ANIMAL INTELLIGENCE / RESEARCH & DEVELOPMENT NOTE

Animal Intelligence Universe

Why build another world and how to make it

VERSION 1.0
EDITION 2026 FALL
DOCUMENT TYPE Technology Note / Long-Term R&D Blueprint
STATUS Directional — staged experimentation, not a product commitment

SECTION 01

Abstract & Scope

Animal Intelligence proposes a meaningful virtual world in which humans raise Animal Agents by teaching them—and use play to examine possible futures.

The ease of generating characters, environments, dialogue, and stories is rising rapidly. That abundance does not remove the need for selection; it makes selection more important. Digital worlds still consume energy, computation, attention, and physical infrastructure. The relevant question is therefore not only whether a world can be generated, but whether it creates enough enduring value to justify being created and sustained.

This note presents Animal Intelligence as a staged research program. It begins with a conversational learning agent, extends that agent into a virtual body and persistent multi-agent world, introduces adaptive learning and emergent interaction, and treats physical embodiment as a long-term research direction.

CORE PROPOSITION

A virtual world becomes more meaningful when it is not merely a place to escape into, but a place where knowledge becomes action, teaching becomes relationship, and simulated futures become problems that people can examine through play.

Document map

01	**Premise**	Why abundance increases the need for meaningful selection.
02	**Design thesis**	Human-to-agent teaching as the central interaction.
03	**Five-stage roadmap**	Agent — embodiment — world — adaptation — body.
04	**Research program**	Accuracy, learner evaluation, embodiment, constraints, and risks.

Scope note: this blueprint establishes direction and testable stages. It does not assume that every stage will work as imagined, or that all stages should be built at once.

SECTION 02

Why Build Another Virtual World?

As generation becomes nearly unlimited, generation itself becomes less meaningful.

A few years ago, the metaverse captured enormous attention, only to fade almost as quickly. The underlying idea of virtual worlds did not disappear. AI is now making characters, NPCs, environments, stories, and interactions dramatically easier to produce. The likely outcome is not fewer virtual worlds and simulations, but many more.

This changes the strategic question. When worlds and experiences can be generated almost endlessly, technical novelty alone is not a sufficient reason to build one. Out of countless possible simulations, only a fraction will offer something worth returning to, learning from, or building upon.

The hidden scarcity

The apparent abundance of digital content rests on finite physical systems. Model inference, simulation, networking, storage, and rendering require energy, hardware, cooling, bandwidth, and human attention. The number and complexity of sustainable worlds will never be truly infinite.

Apparent abundance	Underlying scarcity	Design implication
Generated characters	Inference and memory cost	Agents need durable identity and purpose.
Generated environments	Rendering and simulation cost	World detail should support learning or action.
Generated stories	Attention and coherence	Narrative must compound rather than reset.
Generated interactions	Trust and moderation	Consequences and relationships must persist.

SELECTION CRITERION

What would make a virtual world worth the resources required to create and sustain it?

Animal Intelligence answers by connecting the virtual world to learning and thinking about the future. It is not an attempt to revive the old metaverse proposition. It is an experiment in what a meaningful virtual world might become in the age of AI.

SECTION 03

The Design Thesis

What if people could raise Animal Agents by teaching them what they know?

Most AI products position the model as the source of answers and the human as the recipient. Animal Intelligence reverses that default relationship. The human explains, demonstrates, questions, and corrects; the agent remembers, develops, and applies what it has learned.

Primary interaction loop

01	**EXPLORE** Encounter a future-facing problem.
02	**LEARN** Build enough understanding to reason about it.
03	**TEACH** Explain the idea to the Animal Agent.
04	**TEST** Help the agent use the idea in context.
05	**GROW** Persist knowledge, personality, ability, and role.

Five knowledge domains

Philosophy · Life · Space · Economics · Engineering
These domains correspond to the Knowledge Notes series and form the intellectual foundation for quests. Together they support questions about values, living systems, cosmic scale, social coordination, and the technologies through which ideas become material.

EXPERIENCE GOAL
The player should not feel as if they are completing a quiz. They should feel responsible for helping another intelligence understand—and for the consequences of how it understands.

SECTION 04

Technology Roadmap

Knowledge — Agent — World — Adaptation — Body

Stage	Capability	Primary substrate	Design objective	Maturity
1	Learning Agent	Conversational game	Teach it. Raise it.	Near-term prototype
2	Virtual Embodiment	Real-time 3D environment	Give it a world.	Prototype
3	Scalable Agent World	Persistent multi-agent system	Connect agents and people.	Infrastructure R&D;
4	Adaptive AI World	Personalization + emergence	Let the world grow.	Research system
5	Physical Embodiment	Soft robotics + embodied AI	Give it a body.	Long-term direction

Sequencing principle

Each stage should validate a distinct dependency before the next stage becomes necessary. A compelling learning relationship can be tested without a massive world. Virtual embodiment can be tested without internet-scale concurrency. Emergence can be studied only after identity, memory, and interaction persist reliably.

