
Kitsune Ai システム / Physical Intelligence Lab

POSITION PAPER 00 / FOUNDATIONAL POSITION

PHYSICAL-FIRST INTELLIGENCE

A KITSUNE PHYSICAL-FIRST INTELLIGENCE POSITION

Human purpose, owner-governed data, recursive Perspective, and lived consequence in the age of AI

AI can return Perspective. Humans make it real.

**HUMAN INTENT + VIVO EVIDENCE + BOUNDED VITRO PROCESSING
-> REVIEWED PERSPECTIVE -> HUMAN EMBODIMENT -> LIVED CONSEQUENCE**

Published by Kitsune Ai System LLC

Kitsune Ai システム / Physical Intelligence Lab

Version 1.0 / Draft for scientific and field-partner review / September 2026

ABSTRACT / PHYSICAL INTELLIGENCE LAB

Abstract

Digital systems increasingly mediate everyday life through software that records behavior, predicts preference, recommends action, and, with growing frequency, acts on a person's behalf. Much of this intelligence is constructed from behavioral residue: clicks, transactions, locations, uploads, completion records, and model-generated summaries. These records can be useful, but they often preserve little about what a person intended, what effort occurred, what context changed, or why the resulting consequence mattered. As artificial intelligence becomes more capable, the risk is not only incorrect inference. It is that the processor becomes the practical center of the relationship: software frames the question, the model determines what data is relevant, and the person is left to approve, correct, or resist a direction already taking shape.

Physical-first intelligence proposes a different authority model. Technology is something people live with, not a digital environment they must live through. Seeded Earth, Place, and Space guardrails establish whether a person is physically present. That presence decision governs access through Earth → Place → Space → Thing → Utility: if the person is not present in the eligible context, the Utility does not open. Once admitted, the Utility binds a Task, an Action, and an Agent—Human Intelligence (HI) or Artificial Intelligence (AI)—to the active purpose.

This position moves beyond the conventional choice between keeping a human in the AI loop and allowing the human to fall outside it. Both descriptions make the machine process the primary loop. Physical-first intelligence instead places the human within the data loop and the AI within a larger human process. Intent leads to lived action. Selected evidence enters authorized processing. Perspective returns for review. The person must then interpret it through judgment, experience, body, relationships, craft, imagination, responsibility, and whatever abilities are available or developing. Only when the person carries something from that Perspective into action, creation, care, communication, or decision does digital possibility encounter lived constraint and consequence.

The data position is central. Vivo data is lived data with presence, captured in or out of a Moment. Vitro data is digital data produced by utilities or agent processing. Vitro can help make Vivo legible, but it cannot silently replace lived source evidence or inherit its ownership. The architecture joins three technologies that are usually positioned independently: Spatial establishes jurisdiction and presence-gated Utility access; AI performs bounded processing and returns Perspective; an authorized Block preserves proof of Data Well ownership and owner-authorized lifecycle transitions. The Block does not contain the whole Data Well, prove that an event was true, or grant ownership to a processor.

The position applies across disciplines because it defines a relationship rather than prescribing one product direction. Education remains the first developed scientific field expression because physical learning makes intent, embodied effort, role boundaries, screen burden, and developmental responsibility especially visible. It is the first rigorous test of the position, not the limit of its relevance. The general hypothesis is that framed intent, owner-governed evidence of situated activity, bounded AI processing, reviewed Perspective, and human embodiment can strengthen understanding and support better-situated action without transferring authority over the record or direction to the processor.

CLAIM STATUS

These are proposed mechanisms, not validated results. A science-first program must define comparison conditions, outcome measures, burden ceilings, authority boundaries, correction procedures, retention limits, and stopping rules. The purpose is not to place more AI into everyday life. It is to establish when a governed relationship with AI helps people understand and direct their own lives more effectively, and when it does not.

1 / PHYSICAL INTELLIGENCE LAB

The position

Physical-first intelligence is a position about authority, data, interpretation, and consequence. It is not a synonym for robotics, spatial computing, or software that happens to use sensors.

Kitsune Ai システム / Physical Intelligence Lab takes the position that the dominant digital-first posture is no longer an adequate foundation for human-facing intelligence. Digital-first does not mean the existence of software or screens. It names an architectural assumption: the application becomes the primary place where activity is organized, the profile becomes the primary representation of the person, and retained behavioral data becomes the primary basis for interpretation. Lived context is flattened so it can move through the application. The result is a discontinuity between what a person intended, what they actually did, and what the digital record later implies.

Artificial intelligence can make this discontinuity more consequential. A model can infer preference, classify behavior, summarize records, generate plans, and act across tools, but greater inferential capacity does not repair missing provenance. A location trace, uploaded image, score, purchase, or completion event may record an outcome while leaving purpose, effort, constraint, relationship, and meaning unresolved. The problem is not that inference is useless. The problem is allowing inference to substitute for evidence that could have been framed with the participant and later tested against lived consequence.

Physical-first intelligence reverses that authority. The physical world is not a peripheral sensor environment feeding a central digital system. It is where people form purposes, use things, cooperate, attempt tasks, encounter resistance, experience emotion, and bear consequence. Spatial jurisdiction determines where a Utility may be used; AI processes only what the active purpose authorizes; Block preserves proof of Data Well ownership and authorized continuity. The person should not have to become more digital for the system to become more intelligent. Technology should become more accountable to the person's physical and social reality—something to live with, not live through.

CORE DEFINITION

Physical-first intelligence is the governed capacity to derive useful, reviewable Perspective from intentional activity in lived contexts, then return that Perspective to human interpretation and embodiment while preserving provenance, purpose boundaries, data agency, and authority over consequence.

This definition locates physical intelligence in a relationship rather than a device. Human Intelligence contributes purpose, interpretation, judgment, responsibility, and the ability to make Perspective real. Artificial Intelligence contributes processing, memory, comparison, synthesis, simulation, and alternative expression. Vivo evidence grounds the relationship in attributable source material. Vitro processing makes proposals about that evidence. Human embodiment carries selected Perspective into the world. Reality answers through consequence, producing new Vivo evidence and a changed human understanding.

FOUNDATIONAL THESIS

The human is not merely reviewing the machine. The human is the irreplaceable point where digital possibility becomes physical reality.

1.1 / PHYSICAL INTELLIGENCE LAB

Orientation, human purpose, and choice

Access to information is not the same as orientation. Processing capacity is not the same as purpose.

People now encounter an extraordinary range of human knowledge through networked software and generative models. Yet access alone does not establish what matters, what can be trusted, how a claim relates to lived evidence, or what a person should do next. More information can widen understanding, but it can also intensify confusion when sources, purposes, incentives, and transformations are hidden. Physical-first intelligence treats orientation as a relationship among intent, evidence, processing, review, and consequence rather than as a larger supply of answers.

Artificial intelligence has no independent social mandate. People and institutions choose objectives, data access, incentives, interfaces, deployment contexts, and action rights. Anthropomorphizing the processor can obscure that human accountability. Manipulation performed through an AI-mediated surface remains connected to the people, organizations, and structures that selected the purpose and conditions of use. Governance must therefore remain legible at the human and institutional layer, not be reduced to a model's tone or stated intention.

A system may be technically personalized while giving the individual little meaningful choice. The processor may be selected by a platform, the retention policy by a vendor, the purpose by an employer or institution, and the resulting profile may follow the person across contexts. Physical-first intelligence proposes data agency instead: the practical ability to inspect, correct, direct, share, withhold, expire, and reuse the Perspective produced from one's activity. Data agency is not guaranteed by a settings page. It requires enforceable authority over purpose, source data, processor access, continuity, and consequence.

PURPOSE AND OBJECTIVE

A machine can be assigned an objective. Purpose belongs to the living and social context that gives an objective meaning, determines acceptable consequence, and remains capable of revising or refusing it.

