# Field and Desk Report

How a research organization stays current, what it looks for in people who build, and how it stays relevant without repeating what is already fashionable.

This file is the argument. The dashboard is the working file. The five-slide deck is the ten-minute walkthrough for a hiring conversation. Do not copy the same paragraphs into all three.

This edition keeps the method, the builder criteria, and the people matrix. Live lists (what to review today, feeds, events, capital pipeline, stacks) are in the dashboard only.

**Classification.** Working brief for role, program, and network design. Not a public manuscript. Candidate concept if read against a Programs and Policy job. Not an official mark.

**Method claim.** Judgment is not outsourced. Sources are primary. Models are instruments.

**Status language.** Established / synthesis / hypothesis / proposed construct.

**Evidence classes.** Demonstrated / pre-consensus / hypothesized / capitalized / neglected.

**Date.** 2 September 2026. Expands the 1 September working brief.

**Public / private split.** Author voice, public inquiry surface, engineering organization, and raw working corpus are not the same object. This brief is working material.

---

## Contents

---

## 1. What this brief is for

The question is usually asked as if it were personal: how do you stay ahead of AI. That is the wrong unit. A person can stay informed and an organization can still become fashionable, captured, or late. The live problem is how a network stays current across technical change, institutional change, and human judgment, without mistaking attention for evidence.

This brief answers four linked questions.

**Desk.** What is actually being demonstrated, assumed, funded, standardized, or abandoned.

**Field.** What systems do when inserted into work, institutions, and relationships, including among the people who build them.

**People.** What to look for in builders, not only in artifacts. Quality of questions, revision under pressure, perspectival range, and whether someone can hold authority without collapsing meaning.

**Organization.** How the network differentiates itself by mechanism and constraint rather than by slogan, launch, or proximity to the current stack.

AI is one surface. The object is larger: information architectures as governance, alternative computation, verification, cybernetic feedback, and the conditions under which people remain able to see, contest, and stay in reciprocal relation with realities other than their own.

If this brief is used inside a Programs and Policy role, those four questions still hold. The job adds a fifth constraint: every material observation has to be handable as a decision about instrument (invest, grant, convene, policy, research, field-build, connect, watch, decline), not only as a reading list.

---

## 2. Operating stance

Thinking, priorities, and reasoning stay with the people accountable for them. Critical thinking is not offloaded. The research orientation has been built incrementally across safety-critical engineering, identity infrastructure, organizational systems, and governance, then tested across domains.

The organization does not treat popularity, funding, or consensus as proof. Those are signals about where a field is moving, which assumptions are becoming infrastructure, and what may be important precisely because it is not yet receiving attention.

Commit firmly. Believe provisionally. A position must be strong enough to investigate and loose enough to revise. Updating the structure of the model is the point. Tacking exceptions onto a favorite story is not.

### What this stance refuses

- Trend recitation as strategy.
- Vendor narratives as technical fact.
- Benchmark theater without mechanism.
- Centralized "correct values" as a substitute for contestable institutions.
- Interdisciplinarity as decoration.
- Confusing warmth, fluency, or confidence with epistemic competence.
- Auto-summaries that replace the primary source.
- Scoring people.
- Perpetrator-intervention as the center of the work. Harm, when it appears, is a case. It is not the thesis.

---

## 3. Desk report: what the literature and infrastructure are doing

Desk work is not "keeping up with AI news." It is classifying claims by what they actually are.

## 3.1 Four separations that keep the desk honest

**Demonstrated.** Measured under stated conditions, with known limits.

**Hypothesized.** Coherent and testable, not yet established.

**Pre-consensus.** Serious people are arguing it before it hardens into a literature.

**Capitalized or fashionable.** Money, hiring, or zeitgeist are concentrating here. Useful as a map of lock-in. Not evidence of correctness.

A fifth class is required for the operating system and is already implied by the method:

**Neglected.** Load-bearing in the field or in infrastructure, thin in the rooms the network currently occupies.

## 3.2 Families under watch, not product cycles

Watch families of questions. Product names expire. Mechanisms do not.

| FAMILY | WHAT THE DESK IS ACTUALLY LOOKING AT |
|---|---|
| Reality-construction architectures | Ranking, defaults, labels, compression, persona training. Who governs salience, legitimacy, friction, and whose interiority survives reduction to a token. |
| Beyond-LLM computation | Neuromorphic and event-driven systems, analog and in-memory, photonic and reservoir/physical computing, edge and heterogeneous stacks. Mechanism, energy, latency, robustness, data movement. Not "brain-like" branding. |
| World models and latent prediction | Whether accumulated heuristics against representational collapse can be replaced by a principled constraint. JEPA-family work including LeVJEPA is interesting for that reason, not because it is new. |
| Assurance | Verification, validation, provenance, decision lineage, runtime constraints, human intervention. What can be shown to have happened versus what can be asserted afterward. |
| Second-order governance | The observer, designer, and governor are inside the system. Requisite variety. Good-regulator problems. Homeorhesis rather than a fantasy of fixed equilibrium. |
| Polycentric design | Multiple bounded decision centers, appeal, fork, revision. Coordination without a monopoly over interpretation. Decentralization that creates many small tyrannies is a failure mode, not a victory. |
| Relational side effects | Sycophancy, artificial intimacy, uncertainty communication. Research questions, not catastrophe theater. Aligned to whom. |
| Attention economies | Including gendered antagonism as a possible self-reinforcing product. Benefit in multiple currencies. Designed benefit versus emergent benefit. No conspiracy shortcut. |

These families remain the question map. They are not the top-level filing system for a role that also has to watch who can make something become infrastructure. That filing system is

in sections 13 and 14.

### 3.3 Desk sources, ranked by what they can prove

**Formal evidence.** Papers, preprints, proceedings, dissertations, technical reports, benchmarks, datasets, code and checkpoints, replication, standards, patents, regulatory text, safety cases, incident reports.

**Live scientific signal.** Lectures, workshops, lab talks, GitHub issues and PRs, researcher Discords, technical lists, long-form scientific conversation. Pre-consensus, not settled fact.

**Institutional and capital signal.** Grants, hiring clusters, acquisitions, RFPs, standards activity, compute and energy buildout, permitting, supply chain. Map of which assumptions are being poured into concrete.

**Counter-signal.** Negative results, failed replications, abandoned architectures, companies that die after a clean demo, benchmark gaming, retractions, impressive systems nobody adopts, collapsed investment theses, researcher disagreement, evidence that would kill the house hypothesis.

If counter-signal is not first-class, "we do not follow hype" is a temperament. It is not a method.

### 3.4 Recurring desk loop

Signal → structure → consequence → counterevidence → revision condition.

Record, for every material observation: what was seen, source, evidence status, inference, hypothesis, uncertainty, counterevidence, and what would force a revision. That log is the organization's memory. Without it, the network only has vibes and slides.

---

## 4. Field report: what desk work cannot see

Papers describe design. Benchmarks describe measurement under a protocol. Fieldwork describes what the system does when inserted into reality, including among the people who design, ship, maintain, and live with it.

### 4.1 Where the field is

- Operators, maintainers, on-call engineers, safety reviewers, and people who inherit other people's architecture.
- Users and affected communities, especially those with the least standing to name the frame as a frame.

- Rooms where builders argue before the paper exists: workshops, Discords, lab meetings, unrecorded arguments after the talk.
- Hiring, promotion, and review rituals. These reveal what an organization actually rewards.
- Direct conversation and long-form inquiry, including a podcast used as a research instrument rather than as a content channel.
- Procurement, permitting, interconnection queues, and insurance. These rooms decide what can exist outside a slide.

