
Kitsune Ai システム / Physical Intelligence Lab

POSITION PAPER 01 / LEARNING FIELD EXPRESSION

PHYSICAL-FIRST INTELLIGENCE

A KITSUNE PHYSICAL-FIRST INTELLIGENCE FIELD PAPER

Re-anchoring digital systems in human intent, place, embodied effort, and reviewable consequence

Intelligence should belong to life.

The physical world supplies context. People supply purpose. Technology answers to both.

DECLARED INTENT + DIGITIZED EFFORT -> DATA AGENCY

Published by Kitsune Ai System LLC

Kitsune Ai システム / Physical Intelligence Lab

Version 1.1 / Draft for scientific and field-partner review / August 2026

Abstract

Digital systems often infer purpose from behavioral residue: clicks, scores, uploaded artifacts, location traces, and completion records. Physical-first intelligence proposes a different relationship. Technology is something people live with, not an environment they must live through. Seeded Earth, Place, and Space guardrails establish whether a person is present. That binary admission governs Earth → Place → Space → Thing → Utility: if the person is not present in the eligible context, the Utility does not open. Once eligible, the Utility binds a Task, an Action, and an Agent—Human Intelligence (HI) or Artificial Intelligence (AI)—to the active learning purpose.

Education is the proposed first scientific field expression, but learning is treated as lifelong: from early childhood and formal schooling through apprenticeship, professional practice, doctoral research, and self-directed study. The primary hypothesis is that framed intent plus minimal digitized evidence can support a reviewed Perspective for the learner and distinct, authorized Perspectives for others responsible to the learning context, improving understanding of both the material and the learning process. In an early-learning study, those roles may be the child, parent or caregiver, and teacher. A young child may participate through talk, drawing, construction, movement, or shared reflection while an adult operates the tool; an advanced learner may work directly with the tools while an educator, mentor, advisor, collaborator, or institution receives only the bounded Perspective its role authorizes. Direct child exposure to AI remains an experimental condition, not a default.

The system keeps lived evidence, digital processing, purpose, provenance, review, and later action distinguishable. 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. The architecture joins three technologies normally positioned independently: Spatial governs presence and Utility access; AI performs bounded processing and returns Perspective; an authorized Block preserves proof of Data Well ownership and owner-approved lifecycle transitions. Raw evidence may enter a bounded 72-hour review workspace; processing creates no ownership; and the Block is neither the Data Well nor proof that its contents are true.

CLAIM STATUS

This is a position paper. It defines a research program, design commitments, and testable propositions. It does not claim that the runtime has already improved learning, privacy, agency, or institutional decision-making. Those claims require a published methods specification, independent ethics review where human participants are involved, bounded pilots, and transparent reporting of limitations and negative results.

1. The position

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” here does not mean the existence of software or screens. It names an architectural assumption: the application becomes the primary place where a person’s 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 that it can move through the application. The result is often 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 preferences, classify behavior, summarize records, and act across systems, but greater inferential capacity does not repair missing provenance. A click, location trace, uploaded image, score, or completed form may record an outcome while leaving the reason, effort, constraints, and social meaning unresolved. The problem is not that inference is useless. The problem is allowing inference to substitute for evidence that could have been structured with the participant at the moment of action.

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 intentions, use things, cooperate, attempt tasks, encounter constraints, and experience consequences. Spatial jurisdiction determines whether a Place-bound Utility can be reached; AI processes only the Vivo evidence 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 physical and social reality—something to live with, not live through.

CORE DEFINITION

Physical-first intelligence is the capacity of a digital system to derive useful, reviewable perspective from intentional human activity in real contexts while preserving provenance, purpose boundaries, and human authority over the resulting record.

1.1 Learning in the age of AI: a proposed research direction

The learning problem addressed here is not simply whether schools should accept or prohibit AI. Across classrooms, laboratories, workplaces, and self-directed study, indiscriminate adoption and blanket prohibition both avoid the harder empirical question: under what conditions can AI increase understanding without substituting for physical and social learning, increasing screen time, shifting documentation work to people, or turning a learner’s record into a persistent judgment? A digital record built mainly from completion, clicks, scores, or final

artifacts may show that something happened while preserving little about the learner's intention, effort, explanation, or change in understanding.

The proposed direction starts with a physical learning Moment, not an AI interaction. Intent is framed at an appropriate level: declared by the learner, adopted from an educator's or mentor's mission, derived from a research or professional objective, or co-framed with a parent, caregiver, or collaborator. Only the evidence needed to make effort and consequence legible is digitized. The system returns a purpose-bound Perspective to the learner and to any other participant authorized to use a distinct projection of it. The hypothesis is not that more data or more AI will improve learning. It is that better-framed evidence and role-appropriate feedback may improve understanding while asking less of attention.

PROPOSED LEARNING HYPOTHESIS

FRAMED INTENT + DIGITIZED EFFORT → ROLE-APPROPRIATE PERSPECTIVE → GREATER UNDERSTANDING

Greater understanding is the outcome to be tested, not an automatic property of the technology. Each arrow requires an operational definition, comparison condition, and failure threshold.

A science-first program must define the learning question, intended outcome, comparison condition, exposure level, burden ceiling, privacy boundary, and stopping rule before treating a tool as successful. Educational and research partners are co-designers of the work, not test sites. They contribute curriculum or disciplinary objectives, developmental and domain judgment, participant protection, and the realities of classrooms, laboratories, workplaces, and households. Kitsune contributes purpose-bound utilities, journaling, spatial visualization, and governed information flows. Neither contribution substitutes for the other, and the evidence must be allowed to change the design.

In early-childhood applications, direct child use remains one experimental condition, not the assumed interface. For a kindergarten-age learner, a parent or caregiver may be the primary user of a journaling utility while the child contributes through talk, drawing, construction, movement, or shared reflection. In a classroom, educators and the research protocol should determine whether, when, and how children encounter an AI-mediated surface. The tool may disappear during the activity and return only when reflection adds value.

