

Animal Language Processing in 2026: Computational Foundations, Robotic Epistemology, and the Horizon of Interspecies Translation

The year 2026 represents a critical watershed in our scientific understanding of nonhuman communication. The convergence of deep learning foundation models, high-density bioacoustic telemetry, and computational ethology has established the field of nonhuman Animal Language Processing (ALP). Moving far beyond historical methods of manual audio classification, modern ALP researchers are mapping the acoustic, gestural, and contextual worlds of nonhuman species using unsupervised machine translation (UMT).

Despite these computational breakthroughs, a profound tension remains between semantic "classification" and genuine "translation". Deciphering the vocal signatures of aggression, naming, or alarm does not equate to achieving open-ended, two-way dialogic exchange. This report evaluates the state of ALP in 2026, analyzing the computational and robotic systems driving bioacoustic discovery, the mathematical formulations of unsupervised translation, and the profound epistemic, legal, and ethical paradigms emerging as the boundary between human and nonhuman intelligence begins to dissolve.

Computational Foundations and Taxonomic Progress

The shift from classic bioacoustics to ALP has been accelerated by the development of self-supervised foundation models designed specifically for nonhuman vocalizations. The Earth Species Project (ESP) has pioneered this structural paradigm, releasing NatureLM-audio, the first large-scale audio-language model tailored for bioacoustics, alongside AVES (and its avian counterpart BirdAVES), which functions as a self-supervised transformer-based acoustic representation encoder—effectively acting as a "BERT for animals".

To standardize performance metrics across highly diverse computational ethology tasks, the academic community relies on BEANS (Benchmark of Animal Sounds), which operates as the nonhuman equivalent to human natural language benchmarks.

Parallel to these generalist systems, highly targeted computational initiatives are mapping the communication networks of key species. Table 1 details the primary taxonomic initiatives, computational methodologies, and current milestones achieved across highly communicative species.

Table 1: Key Taxonomic Initiatives in Animal Language Processing (2024–2026)

Target Species	Lead Organization	Computational Architecture	Key Feature Extracted	Empirical Milestone (2025–2026)
Sperm Whale	Project CETI	Transformer	Structured click	Demonstration of a

Target Species	Lead Organization	Computational Architecture	Key Feature Extracted	Empirical Milestone (2025–2026)
(<i>Physeter macrocephalus</i>)		models, Graph-based Clustering, HMMs	sequences ("codas")	phonetic alphabet and human-like coda vowel phonology.
African Elephant (<i>Loxodonta africana</i>)	ElephantVoices / Save the Elephants	Random Forest Classifiers, Variational Autoencoders	Low-frequency rumbles	Verification of non-imitative "vocal labels" corresponding to individual names.
Egyptian Fruit Bat (<i>Rousettus aegyptiacus</i>)	Tel Aviv University (Yossi Yovel Lab)	Convolutional Neural Networks, Multi-class LSTMs	High-frequency aggressive squeaks	Real-time classification of aggressive calls into distinct social conflict categories with 88.5% accuracy.
Marmoset Monkey (<i>Callithrix jacchus</i>)	Hebrew University / Partner Institutions	Random Forest, Leave-one-session-out models	Spontaneous dialogue "phee-calls"	Discovery of vocal labeling of conspecifics and social vocal learning within families.
Bottlenose Dolphin (<i>Tursiops truncatus</i>)	Google DeepMind / Georgia Tech / Wild Dolphin Project	DolphinGemma Large Language Model	Signature whistles and click trains	Unsupervised segmentation of complex aquatic vocalizations at continental scale.

Sperm Whales and the Steganographic Hypothesis

Under the leadership of David Gruber, Project CETI has focused its efforts on decoding the communication structures of sperm whales off the coast of Dominica. Sperm whales communicate through sequences of rapid clicks known as "codas".

A study accepted to the International Conference on Learning Representations (ICLR 2026) by Project CETI Machine Learning Lead Michael Bronstein and Antonio Norelli reimaged cetacean communication through the lens of steganography—the art of concealing a message within a seemingly benign signal.