### 4.2 What field sensing is looking for

Not anecdotes that confirm the thesis. Discriminating observations.

- **Workarounds.** The informal system that makes the official system survivable.
- **What verification cannot see.** Excluded inputs, silent defaults, semantic capture, local vetoes inside "decentralized" structures.
- **Status and belonging economies** inside technical communities. Who can dissent without exile.
- **Transfer.** Whether a builder's stated philosophy survives contact with shipping pressure, metrics, and audience.
- **Missing participants.** Whose consequences are outside the model that is being optimized.

### 4.3 Field method constraints

Fieldwork is not immersion as vibe. It has the same evidence discipline as the desk. Observation is separated from inference. A conversation is not a finding. A single site is not a system. Positionality is logged: which rooms are reachable, which are closed, which of the observer's categories are doing governance work on the material.

The researcher is inside the system. Second-order cybernetics is not a citation here. It is a protocol. If the network only talks to people who already share the frame, the field report will launder confirmation as intelligence.

---

## 5. What the organization looks for in people who build

Artifacts are downstream. The organization that only evaluates what got built will keep hiring people who can produce a demo inside the current reward function. The more consequential question is how a person searches a possibility space, holds authority, and updates when the model breaks.

This is not an IQ screen. It is not ideological conformity named as emotional intelligence. It is a burden that should rise with the authority a role has to construct consequential realities for

other people.

## 5.1 Builder criteria

| LOOK FOR | WHAT IT ACTUALLY MEANS IN PRACTICE |
|---|---|
| Quality of questions | Can they ask questions that could change the model, not only questions that confirm it. Low-dimensional questions presuppose the answer. |
| Perspectival bandwidth | Can they represent, relate, and hold multiple causally relevant perspectives without prematurely collapsing them into one story. Listing viewpoints is not the same as using them. |
| Metacognitive governance competence | Can they inspect their own interpretive process, name incentives shaping the model, and revise the structure of the model when evidence requires it. |
| Mechanism over workaround | Do they reach for a principled constraint, architecture, or objective, or do they accumulate heuristics until the system cannot be explained. |
| Falsification literacy | What would make them change their mind. If they cannot answer, they are defending identity, not investigating a system. |
| Affect tolerance | Can they stay in contact with disagreement, uncertainty, and status threat without domination, dismissal, or sycophancy. Warmth without independence is not competence. |
| Scale and time | Can they move between individual, relational, institutional, and technical scales, and ask what happens if the intervention succeeds. |
| Missing-variable reflex | Who is not in the room. What information was discarded before the record existed. Which population experiences this differently. |
| Authority hygiene | Do they make authority, uncertainty, and revision conditions explicit. Do they confuse provenance with moral truth. |
| Non-heroic shipping | Can they maintain, document, hand off, and live with failure modes. Builders who only want the greenfield problem are incomplete for governance-bearing systems. |

## 5.2 How to assess without creating a priesthood

Who vets the vetters. If this becomes "the smartest people should rule," the organization has reproduced the centralized epistemic hierarchy it claims to oppose.

Assessment, where stakes justify it, should be polycentric and task-based: scenario reasoning, adversarial cases, stakeholder modeling, uncertainty calibration, steelmanning, causal mapping, demonstrated response to being wrong. Not one test. Not one institution's idea of rationality. Not the house politics wearing a competence mask.

Watch what a candidate does with a rival explanation that would, if true, cost them status. That is more informative than their publication list.

Do not score people. Keep examples against the criteria. A profile of a collaborator is evidence, not a rating.

### 5.3 Roles the network is structurally short of

A research organization becomes fashionable when everyone in it can narrate the same stack. Requisite variety in the network is part of the method.

- Control theorists and systems-safety engineers who will ask about observability, controllability, and failure boundaries.
- Formal-methods and assurance people who will refuse a story that cannot be checked.
- Device physicists, photonics and materials people who live at the constraint the software story ignores.
- Operators and maintainers who know the unofficial system.
- Institutional economists, STS, and public-administration scholars who treat rules and feedback as empirical objects.
- Regular counterparties who reject the frame: a measurement-maximalist, a labor-process skeptic, an engineer who thinks the philosophy is epiphenomenal. Without them the network confirms itself.

Section 11 turns this from a reminder into a list. Section 31 names what is still thin after that list.

### 5.4 What we do not select for

- Speed of opinion in public.
- Fluency in the current moral vocabulary of a scene.
- A portfolio of demos with no account of what was measured, discarded, or abandoned.
- Loyalty to a stack, a lab brand, or a funding narrative.
- The ability to make uncertainty sound like a finished story.

---

## 6. How the organization stays current

An organization does not stay current by adding a Slack channel called #ai-news. It stays current by institutionalizing the same loop a careful researcher already runs, then making the output handable to someone else on a clock.

### 6.1 Organizational sensing cadence

| CADENCE | WHAT GETS PRODUCED |
| --- | --- |

| | |
|---|---|
| Weekly | Pre-consensus notes: what serious rooms argued, what shipped, what failed, evidence status attached. No recap theater. |
| Monthly | Capital and infrastructure scan: hiring, grants, RFPs, energy and compute, standards movement. What is being poured into the floor. |
| Quarterly | Model revision review: which house assumptions took a hit, which constructs still lack measurement, which neglected line now has a mechanism. |
| Ad hoc | Incident and anomaly briefs: when a system, a community, or a public case compresses a person or a claim into a token. Use the seven-question inquiry method, not a verdict. |

After an event: same-day observation, one-week synthesis, one-month persistence test. If the signal did not survive 30 days, record that. Fashion often does not.

### 6.2 What "current" means here

Current does not mean first to post. It means the organization can distinguish, on demand:

- what is demonstrated
- what is hypothesized
- what people are beginning to investigate
- what capital is betting on
- what the scientific zeitgeist currently rewards
- what may matter because it is not receiving attention

An organization that can only speak the third and fifth of those is a marketing department. An organization that can only speak the first is late to lock-in. Both failures are common.

### 6.3 How sensing becomes a decision artifact

Field sensing that cannot be handed to another person as a revisable brief will look like taste. Every material note should answer: what changed, what is demonstrated, what is being baked into infrastructure, what is neglected, and what would change our mind this quarter. If it cannot survive that template, it is not ready to steer a program.

Cybernetic primitives are useful here as analytics, not as decoration: feedback, feedforward, requisite variety, the good-regulator problem, homeorhesis, circular causality, observability, controllability. A low-variety model trying to regulate a high-variety world is an organizational failure mode, not only a technical one.

## 7. Differentiation without hype

Relevance is not proximity to the loudest stack. Differentiation is not a unique slogan. The organization stays relevant by being the place that can still see the frame as a frame after the market has named the frame "reality."

### 7.1 What would make this organization interchangeable

- Publishing takes on every model release.
- Borrowing the vocabulary of the current lab and calling it a thesis.
- Growing the lexicon faster than measurement.
- Collapsing Tiffany, the public inquiry surface, the engineering organization, and the raw working corpus into one public object.
- Replacing one centralized optimizer with a morally preferred centralized optimizer and calling that polycentric.