2. Intellectual lineage and the gap this position addresses

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 has described ways in which 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).

Human-computer interaction has long explored 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). These traditions support a physical-first posture, but they do not by themselves specify how intent, provenance, purpose boundaries, model processing, and continuity should be organized as a runtime.

Privacy and data-governance research supply 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 that 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). Human-AI interaction research likewise emphasizes appropriate timing, communication of capability and uncertainty, correction, and user control (Amershi et al., 2019).

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 instead proposes a data-governance architecture for human activity: what evidence enters a system, under whose purpose, with what provenance, how processors may act on it, and how resulting Perspective returns to human use.

Agent research provides another relevant lineage. Wooldridge and Jennings (1995) described autonomy, social ability, reactivity, and proactiveness as central properties in a widely used account of intelligent agents. Contemporary language-model agents can 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 sources establish

the importance of agentic capability and governance; they do not determine where authority over a human record should reside. That unresolved question is the specific gap addressed here.

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 adds physical context, owner-governed continuity, explicit jurisdiction, processor substitutability, and a reviewed return from digital expression to lived consequence. Its additional gap claim is architectural: Spatial, AI, and Block should operate as one environment in which admission, processing, and proof each constrain the others under human purpose.

Research on recursively generated training data motivates a separate 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 the Kitsune runtime. They support only the narrower position that synthetic data is an instrument whose value depends on retained or verified source material, not a self-renewing substitute for lived evidence.

3. Education as the first scientific field expression

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 that process. Its unit of design and study is the physical learning Moment, and its question is whether a bounded digital complement can make that Moment more understandable to the learner and other authorized participants.

Learner is the canonical participant in this position. Child and student name specific developmental or institutional populations; parent, caregiver, teacher, mentor, advisor, collaborator, and institution name context-dependent roles. The first field experiment focuses on young learners because adult mediation, screen exposure, and role boundaries make the

scientific and ethical questions especially visible. That population defines the first test, not the boundary of the thesis.

3.1 The unit of study: a physical learning Moment

A learning Moment begins after the applicable presence gate has admitted the participant to an eligible Utility. Intent is then framed at an appropriate level: declared by the learner, adopted from an educator's or mentor's mission, derived from a research or professional objective, or co-framed with a parent, caregiver, or collaborator. Vivo data is lived data with presence, captured in or out of a Moment; the study should digitize only the minimum observation, artifact, explanation, measurement, decision, or change required by its purpose. Evidence remains subordinate to the learning activity.

A utility can enter briefly Before, During, or After the activity. Before may be a single prompt that confirms the question. During may involve no interface at all, or one deliberate capture of an artifact or observation. After may support reflection, comparison, or spatial visualization of what was tried and what changed. AI may organize, compare, or express authorized evidence, but its output remains Vitro data: a proposal for review rather than a score, diagnosis, or final account. Some experimental conditions should include no direct child-facing AI.

**PHYSICAL LEARNING MOMENT → FRAMED INTENT + DIGITIZED EFFORT → REVIEWED EVIDENCE →
ROLE-APPROPRIATE PERSPECTIVE → NEXT PHYSICAL LEARNING ACTION**

3.2 One Moment, distinct authorized Perspectives

A learner and the people authorized to support, guide, or assess that learning are not audiences for one universal dashboard. Their responsibilities, questions, rights, and exposure models differ by context. In early childhood, the relevant roles may be child, parent or caregiver, and teacher. In advanced study, they may include learner, educator, mentor, advisor, collaborator, or institution. A young child may contribute most meaningfully through the physical activity and a conversation while an adult operates the tool; an older or advanced learner may use an appropriate reflective or spatial interface directly. The exposure model belongs to the learning protocol, not to the novelty of the technology.

For a kindergarten-age child, the parent or caregiver may be the primary journaler: capturing a short observation, inviting the child's explanation, and using the resulting Perspective in a shared conversation. In a classroom, the teacher may receive supplementary depth only when it clarifies the learner's strategy, evidence, or change in explanation without creating a second documentation job. Each role receives a purpose-bound projection, and none inherits the others' full record.

Table 1. In an early-learning context, one governed Moment can support different useful views without becoming a universal child record.

Participant	Useful Perspective	Non-negotiable constraint
Child	Age-appropriate reflection on what was tried, changed, and understood; this may be spatial, visual, spoken, drawn, or revisited with an adult.	Not a score, diagnosis, or permanent profile. Direct AI exposure only when age, context, and protocol support it.
Parent / caregiver	Joint journal and conversation aid connecting home observation to the child's learning Moment; may mediate capture for young children.	Not continuous monitoring or a surveillance report. Sharing remains specific, reviewable, and purpose-bound.
Teacher	Concise, source-linked depth about strategy, explanation, evidence, and change that may inform the next lesson or conversation.	Must reduce or hold administrative burden. No automatic ranking and no whole-record access.

Program-level research is a separate projection, not another personal view. It should receive only the consented and de-identified evidence required by the study, with aggregation rules, retention, and re-identification risk defined in advance. The aim is to build models for learners, not totalizing models of them.

3.3 Malleable information; fixed authority

The information is malleable because the same governed Moment can be expressed differently for each authorized role. In an early-learning context, that may mean a child's spatial journey, a parent's joint journal, or a teacher's source-linked note. In advanced learning, it may mean a learner's research map, an advisor's evidence-linked review, or an institution's deliberately limited program view. These are projections rather than copies of a whole record. Each can select different evidence, scale, language, timing, and visual form while preserving provenance and the purpose that authorized it. Malleability is therefore a property of representation. It is not permission to merge roles, infer new purposes, or expose one participant's stream to another.

