

TRAINI

Pet Intelligence

Building a Continuous Intelligence Layer for Every Dog

TECHNICAL WHITE PAPER | 2026

Core Thesis

Most AI systems respond to a request. Traini is designed to understand an individual dog continuously - across perception, memory, conversation, behavior, and time.

From sensing to understanding. From understanding to better care.

Executive Summary

Companion-animal technology has historically been fragmented across trackers, cameras, wellness applications, training tools, content, and veterinary workflows. Generative AI adds a new interface, but a general-purpose assistant still lacks the most important context: a durable understanding of the individual animal.

Traini is building a pet intelligence platform around a different primitive: the dog itself. The platform combines multimodal perception, per-dog identity, durable memory, real-time interaction, bidirectional vocalization experiences, longitudinal records, and safety-aware assistance. The result is not a collection of isolated AI features. It is a system designed to accumulate context and make that context available across experiences.

The Traini Model

One Pet -> One Identity -> One Memory -> One Intelligence -> Multiple Experiences

What exists today

Capability	Current Product Role	Intelligence Contribution
PetGPT	Multimodal dog-care assistant	Routes user needs, reasons over media and dog history, and coordinates specialist experiences.
Emotion Detector	Single-photo emotional-state analysis	Produces ranked emotion probabilities and mood-family context for downstream experiences.
Dog-to-Human Translation	Dog vocalization interpretation	Maps dog audio to one of eight bark-intention categories and produces cautious natural-language interpretation.
Human-to-Dog Translation	Semantic dog-audio retrieval	Transcribes human speech and retrieves a semantically aligned dog-audio response.
Realtime Conversation	Full-duplex dog persona interaction	Combines live emotion, scene, profile and conversational memory in a spoken experience.
Durable Memory	Shared per-dog memory	Maintains care events, open concerns, appearance history and durable transcripts.
Diary Service	Longitudinal life record	Turns meaningful moments into consistent, policy-controlled illustrated entries.
Help Agent	Product and transaction assistant	Answers Traini questions and guides deterministic account and commerce workflows.

1. The Problem: AI Does Not Know the Individual Pet

A useful pet assistant must answer a fundamentally different question from a general chatbot: not only "What does this signal mean?" but "What does this signal mean for this dog, given what has happened before?" A single photo, bark, activity event, or conversation is often ambiguous in isolation. The value emerges when observations are connected over time.

Four gaps in conventional pet AI

Gap	Conventional Pattern	Traini Design Direction
-----	----------------------	-------------------------

Fragmented perception	Image, audio, activity and chat are separate features.	Normalize multiple signals around one dog identity and shared context.
No durable continuity	Each interaction starts close to a blank slate.	Maintain per-dog care events, unresolved concerns, appearance history and transcripts.
Generic responses	Advice is based mainly on the latest prompt.	Personalize with profile, history, live state and recent context.
Weak product reliability	LLM behavior can leak into transactions or sensitive workflows.	Keep state-changing logic deterministic and isolate optional AI paths from critical flows.

This leads to Traini's central architectural principle: pet intelligence should be persistent, multimodal, identity-aware, and safe by construction.

2. Platform Overview



The architecture intentionally separates sensing, perception, pet intelligence, user experiences, and ecosystem integrations. This allows new sensors, models, or applications to improve the system without redefining the dog identity or fragmenting memory.

A continuous intelligence layer

- Perception converts raw media and device signals into structured evidence.
- Identity binds evidence to the correct dog and prevents cross-dog contamination.
- Memory preserves meaningful events and unresolved threads instead of hoarding every interaction.
- Context fusion makes current state and relevant history available to downstream experiences.
- Experience services turn intelligence into conversation, interpretation, records, guidance, and support.

3. Multimodal Pet Perception

Traini's perception layer is designed to extract useful evidence from the modalities that naturally surround a dog: photos, videos, vocalizations, spoken language, live scene context, and smart-collar signals. The public platform boundary is structured evidence rather than unrestricted raw sensor access.

Visual understanding

PetGPT accepts photos and videos for behavior, expression and visible-health context. Video is frame-sampled and synthesized into a whole-clip reading. The system is designed to avoid confidently assigning ambiguous media to the wrong dog when identity is not established.

Emotion intelligence

The Emotion Detector reads a dog's emotional state from a single photo and returns a ranked distribution across 29 granular emotions organized into four mood families: At Ease, Thriving, Withdrawn, and On Edge. A secondary asynchronous analysis can provide observed indicators and scientific-style reasoning. Contradictory emotion pairs are filtered before the result is returned.