Using their developed Calgacus protocol, the team demonstrated that modest 8-billion-parameter open-source LLMs can successfully hide and extract highly coherent messages within text blocks of identical length. This mathematical insight is being applied to investigate whether sperm whale click variations contain hidden layers of information nested within broader social signals.

In April 2026, CETI researchers published findings in *Proceedings B* titled "*The Phonology of Sperm Whale Coda Vowels*," demonstrating that sperm whale codas possess a phonetic and vocalic architecture that represents one of the closest evolutionary parallels to human

phonology ever analyzed.

Elephants and Non-Imitative Naming

In collaboration with Save the Elephants and ElephantVoices (led by Scientific Director Joyce Poole), researchers have uncovered that elephants address one another using non-imitative vocal labels. Under the guidance of Mickey Pardo and George Wittemyer, a Random Forest machine learning classifier isolated distinct acoustic "signatures" embedded within long-distance contact calls.

Unlike bottlenose dolphins or certain parrot species, which address individuals by imitating the receiver's own signature whistle, elephants construct vocal labels that do not sound like the target. This indicates a capacity for abstract, symbolic thought, as the label operates independently of acoustic mimicry.

During the SMU Hackathon in April 2026, students and engineers developed "Project Tusk," an algorithm that successfully isolated low-frequency elephant rumbles while filtering out overlapping mechanical interferences from aircraft, generators, and land vehicles, thereby expanding the training datasets available to researchers.

Prairie Dogs, Bats, and Primate Dialogues

The granularity of ALP classification is further demonstrated in terrestrial mammals. Prior research led by Con Slobodchikoff revealed that prairie dog alarm calls vary systematically based on the size, shape, speed, and color of approaching predators, showcasing a high degree of semantic specificity.

At Tel Aviv University, Yossi Yovel's continuous monitoring of Egyptian fruit bat colonies has demonstrated that aggressive bat "screaming" is actually structured dialogue. By aligning 91,080 annotated vocal events with continuous video monitoring, his lab proved that bat vocalizations contain distinct features identifying both the calling individual (the emitter) and the targeted individual (the addressee).

Similarly, marmoset monkeys have been shown to use spontaneous "phee-call" dialogues to vocally label their conspecifics. Machine learning models demonstrate that family members utilize shared acoustic features, indicating that marmoset vocalizations serve as a critical evolutionary model for understanding human social language transmission.

The SETI and Exobiology Connection

The computational challenges of ALP have established deep parallels with the Search for Extraterrestrial Intelligence (SETI). Because both fields involve decoding completely unknown, nonhuman communication systems with zero access to bilingual parallel corpora, researchers are sharing methodologies.

This theoretical overlap was demonstrated in a collaborative study involving SETI scientists and marine biologists, who recorded a structured, 20-minute acoustic exchange with a humpback whale named "Twain". By utilizing back-and-forth call-and-response structures, the team established a primitive communication loop, demonstrating the potential for using aquatic mammals as a terrestrial analog for exobiological contact scenarios.

The Robotic and Sensor Revolution in Bioacoustic

Telemetry

The massive datasets required to train ALP foundation models are generated by a parallel revolution in soft robotics, autonomous vehicles, and non-invasive telemetry. Gathering continuous, high-fidelity acoustic and environmental data in extreme environments requires systems that minimize human disturbance. Table 2 details the key technological and robotic innovations developed between 2024 and 2026 to facilitate this data collection.

Table 2: Robotic and Hardware Innovations in Bioacoustic Data Gathering (2024–2026)

Technology	Developer / Instigator	Operational Mechanism	Data Captured / Function	Primary Resource / Reference
"Tap-and-Go" Drone Delivery	Project CETI / Harvard University	FPV racing drones deploy suction-cup sensors onto dorsal skin.	Minimizes tag deployment disturbance.	PLOS One (August 2025)
Passive Vacuum Regeneration	Harvard Robotics / Project CETI	Bioinspired suction cups powered by pressure changes during deep dives.	Extends sensor attachment times on marine mammals.	IEEE Oceans (September 2025)
Multimodal Bio-Logger	Harvard Robotics / Project CETI	Suction-mounted array with three synchronized hydrophones, motion, and environmental sensors.	High-bandwidth audio; maps angle-of-arrival to isolate individual whale voices.	Harvard SEAS Report (2025)
"Backseat Driver" Glider	Project CETI	SeaExplorer glider executing real-time acoustic source separation on a 4-element hydrophone array.	Passively tracks and follows individual sperm whales based on real-time vocalizations.	Scientific Reports (April 2026)
VAE Anomalous Filter	Save the Elephants	Variational Autoencoders deployed on ultra-low-power edge monitoring devices.	Filters redundant acoustic data; retains rare or novel vocalizations at 80-85% accuracy.	Save the Elephants Technical Paper (February 2026)


