Proposal: Deciphering Electrocommunication with MARL and Unsupervised Machine Translation

Satpreet H. Singh Harvard University
Sonja Johnson-Yu Harvard University
Zhouyang Lu Brown University
Aaron Walsman Harvard University

Federico Pedraja Columbia University
Denis Turcu Columbia University

Pratyusha Sharma MIT
Naomi Saphra Harvard University
Nathaniel B. Sawtell Columbia University
Kanaka Rajan Harvard University

Abstract

Unsupervised machine translation (UMT) has recently been proposed as a tool for deciphering animal communication. Previous efforts, however, have attempted to align animal signals directly with human language, introducing large ecological and representational gaps that inevitably limit success. We argue that a more promising path is to generate synthetic corpora through *rich, situated, biologically realistic* multi-agent reinforcement learning (MARL). Such simulations yield emergent communication signals that share statistical and functional properties with real animal data, thereby narrowing the gap that hampers translation. As a case study, we present MARL agents inspired by pulse-type weakly electric fish (WEF), which rely on electric organ discharges (EODs) for both sensing and social communication. WEF provide an ideal test case because their communication signals are tightly coupled to collective behaviors such as foraging, resource sharing, and dominance interactions. Our MARL agents reproduce key features of real WEF behavior and communication, including socially aware foraging strategies, heavy-tailed EOD interval distributions, and context-dependent shifts in EOD rate. These synthetic corpora can be generated at scale, with complete access to both neural and behavioral variables, and allow for mechanistic interpretation and virtual interventions that are expensive or infeasible in vivo. We propose a methodology to combine the MARL-generated emergent communication with UMT techniques to decipher real fish EOD data. This integration opens a path toward AI-assisted deciphering of animal communication, with WEF as a proving ground and strong potential for extension to other species.

1 Introduction

Unsupervised machine translation (UMT) has recently emerged as a tool for deciphering animal communication [1, 2]. Most existing proposals attempt to align animal signals directly with human language [3]. This direct alignment introduces ecological and representational gaps that limit success, since the latent structure of human language differs from that of animal signals [4, 5].

We propose a different strategy. Rather than forcing alignment to human words, we first generate synthetic corpora using *emergent communication* arising among agents trained with multi-agent reinforcement learning (MARL) in *rich, situated, biologically realistic simulations*. When placed in ecologically grounded environments, with realistic constraints on motion, sensing, and energetic

Accepted to the AI for Non-Human Animal Communication workshop (NeurIPS 2025). Do not distribute.

cost, agents generate emergent communication signals that more closely resemble animal signals than human text.

Here we present a case study focused on pulse-type weakly electric fish (WEF) [6, 7, 8, 9]. These fish use electric organ discharges (EODs) for both active sensing and social communication, signals that are central to collective behaviors such as foraging, resource sharing, and dominance [10, 11]. By combining MARL with UMT, we aim to bridge simulation and biology, enabling structured interpretation of animal signals in their ecological context. In the remainder of this proposal, we outline a research design and UMT-based translation framework, present early results from MARL simulations of WEF communication, and discuss key considerations for integrating MARL with UMT.

2 Unsupervised Machine Translation (UMT)

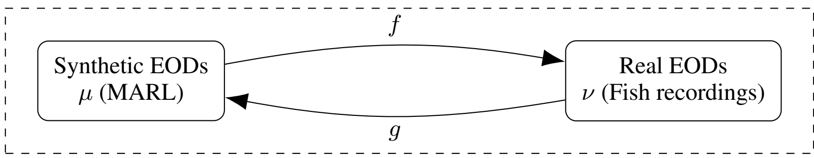

Cycle consistency + Priors ρ_X, ρ_Y constrain translators

Figure 1: Two-domain UMT framework. Synthetic EODs generated by MARL (μ) and real EODs recorded from weakly electric fish (ν) serve as two monolingual corpora. Translators f and g are trained with priors for fluency and cycle-consistency for information preservation.

Unsupervised machine translation (UMT) seeks to learn a mapping between two monolingual corpora in the absence of parallel pairs. In our case, one corpus is composed of synthetic electro-communication signals generated by MARL agents, while the other consists of real EOD recordings from weakly electric fish. We denote the synthetic distribution by μ over alphabet Σ_X and the real distribution by ν over alphabet Σ_Y. The goal is to find a translator $f : \Sigma_X^\star \to \Sigma_Y^\star$ that maps synthetic sequences to real ones (and vice versa via g), such that the translation preserves semantics of ecological context.