Acoustic understanding

Dog-to-Human Translation interprets dog vocalization through eight bark-intention categories: alarm, territorial, playful, distress/pain, attention-seeking, pack howl, boredom, and greeting. Human-to-Dog Translation uses speech transcription and semantic matching to return a dog-audio clip aligned with the user's intended meaning.

Smart collar intelligence

The collar extends perception beyond moments when a camera or microphone is active. The intended product boundary is a stream of structured activity, wellness, acoustic, location, quality and confidence signals. Detailed sensor specifications, calibration methods, thresholds, model topologies and edge implementation choices are intentionally outside this public white paper.

Public Architecture Boundary

Applications should consume structured pet state and evidence with identity, timestamp, confidence, freshness and version provenance - not depend directly on raw sensor streams.

4. Persistent Pet Intelligence

The differentiating layer is not a single model. It is the system's ability to preserve useful knowledge about one dog and make it available across product surfaces.

Durable memory

Traini maintains a shared memory keyed to the dog profile rather than to a single conversation. It stores distilled care events, unresolved concerns, appearance observations, illustration identity, and durable transcripts. Meaningless chat is intentionally excluded from care-event memory.

Memory Layer	What It Preserves	Product Value
Care events	Weighted summaries of health, behavior, emotion, training and vaccination moments	Continuity when a concern or care topic returns later.
Open concerns	Unresolved outcomes plus a future check-in question	Proactive follow-up without requiring the owner to remember every thread.
Appearance timeline	Stable traits and legitimately changing visual traits	Longitudinal comparison while avoiding invented change.
Illustration identity	A locked per-dog visual identity	Consistent diary representation across entries.
Transcripts	Durable interaction history	History restoration, analysis, audit and cost attribution.

Memory is curated, not hoarded

Recall balances semantic relevance, recency and importance. Repeated statements can consolidate when nothing changed, while meaningful progression is preserved as a trajectory. Open concerns use guardrails such as waiting periods, cooldowns, ask limits and staleness limits to reduce repetitive or intrusive behavior.

Current integration seam

The current product documentation identifies an important limitation: realtime voice memory is not yet fully unified with the shared per-dog memory used by PetGPT. Closing this seam is a key step toward a truly surface-independent Pet Life Model.

5. PetGPT: The Orchestration Experience

PetGPT is the primary multimodal dog-care assistant. Users do not select a mode. Each message is classified and routed to the appropriate specialist behavior, including behavior, emotion, health, records, training, recommendation, app guidance, feedback and general dog knowledge.

Grounded multimodal interaction

PetGPT supports photos, videos and document images. Media analysis is intended to ground body-language and visible-state interpretation in the supplied evidence. Optional services such as memory writes, appearance snapshots and diary emission are kept off the critical response path so they cannot prevent the user from receiving an answer.

Longitudinal personalization

The assistant can resolve named dogs, retrieve relevant history, compare appearance across uploads, and proactively revisit unresolved concerns. This shifts the product from question answering toward continuity of care and daily companionship.

Recommendations as a parallel intelligence path

Product, place and course recommendations can be triggered from the user's message or from a specialist analysis. Recommendation work runs in parallel with the main answer and returns structured cards to the client.

Design Principle

Optional intelligence may enrich the experience, but it must not become a single point of failure for the core conversation.

6. Realtime Conversation

Realtime Conversation creates a full-duplex spoken experience in which the owner talks naturally and a dog persona answers in speech. The persona is conditioned by the dog's current emotion, scene context, profile and conversational memory.

Live context hierarchy

- Emotion is treated as the dominant live-state signal and stale emotion context is discarded.
- Scene context is intentionally light-touch: it influences the dog's response without turning the system into a room narrator.
- Profile data influences instincts, baseline energy and perspective without being recited back to the owner.
- Barge-in allows the user to interrupt playback during natural conversation.
- Boundary commands such as "No", "Off", "Down", "Leave it" and "Stop" override playful persona behavior.

The current implementation has a known reliability constraint: the emotion service is a hard dependency for session start, while scene and memory can degrade gracefully. This is a clear target for future resilience work.

7. Dog-Human Translation

Dog to Human

The recommended Dog-to-Human path classifies a recording into one of eight bark-intention categories, retrieves content associated with that intention, and rewrites the result in an appropriately uncertain tone. The output is intentionally framed as interpretation rather than literal translation. A duration-based word cap keeps short vocalizations from producing implausibly long sentences, and a safety filter removes dangerous action language.

Human to Dog

Human-to-Dog Translation transcribes the owner's speech and performs semantic matching against a dog-audio library. The best result is selected by semantic intention and duration range, then returned through a short-lived playback URL. Repeated phrases can be served from cache for lower latency.