This is the practical value of purpose-bound projection. A model may propose a summary, association, or visualization, but the proposal remains attributed Vitro data until the authorized person reviews it. Corrections create a successor account rather than erasing the original evidence. A learner's Perspective can remain revisable; an authorized caregiver, collaborator, or advisor can withhold a contribution; an educator can challenge an interpretation; and a processor can be replaced without taking ownership of the Moment.

3.4 The adoption threshold is part of the evidence

A learning tool fails if its output is not worth the energy required to use it. That threshold should be measured, not assumed. A pilot should record screen minutes, number and timing of

prompts, learner and adult documentation time, abandonment, correction cost, perceived usefulness, and voluntary reuse alongside domain-defined evidence of understanding. Compliance under institutional, educator, caregiver, or supervisory pressure is not the same as desirability. If learners or the people asked to support the experience do not want to use it again when they have a meaningful choice, the design has failed an important field condition.

Educational partners should define the content, age range, acceptable exposure, learning evidence, and failure thresholds. Kitsune can make the tools and projections sufficiently malleable to compare approaches without rebuilding the entire system for each condition. The first experiment should remain small: one population, one learning objective, one Place or Space, one Thing or kit, one narrow utility, one bounded time window, and a predeclared path to stop, expire, or revise the design.

FIELD PRINCIPLE

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

3.5 A bounded first experiment

The first study proposed here should remain a bounded early-learning test that distinguishes the value of framing evidence from the value of AI-supported expression. It 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.

- **Condition A. Physical activity and usual reflection.** Learners complete the activity using the partner's existing reflection or documentation practice. No Kitsune-generated Perspective is introduced.
- **Condition B. Framed intent and minimal evidence.** The same activity adds an age-appropriate statement or co-framing of intent and the minimum deliberate capture needed to make the attempt and result legible. No AI-generated Perspective is returned.
- **Condition C. Reviewed, role-appropriate Perspective.** Condition B is extended with an attributed AI proposal that an authorized adult reviews before the proposal is expressed as a child-facing Perspective, a parent or caregiver Perspective, or a teacher Perspective. For younger learners, the adult operates the capture and review; direct child-facing AI is outside this first experiment.

The primary outcome should be curriculum-defined evidence of understanding, preferably explanation or transfer in a new off-screen task. Secondary outcomes should include screen minutes, adult documentation time, correction cost, perceived usefulness, and voluntary reuse. Privacy events, inappropriate cross-role exposure, uncorrected misinterpretation, and retention

exceptions are guardrails. The design should be narrowed or stopped 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, minimal evidence, or the returned Perspective. The third condition makes that distinction observable.

4. From a digital-first posture to a physical-first posture

The distinction is easiest to see as a change in the unit of authority. An application-centric system begins with a service and recruits people, places, and objects into its data model. A physical-first system begins with seeded Earth, Place, and Space guardrails, admits only a present participant to an eligible Thing and Utility, and activates only the Task, Action, and Agent required by the purpose. This is not a universal description of existing software; it is the contrast that guides the proposed runtime.

Design question	Typical application-centric posture	Proposed physical-first posture
Primary unit	App, account, or persistent profile	Presence-gated Utility in a seeded Earth, Place, Space, and Thing relationship
Basis for interpretation	Behavioral residue and later inference	Vivo lived data with presence, framed intent, and contextual evidence of effort and outcome
Context	Metadata added after capture	Spatial admission before Utility access; context constrains collection and use
Role of AI	Central interpreter and optimizer	Replaceable AI processor operating within a governed stream
Authority over action	Delegated through an app, account, or agent credential	Derived from presence eligibility, participant authorization, the active SOT, and the Utility-bound Task, Action, and Agent
Agent identity	May combine planning, memory, tools, and action in one persistent actor	Observer Agent Node, Personal AI Agent Node, Utility Agent Node, processor, data jurisdiction, and durable record remain distinct
Durable record	Accumulated raw events or platform history	Reviewed Perspective, provenance, participant-controlled continuity, and authorized Block proof of Data Well ownership
Interface	Destination through which activity is performed	Entry, capture, coordination, and visualization in support of activity

Table 2. A design contrast, not a claim that all existing digital systems share one architecture.

The phone remains useful in this posture, but its role changes. It is a tool of entry, authorization, capture, and visualization, not the permanent container for every utility and every context. A chair does not need a universal “chair app,” and it does not need to become an autonomous computer. It can carry an identity-bound utility that is discovered when relevant. A room can maintain a declared relationship to that chair, a table, or a tool. Meaning can then be resolved from deterministic relationships and the active purpose rather than from continuous ambient inference.

5. The physical-first runtime

The runtime is the execution model that makes the position operational. It is not one app, interface, model, or Block implementation. It separates spatial admission, Utility execution, evidence processing, owner review, and proof of continuity so the same rules can take different forms in education, land stewardship, professional work, or public experience. Current prototypes may use Mapbox for geospatial grounding, MindAR for discovery, and Three.js for visualization; those technologies are interfaces. The runtime described here is the governed sequence beneath them.

SPATIAL ACCESS: EARTH → PLACE → SPACE → THING → UTILITY
 UTILITY EXECUTION: TASK + ACTION + AGENT (HI OR AI)
 GOVERNED RETURN: MOMENT → PERSPECTIVE → NEXT PHYSICAL ACTION

PUBLIC-LANGUAGE BRIDGE

Intent → Action → Consequence → Perspective → Next action. This describes the same governed loop in plain language. Perspective remains revisable understanding, not a score, diagnosis, or verdict.