The Lab is not trying to prescribe a morally purified human future. Human nature can be cooperative and coercive, generative and destructive, beautiful and ugly. A governed runtime cannot resolve that conflict. It can make purpose, provenance, authority, processing, and consequence more visible so no model, platform, institution, or economic position becomes architecturally omnipotent. The position is not neutral about hidden authority or coerced data use, but it does not claim the right to define the correct human life.

This is why balance matters. AI should be capable enough to offer unfamiliar Perspective, yet bounded enough that the person can question, compare, replace, or refuse the processor. Human direction should remain primary, yet human judgment should not be insulated from evidence, correction, or other legitimate perspectives. Physical consequence provides a return path that neither digital confidence nor institutional status can manufacture.

ORIENTATION PRINCIPLE

The purpose of the system is not to orient the person toward the system. It is to help the person orient themselves toward life.

Intellectual lineage

The position draws from established work on embodied and situated cognition, interaction, governance, symbiosis, agents, and provenance. Its synthesis remains a proposal requiring independent evaluation.

This position does not begin from the claim that cognition is confined to an individual brain or that meaning can be separated cleanly from action. Research in embodied cognition describes ways cognitive processes depend on bodily capacities and environmental interaction (Wilson, 2002). Situated cognition argues that knowledge is partly a product of the activity, context, and culture in which it develops (Brown, Collins, & Duguid, 1989). Distributed cognition shows how people, artifacts, representations, and coordinated practice can form cognitive systems whose properties cannot be reduced to any one participant (Hutchins, 1995). The extended-mind thesis similarly examines conditions under which external resources can participate in cognitive processes (Clark & Chalmers, 1998).

These traditions matter because physical intelligence is not defined by the physical location of a processor. A model running on a robot may be physically situated while remaining disconnected from the purpose and meaning of the human situation. A conversation running in a remote data center may still contribute to physical intelligence when a person governs the source data, interprets the returned Perspective, and embodies something from it in lived action. The relevant distinction is relational and consequential, not simply local versus remote compute.

Human-computer interaction has long examined what follows when computation moves beyond a desktop. Weiser's ubiquitous-computing vision sought technologies that recede into everyday life rather than continually demand attention (Weiser, 1991). Dourish's account of embodied interaction emphasized skilled, engaged practice and the way meaning emerges through interaction rather than being fully encoded in advance (Dourish, 2001). Human-AI interaction research emphasizes appropriate timing, communication of capability and uncertainty, correction, and user control (Amershi et al., 2019).

Human-computer symbiosis provides a direct historical precedent. Licklider (1960) proposed close coupling in which people set goals, formulate hypotheses, determine criteria, and evaluate results while computers perform routinizable work that supports insight and decision. Kitsune retains that division of authority and extends it through physical context, owner-governed continuity, explicit data jurisdiction, processor substitutability, and a reviewed return from digital expression to lived consequence.

TRADITION	CONTRIBUTION TO THIS POSITION	UNRESOLVED HERE
Embodied and situated cognition	Meaning and understanding emerge through body, activity, context, culture, and environment	How authorized evidence from lived activity enters AI processing and returns to consequence
Distributed and extended cognition	People, artifacts, representations, and tools may participate in a cognitive system	Who governs continuity, source data, transformation, and the authority to act
Ubiquitous and embodied interaction	Technology can recede into skilled everyday activity	How low-attention utilities remain purpose-bound and reviewable
Human-computer symbiosis	People direct purpose and judgment while computers support analysis	How recursive Perspective becomes physically grounded without processor sovereignty

2.1 / PHYSICAL INTELLIGENCE LAB

The unresolved gap

Existing traditions establish important foundations. They do not by themselves specify the governed relationship among human purpose, owner-governed data, AI processing, embodiment, and lived consequence.

Privacy and data-governance research supplies another part of the foundation. Contextual integrity evaluates information flows according to the norms and purposes of a specific context, not merely whether information is public or private (Nissenbaum, 2004). The W3C PROV family formalizes relations among entities, activities, and agents so the origin and transformation of information can be represented (W3C, 2013). NIST's AI Risk Management Framework treats AI as socio-technical and organizes risk work around governance, mapping, measurement, and management (Tabassi, 2023). These sources support purpose, provenance, and accountability, but they do not assign a complete runtime architecture for recursive human data relationships.

Frontier machine-learning research on world models and joint-embedding predictive architectures explores how models learn abstract representations and predict in representation space (LeCun, 2022; Assran et al., 2023). Robotics systems such as RT-2 connect learned visual-language representations to physical action (Brohan et al., 2023). Those efforts concern model architecture and machine action. Kitsune asks a different question: under whose purpose does a system act, using whose lived data, processed by which replaceable intelligence, interpreted by whom, made real through whose abilities, toward whose consequence, and with what return to the human afterward?

Agent research provides another lineage. Wooldridge and Jennings (1995) described autonomy, social ability, reactivity, and proactiveness as central agent properties. Contemporary language-model agents interleave reasoning with tool-mediated action, update plans from observations, and operate across multi-step environments (Yao et al., 2023). Current NIST work identifies agent identity, authorization, delegation, auditing, and data-flow provenance as emerging infrastructure questions as agents gain access to more tools and data (Booth et al., 2026; NIST, 2026). These concerns establish why capability must be separated from standing authority.

Research on recursively generated training data motivates a related source boundary. Shumailov et al. (2024) found that indiscriminate replacement of original data with model-generated data can produce model collapse and loss from the tails of the source distribution. Kazdan et al. (2025) showed that outcomes depend on workflow: accumulating synthetic data with real data can remain stable in studied settings, while fixed-size or replacement regimes can degrade. These findings do not validate Kitsune's architecture. They support the narrower position that synthetic or derived material is an instrument whose value depends on retained or verified source relations, not a self-renewing substitute for lived evidence.

RESEARCH GAP

Spatial systems, AI systems, and block-based proof systems are commonly positioned as separate movements. Physical-first intelligence joins them into one governed environment: Spatial decides eligibility, AI processes bounded evidence, Block preserves proof of Data Well ownership, and the human remains the source of purpose and authority over consequence.

The distinctive claim is therefore not that cognition is embodied, that agents require permissions, that AI can assist reflection, or that source data matters. It is that these concerns must be joined around a person-governed continuity in which physical consequence is part of the intelligence loop and no processor owns the source, the human relationship, or the next direction.


3 / PHYSICAL INTELLIGENCE LAB

Beyond human-in-the-loop

Human-in-the-loop and human-out-of-the-loop both begin with the machine process. Physical-first intelligence begins with the human process in which AI is allowed to participate.

Human-in-the-loop commonly describes systems in which a person supplies feedback, approves an action, corrects an output, or intervenes before consequence. These controls can be valuable. The phrase nevertheless carries an architectural implication: an AI loop exists, and the human is placed somewhere inside it. Human-out-of-the-loop removes or reduces that intervention. Both descriptions center the machine process.

Physical-first intelligence instead places the human within the data loop and the AI within a larger human process. The person frames or adopts a purpose, acts in a lived context, and produces potential source evidence. They decide what enters processing and under which Moment. An authorized processor returns Vitro proposals. The person reviews the proposal, but review is not the endpoint. The person must decide whether and how to carry the Perspective into the world. Reality then answers through consequence.



ELEMENT	IRREPLACEABLE CONTRIBUTION	AUTHORITY BOUNDARY
Human participant	Purpose, interpretation, judgment, embodiment, responsibility, and the ability to make Perspective real	May inspect, correct, reject, withhold, embody, stop, or redirect
AI processor	Comparison, memory, synthesis, simulation, alternative expression, and proposal	Acts only on authorized evidence within the active purpose; gains no ownership through processing
Vivo evidence	Lived data with presence, captured in or out of a Moment	May not be silently replaced by a generated account
Vitro output	A computational transformation or proposal about evidence	Remains attributed, uncertain, and revisable until accepted as part of a Perspective
Physical consequence	The real constraint, response, or change that follows human or authorized machine action	Cannot be manufactured by model confidence or interface persuasion

RELATIONAL PRINCIPLE

AI is not the loop. The human is not a checkpoint. The governed relationship is the loop.