### 7.2 Actual differentiators

- **Mechanism over narrative.** Ask what the system does, what it depends on, what it discards, and how feedback changes behavior.
- **Architecture over patching.** Prefer a constraint that removes a family of failure modes over a policy that promises good conduct.
- **Verification before interpretation.** Show what happened. Do not pretend a hash proves a worldview.
- **Coordination without homogenization.** Enough shared structure to inspect and appeal. Not a Codex that silently owns meaning.
- **Reciprocity as a design objective.** Not "make users believe the correct thing." Preserve the ability to encounter independent subjectivities.
- **Status language in public.** Established, synthesis, hypothesis, proposed construct. If the organization will not say which, it is performing certainty.
- **Case before total theory.** Gender-war political economy, symbolic compression in public controversy, sycophancy in persona-tuned systems: cases. Not a master pipeline from content to harm.

### 7.3 Hype tests

Before the organization amplifies a technical or social claim, run these.

1. What, exactly, was demonstrated, under which conditions, against which baseline.
2. Which heuristics, proxies, or human workarounds is the result still depending on.

3. Who benefits if this becomes the story, and in which currency: money, attention, data, status, authority, certainty, moral legitimacy, constrained exit.
4. What would look the same if the claim were false.
5. What the field does with this once operators, not authors, own it.
6. Whether amplifying it adds information or only adds repetition.

---

## 8. How this is done in the role

In a role, the method has to leave the notebook. The job is to keep the institution's model of the technical and social environment revisable, and to keep the people with authority over that environment from being selected only for fluency inside the current reward function.

If the role is Programs and Policy / Responsible Technology and AI, the unit of work is not a newsletter. It is five motions (section 12) across four domains (section 13), with an instrument attached to every material signal.

### Week-one posture

Map existing sensing: who is already watching formal evidence, capital, operators, and counter-signal. Name the holes. Do not add a newsletter on top of an unowned map.

### Ongoing posture

Hold the four-layer source practice. Keep builder criteria in hiring and collaboration, not only in essays. Force status language in internal memos. Protect the public/private split: author voice, public inquiry surface, engineering organization, and raw corpus are not the same object. Produce quarterly revision conditions, not annual thought leadership.

### What success looks like

Not being early on Twitter. Being able to tell leadership, partners, and builders, without theater: this is demonstrated, this is fashion, this is lock-in, this is neglected, this is what would change the bet, and these are the people whose questions we do not currently have in the room.

Counts that do not count: articles read, events attended, takes published.

Counts that do: signals independently corroborated; developments detected before they were fashionable; relationships that became reciprocal; missing perspectives added; introductions that created useful connections; opportunities originated; weak opportunities rejected early; hypotheses changed; decisions made more reversible; counterevidence surfaced; institutional points of view revised.

## 9. What is still missing

The source practice is already broader than "read papers and follow smart people." The remainders are structural.

- **Named counter-signal.** Give negative results and dead architectures the same care as emerging work.
- **Positionality protocol.** Log which categories are governing the material and where network variety is low.
- **Measurement before lexicon growth.** Proposed constructs need dimensions and discriminability more than they need new names.
- **Cybernetics as analytics, not a sticker.** Feedback, variety, observer-inclusion, multi-level control. Use the primitives. Do not file the program under a dead interdisciplinary slogan.
- **Assurance failure modes.** Provenance is necessary and insufficient. Study what verification cannot see.
- **Counterparties in the language of the work.** People who will reject the frame without being unserious.
- **Role translation.** Sensing that cannot become a revisable brief on a clock will be read as taste.

The sections after this one are the attempt to make those remainders operational: a function-first network, then the operating surfaces that used to live in software.

They do not close the list. A list that claims to be complete is already a coverage failure.

---

## 10. Inquiry method for field cases

Use when the organization studies a public controversy, a product behavior, a community, or an architecture. Do not use as a verdict machine.

1. What happened.
2. What information architecture made some aspects salient and others invisible.
3. What human need, norm, or vulnerability interacted with it.
4. Who benefited, and in which currency.
5. Which rival explanations still survive.
6. What could interrupt the reinforcing loop.
7. What is known, inferred, contested, or speculative.

## 11. Network matrix: people, organizations, labs, and bodies by function

Not a hall of fame. A variety map. Talk to them. Do not outsource judgment.

Nodes are chosen for **function**. Counterparties are load-bearing. Missing seats are named as functions, not filled with invented individuals. Coverage language describes the network's contact with a seat, not the merit of a person.

**Medium.** Person / organization / lab / body.

**Seat.** The job they do in the sensing system: assurance, builder, academic, civil, labor, operator, security, infrastructure, policy, standards, competition, journalist, affected, funder, investor.

A person can sit in more than one seat. The matrix still files them where they are most load-bearing for this brief. Domain tags are the four ON priority domains plus the older question families when those are the reason the node is here.

The first list is the original variety map from the 1 September brief. The second list is the functional matrix: original nodes plus the intermediaries, constraint-setters, and implementation actors that the desk was structurally short of.

### 11.0 Original variety map (kept)

**Philosophy and social epistemology.** Elizabeth Anderson; Miranda Fricker as contrast; Rae Langton as antecedent to distinguish; Karen Yeung; Just and Latzer; C. Thi Nguyen; Emily C. R. Tilton.

**Polycentricity and cybernetic governance.** Elinor and Vincent Ostrom; Paul Lewis and Paul Dragos Aligica on the Ostroms and cybernetics; Shagun Jhaver, Seth Frey, Amy X. Zhang; Michael McGinnis; Stafford Beer readers who still do empirical organizational work; Raul Espejo.

**Observers and meaning.** von Foerster lineage; Bateson lineage; Gordon Pask; Maturana and Varela readers working on autonomy without mysticism; ISA RC51 sociocybernetics.

**Empirical feed, norms, and relational AI.** 2026 feed-ranking registered-report authors; Renström and Bäck; dynamic-norms line; OpenAI/MIT affective-use authors; warm-persona and sycophancy paper authors.

**World models.** Yann LeCun; Randall Balestriero; V-JEPA authors (Assran, Bardes, Ballas and collaborators); Kuhn, Maes, Le Lidec, Buettner on LeVJEPA and LeWorldModel; David Klindt; contrastive anti-collapse work (Boylan, Hokamp).

**Neuromorphic and unconventional compute.** Giacomo Indiveri; Tobi Delbruck; Yulia Sandamirskaya; Chiara Bartolozzi; Gert Cauwenberghs; James Aimone / Sandia; Dhireesha

Kudithipudi and coauthors; Jason Eshraghian; Loihi circle as industrial signal; Telluride / Brains and Machines circuit as a field room.

**Funders as sensors, not oracles.** NSF NeuroAI and CISE/ENG officers; DARPA and IARPA microsystems portfolios; DOE / NNSA lab computing; Wellcome and NIH-adjacent neurotech; Sloan, Simons, Templeton as fashion-and-gap detectors; whoever is actually writing checks for analog, photonic, and in-memory work rather than another wrapper on the same stack.

## 11.1 How to read the functional matrix

Each block is a seat. Inside the seat, nodes are grouped by medium. A seat with many academics and no operators is not "covered." It is academically occupied.

**Strong** in a coverage note means the network can currently name a serious person or body and reach them. **Thin** means a function is named but contact, evidence, or rivalry is weak. **Absent** means the seat exists in the world and is not yet in this map as a reachable node. Whole-seat paint is forbidden: one missing node does not make an occupied seat absent.

Named individuals are starting points. Lineages, desks, and functions matter more than celebrity.