- **Earth and Place.** Earth is the shared planetary reference. Place is a seeded, named environment with authority, purpose, and relationships. Together they establish the first spatial guardrails for presence eligibility rather than serving as incidental coordinates.
- **Space.** A Space is a seeded, bounded, socially meaningful environment such as a classroom, workshop, room, field site, or temporary zone. Presence inside the eligible Space governs access to its related Things and Utilities.
- **Thing.** A Thing is an addressable physical artifact or resource connected to its eligible Place and Space. Identity and declared relationships matter more than embedded compute.
- **Utility.** A Utility is a purpose-bound application packaged for temporary use in context. It opens only when the runtime confirms presence in its eligible seeded context. Inside the Utility, a specific Task, Action, and Agent (HI or AI) are bound to the active purpose; they are execution details, not additional spatial tiers.

- **Moment.** A Moment is the governed evidence window in which intent, Utility execution, Vivo and Vitro data, review, and consequence become specific after spatial admission.
- **Perspective.** Perspective is a derived, reviewable account of what occurred, what changed, and what may matter next. It is not treated as objective truth or a permanent identity claim.

5.1 Intent, effort, and data agency

The runtime begins with intent because action traces are underdetermined. The same observable action can express different goals, constraints, or levels of understanding. Intent may be declared by the participant, adopted from a shared mission, negotiated with a teacher or collaborator, or revised during the activity. It is not assumed to be perfectly articulated. Its value is that it gives later evidence an explicit point of comparison.

Vivo data is lived data with presence, captured in or out of a Moment. Presence is a runtime access condition: a person is either admitted to the eligible Earth → Place → Space → Thing → Utility path or the Utility remains unavailable. Once admitted, Vivo may preserve the relevant artifact, observation, change, explanation, decision, measurement, relationship, or outcome under the defined protocol. Vitro data is digital data produced by Utilities or Agent processing and must remain linked to its Vivo source, Task, Action, Agent (HI or AI), and review history.

WORKING THESIS

DECLARED INTENT + DIGITIZED EFFORT → DATA AGENCY

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

5.2 Moment gates as capture-time jurisdictions

The runtime separates four types of moments: private, personal, professional, and public. These are not cosmetic visibility settings applied after collection. They are intended as operational jurisdictions selected before or at activation, each carrying its own rules for purpose, participants, processors, retention, sharing, and review. The same person may act across all four, but the resulting streams should not be combined merely because a system can technically join them. This design extends the logic of contextual integrity into the architecture of capture itself (Nissenbaum, 2004).

SENSITIVE CAPTURE BOUNDARY

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

- **Private.** A moment intended for the individual or a tightly authorized circle; default disclosure is minimal.
- **Personal.** A moment connected to self-directed life, family, or trusted relationships, with participant-controlled sharing.
- **Professional.** A role-bound moment governed by organizational purpose, responsibility, and applicable policy.
- **Public.** A moment prepared for possible wider release under explicit terms; public does not mean unrestricted reuse. Creating, reviewing, or completing a Public Moment does not publish it. Release is a separate, authorized action.

5.3 Data wells and governed information flows

Private HQ is a signed-in interface, but authentication is not the source of authority. Sign-in may admit a person to their person-centered System of Things only after the relevant identity, admission, and ownership rules are satisfied. Personal, Root, Partner, and Project SOTs remain distinct authority boundaries; access to one does not grant access to another. Cross-boundary use occurs only through an explicit, purpose-bound projection that can be reviewed and revoked; Partner and Project projections are de-identified by default unless an authorized purpose requires otherwise.

A Data Well is not a raw data lake, a comprehensive profile, or an automatic product of login. 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 protocol requires. 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. MER always means Master Engagement Record.

Governed information flow means controlling how one provenance-bearing record is transformed into different authorized representations for different purposes and roles. It does not mean that every participant receives the same record. In an early-learning context, a child may receive a spatial or visual learning journey, a parent may receive a joint journal or conversation prompt, and a teacher may receive a concise, source-linked point of reference. In advanced study, the learner may work directly with a research or practice Perspective while advisors, collaborators, or institutions receive only the bounded projection their roles authorize.

A study receives only the consented, de-identified evidence it requires. Each projection retains its origin, scope, uncertainty, and review history. Representation is malleable; authority is not. W3C PROV offers a relevant foundation for representing how entities, activities, and agents contribute to derived information (W3C, 2013).

5.4 One environment: Spatial, AI, and Block

Spatial computing, artificial intelligence, and block-based proof are commonly positioned as independent movements. Physical-first intelligence treats them as one constrained environment. Spatial seeds Earth, Place, Space, Thing, and Utility relations and enforces the binary presence gate. AI processes authorized Vivo evidence through bounded Utilities or Agents and returns attributed Vitro proposals for review. Block preserves proof of Data Well ownership and owner-authorized transitions after review and promotion.

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

The boundaries are as important as the contributions. Spatial admission does not authorize unrelated capture. AI processing creates no ownership, standing access, or right to determine the next direction. The Block is not the Data Well and does not establish lived truth; it preserves proof that an owner authorized a defined continuity transition.

BLEND-ENVIRONMENT PRINCIPLE

A person cannot remotely use or react through a Place-bound Utility unless they were physically present in the seeded context. Digital participation follows lived participation. The objective is not a more complete online world to live through, but situated intelligence that people and institutions can live with.

5.5 Iterative human-AI work

Iterative work with a language model creates a limited human-AI coupling. A person states an intent, evaluates the model's response, corrects errors or drift, and revises the next instruction or action. This can support an ongoing reasoning process without making the model identical to the person or authoritative over the person's record.

DECLARED INTENT → MODEL RESPONSE → HUMAN REVIEW + CORRECTION → REVISED INTENT OR ACTION ↻

The dialogue alone is insufficient. Context may remain session-bound, digital output may separate from physical consequence, and provider memory may not preserve provenance or participant authority. Kitsune extends the loop by connecting declared intent, selected evidence from physical activity, attributed processing, human review, and a next physical action.

Kitsune uses Symbiont Mind as design shorthand for that broader governed relationship. The term does not claim that a person and an AI form a literal mind, share ownership, or necessarily benefit each other. The person remains the authority, processing creates no automatic ownership, and whether the relationship improves human capability without creating dependency, persuasion, or extraction remains an empirical question.