The canonical UMT formulation balances three ingredients: *(i) Language priors*, ensuring that translations resemble valid samples in the target domain; *(ii) Cycle consistency*, ensuring that $x \to f(x) \to g(f(x)) \approx x$ and $y \to g(y) \to f(g(y)) \approx y$; and *(iii) Optional denoising/back-translation*, where pseudo-parallel pairs are generated by translating monolingual examples and training conditional models on them [12, 13].

Formally, the two-domain UMT objective is

$$\max_{f,g} \; \mathbb{E}_{x\sim\mu}[\log \rho_Y(f(x))] + \mathbb{E}_{y\sim\nu}[\log \rho_X(g(y))] - \lambda_{\text{cyc}}\, \mathcal{L}_{\text{cyc}}(f,g), \tag{1}$$

$$\mathcal{L}_{\text{cyc}}(f,g) = \mathbb{E}_{x\sim\mu} d_X\big(g(f(x)), x\big) \; + \; \mathbb{E}_{y\sim\nu} d_Y\big(f(g(y)), y\big). \tag{2}$$

where ρ_X, ρ_Y are language-model priors trained on synthetic and real corpora, and $d_X : \mathcal{X} \times \mathcal{X} \to \mathbb{R}_{\geq 0}$ and $d_Y : \mathcal{Y} \times \mathcal{Y} \to \mathbb{R}_{\geq 0}$ are distances (or divergences) defined on the synthetic EOD sequence space $\mathcal{X}$ and real EOD sequence space $\mathcal{Y}$ respectively.

The objective in Eq. 1 can be seen as encouraging three things simultaneously: translations should be *fluent* in the target corpus, *information-preserving* under round-trip translation, and *robust* to noise or mismatch via back-translation. Eq. 2 measures how close a round-trip translation is to the original sequence. Similar formulations have been shown to succeed in both word-level alignment and sentence-level UMT [12, 13].

3 MARL for Weakly Electric Fish

Weakly electric fish are an ideal case study for our proposal. They emit pulsatile EODs that are tightly coupled to ecological and social behavior. EODs serve multiple roles, from electrolocation

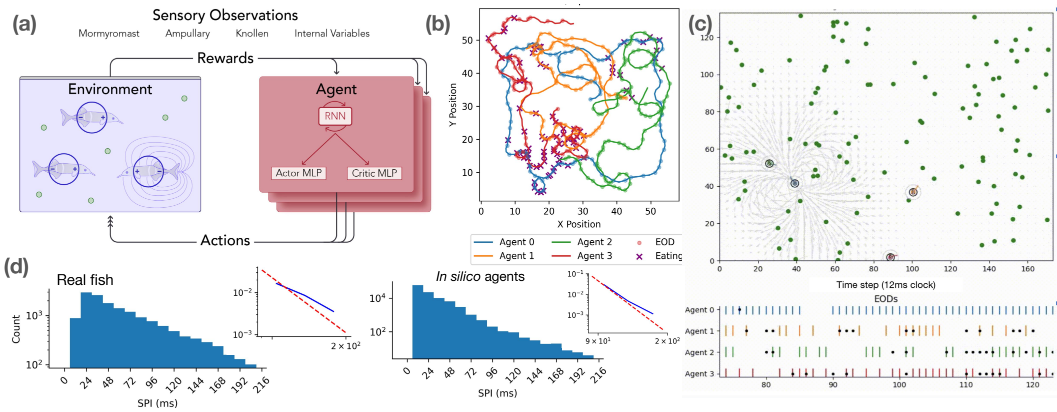

Figure 2: Overview of our MARL framework for modeling weakly electric fish communication. **(a)** Schematic of the training loop, where agents interact with a simulated arena, emitting and sensing electric organ discharges (EODs) through weakly electric fish-inspired sensors. Rewards encourage successful foraging and penalize aggressive encounters. **(b)** Example trajectories from four agents in a single foraging episode, showing exploration and food acquisition. **(c)** Snapshot of the arena (top) showing agents, food sources, and simulated electric fields; bottom shows temporally-structured EOD spike trains across individual agents. **(d)** Sequential Pulse Interval (SPI) distributions from real fish (left) and MARL-trained agents (right), showing that in silico agents reproduce the heavy-tailed statistics observed in biological data. Insets show log-linear curves compared to empirical curve fits.

to communication with conspecifics, and have been extensively studied in both ethological and neurophysiological contexts [14, 15, 6]. The frequency, interval distribution, and context-dependent motifs of EODs vary with foraging, resource sharing, and dominance interactions. Decades of work have shown how temporal patterning in EODs encodes social and ecological information [16, 17, 18]. More recently, studies have revealed collective sensing and group-level coordination in electric fish [11].