This hardware ecosystem has resolved several historical bottlenecks in bioacoustics. The "cocktail party problem"—the challenge of isolating individual voices within highly dense, overlapping animal groups—has been mitigated by both the Harvard-designed tri-hydrophone bio-logger and Earth Species Project’s custom source-separation algorithms. On terrestrial landscapes, the deployment of low-power, unsupervised edge devices running compressed VAE architectures allows researchers to monitor wildlife activity across vast

geographies without exhausting local memory storage or bandwidth, capturing critical behavioral changes and vocal trends in real time.

Information-Theoretic Limits of Unsupervised Machine Translation

The theoretical pathway to true interspecies translation relies on Unsupervised Machine Translation (UMT). In human translation, UMT operates by transforming two distinct languages into high-dimensional geometric shapes (manifolds) where words with similar semantic profiles occupy adjacent spatial coordinates.

Because human experiences of gravity, spatial relations, and emotional states are universal, the geometric shapes of human languages align, allowing direct coordinate-to-coordinate translation.



When translating between human and nonhuman communication systems, however, this spatial alignment breaks down due to divergent cognitive worlds, or *Umwelten*. To address this "domain gap," Shafi Goldwasser, David Gruber, Adam Tauman Kalai, and Orr Paradise formulated an information-theoretic framework for analyzing UMT between languages with highly disparate structural properties and no parallel bilingual data.

The Goldwasser-Gruber-Kalai-Paradise framework establishes that unsupervised translation is mathematically possible under two distinct conditions:

- **Prior Access:** The system has access to an informational "prior" distribution over the target language that can estimate the statistical likelihood that a candidate sentence is a valid translation of a source signal.
- **High-Probability Rejection:** Most incorrect or mutated candidate translation functions output sentences that are deemed highly implausible or "sensical nonsense" by the prior distribution, allowing the model to systematically eliminate invalid translation pathways.

The sample complexity—the volume of unlabeled data required to converge on an accurate translation function—is formally bounded by the complexity of the source language and the degree of shared physical or behavioral "common ground". Let $E(f)$ represent the translation error rate of a candidate translator f , C_L represent the structural complexity of the target language, and G represent the shared contextual ground.

To formalize the necessary sample complexity N required to achieve an error bound of δ with high probability, the framework establishes the following relationship:

where d represents the dimensional complexity of the latent language space (incorporating random sub-tree and knowledge-graph parameters), and κ represents the parameter of shared common ground (the overlap in semantic states or physical interactions).

This inequality demonstrates that if the shared common ground κ approaches zero—as it might when attempting to map human linguistic symbols onto the hyper-acoustic, pressure-sensitive sensory experience of a deep-diving cetacean—the required sample


complexity N approaches infinity.
Consequently, the mathematical feasibility of UMT in animal communication depends not on computing power, but on the physical, ecological, and cognitive overlap shared between the human and nonhuman subjects.


Consumer Commercialization and the Anthropomorphic Illusion

The rapid acceleration of academic ALP has generated a parallel consumer market focused on domestic pet care. By 2035, the AI in Pet Care Market is projected to reach \$15.34 billion, driven by the desire of pet owners to decode the emotional and physical states of their companion animals. This has led to the emergence of consumer devices that claim to offer direct, real-time translation of barks, meows, and body gestures. A prominent example is PettiChat, a \$118, 27-gram AI-powered collar launched on Kickstarter in April 2026 by the Chinese startup Meng Xiaoyi (translating to "cute little translator"). Backed by \$1 million in angel funding and securing over 10,000 pre-orders within weeks of its launch, the device claims to translate dog and cat vocalizations into complete human sentences with a real-time translation accuracy of 94.6%.