5.6 BDA as the laboratory observation process

BDA is the recursive Before-During-After human-AI working loop. A person declares intent, asks the AI to reflect the meaning it understands, reviews developing work, corrects drift, and decides what continues. The runtime makes the evidence surrounding that relationship visible and correctable; it does not redefine BDA or make the model the authority.

BDA · LABORATORY POSTURE
HUMAN INTENT → VIVO / VITRO EVIDENCE → 72-HOUR EPHEMERAL PROCESSING → HUMAN
REVIEW → REVIEWED PERSPECTIVE

- **Observe.** Purpose-bound physical engagement and presence-gated Utilities produce selected signals and multimodal artifacts with Earth, Place, Space, Thing, Utility, Task, Action, Agent (HI or AI), time, source, and Moment provenance. Within the Observatory environment, Vivo and Vitro evidence can be viewed together as a spatiotemporal stream rather than flattened into a chat transcript or event table.
- **Process.** Inside the Ephemeral Intelligence Plane, multiple authorized, substitutable processors or agent nodes may examine the same bounded evidence through different declared processing lenses or viewpoints. Every annotation, association, classification, summary, or candidate remains attributed Vitro output: a proposal for review, not a fact, owner decision, Perspective, or durable record. An explicit DLA controls eligible evidence, transformations, outputs, and stopping conditions; processing never transfers ownership or authority.
- **Review and purge.** Ephemeral is a short-lived laboratory workspace with a governed 72-hour custody cycle. During that window, the owner can compare multimodal evidence and processor proposals, inspect uncertainty and provenance, and correct or reject drift. Unsupported raw material expires or is purged according to the active protocol. The plane may return a provenance-bearing cycle record or a candidate for MER review; it does not mint a final MER or permanent truth.
- **Promote only by authority.** The proposed lifecycle is Signal → 72-hour Ephemeral cycle record → reviewed Moment → accepted Perspective → weekly Observatory Perspective → monthly Master Engagement Record → authorized Block. The Block preserves proof of Data Well ownership and the owner-approved transition; it is not the Data Well. Promotion, release, proof, and action remain separate governed decisions.

The interface below makes the laboratory posture tangible. It shows the Ephemeral view inside the N.01 Observatory environment: a 72-hour lived field, review lenses for presence, events, multimodal evidence, Perspective, collective use, and return; a spatiotemporal Observer plot; a living evidence wall; signal count and time remaining; and a 72-hour scrubber. The surface is designed to let a person see the stream and its transformations before deciding what should endure.




Figure 1. Prototype N.01 Ephemeral Intelligence Plane within the Observatory environment, captured August 25, 2026. The interface exposes event evidence, spatial relationships, temporal custody, review lenses, and the 72-hour time scrubber. It demonstrates the intended laboratory posture; it does not establish that the full governed promotion path through Observatory and MER is complete.

LAB STATUS Recent implementation checkpoints include Ephemeral visualization and agent-flow surfaces. The deterministic end-to-end path from phase closure through weekly Observatory Perspective, monthly MER, and authorized Block proof remains under reconciliation. This figure is interface evidence from an active prototype, not proof of a completed pipeline.

5.7 The System of Things is the person-centered environment

A System of Things 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 the active relationship and purpose rather than reconstructed through continuous observation. Interfaces such as Public HQ and Private HQ provide entry and review, but authentication alone does not create authority.

Program operations belong to the applicable Partner SOT; local participants, learning Places, utilities, and activity belong to the applicable Project or personal SOT. Movement across those boundaries requires an explicit, purpose-bound release or projection.

- **Data-Logic-Action (DLA) Unit.** A DLA Unit is a versioned governance boundary that specifies the data a utility may use, the processing it may perform, the outputs and actions it may produce, and the retention and stopping rules that apply. It operationalizes purpose limitation and least privilege; it cannot grant authority beyond the active SOT and Moment.
- **Goal-Based Tasks (GBT).** GBT narrows the active objective, permitted tasks, and conditions under which the utility should stop, ask, or return. It cannot grant authority beyond the DLA Unit or active Moment.
- **Presence-gated Utility access.** Seeded Earth, Place, and Space guardrails establish whether a person is physically present for an eligible Thing and Utility. The decision is binary at runtime: present opens the eligible Utility; not present denies its use. A remote participant cannot access, react to, or 'like' through that Place-bound Utility. False admission, false denial, and spoofing are security failures to test and correct without weakening the rule.

5.8 Canonical agent-node names

N.01 is the Observer Agent Node; N.02/P.AI is the Personal AI Agent Node; and N.03 is the Utility Agent Node. These names identify node classes only. The active SOT and utility contract separately define behavior, data access, retention, and action authority. A processor remains replaceable and gains no ownership, authority, or control of the governed record through processing.

6. Agentic capability without agentic sovereignty

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

Kitsune's distinction is structural: processing authority is derivative of the active System of Things, presence-gated Utility, and Moment contract, not a standing property of an account, memory, credential, or node name. A Task, permitted Action, and Agent (HI or AI) are bound inside the eligible Utility. None of those capabilities independently determines why data exists, which jurisdiction applies, what may become durable, or who may use the result.

This posture applies established security ideas to the governance of human activity data. Least privilege limits each role and processor to the minimum capability required; complete mediation requires authority to be checked at each protected action rather than assumed from a prior login (Saltzer & Schroeder, 1975). NIST's 2026 work on software and AI agent identity similarly treats identification, authorization, delegated authority, logging, and provenance as open implementation 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 presented as compliance with a completed agent standard.