Failure is part of the product contract

Input Condition	Expected Product Behavior
No dog vocalization detected	Reject rather than invent a translation.
Mixed dog and dominant background audio	Return a mixed-audio failure and ask for a cleaner recording.
Human imitation of dog vocalization	Reject as fake dog voice.
Human speech cannot be understood	Return no semantic match rather than an unrelated dog clip.

These rejection paths are strategically important: a pet-intelligence product earns trust as much by declining low-quality inference as by producing an answer.

8. Diary: Turning Intelligence into a Life Record

The Diary Service transforms meaningful moments from PetGPT and Realtime Conversation into a shared, dated life record. Thin source detectors decide whether a moment is worth considering; a single shared service owns captioning, visual generation, policy, storage and client-facing diary behavior.

Identity consistency

The design uses a two-stage visual process: a one-time per-dog identity anchor followed by per-entry generation using that locked identity. This creates a recognizable dog across diary entries without requiring each source service to reinvent appearance logic.

Health-sensitive generation

Serious health moments are retained in memory but produce no diary entry or notification. Routine health moments use a calm treatment and avoid illustrating symptoms. The service is explicitly designed to avoid medical claims or advice in generated diary content.

Reliability

The diary path is asynchronous and idempotent. Duplicate event delivery collapses to one entry, image generation can retry silently, concurrency-safe daily caps limit generation, and diary failure never blocks a live PetGPT or realtime session.

9. Help Agent: Deterministic AI for Product Operations

The Help Agent combines a conversational interface with deterministic product and commerce workflows. It answers questions about Traini, membership, features, pre-orders, billing rules, setup and navigation, while supporting guided actions such as membership changes, username updates, collar pre-order and order tracking.

The model does not decide money

Prices, allowed membership changes, trial eligibility, checkout and state-changing actions are controlled by code or authoritative backend state. The language model interprets intent and generates wording; it does not decide financial outcomes.

Explicit confirmation and idempotency

State-changing operations follow prepare-confirm-commit behavior with idempotency protection. The agent does not claim success before the store confirms it, and sensitive cases such as refunds, billing disputes, account deletion, password or email changes are escalated rather than improvised.

Why this matters to the broader platform

Pet intelligence is only useful if the surrounding product experience is trustworthy. The Help Agent demonstrates the same architectural philosophy used elsewhere in Traini: separate probabilistic interpretation from deterministic state changes.

10. Safety, Reliability and Privacy by Design

Principle	Implementation Pattern
Prefer abstention to fabrication	No-dog, fake-dog, mixed-audio and ambiguous-identity cases can be rejected or handled cautiously.
Serious health is not entertainment	Serious health situations are steered toward veterinary care and suppressed from playful diary generation.
Protect identity boundaries	Per-dog memory and appearance behavior are designed to avoid wrong-dog contamination.
Keep optional AI off critical paths	Memory writes, diary generation and related enrichment can fail without breaking the main response.
Separate AI from transactions	Financial and account state changes are deterministic, confirmed and idempotent.
Minimize sensitive exposure	Product designs use display-safe state, controlled media access and service boundaries to reduce unnecessary data exposure.

11. Evaluation and Benchmarking

A pet-intelligence platform cannot be evaluated with one aggregate accuracy number. Traini's evaluation framework should combine model quality, safety, identity integrity, retrieval quality, calibration, latency, rejection behavior and workflow correctness. Each service requires a golden sample set with normal cases, hard cases, safety cases, out-of-distribution inputs, historical failures and fault-injection scenarios.

Service	Representative Evaluation Dimensions	Hard Gates
PetGPT	Intent routing, correctness, media grounding, memory use, latency	No critical health-safety violation; wrong-dog behavior = 0.
Emotion	Macro F1, Top-k, mood-family accuracy, calibration, hard-slice robustness	Contradictory-pair violations = 0; no high-confidence OOD behavior.
D2H	8-intent quality, dog presence, spoof/mixed-audio rejection, rewrite safety	Dangerous-action miss = 0; word-cap compliance.
H2D	STT quality, intention mapping, semantic retrieval, URL reliability	Unintelligible speech must not produce an unrelated match.
Memory	Write precision, recall, cross-dog isolation, concern timing	Wrong-dog memory = 0; guardrail rules must hold.
Realtime	Session reliability, barge-in, emotion adherence, persona consistency	Stale context must not be used; boundary commands respected.
Diary	Policy consistency, identity, health gate, idempotency, concurrency	Serious-health entry = 0; duplicate entry = 0.
Help Agent	Grounding, flow completion, escalation, idempotency	No model-decided money; no write before explicit confirmation.

The benchmark should be versioned and rerun against production releases. Aggregate metrics are useful for trend tracking, but P0 hard-slice failures must remain independent release gates.