The Technological Reality of PettiChat

Developed in partnership with Zhejiang University's College of Animal Sciences, PettiChat utilizes a lightweight collar-mounted hardware array consisting of highly sensitive microphones and motion sensors. Rather than performing literal semantic translation, the device captures raw audio and physical motion data, processes it locally, and transmits the feature coordinates to Alibaba Cloud's Qwen large-language-model platform. The data is processed through a proprietary pet-translation classifier trained on five million voiceprint samples from over 1,000 animals. The algorithm operates on a multi-stage classification pipeline, as shown in the following schematic:





As illustrated, the system does not translate syntax; it classifies emotional states across more than 20 distinct categories (such as excitement, frustration, fear, or hunger) and uses a generative LLM to reconstruct these classifications into playful, human-readable sentences.

The Anthropomorphic Projection Feedback Loop

The claim of 94.6% accuracy represents a highly managed corporate metric. Independent bioacoustic validation indicates that the physical ceiling for classifying dog vocalizations from acoustic signals alone under laboratory conditions is approximately 57.3%.

Even when integrating multi-modal telemetry (video, facial expression tracking, and body orientation data, as attempted by startups like Zoolingua), the theoretical classification ceiling remains capped at 89%.

The discrepancy between this scientific baseline and the commercial success of devices like PettiChat is resolved by a psychological phenomenon known as anthropomorphic projection. Pet owners naturally project human-like intentions, thoughts, and complex emotional narratives onto minor behavioral variations (such as a head tilt, meow, or tail wag).

When a generative AI collar outputs a highly contextualized sentence (e.g., "I don't like when you work so late!") that aligns with the owner's pre-existing emotional narrative, the owner experiences the output as highly accurate, validating the device despite the absence of semantic proof.

While these tools are valuable for identifying broad health alerts, stress levels, or veterinary pain indicators, they represent creative, generative reconstructions of emotional states rather than literal language translation.

Epistemic, Legal, and Ethical Frontiers

As ALP technologies transition from experimental laboratories to the public sphere, they are challenging long-standing anthropocentric paradigms in environmental law, animal ethics, and conservation biology.

Taxonomic Bias and Selective Silencing

A 2025 systemic analysis of 14,600 conservation projects spanning a 25-year period revealed profound taxonomic bias in scientific funding: 29% of resources were allocated to species categorized as being of "least concern," whereas only 6% directly supported threatened species.

This funding disparity is mirrored in ALP development, where research is heavily concentrated on "charismatic megafauna" (whales, dolphins, elephants) or human companions (dogs, cats). This selective prioritization creates an ethical dilemma known as selective silencing. While

humanity allocates computational resources to decode the voices of species that do not threaten human economic interests, highly exploited agricultural species (such as pigs, cows, and chickens) are systematically ignored.

Ethicists argue that this dynamic is maintained because acknowledging the linguistically legible preferences, pain states, or active resistance of agricultural animals would threaten the global food industry.

Currently, humans rely on the "moral protection" of human ignorance: the ability to dismiss animal vocalizations of distress as mere instinctual, mechanical reactions rather than conscious suffering. By translating these signals into clear, undeniable communication, ALP removes this ignorance, forcing a direct confrontation with the ethics of confinement and industrial exploitation.

The Interspecies Veil and Epistemic Tiers

To analyze the limits of what humanity can legitimately understand through ALP, researchers classify animal cognitive experiences into three distinct epistemic tiers, illustrating how deep the interspecies veil penetrates:

- **Tier 1: Acute Affective States:** Immediate emotional expressions such as terror, excitement, or acute pain. These states are highly correlated with distinct physiological behaviors and are easily classifiable by modern machine learning models. However, translating these states does not offer deep insight into the animal's unique conceptual world.
- **Tier 2: Continuous Aversions and Preferences:** Expressed states of sustained desire, such as a captive animal communicating a consistent aversion to its cage. Mapping these preferences is technically viable but highly disruptive, as it directly challenges the legal justification for zoological confinement and pet ownership, shifting the human-animal relationship from "ownership" to a model of legal "guardianship".
- **Tier 3: Narrative and Relational Inner States:** Highly complex, internal cognitive processes, such as long-term grief, existential reflection, or abstract familial memories. Because these states lack immediate, acute behavioral markers, they remain opaque behind the interspecies veil, and current machine learning models lack the contextual data to decode or represent them.

