ecosystem will expand greatly.

The principal source of value is likely to be relevance to a particular person, agent, community or decision. The principal source of power is likely to be control over the context, intent and interfaces through which relevance is determined. The principal public challenge is to preserve agency, pluralism, shared knowledge and sustainable information production when those control points offer extremely strong incentives for concentration.

Investment opportunities will emerge in the informational supply chain, in personal and social context, in integrity and assurance, in agentic representation, in market and reward mechanisms, and in the public infrastructure required to prevent a highly personalized ecosystem from becoming a highly centralized one.

This information ecosystem will inevitably be more technical than the information ecosystem that we are leaving, even as AI blurs the distinction between technical and non-technical. It will likely be more abstract, more systemic and less understandable

by people. It will likely be vastly larger, with both supply of and demand for information greatly exceeding what it is now.

There is every possibility that as this new information ecosystem develops, matures and diffuses through society it will provide consumers with informational powers that are currently almost inconceivable, with dramatic benefits for the functioning of communities and societies. There is also every possibility that the transition from our current information ecosystem could be turbulent and perhaps even dangerous.

Perhaps most importantly, we learned that skilled and experienced people are already deeply engaged in this transition, and that they are well aware of the scale of the challenge of ensuring that AI-mediated information will benefit all of us.

What did we learn about the process?

The overall design of the summits held up well over four versions. All summits followed the same structure, in the same way, and only relatively marginal changes were made between summits – typically in merging categories of discussion that became similar.

We learned that by recruiting participants from outside or from the periphery of the existing information ecosystem, and by clearly setting expectations for focus and subject matter, we could cultivate discussions at or near the frontier of possibilities. The tabletop exercise probably also played a role in quickly setting norms for imaginative thinking and consideration of extreme possibilities. We repeatedly heard that the conversation at these summits “began where the conversation in other venues ended”.

The structured approach was valued in essentially all participant feedback, and worked well to keep the discussions focused in the direction of the desired outcomes. It made the summits relatively easy to manage and facilitate, made note-taking and the consolidation of notes across summits much easier, and appeared to provide all the features generally expected by participants, including learning, engaging and networking.

A significant lesson was that the summit design requires the full day and a half to comfortably accommodate the modules, and the one summit that had been shortened (San Francisco) was probably unreasonably intense for its participants.

Beyond the summits themselves we learned that there is clearly significant demand for these conversations. This series of summits was originally to be a single gathering

– at Harvard – and grew first to include New York, then San Francisco and then Copenhagen. Requests to host additional versions at other venues have been many.

What could we have done better?

While the general structure of the summits seems to have worked well, there were many things that in hindsight we could have done better.

On a tactical level we should probably have applied more structure to the individual table discussions within each session, including providing a facilitator for every table. We should probably also have recorded the table-level discussions in order to extract an anonymized transcript and more systematic summary context – a step that we had initially felt might inhibit conversation.

We should probably have substantially redesigned the tabletop exercise, orienting it away from the perspective of wealthy consumers while retaining the ‘radical abundance’ incentive for imagining new products and services. In hindsight the conceit of a billionaire news product clashed too often with our objective of imagining the kinds of radically expanded information access ordinary people might expect from an AI-mediated information ecosystem. Some additional facilitation of the tabletop exercise might also have improved both the experience for participants and the outcomes.

In terms of specific content there were several additional pools of information that, in hindsight, would probably have been beneficial to participants prior to the discussions. A simple glossary of terms and concepts would have benefited some, especially those from disciplines without direct contact with applied AI in informational media. A brief analysis of how other knowledge-producing fields, such as science, government intelligence, law, etc., were approaching these opportunities and challenges might have been useful. We hope that this report, or portions of it, might serve as preparation for other, similar gatherings.

Finally, it was observed at each summit that the people assembled were not representative of typical information consumers. Participants were typically highly educated, highly motivated, highly engaged people in significant professional or entrepreneurial roles in the US, UK or Europe, and were quite comfortable with intellectual risk-taking. While this was somewhat unavoidable for a specialist group of AI experts, in hindsight it would also have been useful to include some kind of input from participants more representative of ordinary information consumers.

Next Steps: The Signals at Scale Studio

The summits described in this report were conceived as an exploratory 'mapping' exercise: an attempt to clarify the approximate shape of the emerging AI-mediated information ecosystem and to identify where within it durable value — consumer, commercial and societal — might be created. That mapping work has produced a concrete next step aimed at continuing and applying that exploration. The organizers of the summits have established the Signals at Scale Studio, a mission-driven venture studio that will fund and support builders of early products, services and infrastructure for the new ecosystem. The studio will launch publicly in September 2026.