For this paper, agentic sovereignty names a failure condition rather than a category applied to every agent: a software actor accumulates enough cross-context memory, delegated credentials, goal-setting power, and action rights to become the de facto authority over a person or domain. The proposed architecture attempts to prevent that condition through separation of roles, jurisdictions, contracts, processor slots, and durable records. It does not guarantee safety. A flawed SOT, an overbroad DLA Unit, a coercive institution, a compromised orchestrator, or an unreviewed model output can still produce harm. The value of governance is that authority and failure become more visible, bounded, and auditable.

7. What the position rejects

A new term is useful only if its boundaries are clear. Physical-first intelligence rejects seven shortcuts that would reproduce the digital-first posture under new language.

- **It is not anti-screen or analog-only.** Screens, maps, augmented reality, voice, and three-dimensional visualizations can be valuable when they return perspective to physical activity rather than replace it.
- **It is not robotics under another name.** Robotics and vision-language-action research concern machines perceiving and acting in the physical world. This position concerns the governance and interpretation of human activity, whether or not a robot is present.
- **It is not ambient surveillance.** Continuous cameras and sensor oversight are not required. The runtime favors activated, purpose-bound capture and declared relationships over generalized observation.
- **It is not a claim that an access gate resolves truth. The presence rule decides whether a Place-bound Utility may be used. Vivo provenance, human review, uncertainty, and lived consequence still determine what can responsibly be said about the resulting evidence.**
- **It is not a totalizing model of a person.** Perspective is domain-bound and revisable. A learning record should support the learner and authorized participants in the learning context; it should not become a portable verdict about the learner.

- **It is not intelligence owned by one model.** Models are processors. They may be selected, compared, constrained, or replaced. The governed stream and the participant's authority are the durable system.
- **It is not a human-data replenishment strategy.** People are not valuable because they produce fresh material for model training. Lived contribution may be authorized for a specific purpose, but no processor receives a standing claim on a person's life, identity, attention, or whole Data Well.

8. Social protocol is part of the runtime

Physical integration changes how agents enter human situations. A model that can speak is not automatically entitled to interrupt, and a completed task does not create a right to demand attention. Timing, permission, role, safety, and the right to decline are part of the system behavior, not decorative sentiment. 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, and whether the participant wants the result (Amershi et al., 2019). Hardware may support these interactions, but the thesis is that agents participate under social and data protocols legible to the humans whose environment they enter.

9. Research agenda: the learning hypothesis first

The position becomes scientific only when its claims can fail. The primary learning proposition and the secondary runtime propositions below should become operational hypotheses in a methods paper; they are not findings. A pilot does not need to test every proposition at once.

9.1 Primary learning hypothesis

PRIMARY HYPOTHESIS · TO BE TESTED

FRAMED INTENT + DIGITIZED EFFORT → ROLE-APPROPRIATE PERSPECTIVE → GREATER UNDERSTANDING

In a bounded physical-learning context, the proposed mechanism should increase role-appropriate understanding without displacing physical activity, increasing unnecessary screen time, or creating disproportionate work for the learner or other authorized participants.

- **P1. Understanding.** Compared with a final artifact, completion record, or generic summary alone, an authorized role-appropriate Perspective will be associated with stronger curriculum-defined evidence of understanding, such as explanation, recall, revision, or transfer in a new off-screen task. The educational partner must choose the measure and meaningful effect before the study.

- **P2. Evidence of change.** Framed intent plus minimal digitized effort will reduce disagreement about what the learner attempted and what changed compared with interpretation of the outcome artifact alone, while leaving measurable uncertainty and room for correction.
- **P3. Role relevance.** The learner and other authorized participants in the learning context will find separately designed Perspectives more useful for their respective next actions than one shared summary, without increasing inappropriate cross-role exposure. In the first early-learning experiment, those roles are the child, parent or caregiver, and teacher.
- **P4. Burden and desirability.** The experience will meet predeclared ceilings for screen time, prompts, adult documentation, and correction cost, and will achieve a predeclared voluntary-reuse threshold. Failure on burden or reuse is evidence against the design even if a learning measure improves.
- **P5. Developmental exposure.** For younger learners, adult-mediated or joint journaling will be compared with any direct child-facing condition. The study will not assume that more direct AI exposure is better; equal or stronger learning evidence with less exposure should favor the lower-exposure design.
- **P6. Governance and correction.** Purpose-bound projections, provenance, review, and withholding controls will reduce unnecessary raw-data exposure and cross-role leakage compared with a whole-record view. Participants or authorized adults must also be able to identify and correct material misinterpretations.

9.2 Secondary runtime propositions

The following propositions concern the general runtime. They remain relevant to education but should not be made prerequisites for demonstrating the narrower learning hypothesis.

- **R1. Spatial admission, intent, and provenance.** Presence-gated Earth → Place → Space → Thing → Utility access will prevent remote use outside the seeded context while meeting predeclared ceilings for false admission, false denial, spoofing, accessibility burden, and recovery. Framed intent and provenance should then reduce disagreement compared with action logs alone.
- **R2. Bounded review.** A 72-hour Ephemeral review of attributed Vivo and Vitro evidence will help authorized reviewers identify provenance gaps and contest derived interpretation more effectively than a flat log or single unreviewable summary; protocol-bound expiry should reduce retention exposure.
- **R3. Purpose and action boundaries.** Spatial admission, capture-time Moment gates, and explicit Task, Action, Agent, DLA Unit, and GBT constraints will reduce inappropriate cross-context reuse and unauthorized Utility or data use compared with functionally similar broad standing authority.
- **R4. Processor substitutability.** Replacing a processor while holding the owning System of Things, Moment contract, and accepted behavior constant will preserve authorization and provenance semantics, even when the quality of derived Perspective changes.

9.3 Methods threshold

Before recruitment or field use, a study should declare the population and owning SOT, educational objective, physical context, comparison conditions, exposure levels, unit of analysis, outcome instruments, burden ceilings, privacy and adverse-event definitions, retention and expiry rules, stopping criteria, and treatment of missing data. Understanding should not be reduced to one dashboard metric. Results should report learning evidence, burden, uptake, correction, privacy events, null effects, and negative outcomes together.