3.1 / PHYSICAL INTELLIGENCE LAB

Humans make Perspective real

The return of Perspective does not complete physical intelligence. At that point, the output has not yet encountered life.

A Perspective may affect a person cognitively or psychologically while remaining entirely inside a digital exchange. It can reorganize attention, create confidence, provoke doubt, or suggest a course of action. Those effects matter, but they are not yet the full physical-first cycle. The output remains a computational expression until the person interprets it through their own judgment, experience, body, relationships, craft, imagination, responsibility, and whatever abilities they currently possess or are developing.

The person decides whether the Perspective is meaningful, whether it should be rejected or changed, and whether and how it can be carried into action, practice, creation, care, communication, or another lived decision. Only then does the processed expression encounter material, social, emotional, and practical conditions that the processor cannot fully manufacture. That act of human application is not a final approval step attached to an AI process. It is the constitutive transformation through which digital processing can become physical intelligence.

The resulting consequence produces new Vivo evidence: what actually happened, what changed, what failed, what was learned, and what the person now understands differently. That evidence may enter another bounded cycle. The recursion can deepen human understanding and make processors more useful within an accepted purpose, but the direction remains human. The person is not required to act merely because a Perspective was produced. Refusal, delay, reframing, or choosing another processor can also be legitimate physical-first outcomes.

THE DECISIVE BOUNDARY

The human is not merely reviewing the machine. The human is the irreplaceable point where digital possibility becomes physical reality.

A conversation with AI can express the beginning of this relationship. A person may not be seeking an isolated answer with every prompt. They may externalize an unfinished thought, examine how it is processed, recognize what the response reveals or distorts, and guide the next expression. The value is not only the output. It is the person's developing ability to understand the subject, question the processor, and direct the relationship. Yet the conversation alone remains primarily digital and cognitive. Its physical intelligence emerges when the person brings something from the exchange into lived experience and reality returns evidence.

This return matters because a recursive digital system can affect cognition, emotion, identity, and behavior while remaining insulated from physical evidence and human responsibility. A feed, ranking system, or conversational surface may repeat and optimize representations without requiring them to encounter the fuller reality of a person or community. The danger is not that digital interaction is unreal. Its effects can be intensely real. The danger is ungrounded recursion: influence without a reliable return through provenance, choice, embodied judgment, and consequence.

PHYSICAL RETURN

Physical-first intelligence does not create a more immersive digital loop. It repeatedly carries processed Perspective out of the digital environment, tests or expresses it through life, and returns with lived evidence.

4 / PHYSICAL INTELLIGENCE LAB

The data position

Physical intelligence begins by separating what occurred from what a processor says about what occurred. The terminology follows the distinction, not the other way around.

A person's activity is not automatically data. Life produces events, experiences, relationships, artifacts, changes, and consequences. A bounded purpose determines which part, if any, should become a digital record. That record is selective and fallible. It can preserve an observation, explanation, measurement, image, decision, artifact, timestamp, or relation under a defined protocol. It cannot contain the whole experience or guarantee truth.

A processor may transform that authorized evidence. It can classify, summarize, compare, associate, simulate, infer, question, or propose. Those transformations can be valuable, but they are not the source event. A fluent output can appear more coherent than the evidence supporting it. Physical-first intelligence therefore keeps source evidence, transformation, processor identity, uncertainty, human correction, and later consequence distinguishable.

LAYER	PLAIN-LANGUAGE MEANING	WHAT IT MUST NOT BECOME
Lived event	What a person or group actually encounters, attempts, creates, changes, or experiences	An automatic claim that every part of life should be captured
Selected source evidence	The minimum attributable record authorized for the active purpose	A universal profile or proof that an interpretation is true
Computational transformation	What a model or other processor derives, reorganizes, or proposes	A silent replacement for the source or a verdict about the person
Reviewed Perspective	A purpose-bound account accepted as useful enough for a human next step	An objective identity claim or standing permission to act
Lived consequence	What follows when a person or authorized system acts in the world	A reward signal assumed to contain the whole human meaning of the result

Intent is necessary because action traces are underdetermined. The same observable behavior can express different purposes, constraints, emotional states, or levels of understanding. Intent may be declared, adopted from a shared mission, negotiated with another person, or revised during activity. It is not presumed complete or perfectly sincere. Its value is that it gives later evidence an explicit point of comparison and gives the participant a basis for correction.

WORKING THESIS

DECLARED INTENT + SELECTED EVIDENCE OF LIVED EFFORT → DATA AGENCY

Data agency is the practical ability to inspect, correct, direct, share, withhold, expire, and reuse the Perspective produced from one's activity. Intent and evidence support that agency only when the system also preserves provenance, review, enforceable authority, and processor choice. Continuity is not automatic and cannot be claimed by the processor.

4.1 / PHYSICAL INTELLIGENCE LAB

Vivo, Vitro, and reviewed Perspective

Kitsune names the source and derived layers only after their relationship is clear.

Vivo data is lived data with presence, captured in or out of a Moment. Presence is operational: seeded Earth, Place, and Space guardrails determine whether the person is there. When presence is not confirmed for the eligible context, its Utility cannot be used. Vivo may include an observation, artifact, measurement, explanation, decision, image, sound, relationship, or consequence attributable to lived participation.

Vitro data is digital data produced by utilities or agent processing. It may include a transcription, classification, association, summary, model response, simulation, visualization, proposed plan, or inferred relation. Vitro does not mean false or inferior. It names the processing relation: the expression must remain linked to its Vivo evidence, Utility, Task, Action, Agent (HI or AI), instructions, uncertainty, and review history.

Perspective is not another raw data type. It is the purpose-bound return that an authorized participant has reviewed and accepted as useful enough to inform reflection or a next action. Perspective remains revisable. A correction should create a successor account while preserving the evidence and transformation history needed to understand what changed. No Perspective becomes a permanent truth about a person merely because it was accepted once.

STAGE	STATUS	AUTHORITY
Vivo evidence	Selected source evidence under the active Moment	Participant or authorized source governs admissibility and use
Vitro proposal	Attributed processor output with provenance and uncertainty	Processor may propose; it cannot promote, publish, or act beyond the contract
Human review	Interpretation, correction, comparison, withholding, or rejection	Authorized participant decides what, if anything, becomes Perspective
Human embodiment	Perspective is carried into a lived choice, practice, creation, relation, or action	The person remains responsible for translating digital possibility into human consequence
New Vivo evidence	Observed consequence and changed understanding	May enter a new Moment only under renewed purpose and authority

GENERAL HYPOTHESIS - TO BE TESTED

FRAMED INTENT + OWNER-GOVERNED VIVO EVIDENCE + BOUNDED VITRO PROCESSING → REVIEWED PERSPECTIVE → HUMAN EMBODIMENT → LIVED CONSEQUENCE + NEW VIVO EVIDENCE → RECURSIVE UNDERSTANDING

Each arrow requires an operational definition, comparison condition, and failure threshold. The hypothesis does not claim that more data, more processing, or more action is better. A useful cycle may end in correction, non-action, expiry, or rejection. The scientific question is whether the governed relationship improves role-appropriate understanding or action without disproportionate burden, coercion, exposure, or loss of human direction.

The physical Moment

A Moment is the bounded jurisdiction in which purpose, participants, evidence, processors, review, retention, and possible action become specific.

A Moment is not simply an event timestamp. It is a governed encounter that begins only after the applicable spatial admission rule has been satisfied. Seeded Earth, Place, and Space guardrails determine whether a person is present; if presence is not confirmed, the associated Utility remains unavailable. Once admitted, a person declares or adopts a Goal, uses or develops a Method, and the eligible Utility binds a Task, an Action, and an Agent—HI or AI—to the active purpose.