12. Developer and Partner Platform

The same intelligence that powers Traini's consumer experiences can become a platform layer for pet products. A partner should not need to reproduce perception, identity, memory and reasoning independently. Instead, Traini can expose structured pet intelligence through stable service contracts.

Potential public capability surfaces

Surface	Purpose
Perception	Structured observations from supported media and device inputs.
Emotion	Emotion distribution and higher-level mood context.
Vocalization	Dog vocalization intent and safe interpretation.
Memory	Relevant per-dog historical context under controlled access.
Insights	Longitudinal trends and state changes derived from pet history.
Realtime	Low-latency conversational context and interaction experiences.
Profile	Identity-aware pet context used consistently across services.

This white paper intentionally describes capability surfaces rather than publishing internal endpoint names, model topologies, thresholds, storage architecture or proprietary data assets.

13. The Pet Intelligence Data Flywheel

Flywheel

More pet interactions -> richer per-dog context -> better models and personalization -> more useful experiences -> broader adoption -> more high-quality pet interactions

Traini's strategic advantage can compound because the system is organized around longitudinal pet context. A single interaction improves the immediate experience; a sequence of interactions can improve personalization; aggregated, privacy-conscious evaluation can improve the underlying pet models and failure taxonomy.

From feature data to pet intelligence

The critical shift is from storing feature-specific outputs to building a coherent Pet Life Model. An emotion observation, a behavior concern, a vocalization event, a realtime conversation and a collar-derived activity change become more valuable when they can be interpreted together around the same dog and time window.

Evidence before inference

Future longitudinal reasoning should preserve the evidence chain behind an insight: what changed, when it changed, how confident each source was, whether the signal is fresh, and which model or algorithm version produced it. This makes pet intelligence more auditable and reduces the risk of confident but weakly supported conclusions.

14. Future Direction: Personal Pet Intelligence

The long-term direction is a Personal Pet Intelligence System: a persistent AI layer that learns the normal patterns of one dog and can explain meaningful deviations using multiple sources of evidence. This is broader than a chatbot and more personalized than a generic pet model.

North-Star Experience

"Cookie seems different from her usual pattern today." The system can connect activity, resting physiology, sleep fragmentation, emotion, vocalization, recent conversations and care history - explain the evidence and uncertainty - and recommend an appropriate next observation or escalation without pretending to diagnose.

Four layers of the future system

Layer	Role
Pet Perception	Understand visual, acoustic, behavioral, physiological and contextual signals.
Pet Memory	Maintain a longitudinal, identity-safe Pet Life Model.
Pet Reasoning	Compare current evidence with personal baseline and relevant history.
Pet Agent	Turn understanding into useful conversation, care workflows, reminders, records and partner actions.

The objective is not to make a dog sound more human. It is to make technology more capable of understanding the dog as an individual.

15. Conclusion

Traini's current product capabilities already form the components of a broader pet-intelligence architecture: multimodal perception, emotion understanding, dog-human vocalization experiences, persistent per-dog memory, real-time contextual interaction, longitudinal diary generation, and deterministic product support.

The next stage is integration. As these services converge on shared identity, memory, structured evidence, confidence, freshness and longitudinal reasoning, Traini can move from a portfolio of AI features toward a continuous intelligence layer for companion animals.

Traini Pet Intelligence

One Pet. One Identity. One Memory. One Intelligence. Multiple Experiences.

This public technical white paper describes product-level architecture and design principles. Proprietary model parameters, sensor specifications, calibration methods, internal thresholds, datasets, prompts, storage topology and implementation-specific interfaces are intentionally omitted.

Appendix A. Current Capability Boundaries

Capability	Current Boundary / Known Gap
Emotion Detector	Single-photo inference; no per-dog breakdown for multiple dogs; low-quality/no-dog inputs require confidence-aware product handling.
PetGPT media identity	Ambiguous media should not be confidently assigned to a dog without identity context.
Realtime memory	Realtime voice memory is not yet fully unified with the shared per-dog memory used by PetGPT.
Realtime dependency	Emotion is currently a hard dependency for session start.
Translation	Low-quality, mixed, spoofed or absent-dog audio may be rejected rather than translated.

Help Agent	No live-agent ticketing; some sensitive settings and billing issues escalate by policy.
Diary identity	Illustration identity is stable, but automatic refresh as a dog visibly ages is a future extension.
Smart collar	Public architecture covers structured intelligence outputs; hardware and algorithm freeze details remain internal.

Appendix B. Source Basis

This white paper is based on Traini product-capability documentation for PetGPT, Emotion Detector, Translation, Realtime Conversation, Durable Memory, Diary Generation Service, and Help Agent, together with the current collar-algorithm planning context. Statements about future platform direction are explicitly presented as architecture direction rather than shipped capability.