10. Failure modes and limitations

The runtime cannot be evaluated only when it works as intended. Fourteen failure modes must remain visible in research and product decisions.

- **Declared intent can be incomplete or performative.** Participants may not know their intent, may revise it, or may state what they believe 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 condition.**
- **Capture changes behavior.** Journaling, prompts, cameras, or agent interventions can burden the activity and alter the phenomenon being studied. A high-friction tool can steal attention from physical learning, shift work onto learners, educators, caregivers, mentors, or collaborators, and produce compliance that is mistaken for value.
- **Derived perspective can distort.** Summaries, model outputs, and compelling learner-facing visualizations may omit context, overstate coherence, or introduce bias. Supporting evidence, uncertainty, and correction paths are essential.
- **Domain structures may not transfer.** A SOT designed for education may encode relationships and values that are inappropriate in health, work, land, or public contexts.
- **Access is uneven. Presence-gated discovery and visual wrappers can exclude participants because of device, connectivity, mobility, vision, hearing, language, time, developmental stage, adult-support requirements, or false denial. Accessibility and recovery paths must be designed without converting physical presence into remote Utility access.**
- **Continuity can become surveillance.** A useful record can become coercive when purposes expand, retention becomes indefinite, or an institution treats perspective as a fixed identity.
- **Governance structures can be wrong.** A DLA Unit, GBT structure, or SOT can encode an overbroad purpose, omit a harmed participant, or authorize an action that should never have been allowed.

- **Agent-node separation can be nominal.** The canonical N.01, N.02/P.AI, and N.03 names provide no protection if one credential, processor, or operator can silently bypass their accepted boundaries.
- **Iterative confirmation can amplify error.** Repeated confirmation between a person and a model can reinforce a mistaken premise, narrow alternatives, or create false confidence rather than improve thought.
- **Coupled continuity can create dependency or persuasion.** A continuity layer can make one processor, style, or recommendation feel inseparable from the person, increasing lock-in or reducing independent judgment.
- **Method or Block proof can be overclaimed. Sources may select favorable evidence or obscure failed conditions, and a Block that proves an ownership transition may be misrepresented as proof that the underlying event or interpretation is true.**
- **Synthetic output or inferred intent can be misattributed.** Vitro material can enter a record without retaining its source, be represented as Vivo evidence, or be attributed to a participant as though the participant expressed or authorized it. Source, authorship, and transformation history are required for later review.
- **Laboratory visibility can be mistaken for truth.** A compelling spatiotemporal display, a dense multimodal record, or agreement among several processors can create false confidence. Every artifact and proposal still requires provenance, uncertainty, owner review, and authority outside the visualization.

These are not edge cases to be solved after deployment. They define the conditions under which the thesis should be rejected, narrowed, or redesigned. In child-facing work, independent ethics review, developmentally appropriate assent, parental permission where required, data minimization, retention limits, and educator oversight are prerequisites, not claims of safety.

11. Research and governance commitments

Kitsune proposes the following commitments for systems and studies developed under this position. The next methods paper must translate each commitment into data structures, operational controls, audit procedures, and measurable criteria.

- **Purpose before capture.** A utility and moment gate must identify why data is needed before collection begins.
- **Presence before Utility access.** Enforce seeded Earth, Place, and Space eligibility before a person can reach a Place-bound Utility; bind Task, Action, and Agent only after admission.
- **Minimum necessary evidence.** Collect only what the moment requires; do not treat future model usefulness as unlimited justification.

- **Provenance with uncertainty.** Record how evidence was generated and transformed, and preserve confidence, missingness, and contestation.
- **Separation by jurisdiction.** Private, personal, professional, and public streams do not merge by default.
- **Participant review, correction, and stop authority.** People must be able to inspect and challenge derived Perspective before it becomes durable or consequential, and to pause or stop before an irreversible action.
- **Processor substitutability and portable continuity.** No model is the permanent owner of the record or sole authority on its meaning. Authorized intent and Perspective must remain separable from any one model, vendor, session, or interface.
- **Authority outside probabilistic output.** Permissions, jurisdiction, retention, and consequential state transitions must be enforced by the governed runtime rather than inferred from model text.
- **Least privilege and revocation.** Every agent role, processor, tool, and data projection must receive only the authority required for the active purpose, with an explicit path to suspension or revocation.
- **Agent identity and action provenance.** The system must distinguish participant expression from processor inference and identify every participant, canonical agent node, processor, tool, and operator involved in consequential action. Inference is not participant authorization.
- **Bounded retention.** Raw evidence should expire according to protocol; durable continuity should be intentional and proportionate.
- **Block proof without data custody.** Use authorized Block records to preserve proof of Data Well ownership and owner-approved lifecycle transitions, never as a substitute for the Data Well, its evidence, or human review.
- **Transparent field evidence.** Educational partners must help define learning measures, exposure, burden thresholds, and adverse events. Pilots must report limitations, screen and documentation burden, abandonment, privacy events, null effects, and negative results alongside positive findings.
- **Special protection for children.** No child-facing SOT should rank, diagnose, manipulate, or define a child beyond the explicit educational purpose and independently reviewed protocol.
- **Physical and social primacy.** AI must remain a bounded complement to the learning environment. A study must declare the maximum screen and documentation burden it will tolerate and must not treat more interaction as success by itself.
- **Developmentally appropriate exposure and role boundaries.** Direct child use is not a default. Child, parent or caregiver, teacher, and program-research projections require separate purpose, authority, and minimum-necessary evidence; no role inherits a whole record.

- **Source labeling.** Vivo and Vitro evidence must remain distinguishable, and every transformation must preserve the source relations needed for later review.

12. Conclusion

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 be reduced to 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 learners, families, educators, and institutions live with, not an environment they are required to live through.