## 11.2 Assurance

The people and bodies who will eventually say a system is safe enough, compliant enough, insurable enough, or auditable enough, and the researchers who will refuse that sentence if it cannot be checked. Provenance is necessary and insufficient. Who audits the auditors is an open question, not a slogan.

| NODE | MEDIUM | DOMAIN | WHY THEY ARE HERE |
|---|---|---|---|
| Inioluwa Deborah Raji | person | Frontier AI | Evaluation, documentation, and the politics of what gets measured. |
| Beth Barnes / METR | person / org | Frontier AI | Independent evals of frontier models. Pre-consensus on what "evals" can prove. |
| Marius Hobbhahn / Apollo Research | person / org | Frontier AI | Model organisms, deceptive-behavior evals. Treat as a research program, not as settled risk. |
| Rumman Chowdhury / Humane Intelligence | person / org | Frontier AI | Red-teaming as a practice with community and labor inside it. |
| NIST TEVV researchers | body | Frontier AI | Test, evaluation, validation, verification. Draft TEVV-Athlon (August 2026) is a standards-and-method signal, not a finding that TEVV is solved. |
| MITRE ATLAS contributors | org / body | Frontier AI / security | ATLAS map of claimed AI attack classes. Use as a map. Not as a completed threat model. |

| | | | |
|---|---|---|---|
| Independent safety-case practitioners | function | all | The people who will write, or refuse, a safety case. Seat is named. Individuals still thin. |
| AI incident databases | org | Frontier AI | Post-hoc record. Useful as counter-signal. Incident databases do not equal causal accounts. |
| Measurement-maximalist psychologist | person (counterparty) | Adjacent | Frame-rejecter: if it cannot be measured as specified, it is not yet a finding. Required. |
| Measurement-maximalist nuclear safety-case expert | person (counterparty) | Nuclear | Frame-rejecter from a domain that already has a safety-case culture. Required. |

**Thin.** Insurers, reinsurers, actuarial shops, certification bodies, corporate internal audit, rating startups. The assurance *economy* is not the same as assurance *research*. This seat currently overweights researchers and underweights people who will sign.

### 11.3 Builder

People who search a possibility space and ship. Artifacts age. Selection is for the quality of their model of reality, and whether that model survives contact with metrics.

| NODE | MEDIUM | DOMAIN | WHY THEY ARE HERE |
|---|---|---|---|
| Yann LeCun | person | Frontier AI | World-model / JEPA program as a live alternative to next-token hegemony. Watch the mechanism, not the brand. |
| Randall Balestriero | person | Frontier AI | Geometry of representations; collaborator on JEPA-family constraints. |
| Mido Assran, Adrien Bardes, Nicolas Ballas and collaborators | person | Frontier AI | V-JEPA authors. Treat as a research line. |
| Lukas Kuhn, Lucas Maes, Quentin Le Lidec, Buettner | person | Frontier AI | LeVJEPA / LeWorldModel. Interesting if a principled constraint replaces collapse-prevention heuristics. Hypothesis until measured. |
| David Klindt | person | Frontier AI | Representation structure; useful as a check on slogan-level "world models." |
| Boylan, Hokamp | person | Frontier AI | Contrastive anti-collapse work. Mechanism question, not a product. |
| Christopher Langton | person | Adjacent | Artificial life as a reminder that computation-as-life is an old claim with a specific history. |
| Carver Mead | person | Adjacent | Neuromorphic and analog lineage. Constraint-first hardware thinking. |

| | | | |
|---|---|---|---|
| Andra Keay / SV Robotics | person / org | Adjacent | Robotics field room, not a lab demo. |
| Open-source model maintainers | function | Frontier AI | People who keep weights, evals, and forks alive after the launch post. Thin as named individuals. |
| An engineer who thinks the philosophy is epiphenomenal | function (counterparty) | all | Required. If this person cannot be named, the network is a seminar. |

**Thin.** Device physicists, photonics and materials people, analog and in-memory builders who are not on the conference circuit. The software story still occupies this seat.

### 11.4 Academic (philosophy, governance, cybernetics, empirical social, neurotech)

The people who treat rules, meaning, and feedback as empirical objects, plus the labs that still argue before the paper exists.

| NODE | MEDIUM | DOMAIN | WHY THEY ARE HERE |
|---|---|---|---|
| Elizabeth Anderson | person | Adjacent | Relational equality, epistemic justice, and the democratic standing of people subject to others' engineering choices. Highest-authority correspondence in the underlying research. |
| Miranda Fricker | person (contrast) | Adjacent | Epistemic injustice as a contrast, not as a transplant. Keep the distinction. |
| Rae Langton | person (antecedent) | Adjacent | Speech-act / silencing lineage. Distinguish rather than absorb. |
| Karen Yeung | person | Frontier AI / policy | Algorithmic regulation. Legal-institutional, not vibes. |
| Just and Latzer | person | Adjacent | Governance by information architecture. Ranking and defaults as institutional facts. |
| C. Thi Nguyen | person | Adjacent | Games, echo chambers, opacity of expert systems. Useful on value capture. Hypothesis-adjacent; do not over-apply. |
| Emily C. R. Tilton | person | Adjacent | Epistemic justice extensions. Keep as contrast literature. |
| Elinor Ostrom | person | Adjacent | Polycentric governance as empirical institutional work, not a slogan. |
| Vincent Ostrom | person | Adjacent | Democratic administration and polycentricity as a pair with Elinor, not a footnote. |
| Paul Lewis and Paul Dragos Aligica | person | Adjacent | Ostroms and cybernetics. One of the few explicit bridges. Read them. Do not paraphrase them into "complexity." |

| | | | |
|---|---|---|---|
| Michael McGinnis | person | Adjacent | Institutional analysis. Polycentricity as a research program. |
| Shagun Jhaver, Seth Frey, Amy X. Zhang | person | Adjacent | Community governance of platforms as empirical polycentricity / failure. |
| Stafford Beer readers who still do empirical organizational work | function | Adjacent | Viable System Model as analytics. Skip the sticker. |
| Raul Espejo | person | Adjacent | Organizational cybernetics still attached to actual organizations. |
| von Foerster lineage | function | Adjacent | Observer included. Second-order as protocol. |
| Bateson lineage | function | Adjacent | Levels of learning, schismogenesis. Use with restraint. |
| Gordon Pask | person | Adjacent | Conversation theory. Interaction as a control problem. |
| Maturana and Varela readers working on autonomy without mysticism | function | Adjacent | Autopoiesis is easy to abuse. Keep the readers who will not. |
| ISA RC51 sociocybernetics | body | Adjacent | A room, not a theory. |
| 2026 feed-ranking registered-report authors | function | Frontier AI | Ranking as an empirical object. Registered reports beat takes. |
| Renström and Bäck | person | Adjacent | Gender and political affect. Case literature, not a master pipeline. |
| Dynamic-norms line | function | Adjacent | Perceived norms as a lever. Status: hypothesized / pre-consensus depending on the paper. Do not moralize. |
| OpenAI/MIT affective-use authors | function | Frontier AI | Companion / affective use. Read methods. Do not cite the press release. |
| Warm-persona and sycophancy paper authors | function | Frontier AI | Relational side effects as a research question. Aligned to whom. |
| Giacomo Indiveri | person | Adjacent | Neuromorphic circuits as a living research program. |
| Tobi Delbruck | person | Adjacent | Event-based sensing. Mechanism, energy, latency. |
| Yulia Sandamirskaya | person | Adjacent | Neuromorphic for robotics / control. |
| Chiara Bartolozzi | person | Adjacent | Event-driven sensing and embodiment. |
| Gert Cauwenberghs | person | Adjacent | Analog / mixed-signal neuromorphic. |
| James Aimone / Sandia | person / lab | Adjacent | National-lab neuromorphic and computing. |