Our MARL environment instantiates these ecological and biological priors (Figure 2). Agents are recurrent actor–critic networks that receive egocentric observations from simulated electrosensory modalities and generate actions including movement, turning, EOD emission, and biting [7]. EODs induce fields that interact with food objects, walls, and other agents, which are then sensed by different receptor types on each agent's body. Agents are trained using Multi-Agent Proximal Policy Optimization [19, 20, 21] with rewards that encourage successful foraging and provide asymmetric penalties during aggressive encounters between fish of different dominance levels. Trained MARL agents develop socially aware foraging strategies, dominance displays, and context-dependent EOD modulation. The resulting synthetic corpora reproduce power-law–like SPI distributions and other temporal motifs observed in real fish, while providing complete access to neural, behavioral, and environmental variables. Full observability and controllability of the artificial neural network enables mechanistic interpretation, *in silico* ablations and inter-agent comparisons [22, 23, 24, 25, 26, 27, 28].

4 Key Considerations

A key theoretical insight is that UMT performance depends on two properties of the signal distributions [1]. *Complexity* (d) captures the richness of statistical structure: context-dependent signals restrict the set of valid translators. *Common ground* (α) quantifies structural overlap between synthetic and real signals. Translation error decreases with both d and α: if signals are too simple or the domains share little overlap, degenerate translators can satisfy the UMT objective without preserving semantics.

Our design follows directly from the above. Generate MARL corpora under ecologically rich conditions (to raise d), and embed in the simulation constraints known from fish ethology so that synthetic motifs share structure with real EODs (to increase α). By doing so, the UMT alignment problem becomes not only feasible but biologically interpretable.

Several research questions guide our proposal: How do emergent motifs in MARL align with those observed in real fish? Can UMT trained on MARL corpora decipher real EOD recordings into structured behavioral descriptors? How does task context, e.g. competition versus cooperation, or group size, affect translatability?

5 Discussion

Traditional UMT approaches have struggled with the mismatch between animal signals and human text [3]. We propose bridging this gap with *rich, situated MARL*, using weakly electric fish as a case study. Our simulations generate communication corpora closer to real biology, enabling UMT to map signals into structured, situated descriptors.

The approach offers several contributions. First, it demonstrates how MARL can create interpretable corpora that reproduce known features of fish behavior and communication while allowing controlled virtual interventions. Second, it outlines a principled UMT framework to align synthetic and real signals into structured descriptors of ecological context. Third, it contributes datasets, both simulated and empirical, that can serve as benchmarks for future work on AI-assisted animal communication [29, 30].

The implications extend beyond weakly electric fish. Any species with structured signals embedded in social contexts, such as birdsong, primate calls, dolphin whistles, or sperm whale clicks, could be studied with this framework [31, 32, 33, 34, 35, 36]. By embedding ecological realism into MARL simulations, we generate corpora that are both interpretable and aligned with biology. Combined with UMT, this approach opens a general path toward AI-assisted translation of animal communication [37]. While our case study focuses on EODs from weakly electric fish, the same MARL–UMT framework can generalize to other modalities of animal communication.

Ethical considerations

Our framework reduces invasive experimentation on live animals by generating large synthetic corpora through simulation. By grounding AI models in ecological realism, we also reduce risks of anthropomorphizing or over-interpreting signals, instead situating interpretation in the behavioral context of the species. All experimental data used in this paper were collected by collaborators for previous neuroscientific studies of weakly electric fish.