Learning by building

The summits demonstrated both the value and the limits of anticipatory discussion. These discussions were productive, and much of this report records that progress, but the summits also ran up against a boundary that became more apparent at each event: many of the most important questions about the emerging ecosystem are empirical and practical, and cannot be resolved by further discussion, however well-informed and imaginative.

Three observations from the summits point in the direction of exploration by building. First, reasoning about a radically unfamiliar information ecosystem is cognitively difficult and even exhausting, and the cultural distance between most participants and mainstream information consumers was significant. Practical returns to additional rounds of speculative discussion are therefore likely to be diminishing. Second, the acceleration in AI capability between the first summit in December 2025 and the last in May 2026 was palpable; the object of study is changing faster than convenings can be organized to study it. Third, the uncertainties identified in this report — about where value will settle, which functions the ecosystem will actually demand, and how its layers will interact — are precisely the kind that resolve only through contact with reality: through building working systems, products and processes, placing them in the ecosystem, and observing what happens.

There is also a question of timing. The foundational layers of the new ecosystem are not yet stable, which makes this a moment of unusual leverage: choices made in the next two to three years will plausibly determine which information reaches whom, in what form, and on what terms, for a generation or more. Building is therefore not only

the best available way to learn; it is also the best way to participate in shaping the outcome. The natural continuation of the summits is a shift from learning by discussing to learning by building — with the understanding produced by building and testing becoming evidence of a kind that no further summit could generate.

The studio

The Signals at Scale Studio is a mission-driven venture studio, incorporated in 2026 and co-founded by the organizers of the summits, Shuwei Fang and David Caswell. Its mission is to enable the creation of information infrastructure for the emerging AI-mediated information ecosystem. The studio operates for public benefit, with a capital architecture designed to align that mission with long-term financial durability.

The studio's model is challenge-based. It defines challenges — specific components and infrastructure that it believes should exist in the next ecosystem — and then recruits and funds individual builders to pursue them. Alongside its cohorts of builders, the studio will continue the convening and synthesis work begun by the summits: further summits and workshops, synthesis publications, and field briefings that draw on learnings across multiple cohorts of builders.

The studio is best understood as the summits' framing put to the test. The summits produced an early map; the studio now sends builders into the territory. Because those builders will work natively inside the emerging ecosystem, their experience will produce the clearest available picture of how AI-mediated information actually evolves — a learning loop that continues, by other means, the work these summits began.

Acknowledgements

Collaborators

The organizers would like to thank the following funders and hosts who made the Signals at Scale summits possible.

- The Institute for Humane Studies
- The Soros Economic Development Fund
- The Tinius Trust
- The Shorenstein Center for Media, Politics and Public Policy at Harvard University
- The International Fund for Public Interest Media
- The Financial Times
- Holtzbrinck Publishing Group
- The Tow-Knight Center for Journalism Futures at City University of New York
- JP/Politikens
- The Copenhagen Institute for Futures Studies
- The University of Copenhagen

Participants

The Signals at Scale summits and the insights in this report were the combined contributions of the people who participated in them, who earnestly interacted with each other and who contributed their time, energy, trust, wisdom and experience to an uncertain cause. Many endured airports, email backlogs, appointment juggling and disrupted routines in order to attend, and all energetically engaged in discussions that were sometimes uncomfortable and occasionally unsettling. The organizers would like to extend their heartfelt gratitude to everyone who participated in this journey. Thank you!

About the authors

David Caswell

Co-founder and Chief Strategy Officer of Signals at Scale. David has been working at the intersection of AI and news since 2010, including as a product innovation executive at BBC News and as Director of Content Personalization at Yahoo. He has published extensive peer-reviewed research on automated journalism and the semantic units of news, has architected systems and workflows for automated media, and is a frequent keynote speaker at major AI and media conferences.

Shuwei Fang

Co-founder and Chief Executive Officer of Signals at Scale. She spent more than a decade in media and technology philanthropy and investment, and is a fellow at Harvard Kennedy School's Shorenstein Center, where her work has focused on the economics of the AI-mediated information ecosystem. She will be a senior fellow at the Newmark Graduate School of Journalism at CUNY.

About this report

Publication

Published by Signals at Scale, September 2026.

DOI: <https://doi.org/10.5281/zenodo.22894513>

How to cite

Caswell, D. and Fang, S. (2026). Value in the emerging AI-mediated information ecosystem. Signals at Scale. <https://doi.org/10.5281/zenodo.22894513>

Licence

This report is published under a Creative Commons Attribution 4.0 International licence (CC BY 4.0). You may share and adapt it for any purpose, with attribution.