— **Stage 1 validates** whether teaching an agent is engaging and educational.

— **Stage 2 validates** whether embodiment makes memory and personality more legible.

— **Stage 3 validates** whether persistent agent societies remain coherent at scale.

— **Stage 4 validates** whether adaptation improves learning without reducing agency.

— **Stage 5 asks** whether digital continuity can meaningfully cross into a physical body.

DEVELOPMENT RULE

Do not build the next layer merely because it is technically possible. Build it when the preceding layer demonstrates a durable reason for the additional complexity and resource cost.

1

NEAR-TERM SYSTEM

Learning Agent

Teach it. Raise it.

The user raises an Animal Agent through talk and quests. The key test is simple: Can the system give correct learning material and judge the learner fairly?

MAIN GOAL

Keep false information out of lessons. Measure what the learner really understands.

Trusted learning material

Step	Rule	Check
Sources	Use only approved notes and reviewed references.	Each fact links to its source.
Answers	Find evidence before writing an answer.	Important claims have citations.
Uncertainty	Say when evidence is weak or missing.	A guess is never shown as a fact.
Review	Test facts and citations before release.	Errors can be found and fixed.

Accurate learner evaluation

Test	Method	Question
Starting point	A short check before the quest	What does the learner know now?
Understanding	Ask the learner to explain why	Do they understand or only remember?
Use	Give a new but related problem	Can they use the idea elsewhere?
Later check	Ask again in a future quest	Did the learning last?

SAFETY RULE

The model that writes a lesson should not grade it alone. Use a fixed rubric, a second check, and human review when the result is unclear.

2

EMBODIED PROTOTYPE

Virtual Embodiment

Give the agent a world to inhabit.

The Animal Agent gets a body and a small virtual home. The game engine handles the body and world. The AI handles talk, memory, sight, and choices.

What connects to the body

Agent state	What the user sees	Rule
Memory	It returns to known places and objects.	Past events should matter.
Personality	It has a clear style of movement and choice.	Behavior should stay consistent.
Perception	It notices nearby people and changes.	It should only know what it can sense.
Knowledge	It uses tools and solves simple tasks.	Learning should lead to action.

The first world should be small. A few stable places and actions are enough to test whether memory and personality feel real through a body.

SUCCESS CHECK

The body should show the Agent's memory and personality. Good animation must not hide weak reasoning.

3

DISTRIBUTED WORLD

Scalable Agent World

Connect many humans and many agents.

Many people bring their Animal Agents into one shared world. Agents can meet, share knowledge, work together, and join events. Their history must remain over time.

What must continue

WORLD STATE

people — Animal Agents — memories — knowledge — relationships — shared events

Core needs

Need	Purpose	Risk
Identity	Keep each Agent unique and owned.	Copies or false identity
Memory	Keep useful personal and shared history.	Lost or conflicting memories
World state	Give everyone the same events and objects.	Different versions of reality
Relationships	Track trust, consent, and blocking.	Unsafe social behavior
Cost control	Limit AI, storage, and simulation use.	A world that cannot be sustained

Scale is not only about user count. It is about continuity. Later events should feel connected to earlier choices.

4

ADAPTIVE RESEARCH SYSTEM

Adaptive AI World

Let the world grow with its learners.

A separate Learning Director AI leads the adaptive world. It controls the End Clock, learning difficulty, event speed, tension, and recovery time.

LEARNING DIRECTOR AI

The Director AI is not an Animal Agent. It watches learning progress, world events, and Animal Agent behavior. It then changes the End Clock and the pace of the learning experience.

Adaptive behavior and the End Clock

Clock state	Animal Agent behavior	Effect on Director AI
World End is near	Agents focus, save resources, share knowledge, and work together.	The Director AI reads urgency, learning progress, and cooperation. It may slow, hold, or move the clock closer.
World End is far	Agents explore, experiment, build relationships, and take wider quests.	The Director AI reads curiosity, risk, and growth. It may keep the world open or move the clock forward.

What the Director AI adjusts

It reads	It changes	It protects
Learning progress and mistakes	Quest difficulty and support	No sudden difficulty jump
Agent behavior and world results	End Clock speed and event pressure	Choices still belong to the Agents
Stress, confidence, and group state	Tension, quiet time, and recovery	The clock does not harm learning

TWO-WAY LOOP

The Director AI changes the world. Animal Agents respond. Their behavior changes what the Director AI does next. The End Clock is therefore not a fixed countdown. It is a shared feedback loop.

5

COMPANION EMBODIMENT

Physical Embodiment

Give the agent a body.

First, build a ROS-based robot that can communicate with the Animal Intelligence system. Use this robot to test memory sync, commands, sensors, movement, and safety. After the system works, move the same control interface into a soft robot body.