Acknowledgements:

We thank Yonatan Belinkov, Eugene Vinitsky, Roy Harpaz, Daphne Cornelisse, Wilka Carvalho, Thomas Fel, and members of the Rajan, Sawtell and Gershman labs for helpful discussions. Funded by NIH (RF1DA056403), James S. McDonnell Foundation (220020466), Simons Foundation (Pilot Extension-00003332-02), McKnight Endowment Fund, CIFAR Azrieli Global Scholar Program, NSF (2046583), Harvard Medical School Dean's Innovation Award, Harvard Medical School Neurobiology Lefler Small Grant Award, Alice and Joseph Brooks Fund Postdoctoral Fellowship (S.H.S.), Shanahan Family Foundation Fellowship at the Interface of Data and Neuroscience at the Allen Institute and University of Washington, supported in part by the Allen Institute (D.T.).

References

[1] Shafi Goldwasser, David Gruber, Adam Tauman Kalai, and Orr Paradise. A theory of unsupervised translation motivated by understanding animal communication. *Advances in Neural Information Processing Systems*, 36:37286–37320, 2023.

[2] Ido Levy, Orr Paradise, Boaz Carmeli, Ron Meir, Shafi Goldwasser, and Yonatan Belinkov. Unsupervised translation of emergent communication. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 39, pages 23231–23239, 2025.

[3] Emily Anthes. The animal translators. *The New York Times*, Aug 2022.

[4] Mélissa Berthet, Camille Coye, Guillaume Dezecache, and Jeremy Kuhn. Animal linguistics: a primer. *Biological Reviews*, n/a(n/a), 2022.

[5] W Fitch. The evolution of language: a comparative review. *Biology and philosophy*, 20(2):193–203, 2005.

[6] Nathaniel B Sawtell, Alan Williams, and Curtis C Bell. From sparks to spikes: information processing in the electrosensory systems of fish. *Current opinion in neurobiology*, 15(4):437–443, 2005.

[7] Sonja Johnson-Yu, Satpreet Harcharan Singh, Federico Pedraja, Denis Turcu, Pratyusha Sharma, Naomi Saphra, Nathaniel Sawtell, and Kanaka Rajan. Understanding biological active sensing behaviors by interpreting learned artificial agent policies. In *Workshop on Interpretable Policies in Reinforcement Learning at RLC-2024*, 2024.

[8] Satpreet Harcharan Singh, Sonja Johnson-Yu, Zhouyang Lu, Aaron Walsman, Federico Pedraja, Denis Turcu, Pratyusha Sharma, Naomi Saphra, Nathaniel Sawtell, and Kanaka Rajan. Understanding electro-communication and electro-sensing in weakly electric fish using multi-agent deep reinforcement learning. In *The Thirty-Ninth Annual Conference on Neural Information Processing Systems workshop: AI for non-human animal communication*, 2025.

[9] Kaden Zheng, Sonja Johnson-Yu, Satpreet Harcharan Singh, Denis Turcu, Federico Pedraja, Pratyusha Sharma, Naomi Saphra, Nathaniel Sawtell, and Kanaka Rajan. Keypoint annotation for electrocommunication source separation with pikachu and raichu. In *The Thirty-Ninth Annual Conference on Neural Information Processing Systems workshop: AI for non-human animal communication*, 2025.

[10] Jan Benda. The Physics of Electrosensory Worlds. In *The Senses: A Comprehensive Reference*, pages 228–254. Elsevier, 2020.

[11] Federico Pedraja and Nathaniel B. Sawtell. Collective sensing in electric fish. *Nature*, 628(8006):139–144, April 2024.

[12] Guillaume Lample, Alexis Conneau, Ludovic Denoyer, and Marc'Aurelio Ranzato. Unsupervised machine translation using monolingual corpora only. In *6th International Conference on Learning Representations, ICLR 2018, Vancouver, BC, Canada, April 30 - May 3, 2018, Conference Track Proceedings*. OpenReview.net, 2018.

[13] Mikel Artetxe, Gorka Labaka, and Eneko Agirre. A robust self-learning method for fully unsupervised cross-lingual mappings of word embeddings. In *Proceedings of the 56th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 789–798, Melbourne, Australia, July 2018. Association for Computational Linguistics.

[14] Gerhard Von der Emde. Active electrolocation of objects in weakly electric fish. *Journal of experimental biology*, 202(10):1205–1215, 1999.

[15] Te K. Jones, Kathryne M. Allen, and Cynthia F. Moss. Communication with self, friends and foes in active-sensing animals. *Journal of Experimental Biology*, 224(22):jeb242637, November 2021.

[16] Bruce A. Carlson and Carl D. Hopkins. Stereotyped temporal patterns in electrical communication. *Animal Behaviour*, 68(4):867–878, October 2004.