| | | | |
|---|---|---|---|
| Dhireesha Kudithipudi and coauthors | person | Adjacent | Lifelong learning on neuromorphic substrates. |
| Jason Eshraghian | person | Adjacent | Spiking networks as an engineering object. |
| Telluride / Institute of Neuromorphic Engineering | lab / room | Adjacent | Field room. Go. Talk to the unfriendly questioner. |
| Stanford HAI | org | Frontier AI | Academic-industry seam. Fashion-and-gap detector as much as a lab. |
| Berkeley BCLT / CITRIS | org | Frontier AI / policy | Law and technology, and a campus engineering room. Different functions; do not collapse. |
| NSF NeuroAI | body / program | Adjacent | Funder and research program. Sensor, not oracle. |
| Nita Farahany | person | BCI | Cognitive liberty / neurorights. Legal-ethical. Keep distinct from device engineering. |
| Sara Goering | person | BCI | Neuroethics with disability and agency inside it. |
| David Winickoff | person | BCI | OECD-adjacent anticipatory governance of neurotechnology. |
| Karen Rommelfanger | person | BCI | Neuroethics as practice. |
| Loihi circle | function / industrial signal | Adjacent | Intel neuromorphic as capital-and-hardware signal. Not proof of a brain-like future. |

**Thin.** Public-administration scholars who do not already speak "AI." Control theorists. Formal-methods people outside the safety-eval circuit.

## 11.5 Civil society and institutional power

Necessary counterweight to lab-centered governance discourse. Do not adopt a worldview because it is the available opposition. Do not ignore it because it is inconvenient.

| NODE | MEDIUM | DOMAIN | WHY THEY ARE HERE |
|---|---|---|---|
| Amba Kak / AI Now | person / org | Frontier AI | Institutional power, compute, accountability. |
| Sarah Myers West / AI Now | person / org | Frontier AI | Same shop, different emphasis. Keep both. |
| Timnit Gebru / DAIR | person / org | Frontier AI | Independent research outside lab capture. Disagreement with a lab is not a finding against the work, or for it. |
| Meredith Whittaker / Signal | person / org | Frontier AI | Surveillance, concentration, infrastructure politics. |

| | | | |
|---|---|---|---|
| Arvind Narayanan and Sayash Kapoor | person | Frontier AI | AI snake oil as a method: ask what was measured. |
| Safiya Umoja Noble | person | Adjacent | Ranking and racialized information architecture. Case and theory. Do not flatten into a slogan. |
| Ruha Benjamin | person | Adjacent | Coded inequity. Same caution. |
| danah boyd | person | Adjacent | Fieldwork on networked publics. Useful against desktop-only governance. |
| Tarleton Gillespie | person | Adjacent | Platforms as governance. Moderation as institutional fact. |
| Data & Society | org | Frontier AI / labor | Empirical work on labor, data, and institutions. Sits on the civil/labor seam. |
| Anti-nuclear civil-society counterparties | function | Nuclear | Required. A nuclear map with only labs and vendors is a vendor map. |

## 11.6 Labor

Not "labor" as a value. People who do the work that makes the official system run, and people who organize around insertion of AI into work.

| NODE | MEDIUM | DOMAIN | WHY THEY ARE HERE |
|---|---|---|---|
| National Employment Law Project | org | Frontier AI | Worker classification, surveillance, enforcement. |
| TechEquity Collaborative | org | Frontier AI | Bay Area labor and tech-power work. |
| Coworker.org | org | Frontier AI | Worker voice inside platforms. |
| Labor economists on AI insertion | function | Frontier AI | Task, wage, and bargaining evidence. Named individuals still thin. |
| Worker centers | function | Frontier AI | Field contact. Not interchangeable with think-tank labor. |
| On-call, annotation, and moderation workers | function | Frontier AI | The unofficial system. Seat is structurally absent as named, reachable nodes. |
| Data annotators, RLHF workers, content moderators | function | Frontier AI | Same. Do not write "labor" and then only meet directors. |
| A labor-process skeptic | function (counterparty) | all | Required. Deskilling, pacing, and managerial control as rival explanations. |

## 11.7 Operator and inheritor

People who keep a system running when the demo is over, and institutions that buy it.

| NODE | MEDIUM | DOMAIN | WHY THEY ARE HERE |
|---|---|---|---|
| Enterprise AI implementers | function | Frontier AI | What actually got deployed, patched, or quietly abandoned. |
| Teachers and clinicians inheriting AI systems | function | Frontier AI | High-consequence inheritors. Thin. |
| Public servants using AI systems | function | Frontier AI | Procurement meets daily work. |
| GSA and federal procurement officers | body / function | Frontier AI | An RFP can tell you more than a keynote. |
| Health-system and school-district buyers | function | Frontier AI | Institutional adoption. Quiet lock-in. |
| Nuclear operators | function | Nuclear | A safety culture with a memory. Do not visit only the advanced-reactor slide. |
| Support and on-call engineers | function | all | Unofficial authority. Field, not LinkedIn. |
| Safety reviewers | function | all | The people who can stop a ship. Ask what they are not allowed to stop. |
| Warehouse and data-center workers; electrical trades; chip-fab workers | function | Infrastructure | Physical layer of "AI." Currently almost absent as named contact. |

## 11.8 Security

A different threat model than public-interest technology. Required as a counter-discourse, not as a conversion.

| NODE | MEDIUM | DOMAIN | WHY THEY ARE HERE |
|---|---|---|---|
| Bruce Schneier | person | Frontier AI / security | Security as economics and institutions. |
| Ross Anderson's security-economics lineage | function | security | Incentives, not only bugs. |
| CISA | body | security | Operational federal cyber. AI-security teams as a watch, not a source of truth. |
| MITRE AI-security researchers | org / body | security | ATLAS and related. |
| Frontier-lab model-security teams | function | Frontier AI | Weight security, insider threat, evals under a national-security frame. Reach is limited and that fact should be logged. |

| | | | |
|---|---|---|---|
| Critical-infrastructure security researchers | function | Infrastructure | Grid, water, industrial control. Thin. |

Also watch, as questions not as a beat: model-weight security, supply-chain attacks, AI-enabled cyber operations, dual-use evaluation, biosecurity, export controls, defense procurement, classified/unclassified research interfaces.

## 11.9 Infrastructure (compute, grid, energy markets)

Where capital turns assumptions into physical lock-in. AI governance and energy governance are becoming the same problem in places. That sentence is a hypothesis about coupling, with some demonstrated institutional movement behind it. It is not a completed finding.

| NODE | MEDIUM | DOMAIN | WHY THEY ARE HERE |
|---|---|---|---|
| FERC | body | Nuclear + energy | Interstate electricity. Rulemaking as infrastructure. |
| NERC | body | Nuclear + energy | Reliability standards. |
| CAISO | body | Nuclear + energy | California market and interconnection reality. |
| PJM | body | Nuclear + energy | Largest RTO. Queue, capacity, reliability. |
| ERCOT | body | Nuclear + energy | Different market design. Do not generalize from CAISO. |
| DOE Office of Nuclear Energy | body | Nuclear | Civilian nuclear R&D and policy. |
| Idaho National Laboratory | lab | Nuclear | Nuclear systems, fuels, and (increasingly) computing overlays. |
| Argonne | lab | Nuclear / energy | National-lab computing and energy systems. |
| Oak Ridge | lab | Nuclear / energy | Same, with a different computing history. |
| Utility planners and ratepayer advocates | function | energy | Who pays, who waits, who is told the data center is inevitable. |
| Data-center developers | function | Frontier AI / energy | Land, water, interconnection, offtake. Capital sensor. |
| Semiconductor supply-chain analysts | function | Frontier AI | HBM, packaging, foundries, export controls. |
| Energy-market economists | function | energy | Capacity markets, PPAs, who subsidizes whom. |
| Transmission developers | function | energy | The constraint the software story ignores. |

Watch, as objects not as people: hyperscaler capacity, GPU/accelerator supply, advanced packaging, HBM, networking/interconnect, foundries, sovereign compute, cloud concentration, training partnerships, inference economics, chip export controls, data-center land, cooling/water, utility interconnection, permitting, power-purchase agreements, data-center financing, military/intelligence compute.