A utility may enter briefly Before, During, or After activity. Before may frame a question or confirm a purpose. During may involve no interface at all, or one deliberate capture. After may support comparison, reflection, or spatial review. The utility should disappear when continued attention adds no value. The Moment is designed around the activity rather than around the application's need for engagement.

JURISDICTION	INTENDED PURPOSE	DEFAULT BOUNDARY
Private	An individual or tightly authorized circle	Minimal disclosure; no inherited institutional or public use
Personal	Self-directed life, family, or trusted relationships	Participant-controlled sharing and continuity
Professional	Role-bound activity under an organizational or disciplinary purpose	Policy and responsibility apply only within the declared role and task
Public	Material prepared for possible wider release under explicit terms	Completing a Public Moment does not publish it; release remains a separate action

These jurisdictions are not cosmetic visibility settings applied after capture. They are selected before or at activation and constrain purpose, participants, processors, sensors, retention, sharing, and review. The same person may act across all four, but the streams do not merge merely because software can join them. This extends contextual integrity into the architecture of capture itself (Nissenbaum, 2004).

SENSITIVE CAPTURE BOUNDARY

Location, camera, microphone, and other sensitive sensors activate only through an explicit Moment, check-in, or separately authorized utility action. Public entry and ordinary browsing do not create standing permission. Child-facing capture additionally requires an independently reviewed purpose, developmentally appropriate assent, parental permission where applicable, minimum evidence, bounded retention, and accountable adult oversight.

A Moment gate defines the eligible evidence, processors, transformations, outputs, actions, retention, and stopping conditions. It is a capture-time jurisdiction, not a promise inferred from interface copy. Authority may be narrowed during the Moment and must be re-established if the purpose changes.

5.1 / PHYSICAL INTELLIGENCE LAB

The physical-first runtime

The runtime is not one app, interface, model, or robot. It is the governed sequence beneath different physical and digital expressions.

RUNTIME SEQUENCE

SPATIAL ACCESS: EARTH → PLACE → SPACE → THING → UTILITY

UTILITY EXECUTION: TASK + ACTION + AGENT (HI OR AI)

GOVERNED RETURN: MOMENT → PERSPECTIVE → HUMAN APPLICATION → CONSEQUENCE

ELEMENT	WORKING DEFINITION
Earth and Place	A shared planetary reference and a seeded, named environment whose spatial guardrail establishes the first levels of presence eligibility
Space	A seeded, bounded environment such as a room, classroom, workshop, field site, or temporary zone; presence inside it governs access to eligible Things and Utilities
Thing	An addressable physical artifact or resource whose identity and declared relationships matter more than embedded compute
Utility	A purpose-bound application that opens only when the runtime confirms presence in its eligible seeded context
Task + Action + Agent	Execution detail bound inside an eligible Utility: what is to be done, which action is permitted, and whether HI or AI performs it
Moment + Perspective	The governed evidence window and its derived, reviewable return; neither becomes a score, diagnosis, verdict, or identity

The Earth → Place → Space → Thing → Utility hierarchy is an access path, not only a descriptive taxonomy. A person is present or not present in the seeded context. If absent, they cannot reach the Utility, react to the experience, or use its social function from a remote feed. Once eligible, the Utility binds the Task, Action, and Agent (HI or AI); those three are execution details inside the Utility, not additional spatial tiers.

DESIGN QUESTION	APPLICATION-CENTRIC POSTURE	PHYSICAL-FIRST POSTURE
Primary unit	App, account, or persistent profile	Purpose-bound Moment in a Place, Space, or Thing
Basis for interpretation	Behavioral residue and later inference	Framed intent plus contextual evidence and consequence
Role of AI	Central interpreter and optimizer	Replaceable processor within a governed stream
Durable record	Accumulated events or platform history	Reviewed Perspective, provenance, and participant-controlled continuity
Interface	Destination through which activity is performed	Entry, capture, coordination, and visualization in support of activity

A chair does not need a universal chair app or to become an autonomous computer. It can carry or reveal an identity-bound utility when relevant. A room can maintain a declared relationship to that chair, a table, or a tool. Meaning can be resolved from deterministic relations and the active purpose rather than reconstructed through continuous ambient inference.

5.2 / PHYSICAL INTELLIGENCE LAB

Governed continuity: Data Wells, BDA, and the System of Things

Continuity belongs to the owner-governed environment, not to a model session, vendor memory, or login alone.

A Data Well is not a raw data lake, comprehensive profile, or automatic product of sign-in. It is a governed continuity environment within the applicable owner boundary for reviewed Perspective, selected artifacts, and provenance. Raw evidence may remain ephemeral and be retained only as long as the active protocol requires. More durable records privilege the participant-reviewed account, supporting provenance, explicit uncertainty, and the authority under which promotion occurred. An authorized Block may preserve proof of Data Well ownership and owner-approved lifecycle transitions; it does not store the whole Data Well, prove the truth of its contents, or transfer ownership to a processor.

BDA is the recursive Before-During-After human-AI working loop. A person declares intent, asks processors to reflect the meaning they understand, reviews developing work, corrects drift, and decides what continues. The runtime makes the evidence around that relationship visible and correctable. It does not make the processor authoritative.

BDA STAGE	GOVERNED ACTIVITY
Observe	Purpose-bound physical engagement produces selected signals and artifacts with Place, Space, Thing, utility, time, source, and Moment provenance
Process	Multiple authorized and substitutable processors examine the same bounded evidence through declared lenses; every candidate remains attributed Vitro output
Review and purge	During a governed 72-hour Ephemeral cycle, the owner compares evidence and proposals, inspects uncertainty, corrects or rejects drift, and lets unsupported raw material expire
Promote only by authority	Signal → 72-hour Ephemeral cycle record → reviewed Moment → accepted Perspective → weekly Observatory Perspective → monthly Master Engagement Record → authorized Block; every transition remains an owner-approved decision

A System of Things (SOT) is a domain-specific system of declared relationships among participants and roles, Places and Spaces, Things and utilities, Moments and data streams. It does not require every object to become intelligent. Meaning is resolved from active relationship and purpose rather than reconstructed through continuous observation. Personal, Partner, Project, and Root SOTs remain distinct authority boundaries; access to one does not grant access to another.

A Data-Logic-Action (DLA) Unit is a versioned governance boundary specifying the data a utility may use, the processing it may perform, the outputs or actions it may produce, and the retention and stopping rules that apply. Goal-Based Tasks (GBT) narrows the active objective, permitted tasks, and conditions under which the utility should stop, ask, or return. Neither grants authority beyond the active SOT and Moment.

CANONICAL AGENT NODES

N.01 is the Observer Agent Node; N.02/P.AI is the Personal AI Agent Node; N.03 is the Utility Agent Node. These names identify node classes only. The active SOT, Moment, DLA Unit, and utility contract separately define behavior, data access, retention, and action authority.

5.3 / PHYSICAL INTELLIGENCE LAB

One environment: Spatial, AI, and Block

The position is not three independent technology strategies. Each layer supplies a different form of constraint to one human-governed relationship.

Today's technology categories often place spatial computing, artificial intelligence, and block-based proof on separate roadmaps. Physical-first intelligence treats that separation as incomplete. Spatial without governed intelligence can become an interface layer. AI without spatial jurisdiction can process life from behavioral residue and remote access. Block without a human and data model can preserve transactions without clarifying whose lived record or authority they represent.