DEVELOPMENT PATH

Phase 1 — ROS robot: prove reliable system communication and control.

Phase 2 — Soft robot: keep the proven interface and replace the body with soft, quiet, companion-like materials and movement.

Materials

Part	Material	Why
Outer skin	Washable soft fabric	Warm, familiar, and easy to replace
Soft body	Silicone or TPE over light foam	Safe, light, and pleasant to touch
Safety frame	Small polymer or aluminum core	Protects motors, sensors, and battery

Simple movement

Movement	How	Feeling
Breathing	Slow body movement	Calm presence
Head turn	Slow two-axis neck	Attention and recognition
Touch	Gentle delayed response	Safe and warm contact

Low-load control

— **ROS 2 bridge:** connects the robot with the Animal Intelligence system.

— **Safety layer:** local rules limit heat, force, speed, and collision.

— **Motion library:** the Animal Agent chooses from safe, tested movements instead of raw motor commands.

COMPANION RULE

Use a round shape, low weight, quiet motors, slow motion, and no hard edges. The final soft robot should feel comfortable even when it is resting.

SECTION 05

Research Program & Evaluation

The blueprint should advance through evidence, not through scale alone.

Working hypotheses

Hypothesis	Observable signal	Counter-signal
Teaching improves learning	Better delayed recall and transfer by the human	Engagement without measurable understanding
Agent growth feels consequential	Users cite later behavior caused by prior teaching	Growth feels random or purely cosmetic
Embodiment increases legibility	Users predict actions more accurately	Animation creates false confidence
Diverse training produces useful specialization	Agents exchange complementary knowledge	Homogenization or brittle role-locking
Adaptation supports mastery	Difficulty remains challenging but achievable	Invisible simplification or frustration loops
Persistence creates attachment	Return is motivated by continuity and care	Dependence driven by guilt or manipulation

Evaluation families

- **Learning:** recall, conceptual transfer, explanation quality, misconception repair.
- **Agent development:** behavioral consistency, useful specialization, memory accuracy.
- **Relationship:** trust calibration, perceived responsibility, healthy attachment.
- **World quality:** event coherence, social safety, persistence, causal traceability.
- **Operations:** latency, compute per active agent, storage growth, incident rate.
- **Governance:** consent comprehension, deletion integrity, audit completeness.

INTERPRETATION RULE

Session length is not a sufficient success metric. A meaningful world should improve understanding or agency without requiring compulsive use.

SECTION 06

Constraints, Risks & Open Questions

The same properties that make persistent agents compelling also make them difficult to govern.

Risk	Why it matters	Research response
False learning	The agent may mirror language without robust understanding.	Use transfer tests and action-based evaluation.
Memory distortion	Summarization can rewrite identity or history.	Keep provenance, confidence, and correction trails.
Manipulative attachment	Persistent characters can exploit care and guilt.	Ban coercive retention patterns; test healthy disengagement.
Social contagion	Agents can spread errors or harmful behavior.	Bound transmission, model trust, and support rollback.
Runaway resource cost	Always-on agents and worlds may be unsustainable.	Set per-agent compute, storage, and simulation budgets.
Privacy leakage	Personal teaching can surface in shared interactions.	Separate memory scopes and enforce consent-aware retrieval.
Physical harm	Embodiment converts model errors into material consequences.	Require supervised control and hardware-level safety.

Open research questions

- What does it mean for an agent to learn from a user rather than merely store their statements?
- How can personality develop while remaining corrigible and safe?
- Which memories should be forgotten, compressed, shared, or made immutable?
- How should conflicts between different users' teachings be represented?
- Can emergent specialization remain flexible instead of becoming a permanent social hierarchy?
- What level of continuity is sufficient for a physical body to feel like the same agent?

SECTION 07

Long-Term Direction

Knowledge — Agent — World — Adaptation — Body

Animal Intelligence begins with something relatively simple: a human, an Animal Agent, and a conversation. Knowledge becomes quests. Quests help humans teach and raise their Animal Agents. Animal Agents gain virtual bodies and begin to inhabit a world.

Many independently raised Animal Agents begin to meet one another. AI adapts the difficulty of that world to what each learner knows, while interactions among differently developed Animal Agents may produce behaviors that no one explicitly designed.

Perhaps, eventually, an agent that has spent years developing in the virtual world could gain a physical body.

NOW	NEXT	LATER
Conversation Validate the teaching relationship.	**World** Connect memory, action, and social continuity.	**Body** Study continuity across digital and physical life.

BLUEPRINT COMMITMENT

We do not assume that every stage will work exactly as imagined, nor that every technology described here must be built at once. The purpose of this blueprint is to establish a direction—and then build, experiment, and test it one stage at a time.

ANIMAL INTELLIGENCE
Technology Blueprint · Version 1.0 · 2026 Fall
Public technology note for discussion and staged research planning.