Jurisprudential Transformations: The MOTH Paradigm

The scientific validation of nonhuman communication is actively transforming legal frameworks. Historically, Western law has treated nonhuman animals as property.

However, in October 2025, researchers from Project CETI and the NYU More-Than-Human Life (MOTH) Program published a paper in *Ecology Law Quarterly* titled "*What if We Understood What Animals Are Saying?: The Legal Impact of AI-Assisted Studies of Animal Communication*". The authors argue that empirical proof of phonetic alphabets, non-imitative naming, and social learning undermines the legal justification for treating animals as property, strengthening arguments for:

- **Legal Personhood:** Granting intelligent nonhumans the legal status of persons, affording them representation and the standing to sue in court. This builds on international precedents, such as Ecuador's 2008 constitutional recognition of the Rights of Nature, the landmark Los Cedros ruling, and Argentina's habeas corpus case that released Cecilia the chimp from a zoo.

- **Right to Culture:** Legally protecting the shared social structures, parenting protocols, and distinct communication dialects of animal communities from human disruption.
- **Acoustic Sanctuary:** Regulating undersea noise pollution, industrial shipping lanes, and military sonar as direct violations of marine mammals' right to communicate and navigate.

The PEPP Framework: Ethical Standards for NACTs

To address the regulatory vacuum surrounding Nonhuman Animal Communication Technologies (NACTs), the NYU MOTH Program released the **PEPP Framework** in November 2025. The framework outlines twelve principles organized across four core pillars, designed for voluntary adoption by researchers, engineers, and funders. Table 3 details the architecture of this ethical framework.

Table 3: The PEPP Framework for NACT Governance (November 2025)

Pillar	Principle	Core Obligation and Ethical Standard	Operational Impact
Prepare	Research Design & 3Rs	Adhere to rigorous experimental design, prioritizing the reduction, replacement, and refinement of animal subjects.	Forces institutional review before NACT deployment.
	Transparency & Governance	Openly publish ethics protocols, share data risk assessments, and report all unintended harms or model shortcomings.	Minimizes proprietary, closed-source animal testing.
Engage	Stakeholder Inclusion	Consult with independent scientific experts and traditional Indigenous knowledge-holders sharing lands/waters with subjects.	Integrates Western ethology with ancestral ecological knowledge.
	Non-Coercion	Respect animal autonomy; facilitate only non-coerced participation and disengage if signs of avoidance appear.	Eliminates invasive restraint during acoustic recording.
Prevent	Precautionary Principle	Cease any planned NACT activity if there are threats of serious, irreversible harm to individual animals or	Shifts burden of proof of safety onto technology developers.

Pillar	Principle	Core Obligation and Ethical Standard	Operational Impact
		ecosystems.	
	Misuse Mitigation	Anticipate and block potential commercial, military, or recreational exploitation of synthesized animal signals.	Limits the development of weaponized or poaching-related bioacoustics.
	Protect	Best Interests of the Animal	Let physical, mental, relational, and environmental best interests guide every computational design decision.
	Remediation of Harms	Urgently address and remediate any direct or indirect ecological harms caused by NACT deployment.	Establishes legal accountability for computational ethologists.

Ethicists warn that the development of generative AI capable of "talking back" represents a profound ecological danger. Emitting synthetic, AI-generated calls into the wild could cause severe cognitive distress, disrupt natural mating patterns, or fracture delicate social bonds within wild populations.

Furthermore, if these generative tools are commercialized, bad actors could easily exploit them. Poachers could use highly optimized, synthesized contact calls to lure endangered species, and commercial fishing fleets could deploy simulated dolphin signals to herd marine life into nets, transforming tools designed for conservation into instruments of highly efficient exploitation.