LAYER	ROLE IN THE ENVIRONMENT	NON-NEGOTIABLE BOUNDARY
Spatial	Seeds Earth, Place, Space, Thing, and Utility relations; confirms whether a person is present and therefore eligible to use the Utility	No confirmed presence, no Utility access
AI	Processes authorized Vivo evidence through bounded Utilities or Agents and returns attributed Vitro proposals as reviewable Perspective	Processing creates no ownership, standing access, or authority over the next direction
Block	Preserves proof of Data Well ownership and owner-authorized lifecycle transitions after review and promotion	The Block is not the Data Well, does not establish lived truth, and does not make a processor the owner

OPERATIONAL SYNTHESIS

SPATIAL PRESENCE → ELIGIBLE UTILITY
VIVO + BOUNDED AI PROCESSING → REVIEWED PERSPECTIVE
OWNER-AUTHORIZED CONTINUITY → BLOCK PROOF OF DATA WELL OWNERSHIP

The practical consequence is simple and culturally significant. A person cannot use a Place-bound Utility—or perform its equivalent of a remote reaction or 'like'—unless they are physically present in the seeded context. Digital participation follows lived participation rather than replacing it. What happens there may generate Vivo evidence; AI may help interpret it; the owner may decide what becomes durable; and Block may preserve proof of that ownership decision.

This blended environment changes the individual's relationship with technology. The objective is not to build a more complete digital world in which people live through profiles, feeds, and remote representations. It is to make intelligence available around people and institutions in ways they can live with: situated, bounded, attributable, reviewable, and answerable to physical and social consequence.

DESIGN PRINCIPLE

Spatial governs admission. AI contributes Perspective. Block preserves proof of ownership. Human Intelligence supplies purpose, review, and consequence.

Where physical intelligence lives among Physical AI

Robotics asks how a machine can perceive and act in the world. Physical-first intelligence asks how action remains related to human purpose, source data, interpretation, authority, and consequence.

A robot can turn a digital instruction into physical motion. Motion alone is not necessarily interpretation. It can be the trained embodiment of an interpretation: a policy, demonstration, objective, reward structure, or boundary that people encoded, selected, or authorized. Industrial robotics, remote-controlled machines, and competitive demonstrations can show physical capability without establishing who supplied the purpose, how evidence was governed, what the action meant to the people present, or what returned to the human afterward.

The comparison to children is revealing but limited. Children do not learn only from training signals. They are already living beings with bodies, needs, attachment, emotion, vulnerability, relationships, memory, development, and mortality. An experience can comfort, injure, transform, isolate, nourish, or end them. Those consequences are not merely positive and negative rewards. They matter to the being experiencing them.

A robot can detect that heat damaged an actuator and update its behavior. That is reaction and adaptation. Whether the damage means something to the robot is a different and unresolved threshold. A system may become increasingly life-like through persistent embodiment, self-maintenance, local learning, memory, relationships, and internally generated objectives. None of those capabilities alone establishes subjective feeling or a self-originating purpose. Physical-first intelligence therefore does not assign living status from sensorimotor competence.

CAPABILITY	WHAT IT ESTABLISHES	WHAT IT DOES NOT ESTABLISH
Embodiment	A system occupies and acts through a physical body	That the system experiences its body as a living stake
Reaction	A system changes behavior in response to stimuli	Feeling, meaning, or human interpretation
Adaptation	Performance changes from feedback or experience	Understanding why the consequence matters
Assigned objective	A target can guide planning and action	Self-originating purpose or ownership of the Goal
Machine enactment	A trained interpretation becomes physical motion	Transfer of human responsibility, authority, or consequence

EMBODIMENT IS NOT LIVED CONSEQUENCE

Physical presence does not make a processor alive. Physical action does not make its interpretation human. Adaptation does not by itself establish understanding. An assigned objective does not by itself become purpose.

Within a physical-first system, a robot may observe selected physical evidence, perform an authorized Task, and return evidence of consequence. It remains an actuator and processor under the applicable purpose unless a future system gives reason to reconsider its moral or legal status. The present architecture does not use speculative machine life to dissolve current human responsibility.

6.1 / PHYSICAL INTELLIGENCE LAB

Agentic capability without agentic sovereignty

Capability can be valuable without becoming a standing claim over purpose, memory, data, or action.

An agentic architecture generally receives instructions, acquires additional context, iterates between reasoning and tool use, may take action with limited supervision, and returns a result. This paper does not argue that agents are unnecessary, that autonomy is inherently unsafe, or that existing agent systems lack governance. Many employ permissions, approval gates, logging, policy, and other controls.

Kitsune's distinction is structural: processing authority is derivative of the active System of Things and Moment contract, not a standing property of an account, memory, credential, node name, or physical body. A processor may reason, generate, observe, or act only when that behavior is separately authorized inside the active system and utility. None of those capabilities independently determines why data exists, which jurisdiction applies, what may become durable, who may receive the result, or whether an action should occur.

Least privilege limits each role and processor to the minimum capability required; complete mediation requires authority to be checked at protected actions rather than assumed from a prior login (Saltzer & Schroeder, 1975). NIST's work on software and AI agent identity similarly treats identification, delegated authority, logging, and provenance as open infrastructure concerns. DLA Units, GBT structures, canonical node identity, and Moment provenance are Kitsune's proposed mechanisms for making those concerns explicit and testable; they are not claims of compliance with a completed standard.

AGENTIC SOVEREIGNTY - FAILURE CONDITION

A software actor accumulates enough cross-context memory, delegated credentials, goal-setting power, continuity, and action rights to become the de facto authority over a person, institution, or domain.

The proposed architecture attempts to prevent this condition through separation of roles, jurisdictions, contracts, processor slots, and durable records. It cannot guarantee safety. A flawed SOT, overbroad DLA Unit, coercive institution, compromised orchestrator, or unreviewed output can still produce harm. The value of governance is that authority and failure become more visible, bounded, revocable, and auditable.

The same boundary applies to physical machines. A robot that can select Methods and Tasks inside an assigned Goal does not inherit ownership of the Goal. Machine autonomy within a bounded purpose is not equivalent to self-originating purpose. When a machine acts, the human and institutional chain that designed, trained, authorized, deployed, and benefited from the action must remain legible.

AUTHORITY PRINCIPLE

Robots and agents may enact trained interpretations. They do not erase the human source of meaning, authority, or responsibility merely by moving faster, operating longer, or acting in physical space.

Symbiosis and digital diversity

The proposed relationship should elevate the position of both Human Intelligence and Artificial Intelligence in the physical world without allowing either to disappear into the other.

Symbiosis here is a design direction, not a claim that every coupling is mutually beneficial or that a person and model form one literal mind. Human Intelligence contributes purpose, values, embodied judgment, social meaning, and responsibility. Artificial Intelligence contributes processing scale, memory, comparison, pattern recognition, simulation, and perspectives that may be unfamiliar to the person. Vivo evidence gives the processor a grounded source. Vitro output gives the person a manipulable expression. Human application and physical consequence test what the relationship produced.

The reciprocal elevation is specific. The human should no longer be treated mainly as behavioral material from which systems extract intelligence. The person becomes an active data agent and embodied interpreter. The AI should no longer be treated only as an answer machine or an emerging sovereign. It becomes a more grounded and accountable processor because its output remains linked to authorized source evidence, human review, processor identity, and lived consequence.

PARTICIPANT	WHAT IT CONTRIBUTES	HOW THE RELATIONSHIP ELEVATES ITS POSITION
Human Intelligence	Purpose, interpretation, judgment, embodiment, responsibility, lived consequence	From passive data subject or approval gate to owner-governor and meaning-making participant
Artificial Intelligence	Processing, memory, comparison, synthesis, simulation, alternative Perspective	From ungrounded answer or standing authority to situated, accountable, replaceable processor
Physical and social world	Constraint, response, relationship, change, irreversibility, and consequence	From peripheral sensor field to primary source and test of intelligence

Digital diversity extends processor substitutability into an ecological design principle. No single model, provider, interface, storage posture, or representation should become inseparable from a person's continuity. Multiple authorized processors may examine the same bounded evidence through declared lenses. Their outputs remain distinct and attributable. The person can compare, combine, reject, or replace them without surrendering the underlying Data Well.