**Demonstrated institutional signal, not a completed technical finding.** DOE/NNSA announced a proposed 1 GW AI data center paired with on-site generation at Savannah River (July 2026). DOE also demonstrated AI-assisted translation of nuclear safety documentation into NRC licensing materials (March 2026). File both as infrastructure-plus-information-architecture signals. Ask what verification cannot see.

### 11.10 Nuclear (beyond the grid bodies)

| NODE | MEDIUM | DOMAIN | WHY THEY ARE HERE |
|---|---|---|---|
| NRC staff and commissioners | body | Nuclear | Licensing is the governance surface. |
| Advanced-reactor developers | function | Nuclear | Capitalized. Treat claims as claims. |
| Fuel-cycle researchers | function | Nuclear | HALEU and the rest of the unglamorous path. |
| Nuclear safety-case experts | function | Nuclear | Including the measurement-maximalist counterparty above. |
| Anti-nuclear civil-society counterparties | function | Nuclear | Repeated on purpose. |

Fusion is a separate speculative class. Do not file it under "advanced nuclear" to borrow seriousness.

### 11.11 Policy, law, procurement, liability

Legislation, regulation, enforcement, litigation, liability doctrine, contracts, insurance, antitrust, IP, privacy, consumer protection, and labor law create different incentives. Do not file them all under "policy."

| NODE | MEDIUM | DOMAIN | WHY THEY ARE HERE |
|---|---|---|---|
| FDA CDRH | body | BCI | Device regulation. Clinical BCI lives or dies here more than in a neurorights essay. |
| OECD AI and neurotech desks | body | Frontier AI / BCI | Anticipatory-governance work. International, slow, sometimes load-bearing. |
| DOJ Antitrust Division | body | Frontier AI | Market structure as a live federal practice. |
| FTC technologists | body | Frontier AI | Technical staff inside consumer-protection and competition. |

| | | | |
|---|---|---|---|
| GSA and federal procurement | body | Frontier AI | Repeated because procurement quietly determines deployment. |

Track as streams, even when the named node is still a function: state and municipal procurement, school districts, health systems, defense procurement, government contractors, approved-vendor lists, contractual AI requirements, indemnification, audit rights.

## 11.12 Standards

What is becoming governable through standards before it is governable through law.

| NODE | MEDIUM | DOMAIN | WHY THEY ARE HERE |
|---|---|---|---|
| NIST | body | Frontier AI | TEVV, AI standards, Zero Drafts, secure-by-design direction. Primary desk. |
| ISO/IEC JTC 1/SC 42 | body | Frontier AI | ISO/IEC 42001 (AI management), 42005 (impact assessment), TS 42119-2:2025 (risk-based testing). Watch participation and absence. |
| IEEE | body | all | Sectoral and AI standards. Also a Bay Area room. |
| IETF | body | Adjacent | Internet infrastructure standards. Easy to forget from an AI desk. |
| W3C | body | Adjacent | Identity, provenance, web platform. |
| C2PA | body | Frontier AI | Content provenance. Necessary and insufficient. |
| MITRE | org / body | Frontier AI / security | ATLAS and related. |

Each standard under watch should eventually carry: body, committee, stage, public-comment deadline, scope, technical mechanism, governance implications, interoperability impact, who is participating, who appears absent, instrument, review status.

**Named live objects (desk, 2025 to 2026).** NIST draft TEVV-Athlon (August 2026). NIST AI Standards Zero Drafts, including public-facing AI documentation. ISO/IEC TS 42119-2:2025. ISO 42001 paired with 42005. File as standards movement. Do not file as "AI is now governed."

## 11.13 Competition and market structure

| NODE | MEDIUM | DOMAIN | WHY THEY ARE HERE |
|---|---|---|---|
| Open Markets Institute | org | Frontier AI | Anti-concentration. One school. Keep rivals. |
| American Economic Liberties Project | org | Frontier AI | Same seam, different shop. |

| | | | |
|---|---|---|---|
| Platform-economics researchers | function | Frontier AI | Switching costs, vertical integration, what looks decentralized and is not. |
| Cloud and semiconductor market analysts | function | Frontier AI | Concentration at compute. |
| Antitrust economists (multiple schools) | function | Frontier AI | Deliberately mix. A single school is a faction. |

Questions that belong on this seat, not as a vibe: Who can enter? Who can exit? What requires permission? Who subsidizes whom? Where are switching costs appearing? Where are vertical integrations forming? What looks decentralized but depends on centralized infrastructure?

Watch concentration at: model, cloud, compute, chips, data, distribution, app stores, identity, payments, advertising, enterprise software, talent, research, energy.

## 11.14 Journalist

| NODE | MEDIUM | DOMAIN | WHY THEY ARE HERE |
|---|---|---|---|
| Investigative AI-beat journalists | function | all | Primary field sensors when they still do documents and not recap. Seat is almost absent as named nodes. |

People-calendars of serious reporters (talks, hearings, FOIAs) are a different source than an RSS publisher. Subscribe to the person, not only the outlet.

## 11.15 Affected parties

| NODE | MEDIUM | DOMAIN | WHY THEY ARE HERE |
|---|---|---|---|
| BCI trial participants and disability advocates | function | BCI | Lived consequence of implants, decoding, and abandonment. Not a mascot seat. |
| Rehabilitation specialists | function | BCI | What happens after the paper. |
| Neural-data / privacy scholars | function | BCI | Ownership, portability, secondary use. |
| Users with the least standing to name the frame as a frame | function | all | Methodological requirement from section 4. Not yet a named network. |

GAO has identified unresolved issues around brain-data ownership, long-term implant support, and insurance coverage. File as a governance gap with a citation, not as a completed ethics position.

## 11.16 Funder

Funders are sensors, not oracles. Track thesis, portfolio, repeated recipients, abandoned priorities, language shifts, and the people who passed on the fashionable category.

| NODE | MEDIUM | DOMAIN | WHY THEY ARE HERE |
|---|---|---|---|
| Open Philanthropy | org | Frontier AI | Large, opinionated, method-heavy. Map the thesis. Do not inherit it. |
| Good Ventures | org | Frontier AI | Adjacent capital. |
| Schmidt Sciences | org | Frontier AI / Adjacent | Academic-science scale. Fashion-and-gap detector. |
| Patrick J. McGovern Foundation | org | Frontier AI | Neuroscience / AI philanthropy seam. |
| Kapor Foundation | org | Frontier AI | Equity-frame tech philanthropy. Counterweight and also a zeitgeist. |
| Ford Foundation tech-and-society | org | Frontier AI | Institutional philanthropy with a longer memory. |
| MacArthur | org | Adjacent | Same. Watch abandoned priorities. |
| Mozilla Foundation | org | Frontier AI | Open-source and public-interest tech. |
| Siegel Family Endowment | org | Adjacent | Infrastructure / systems philanthropy. |
| NSF NeuroAI and CISE/ENG officers | body | Adjacent | Public science funding. |
| DARPA and IARPA microsystems portfolios | body | Adjacent | Unconventional compute as a defense problem. |
| DOE / NNSA lab computing | body | energy / Frontier AI | Compute and energy as one object. |
| Wellcome and NIH-adjacent neurotech | body | BCI | Biomedical frame, different constraints. |
| Sloan, Simons, Templeton | org | Adjacent | Fashion-and-gap detectors. Different metaphysics. Log that. |
| A funder who passed on the fashionable category | function (counterparty) | all | Often more informative than the people writing the current checks. Required. |

Philanthropy has a zeitgeist. Track what civil society currently thinks the problem is, as revealed by where money is actually going. Humanity AI's announced grants around democracy, worker rights, journalism, education, and AI civics are a portfolio signal about problem-framing. They are not evidence that those are the correct objects.