Diagnostic Synthesis: The Horizon of Interspecies Translation

To evaluate how close we are to being able to "discuss" with animals using AI translations in 2026, the scientific consensus requires a clear, bipartite distinction:

Structural Decoding and Intentional Classification: Extremely Close

If "discussion" is defined as the automated, high-fidelity classification of emotional states, context-specific intentions, and individual identities (such as detecting whether a bat is arguing over food, or recognizing which elephant is being addressed by name), **humanity is remarkably close**. State-of-the-art models like NatureLM-audio and DolphinGemma have successfully mapped the baseline structural units and acoustic alphabets of several species, transforming bioacoustics from passive listening into systematic decoding.

Reciprocal, Open-Ended Semantic Dialogue: Profoundly Distant

If "discussion" is defined as a mutual, open-ended exchange of abstract ideas—a literal "Google Translate" that converts a whale's inner thoughts into English sentences—**humanity remains far from this goal**. The primary bottlenecks are not merely computational, but are fundamentally limited by data availability and the vast divergence of species' *Umwelten* :

- **The Data Asymmetry:** Modern human LLMs are trained on trillions of words of written text. Nonhuman animal datasets, even the most extensive, are many orders of magnitude smaller. Because animal communication is continuous, lacks a written form, and is deeply embedded in physical behavior, gathering the necessary scale of multimodal data represents a slow, labor-intensive process that cannot match the scaling velocity of human AI.
- **The Geometry of Subjective Experience:** Mathematically, unsupervised machine translation requires that the semantic structures of the two languages share a congruent geometric shape. However, because a sperm whale's or a bat's sensory experience is built on physical inputs completely alien to human biology (such as echolocation or ultrasonic tracking), their cognitive manifolds may share very little "common ground" with human language. If the geometric shapes of our respective cognitive worlds do not align, unsupervised translation models will fail to find a coherent mathematical mapping function, regardless of computational power.

The future of interspecies communication lies not in forcing animals to speak human languages through creative AI generation, but in using computational tools to appreciate their unique cognitive structures, establishing rigorous ethical standards to protect their habitats, and allowing their decoded voices to reshape environmental law and human ecological responsibility.

Works cited

1. THE ORGANIZATION USING ARTIFICIAL INTELLIGENCE TO UNLOCK COMMUNICATION WITH THE ANIMAL KINGDOM - Parley, <https://parley.tv/journal/earth-species> 2. Masato Hagiwara's Page, <https://masatohagiwara.net/> 3. Earth Species Project – More than 8 million species share our planet ..., <https://earthspecies.org/> 4. A Theory of Unsupervised Translation Motivated by Understanding Animal Communication, https://proceedings.neurips.cc/paper_files/paper/2023/hash/7571c9d44179c7988178593c5b62a9b6-Abstract-Conference.html 5. AI-assisted animal language translation - Wild Animal Initiative, <https://www.wildanimalinitiative.org/blog/ai-animal-translation> 6. Decoding communication of non-human species - Diva-Portal.org, <http://www.diva-portal.org/smash/get/diva2:1811459/FULLTEXT01.pdf> 7. Are We Ready to Hear a Non-Human 'No'? — EA Forum, <https://forum.effectivealtruism.org/posts/G7zvBZoJsW8KnNY5T/are-we-ready-to-hear-a-non-human-no> 8. A four-month-old Chinese startup just launched a \$118 AI collar that claims to translate dog and cat vocalizations into human sentences with 95% accuracy — an extraordinary consumer device that has secured \$1 million in funding despite zero independent scientific proof that it actually works - Make Tech Easier, <https://maketechesasier.com/a-four-month-old-chinese-startup-just-launched-a-118-ai-collar-that-claims-to-translate-dog-and-cat-vocalizations-into-human-sentences-with-95-accuracy-an-extraordinary-consumer-device-that/> 9. Bridging Species: Legal and Ethical Implications of AI-Assisted Animal Communication, https://www.americanbar.org/groups/tort_trial_insurance_practice/resources/brief/2025-winter/legal-ethical-implications-ai-assisted-animal-communication/ 10. AI-enabled decoding of whale