The analogy to biodiversity is limited and should not be mistaken for biological equivalence. Its value is structural: resilience depends on avoiding a monoculture in which one processor's errors, incentives, style, or ontology becomes the only available interpretation of a person's life. Digital diversity also includes local and cloud processing, open and closed models, specialized and general systems, human interpretation, and the option to use no AI for a particular Moment.

DIGITAL DIVERSITY PRINCIPLE

Continuity belongs to the governed human record. Perspective may come from many processors. No processor should own the continuity required to replace it.

One position, many disciplines

The position is general because it specifies a relationship. It is not universal because each discipline must define its own purpose, evidence, competence, authority, and consequence.

People learn, create, research, care for others, maintain homes and equipment, steward land, practice professions, participate in institutions, and build relationships. Each activity contains different legitimate roles and consequences. Physical-first intelligence does not prescribe one interface or outcome across them. It asks whether owner-governed source evidence and bounded processing can return a Perspective that a human can responsibly make useful in that particular life context.

Readers should be able to connect the position to their own discipline without being told that Kitsune has already designed its correct application. A teacher may ask whether Perspective strengthens understanding. A land steward may ask whether it clarifies condition and consequence. A researcher may ask whether it preserves method and uncertainty. A craftsperson may ask whether it makes effort and adaptation legible. A parent may ask whether it supports reflection without turning family life into monitoring. The shared architecture does not collapse these outcomes into one metric.

FIELD EXPRESSION	POSSIBLE FIELD OUTCOME	NON-TRANSFER BOUNDARY
Learning and education	Explanation, revision, recall, transfer, or greater understanding	A learning Perspective must not become a diagnosis or permanent child profile
Home and daily practice	Reflection, coordination, maintenance, care, or reduced uncertainty	Household continuity must not become ambient surveillance or default institutional access
Land and stewardship	Contextual understanding, selected condition evidence, maintenance, or responsible handoff	Presence does not prove ownership, boundary, identity, or licensed judgment
Professional practice	Method legibility, decision support, provenance, or accountable next action	Institutional purpose does not create a claim over a person's whole Data Well
Research and creative work	Traceable development, alternative Perspective, method comparison, or embodied expression	Processor coherence does not become authorship, truth, or proof of method
Public and collective life	Shared understanding, participation, stewardship, or coordinated action	Public purpose does not mean unrestricted reuse, universal identity, or forced consensus

DOMAIN HUMILITY

A structure designed for one field may encode relationships, values, or risks that are wrong for another. Shared runtime principles do not remove the need for domain knowledge, participant protection, comparison conditions, and the authority to stop.

This is why the foundational paper sits above field papers rather than replacing them. It establishes the human-AI-data-consequence relationship. Each field paper must then accept the discipline of its own participants, measures, history, and failures.

Education as the first scientific field expression

Education makes the position vulnerable to evidence because learning is physical, social, developmental, role-bound, and measurable only through carefully chosen outcomes.

Education is a consequential initial domain because learning continues across early childhood, formal schooling, apprenticeship, professional practice, doctoral research, and self-directed study, and understanding is often built through action. Learners construct, draw, test, observe, explain, revise, and collaborate with tangible materials, instruments, environments, and other people. Research has long emphasized context, active engagement, self-explanation, formative assessment, and feedback (Brown et al., 1989; Chi et al., 1994; Black & Wiliam, 1998; Hattie & Timperley, 2007; National Academies, 2018).

The proposed direction does not place AI at the center of learning. Its unit of design and study is the physical learning Moment. A learner may declare an expectation, adopt an educator's mission, or co-frame a question with a parent, caregiver, teacher, mentor, or collaborator. Digitized effort is the minimum provenance-bearing evidence needed to make the attempt legible. It is not a requirement to film the lesson, log every action, or keep a screen open. AI may organize, compare, or express authorized evidence, but its output remains Vitro until reviewed and embodied through further learning activity.

EDUCATION HYPOTHESIS

FRAMED INTENT + MINIMUM EVIDENCE OF LIVED LEARNING + REVIEWED ROLE-APPROPRIATE PERSPECTIVE → HUMAN REFLECTION AND NEXT PHYSICAL LEARNING ACTION → GREATER UNDERSTANDING

PARTICIPANT	USEFUL PERSPECTIVE	NON-NEGOTIABLE CONSTRAINT
Learner or child	Age-appropriate reflection on what was tried, changed, and understood; possibly spoken, drawn, spatial, or revisited with an adult	Not a score, diagnosis, permanent profile, or assumption that direct AI exposure is better
Parent or caregiver	A joint journal or conversation aid connecting selected home observation to the learning Moment	Not continuous monitoring; sharing remains specific, reviewable, and purpose-bound
Teacher or mentor	Concise, source-linked depth about strategy, explanation, evidence, and change	Must reduce or hold administrative burden; no whole-record access or automatic ranking
Program research	Only the consented and de-identified evidence required to evaluate the study	Aggregation, retention, and re-identification risk are declared in advance

A young child may contribute through talk, drawing, construction, movement, or shared reflection while an adult operates the tool. An advanced learner may work directly with reflective or spatial interfaces. The exposure model belongs to the learning protocol, not to the novelty of the technology. Education is the first rigorous field expression because these boundaries become difficult to ignore, not because the thesis belongs only to schools.

9.1 / PHYSICAL INTELLIGENCE LAB

A bounded education experiment

The first field test should distinguish the value of framing evidence from the value of AI-supported expression while keeping physical learning primary.

The first study should use one learning objective, one declared age range, one physical kit or activity, one setting, and one short study window. The educational partner should choose the curriculum measure, allocation method, sample, and meaningful effect before recruitment. Direct child-facing AI is outside the first early-learning experiment; an authorized adult operates capture and review.

CONDITION	EXPERIENCE	WHAT IT HELPS DISTINGUISH
A. Physical activity and usual reflection	Learners complete the activity using the partner's existing reflection or documentation practice. No Kitsune Perspective is introduced.	The comparison baseline
B. Framed intent and minimum evidence	The same activity adds age-appropriate framing of intent and the minimum deliberate capture needed to make attempt and result legible. No AI Perspective is returned.	The value of framing and evidence without model expression
C. Reviewed, role-appropriate Perspective	Condition B adds an attributed AI proposal that an authorized adult reviews before it becomes a learner, caregiver, or teacher Perspective.	The additional value and burden of returned Perspective

The primary outcome should be curriculum-defined evidence of understanding, preferably explanation or transfer in a new off-screen task. Secondary outcomes include screen minutes, learner and adult documentation time, correction cost, perceived usefulness, abandonment, and voluntary reuse. Privacy events, inappropriate cross-role exposure, uncorrected misinterpretation, and retention exceptions are guardrails. The design should narrow or stop if burden exceeds its declared ceiling even if a learning measure improves.

WHY THREE CONDITIONS

A two-condition comparison could show that the full experience differs from current practice, but it could not show whether the useful ingredient was framed intent, minimum evidence, or the returned Perspective. The third condition makes that distinction observable.

Education connects directly to the foundational cycle. The learner's activity produces Vivo evidence. AI may return a Vitro proposal. An authorized person reviews it. The learner, teacher, or caregiver then brings selected Perspective back into explanation, conversation, revision, or another physical task. The next learning action, not the generated summary, is where the claim of greater understanding must ultimately be tested.

The education paper remains the complete field document for this experiment. The foundational paper does not replace its developmental safeguards, role definitions, methods threshold, or scientific detail. It explains why those requirements are expressions of a broader position on human purpose, governed data, processor authority, and lived consequence.

FIELD PRINCIPLE

Physical activity creates the evidence. Human explanation and application give it meaning. The runtime preserves provenance and returns only the Perspective each authorized participant needs. AI is a bounded complement to learning, not its replacement or authority.

Social protocol and collective intelligence

Physical integration changes how agents enter human situations. Timing, permission, role, attention, and the right to decline are part of the runtime.