Education makes the claim testable. A bounded learning Moment can preserve what a learner intended, attempted, created, discovered, and understood, then return a reviewed Perspective to the learner and distinct authorized Perspectives to others involved in that learning context. In the first early-learning experiment, those participants may be the child, parent or caregiver, and teacher. The broader approach merits further study only if those Perspectives improve understanding while keeping physical activity primary, screen time proportionate, participant burden acceptable, and authority legible.

Educational and research partners bring curriculum or disciplinary objectives, educators, mentors, families, developmental and domain judgment, and the realities of physical learning. Kitsune brings malleable tools for presence-gated spatial access, journaling, attributed AI processing, governed information flow, and Block proof of Data Well ownership. Together, they can compare approaches rather than deploy a fixed answer. The appropriate next step is one reversible experiment with predeclared measures and a clear path to stop. The goal is not to prove that AI belongs in education. It is to establish when a carefully bounded, blended environment helps learners understand more—and when it does not.

Appendix A. Secondary hypotheses: collective method and value

The individualized-collective model is secondary to the learning hypothesis. It proposes that a method source and a situated implementation remain distinct, attributable sources: the first contributes a model of action; the second contributes the model in action; reviewed Perspective may connect them as method proof. This construct is not required for a learner or another authorized participant to receive a useful Perspective, and it should not be made a prerequisite for an educational pilot.

Future studies may test whether dual provenance helps an evaluator distinguish method effect, local adaptation, and contextual failure, and whether selective evidence can remain useful with less raw personal exposure than whole-record access. These claims are unproven. They create no automatic authorization for publication, payment, licensing, identity disclosure, or access to a whole Data Well.

References

The external references below establish adjacent scientific and technical foundations. Kitsune's formula, runtime synthesis, Moment gates, System of Things architecture, and learning mechanism are proposed constructs that require independent evaluation.

- Amershi, S., Weld, D., Vorvoreanu, M., Fourney, A., Nushi, B., Collisson, P., et al. (2019). Guidelines for human-AI interaction. Proceedings of the 2019 CHI Conference on Human Factors in Computing Systems, Paper 3, 1–13. [Source](#)
- Assran, M., Duval, Q., Misra, I., Bojanowski, P., Vincent, P., Rabbat, M., LeCun, Y., & Ballas, N. (2023). Self-supervised learning from images with a joint-embedding predictive architecture. Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition, 15619–15629. [Source](#)
- Black, P., & Wiliam, D. (1998). Assessment and classroom learning. Assessment in Education: Principles, Policy & Practice, 5(1), 7–74. <https://doi.org/10.1080/0969595980050102>
- Booth, H., Fisher, B., Galluzzo, R., & Roberts, J. (2026). Accelerating the adoption of software and AI agent identity and authorization (Draft concept paper). National Institute of Standards and Technology, National Cybersecurity Center of Excellence. [Source](#)
- 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. [Source](#)
- Brown, J. S., Collins, A., & Duguid, P. (1989). Situated cognition and the culture of learning. Educational Researcher, 18(1), 32–42. [Source](#)
- Chi, M. T. H., de Leeuw, N., Chiu, M. H., & LaVancher, C. (1994). Eliciting self-explanations improves understanding. Cognitive Science, 18(3), 439–477. [https://doi.org/10.1016/0364-0213\(94\)90016-7](https://doi.org/10.1016/0364-0213(94)90016-7)
- Clark, A., & Chalmers, D. (1998). The extended mind. Analysis, 58(1), 7–19. [Source](#)
- Dourish, P. (2001). Where the action is: The foundations of embodied interaction. MIT Press. [Source](#)
- Hattie, J., & Timperley, H. (2007). The power of feedback. Review of Educational Research, 77(1), 81–112. <https://doi.org/10.3102/003465430298487>
- Hutchins, E. (1995). Cognition in the wild. MIT Press. [Source](#)
- 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. [Source](#)
- LeCun, Y. (2022). A path towards autonomous machine intelligence (Version 0.9.2). OpenReview. [Source](#)
- Licklider, J. C. R. (1960). Man-computer symbiosis. IRE Transactions on Human Factors in Electronics, HFE-1(1), 4–11. [Source](#)
- National Academies of Sciences, Engineering, and Medicine. (2018). How people learn II: Learners, contexts, and cultures. The National Academies Press. [Source](#)
- National Institute of Standards and Technology. (2026). AI Agent Standards Initiative. Updated August 14, 2026. [Source](#)

- Nissenbaum, H. (2004). Privacy as contextual integrity. *Washington Law Review*, 79(1), 119–157. [Source](#)
- Saltzer, J. H., & Schroeder, M. D. (1975). The protection of information in computer systems. *Proceedings of the IEEE*, 63(9), 1278–1308. [Source](#)
- 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. [Source](#)
- Tabassi, E. (2023). Artificial Intelligence Risk Management Framework (AI RMF 1.0) (NIST AI 100-1). National Institute of Standards and Technology. [Source](#)
- W3C. (2013). PROV-O: The PROV ontology. W3C Recommendation 30 April 2013. [Source](#)
- Weiser, M. (1991). The computer for the 21st century. *Scientific American*, 265(3), 94–104. [Source](#)
- Wilson, M. (2002). Six views of embodied cognition. *Psychonomic Bulletin & Review*, 9(4), 625–636. [Source](#)
- Wooldridge, M., & Jennings, N. R. (1995). Intelligent agents: Theory and practice. *The Knowledge Engineering Review*, 10(2), 115–152. [Source](#)
- 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*. [Source](#)

Internal design source

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 PHYSICAL-FIRST INTELLIGENCE

POSITION PAPER 01 / LEARNING FIELD EXPRESSION

Position Paper 01 applies the foundational position to learning and education as the first scientific field expression. It should be read with Position Paper 00, which defines the broader human-AI-data-consequence relationship.

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>