communication could bolster animal ..., <https://www.law.nyu.edu/news/ideas/cesar-rodriguez-garavito-animal-communication-ai> 11. What are animals saying? 3 AI tools that could soon tell us their thoughts, <https://m.economictimes.com/magazines/panache/what-are-animals-saying-3-ai-tools-that-could-soon-tell-us-their-thoughts/articleshow/122180814.cms> 12. Project CETI •-- Research, <https://www.projectceti.org/research/index> 13. Project CETI •-- News & Insights, <https://www.projectceti.org/news-insights> 14. Like Us, Elephants Have Names Too! - ElephantVoices, <https://elephantvoices.org/latest-news-updates/enewsletters/like-us-elephants-have-names-too> 15. iMasons, Grundfos, ElephantVoices and the Open Compute Project Deliver Real-World Impact at SMU Hackathon – Infrastructure Masons, <https://imasons.org/blog/imasons-grundfos-elephantvoices-and-the-open-compute-project-deliver-real-world-impact-at-smu-hackathon/> 16. How AI Could Help Us Understand Animal Language - Devstyler.io, <https://devstyler.io/blog/2025/01/02/how-ai-could-help-us-understand-animal-language/> 17. Vocal labeling of others by nonhuman primates | Request PDF - ResearchGate, https://www.researchgate.net/publication/383562551_Vocal_labeling_of_others_by_nonhuman_primates 18. April 28 - 30 | Santa Fe Institute, <https://www.santafe.edu/info/interspecies/about> 19. Researchers translate bat language. Turns out, they argue a lot - ZME Science, <https://www.zmescience.com/ecology/animals-ecology/fruit-bat-language-02012017/> 20. Can AI build 'Dr. Dolittle machine' to communicate with animals? - IndiaAI, <https://indiaai.gov.in/article/can-ai-build-dr-dolittle-machine-to-communicate-with-animals> 21. (PDF) AI-Driven Communication Analysis for Biodiversity Conservation: A Fruit Bat Prototype, https://www.researchgate.net/publication/397208574_AI-Driven_Communication_Analysis_for_Biodiversity_Conservation_A_Fruit_Bat_Prototype 22. Animal-Borne Adaptive Acoustic Monitoring - Save the Elephants, <https://savetheelephants.org/wp-content/uploads/2026/02/jsan-14-00066-v3.pdf> 23. A Theory of Unsupervised Translation for Understanding Animal Communication, <https://neurips.cc/virtual/2022/59460> 24. Decoding Your Pet's Behavior: How AI Is Revolutionizing Pet Communication, <https://samscatsanddogs.com/blog/decoding-your-pets-behavior-how-ai-is-revolutionizing-pet-communication/> 25. China's AI pet translator sells a dream that animal science cannot yet prove, <https://webiano.digital/chinas-ai-pet-translator-sells-a-dream-that-animal-science-cannot-yet-prove/> 26. AI Pet Translator From China Claims To Understand Pets With Almost 94% Accuracy - 8days, <https://www.8days.sg/entertainment/asian/ai-pet-translator-understand-pets-94-cent-accuracy-858441> 27. Legal and Ethical Principles for Nonhuman Animal Communication Technologies - Moth, https://mothlife.org/wp-content/uploads/2025/11/Listening-to-Our-Animal-Kin_MOTH-Report.pdf 28. Google Translate is getting closer to understanding animal language - Talk Android, <https://www.talkandroid.com/511994-google-translate-is-getting-closer-to-understanding-animal-language/> 29. The PEPP Framework: Ethical Standards for Animal Communication Technologies (NACTs), <https://www.ouranimalkin.com/pepp-framework> 30. PEPP Framework: Ethical Principles for Animal Communication Technology (NACTs), <https://www.ouranimalkin.com/> 31. Listening Is Not Neutral: New Ethical Guardrails for Animal Communication Technologies, <https://bioneers.org/listening-is-not-neutral-new-ethical-guardrails-for-animal-communication-technologies-zmaz2512/> 32. AI for decoding non-human languages - Gemma's Open Research,

<https://open-research.gemmadanks.com/notes/ai-for-decoding-non-human-languages/>