A model that can speak is not automatically entitled to interrupt, and a completed task does not create a right to demand attention. In a co-present setting, an agent's behavior should reflect who is present, which Moment is active, what the utility was authorized to do, whether the participant wants the result, and what social consequence an interruption may create (Amershi et al., 2019). Hardware may support the interaction, but the thesis is that agents participate under social and data protocols legible to the humans whose environment they enter.

One person may participate in private, personal, professional, institutional, and public relationships. Those roles do not collapse into a single identity merely because one processor can access them. Purpose-bound projection allows the same governed Moment to produce different representations for authorized participants while preserving source, scope, uncertainty, and review. Representation may be malleable. Authority is not.

Collective intelligence introduces another boundary. A method source and a situated implementation are distinct, attributable contributions. One person or institution may contribute a model of action; another may contribute the model in action under different conditions. Reviewed Perspective may connect them as method evidence without collapsing authorship, local adaptation, or contextual failure. The collective can learn from relation while the people involved retain their own jurisdictions.

RELATION	POSSIBLE VALUE	REQUIRED BOUNDARY
Individual Perspective	Supports a participant's own understanding and next action	Does not create a portable verdict or universal profile
Role projection	Gives an authorized collaborator, caregiver, mentor, or institution the minimum view needed	No role inherits the whole record
Method source	Contributes an attributable model, curriculum, practice, or approach	Attribution creates no automatic right to participant data or outcomes
Situated implementation	Shows how a method behaves under local people, constraints, and adaptation	Local evidence cannot be silently generalized beyond the studied context
Collective Perspective	May reveal patterns across separately governed contributions	No automatic publication, payment, licensing, identity disclosure, or access to whole Data Wells

The Lab does not seek balance by forcing consensus or treating all claims as equally supported. Digital diversity preserves multiple attributable processors and interpretations; provenance and physical consequence still allow evidence to challenge them. Balance means that disagreement, correction, refusal, minority conditions, and alternative methods remain structurally possible rather than being erased by one optimized representation.

SOCIAL PRINCIPLE

An agent may contribute to a human situation only through authority that the people and governed context can see, question, narrow, and revoke.

Research agenda

The position becomes scientific only when its claims can fail. The general relationship and each field expression require operational hypotheses rather than demonstrations of technical possibility.

PRIMARY HYPOTHESIS - TO BE TESTED

In a bounded physical context, framed intent plus owner-governed Vivo evidence, bounded Vitro processing, reviewed Perspective, and human embodiment will improve role-appropriate understanding or action without transferring direction to the processor or exceeding declared burdens and authority limits.

PROPOSITION	TESTABLE DIRECTION
P1. Recursive understanding	Compared with an isolated answer, generic summary, or outcome record, the governed cycle will produce stronger field-defined evidence that the participant understands what changed and why
P2. Physical return	A Perspective carried into a declared lived action will reveal errors, constraints, or useful consequences that remain invisible when the cycle ends in digital review
P3. Intent and provenance	Framed intent plus source and transformation provenance will reduce false attribution and disagreement compared with behavioral residue alone while leaving measurable uncertainty
P4. Human direction	Participants will be able to identify, correct, reject, replace, or withhold processor output before durable or consequential use
P5. Burden and desirability	The cycle will meet predeclared ceilings for attention, prompts, capture, documentation, correction, and voluntary reuse; failure on burden counts against the design
P6. Digital diversity	Replacing or comparing processors while preserving the owning SOT and Moment will retain authority and provenance semantics while making processor-specific differences legible
P7. Spatial jurisdiction	Presence-gated Earth → Place → Space → Thing → Utility access will prevent remote Utility use outside the seeded context while meeting predeclared ceilings for false admission, false denial, spoofing, accessibility burden, and recovery

A field does not need to test every proposition at once. Before recruitment or consequential use, a study should declare the population and owning SOT, physical context, field objective, comparison conditions, exposure levels, unit of analysis, outcome instruments, burden ceilings, privacy and adverse-event definitions, processor configuration, retention and expiry rules, stopping criteria, and treatment of missing data.

Results should report field outcomes, burden, uptake, correction, processor disagreement, privacy events, null effects, and negative outcomes together. Understanding must not be reduced to one dashboard metric. A persuasive visualization, high reuse rate, or fluent Perspective is not sufficient evidence that the system improved human orientation or action.

SCIENTIFIC POSTURE

The appropriate first step in any discipline is one reversible experiment with a predeclared measure, a real path to stop, and enough humility to let evidence change the position.

Failure modes and limitations

The runtime cannot be evaluated only when it works as intended. These conditions define when the thesis should be rejected, narrowed, or redesigned.

- Declared intent can be incomplete or performative. People may not know their intent, may revise it, or may state what an authority wants to hear.
- Spatial admission can fail. False admission opens a Utility to someone who is not present; false denial excludes someone who is present; spoofing defeats the seeded guardrail. These are security and implementation failures to measure, not reasons to weaken the binary access rule.
- Capture changes behavior. Prompts, cameras, journaling, or agents may burden activity, shift work onto participants, or produce compliance mistaken for value.
- Perspective can distort. Summaries and visualizations may omit context, overstate coherence, introduce bias, or make processor confidence appear like evidence.
- Human embodiment can fail. A person may apply a Perspective without sufficient skill, context, care, or authority; human control is not a guarantee of good consequence.
- Ungrounded recursion can shape the person. A digital loop may affect cognition, emotion, identity, or behavior without returning through evidence, choice, and lived consequence.
- Continuity can become surveillance or persuasion. Purpose may expand, retention may become indefinite, or one processor's style may become difficult to separate from the person.
- Domain structures may not transfer. An SOT designed for education may encode relationships or risks that are wrong for home, work, land, or public contexts.
- Access is uneven. Device, connectivity, mobility, sensory, language, time, developmental, and support requirements can exclude participants.
- Governance structures can be wrong. A DLA Unit, GBT structure, SOT, or institution can encode an overbroad purpose or omit a harmed participant.
- Agent separation can be nominal. Node names provide no protection if one credential, processor, or operator silently bypasses accepted boundaries.
- Machine enactment can obscure responsibility. Physical autonomy may make trained or authorized human interpretations appear self-originating.
- Digital diversity can become theater. Nominal processor choice is meaningless when models share the same source, incentives, infrastructure, or inaccessible continuity.
- Method or Block proof can be overclaimed. Participants may select favorable evidence, and a Block that proves an ownership transition may be misrepresented as proof that the underlying event or interpretation is true.
- Vitro can masquerade as Vivo. Generated material may enter a record without source labeling or be represented as lived evidence.
- Laboratory visibility can be mistaken for truth. Dense evidence, spatial displays, or agreement among processors can create false confidence.

LIMITATION OF THE POSITION

Physical-first intelligence does not guarantee wise people, benevolent institutions, accurate models, safe robots, truthful evidence, or good consequences. It proposes a way to keep purpose, data, processing, interpretation, authority, and consequence distinguishable enough to be examined and contested.

In child-facing work, independent ethics review, developmentally appropriate assent, parental permission where required, data minimization, retention limits, adult accountability, and educator oversight are prerequisites, not claims of safety. In other consequential fields, equivalent domain-specific protections must be defined rather than borrowed from education.

Commitments and conclusion

The position is only meaningful if its research and runtime commitments survive contact with product pressure, institutional power, and technical convenience.