## 11.17 Investor

| NODE | MEDIUM | DOMAIN | WHY THEY ARE HERE |
|---|---|---|---|

| | | | |
|---|---|---|---|
| Frontier-tech VCs | function | Frontier AI | Capital formation around labs and tooling. Strong as a scene, thin as named, interrogable nodes in this map. |
| Deep-tech and energy investors | function | Nuclear + energy | Different time horizon. Do not collapse into "AI VC." |
| Generalist frontier-tech VCs | function | all | Useful for what they will not touch. |

## 11.18 Rooms that are not a person

These are field sites. Attend with one question. Talk to the unfriendly questioner, the postdoc, and the maintainer.

| NODE | MEDIUM | DOMAIN | ROOM QUALITY TO CHECK |
|---|---|---|---|
| Foresight Institute | org / room | Adjacent | Mechanism-heavy in places, capital-adjacent in others. Log which night it was. |
| Long Now Foundation | org / room | Adjacent | Long horizon. Can be serious or atmospheric. |
| SF Bay ACM | org / room | Frontier AI | Working engineers. Often more operator-rich than a "summit." |
| Telluride / Brains and Machines | lab / room | Adjacent | Neuromorphic field room. |
| Stanford and Berkeley seminar calendars | room | Frontier AI | Pre-consensus. People-calendars as well as org-calendars. |
| IEEE and engineering-society chapters | room | all | Operator and standards seam. |

Room quality tags, used as metadata not as scores: mechanism-heavy, capital-heavy, operator-rich, research-room, demo-surface, counterparty-rich. The point is to stop Luma from becoming the research agenda.

## 11.19 How not to use this matrix

- Do not treat it as a guest list for a conference that already agrees with itself.
- Do not collapse Tiffany's private correspondence, a public inquiry surface, and a funder's network into one graph.
- Do not add names to look complete. An honest function with no individual is better than a decorative individual.
- Do not score coverage as a single number. Report worst-in-seat as "N absent" or "thin on operators," not as a red paint over an occupied category.

- Talk to them. Do not outsource judgment.

---

## 12. Five kinds of work

If this method is used in a Programs and Policy role, the job is not "cover AI." The job is five motions. Everything else in this document should support one of them.

**Strategy.** Translate a responsible-technology thesis into action across the four domains, and revise the thesis when the field moves. Strategy that cannot name a revision condition is branding.

**Ecosystem.** Originate and maintain relationships across labs, companies, venture, philanthropy, academia, civil society, operators, and the people who will reject the frame. Many of the best opportunities arrive through this network. That is a fact about origination, not a reason to confuse a network with evidence.

**Convening.** Put constituencies that do not typically share a room into contact, with a question, an attribution rule, and a follow-through owner. A listening event with no artifact and no owner is hospitality.

**Point of view.** Write internal memos and, when warranted, public pieces. Some of this informs investment committees, board materials, and partner conversations. Status language is mandatory. Auto-generated institutional voice is forbidden.

**Capital.** Diligence, structure, memos, and the operational details of deploying money: equity, SAFE, grant, or a decision to wait. Capital is one instrument. It is not the answer to every signal.

The home question, every day, is not "what is trending." It is:

- What changed?
- What requires attention?
- Who should be known or talked to?
- What decision could this affect?
- Where are we under-covered?
- What is becoming infrastructure?
- Which instrument, if any?

---

## 13. Four subject areas

These are the priority domains of the role. They get permanent top-level treatment. The eight question families from section 3.2 remain the analytic layer inside them. Do not replace mechanism families with brand names.

Every domain page, or every domain section of a brief, answers the same questions:

- **Technology.** What changed technically?
- **Capital.** Who is financing it?
- **Governance.** What institutions or rules constrain it?
- **Power.** Who gains capability or authority?
- **Public interest.** Who bears consequences?
- **Network.** Who should be known?
- **Instrument.** Invest, grant, convene, policy, research, field-build, connect, watch, or decline?

### 13.1 Frontier AI

Models, architectures, agents, world models, evals, compute, open and closed ecosystems, safety, assurance, deployment, economic structure.

**Watch (mechanism, not product).** Next-token systems as the current lock-in; JEPA-family and other latent-prediction work as a hypothesized replacement of collapse-prevention heuristics; agents and computer-use as a deployment path that changes the assurance object; evals as a market and as a method; open-weight maintainers as a different political economy than lab APIs.

**House questions.** What can be shown to have happened. Aligned to whom. Who can enter. What becomes a default in enterprise and procurement.

**Typical counterparties.** Measurement-maximalist evaluator; labor-process skeptic; lab builder who thinks the philosophy is epiphenomenal; civil-society shop that will not accept "safety" as the only frame.

### 13.2 Nuclear and advanced energy

Nuclear, geothermal, grid, storage, generation, data-center power, transmission, interconnection, permitting, fuel, supply chains.

**Watch.** Interconnection queues; capacity markets; HALEU and fuel-cycle constraints; advanced-reactor licensing; data-center offtake; water and land; ratepayer incidence; AI used inside licensing and safety documentation.

**House questions.** Which energy assumptions are being poured into concrete for compute. What a safety case can see and what it cannot. Who pays when a campus or a county becomes a node in someone else's training run.

**Fusion.** Separate speculative class. Do not borrow nuclear's institutional seriousness by filing fusion under it.

**Geothermal, storage, gas, renewables.** First-class subdomains. An "advanced energy" desk that only talks SMRs is a fashion desk.

### 13.3 Brain-computer interfaces and neurotechnology

Devices, electrodes, neural decoding, clinical development, data, consent, ownership, rehabilitation, augmentation, long-term stewardship.

Break the stack rather than saying "BCI":

- Hardware: electrodes, implants, sensors, materials, packaging.
- Signal acquisition: EEG, ECoG, intracortical, other modalities.
- Decoding: signal processing, ML, representation.
- Actuation: cursor, speech, prosthetics, robotics.
- Clinical: trials, surgical practice, rehabilitation.
- Data: neural-data ownership, storage, portability.
- Long-term stewardship: who supports an implant after acquisition or bankruptcy.
- Governance: consent, privacy, neurorights, agency, identity, augmentation.
- Capital: who is financing what.

OECD's 2026 neurotechnology work is aligned as a process (anticipatory governance involving neuroscience, technology, policy, investment, and lived experience) rather than as a completed ethic. GAO's unresolved issues on brain-data ownership, implant support, and insurance are a demonstrated governance gap.