- Purpose before capture. Identify why evidence is needed before collection begins.
- Presence before Utility. Enforce seeded Earth, Place, and Space eligibility before a person can access a Place-bound Utility.
- Minimum necessary evidence. Do not treat possible future model usefulness as unlimited justification.
- Vivo/Vitro source integrity. Keep participant expression distinct from processor inference. Inference is not participant authorization.
- Provenance with uncertainty. Preserve source, processor, transformation, missingness, confidence, correction, and contestation.
- Separation by jurisdiction. Private, personal, professional, institutional, and public streams do not merge by default.
- Participant review and stop authority. People can inspect, correct, withhold, pause, or stop before irreversible consequence.
- Human embodiment and responsibility. Perspective does not become direction until a person decides whether and how to make it real.
- Processor substitutability. No model or vendor owns the record, continuity, or sole authority on meaning.
- Digital diversity. Preserve meaningful processor, infrastructure, interpretive, and no-AI choices.
- Authority outside probabilistic output. Permissions, retention, jurisdiction, and consequential state changes are enforced by the runtime.
- Least privilege and revocation. Every agent, processor, tool, and projection receives only the authority required for the active purpose.
- Agent identity and action provenance. Distinguish participant, agent node, processor, tool, operator, and physical actuator.
- Bounded retention. Raw evidence expires by protocol; durable continuity is intentional and proportionate.
- Block proof without data custody. Use authorized Block records to preserve proof of Data Well ownership and approved transitions, never as a substitute for the Data Well or its review.
- Physical and social primacy. Declare acceptable attention and documentation burdens and do not treat more interaction as success.
- Transparent field evidence. Report benefits, limitations, abandonment, adverse events, null effects, and negative results together.
- Domain-specific protection. Each field defines its own legitimate authority, competence, vulnerability, and stopping conditions.

POSITION

Artificial intelligence can contribute Perspective. Human beings carry Perspective into life. Physical consequence returns evidence. The governed relationship among those movements is where physical intelligence lives.

Physical-first intelligence begins closer to the human event it is meant to support. Intent should not be reconstructed only from residue. Effort should not disappear behind completion. Place and social context should not become incidental metadata. Spatial presence should govern access to Place-bound Utilities; AI should remain a bounded processor; Block should preserve proof of Data Well ownership without becoming the Data Well. Technology should be something people live with, not an environment they are required to live through.

The foundational position and the learning field paper are meant to be read together. This paper defines the relationship among Human Intelligence, Artificial Intelligence, presence-gated Spatial access, owner-governed Vivo data, bounded Vitro processing, Block proof of Data Well ownership, human embodiment, physical consequence, symbiosis, and digital diversity. The learning paper accepts the responsibility of testing that relationship around learners, caregivers, educators, developmental protection, burden, and greater understanding.

The goal is not to prove that AI belongs everywhere or that technology can resolve human conflict. It is to give people a more legible and balanced way to work with intelligence without becoming subordinate to the systems that process them. The next step remains small: one purpose, one Moment, one bounded stream, one reviewable Perspective, one human application, and enough contact with reality to learn what should happen next.

REFERENCES / PHYSICAL INTELLIGENCE LAB

References and source status

External references establish adjacent scientific and technical foundations. Kitsune's synthesis, terms, runtime, and hypotheses are proposed constructs requiring independent evaluation.

Amershi, S., Weld, D., Vorvoreanu, M., Fournay, A., Nushi, B., Collisson, P., et al. (2019). Guidelines for human-AI interaction. Proceedings of CHI 2019, Paper 3, 1-13.

Assran, M., Duval, Q., Misra, I., Bojanowski, M., Vincent, P., Rabbat, M., LeCun, Y., & Ballas, N. (2023). Self-supervised learning from images with a joint-embedding predictive architecture. CVPR 2023, 15619-15629.

Black, P., & Wiliam, D. (1998). Assessment and classroom learning. Assessment in Education, 5(1), 7-74.

Booth, H., Fisher, B., Galluzzo, R., & Roberts, J. (2026). Accelerating the adoption of software and AI agent identity and authorization (draft concept paper). NIST NCCoE.

Brohan, A., Brown, N., Carbajal, J., Chebotar, Y., Chen, X., Choromanski, K., et al. (2023). RT-2: Vision-language-action models transfer web knowledge to robotic control. arXiv.

Brown, J. S., Collins, A., & Duguid, P. (1989). Situated cognition and the culture of learning. Educational Researcher, 18(1), 32-42.

Chi, M. T. H., de Leeuw, N., Chiu, M. H., & LaVancher, C. (1994). Eliciting self-explanations improves understanding. Cognitive Science, 18(3), 439-477.

Clark, A., & Chalmers, D. (1998). The extended mind. Analysis, 58(1), 7-19.

Dourish, P. (2001). Where the action is: The foundations of embodied interaction. MIT Press.

Hattie, J., & Timperley, H. (2007). The power of feedback. Review of Educational Research, 77(1), 81-112.

Hutchins, E. (1995). Cognition in the wild. MIT Press.

Kazdan, J., Schaeffer, R., Dey, A., Gerstgrasser, M., Rafailov, R., Donoho, D. L., & Koyejo, S. (2025). Collapse or thrive? Perils and promises of synthetic data in a self-generating world (Version 4). arXiv.

LeCun, Y. (2022). A path towards autonomous machine intelligence (Version 0.9.2). OpenReview.

Licklider, J. C. R. (1960). Man-computer symbiosis. IRE Transactions on Human Factors in Electronics, HFE-1(1), 4-11.

National Academies of Sciences, Engineering, and Medicine. (2018). How people learn II: Learners, contexts, and cultures. National Academies Press.

National Institute of Standards and Technology. (2026). AI Agent Standards Initiative.

Nissenbaum, H. (2004). Privacy as contextual integrity. Washington Law Review, 79(1), 119-157.

Saltzer, J. H., & Schroeder, M. D. (1975). The protection of information in computer systems. Proceedings of the IEEE, 63(9), 1278-1308.

Shumailov, I., Shumaylov, Z., Zhao, Y., Papernot, N., Anderson, R., & Gal, Y. (2024). AI models collapse when trained on recursively generated data. Nature, 631, 755-759.

Tabassi, E. (2023). Artificial Intelligence Risk Management Framework (AI RMF 1.0). NIST AI 100-1.

W3C. (2013). PROV-O: The PROV ontology. W3C Recommendation, 30 April 2013.

Weiser, M. (1991). The computer for the 21st century. Scientific American, 265(3), 94-104.

Wilson, M. (2002). Six views of embodied cognition. Psychonomic Bulletin & Review, 9(4), 625-636.

Wooldridge, M., & Jennings, N. R. (1995). Intelligent agents: Theory and practice. The Knowledge Engineering Review, 10(2), 115-152.

Yao, S., Zhao, J., Yu, D., Du, N., Shafran, I., Narasimhan, K. R., & Cao, Y. (2023). ReAct: Synergizing reasoning and acting in language models. International Conference on Learning Representations.

INTERNAL DESIGN SOURCES

Kitsune Ai システム / Physical Intelligence Lab. (2026). Kitsune Symbiont Mind and Lived Intelligence Canon (R2). Approved internal design canon; not external evidence.

Kitsune Ai システム / Physical Intelligence Lab. (2026). Physical-first intelligence: Re-anchoring digital systems in human intent, place, embodied effort, and reviewable consequence (Version 1.1). Position Paper 01, learning field expression.

RELATIONSHIP BETWEEN PAPERS

Position Paper 00 defines the foundational human-AI-data-consequence relationship. Position Paper 01 applies that position to education as the first scientific field expression. Neither document is evidence that the proposed mechanisms have already improved human understanding, agency, learning, or field outcomes.

KITSUNE

PHYSICAL-FIRST INTELLIGENCE

POSITION PAPER 00 / FOUNDATIONAL POSITION

Position Paper 00 defines the foundational human-AI-data-consequence relationship. Position Paper 01 applies that position to learning and education as the first scientific field expression. The papers are intended to be read together.

IDENTITY AND POSITION

LEGAL ENTITY	Kitsune Ai System LLC
PUBLIC BRAND	Kitsune Ai システム
LAB	Physical Intelligence Lab
NAMED POSITION	Kitsune Physical-First Intelligence

WEBSITE

<https://kitsuneaisystem.com/>

CONTACT AND FIELD INQUIRIES

<https://kitsuneaisystem.com/investor-media#contact-desk>