**Typical counterparties.** Device engineer; neurosurgeon; rehabilitation specialist; trial participant; privacy scholar; CDRH reviewer.

### 13.4 Adjacent and emerging

Neuromorphic, analog, photonic, in-memory, robotics, biological computing, quantum where relevant, semiconductor changes, novel interfaces, technologies that do not yet fit an established category.

This is where the original beyond-LLM family lives. It is also where a desk goes to die if it treats "adjacent" as a junk drawer. Each line still needs mechanism, energy, latency, robustness, and data movement. "Brain-like" is branding.

Quantum is not a default member. Include it when a mechanism, a capital movement, or a standards process makes it load-bearing. Otherwise it is zeitgeist.

---

## 14. Two separate lists that should not be merged

Do not make the top-level categories correspond only to intellectual interests. Three axes, kept separate.

### Domain

What is changing. The four above, plus, when needed as subdomains rather than as new religions: compute architectures, agents, world models, robotics, chips, data infrastructure, cyber and AI security, identity, information architectures, social platforms, companion AI, labor and workplace AI, education and youth, journalism and information ecosystems, public-sector AI.

### Governance surface

Where trajectory can actually be altered. Research, engineering, standards, regulation, procurement, litigation, liability, insurance, investment, grantmaking, labor organization, auditing and assurance, infrastructure permitting, competition policy, civil society, journalism, public opinion, community governance.

### Signal plane

What kind of evidence this is. Formal evidence, live scientific / pre-consensus, field / operator, capital, infrastructure, policy, cultural, counter-signal.

A capitalized claim with no formal evidence is a capital signal. It may still matter. It is not a demonstration.

---

## 15. Status language, next actions, and the five-step loop

### 15.1 Status of a claim

Established. Synthesis. Hypothesis. Proposed construct.

If the organization will not say which, it is performing certainty.

## 15.2 Evidence class of a signal

Demonstrated. Pre-consensus. Hypothesized. Capitalized. Neglected.

These can stack. A result can be demonstrated under a protocol and capitalized beyond what the protocol can support.

## 15.3 Instrument

Every material signal gets one, including "watch" and "decline."

| INSTRUMENT | USE WHEN |
|---|---|
| Watch | Too early, too thin, or the right action is to keep the object in view. Default. |
| Research | A mechanism or gap needs desk or field work before any other instrument. |
| Field | The question cannot be answered from papers. Go there. |
| Convene | The missing ingredient is the room, not more capital and not another memo. |
| Connect | Two nodes should know each other and currently do not. |
| Grant | Public-interest work that should not be forced through an equity story. |
| Invest | A for-profit path is the actual lever, and diligence can name what would kill it. |
| Policy | The lever is a rule, a docket, a procurement text, a standard, or an enforcement posture. |
| Decline | The fashionable option. Use it. Record why. |

Then ask: why this instrument? If that sentence cannot be written, the instrument is a reflex.

## 15.4 The loop

Most intelligence systems stop at act. This one does not.

**01 Observe.** RSS, papers, GitHub, conversations, events, capital, standards, policy, infrastructure. Capture: observed / inferred / heard / question / counterevidence.

**02 Structure.** Entities, domains, stacks, mechanisms, relationships, incentives, authority.

**03 Judge.** Evidence status, rival explanations, consequences, missing actors, counterevidence, confidence.

**04 Act.** Read, contact, convene, diligence, grant, invest, write, investigate, decline, wait.

**05 Revise.** What did we believe? What happened? What changed our mind? What did we fail to see?

Revision is the differentiator. A network that cannot show a changed mind is a content operation.

## 16. What lives in the dashboard instead of this brief

Sections that used to continue here (signal records, feeds, today's review list, opportunities, diligence checklists, convening logs, point-of-view queue, capital pipeline, standards watch, infrastructure watch, stacks, graveyard, events calendar, capture fields, review cadence) are files. They change when a document arrives or a conversation happens. They do not belong as frozen chapters.

Open `research-ops-dashboard.html`. Mark what has been reviewed. Do not treat a checked box as a finding.

What stays in this brief: the questions, the status language, who to look for in builders, the people matrix, the holes, and what this work will not become.

---

## 17. Coverage holes named as holes

The matrix in section 11 is occupied in the seats this practice already knew how to see: academic, funder, some civil, some infra bodies. It is thin where implementation actually happens.

**Structurally thin or absent, as of this edition:**

- Journalists as named, reachable nodes.
- Standards as a living desk (NIST is present; committees, comment calendars, and "who is absent from the room" are not yet a practice).
- Affected parties as people one can actually talk to, especially BCI trial participants and inheritors of workplace AI.
- Investors as interrogable nodes rather than a scene.
- Builders outside the JEPA / neuromorphic circuit: open-source maintainers, device physicists, analog and photonic people.
- Operators: on-call, annotation, moderation, teachers, clinicians, nuclear operators, trades on the data-center side.
- Assurance *economy*: insurers, certification bodies, internal audit, the people who will sign.
- Procurement officers as contacts, not as a category label.
- International governance beyond OECD-as-a-citation.
- A named funder who passed on the fashionable category.
- A named engineer who thinks the philosophy is epiphenomenal, in the language of this work.

Report coverage as "N absent in seat" or "thin on operators." Do not paint a whole seat absent because one node is missing. Do not paint a seat strong because it has many academics.

**What we are not seeing** should stay a standing section of any weekly note:

- missing counterparties
- domains without counter-signal
- domains without field evidence
- hypotheses without measures
- networks with low perspectival variety
- claims based primarily on capital signal
- high-confidence records sourced from one community
- decisions without affected-party representation

That is the missing-seat list. It is not a radar. It is not a heat map.

---

## 18. What this system will not become

- Generic KPI cards.
- An "AI trends" page.
- A red / yellow / green risk score.
- Opaque model-written summaries replacing primary sources.
- A forced single score for organizations or people.
- A graph visualization whose only function is to look like a network.
- An auto-generated institutional point of view.
- A giant ontology before distinctions have been shown to discriminate anything.
- A Polycentria product, mark, or lockup. If a design system is reused, it is ceramic, type, and constraint. It is not the identity of the work.
- An official Omidyar Network product. Candidate concept. No borrowed authority.
- A public object that collapses author voice, inquiry surface, engineering organization, and raw corpus.
- A perpetrator-intervention machine. Cases of harm are cases. They do not become the center of gravity.

Measurement before lexicon growth still applies to this document. If a new label in here cannot be used to reject an observation, cut it.

---

## 19. Closing

Staying ahead of technology is not a consumption habit. It is a governance problem inside the organization that claims to study governance. The network either selects for people who can hold plurality, friction, and revision, or it selects for people who can perform currency. Those two selection systems do not produce the same institution.

The work is to keep the second system from winning while still knowing, in detail, what the first system is up against: architectures that allocate attention, builders under metric pressure, capital pouring assumptions into infrastructure, and publics that reward verdicts more cheaply than inquiry.

That is the field. That is the desk. That is the hiring problem. That is the differentiation problem. They are the same problem at different scales.

The dashboard files the work. This brief is the argument. If inner work, research work, or organizational work never changes how people are treated, what authority is structured, or what would count as being wrong, it is not doing the job it describes. Method is only method if it leaves a revision mark.

This brief should itself be treated as provisional. The next field contact, failed replication, or infrastructure decision can require changing the map. That is the point.

Who governs the conditions under which we construct reality and one another, who benefits from those conditions, and how might we build institutions that preserve the capacity to see, contest, and remain in reciprocal relation with realities other than our own.

End of brief.