[17] Matthew E Arnegard and Bruce A Carlson. Electric organ discharge patterns during group hunting by a mormyrid fish. *Proceedings of the Royal Society B: Biological Sciences*, 272(1570):1305–1314, July 2005.

[18] Angel Ariel Caputi. The electric organ discharge of pulse gymnotiforms: the transformation of a simple impulse into a complex spatio-temporal electromotor pattern. *Journal of Experimental Biology*, 202(10):1229–1241, May 1999.

[19] John Schulman, Filip Wolski, Prafulla Dhariwal, Alec Radford, and Oleg Klimov. Proximal policy optimization algorithms. *arXiv preprint arXiv:1707.06347*, 2017.

[20] Tianwei Ni, Benjamin Eysenbach, and Ruslan Salakhutdinov. Recurrent model-free rl is a strong baseline for many POMDPs. *arXiv preprint arXiv:2110.05038*, 2021.

[21] Chao Yu, Akash Velu, Eugene Vinitsky, Jiaxuan Gao, Yu Wang, Alexandre Bayen, and Yi Wu. The surprising effectiveness of PPO in cooperative multi-agent games. *Advances in Neural Information Processing Systems*, 35:24611–24624, 2022.

[22] Chenguang Li, Gabriel Kreiman, and Sharad Ramanathan. Discovering neural policies to drive behaviour by integrating deep reinforcement learning agents with biological neural networks. *Nature Machine Intelligence*, 6(6):726–738, 2024.

[23] Satpreet H Singh, Floris van Breugel, Rajesh PN Rao, and Bingni W Brunton. Emergent behaviour and neural dynamics in artificial agents tracking odour plumes. *Nature Machine Intelligence*, 5(1):58–70, 2023.

[24] Satpreet Harcharan Singh. Neuroprospecting with DeepRL agents. NeurIPS 2021 Workshop on AI for Science, 2021.

[25] Ann Huang, Satpreet Harcharan Singh, and Kanaka Rajan. Learning dynamics and the geometry of neural dynamics in recurrent neural controllers. In *Workshop on Interpretable Policies in Reinforcement Learning RLC-2024*, 2024.

[26] Raaghav Malik, Satpreet H Singh, Sonja Johnson-Yu, Nathan Wu, Roy Harpaz, Florian Engert, and Kanaka Rajan. Dissecting larval zebrafish hunting using deep reinforcement learning trained rnn agents. *arXiv preprint arXiv:2510.03699*, 2025.

[27] Ann Huang, Mitchell Ostrow, Satpreet H Singh, Leo Kozachkov, Ila Fiete, and Kanaka Rajan. Inputdsa: Demixing then comparing recurrent and externally driven dynamics. *arXiv preprint arXiv:2510.25943*, 2025.

[28] Ann Huang, Satpreet H Singh, Flavio Martinelli, and Kanaka Rajan. Measuring and controlling solution degeneracy across task-trained recurrent neural networks. *ArXiv*, pages arXiv–2410, 2025.

[29] Peter C Bermant, Michael M Bronstein, Robert J Wood, Shane Gero, and David F Gruber. Deep machine learning techniques for the detection and classification of sperm whale bioacoustics. *Scientific reports*, 9(1):1–10, 2019.

[30] Shane Gero, Hal Whitehead, and Luke Rendell. Individual, unit and vocal clan level identity cues in sperm whale codas. *Royal Society Open Science*, 3(1):150372, 2016.

[31] Pratyusha Sharma, Shane Gero, Roger Payne, David F Gruber, Daniela Rus, Antonio Torralba, and Jacob Andreas. Contextual and combinatorial structure in sperm whale vocalisations. *Nature Communications*, 15(1):3617, 2024.

[32] Volker B Deecke and Vincent M Janik. Automated categorization of bioacoustic signals: Avoiding perceptual pitfalls. *The Journal of the Acoustical Society of America*, 117(4):2470–2470, 2005.

[33] Heather M Hill, Sarah Dietrich, and Briana Cappiello. Learning to play: A review and theoretical investigation of the developmental mechanisms and functions of cetacean play. *Learning & Behavior*, 45(4):335–354, 2017.

[34] Alison J. Barker, Grigorii Veviurko, Nigel C. Bennett, Daniel W. Hart, Lina Mograby, and Gary R. Lewin. Cultural transmission of vocal dialect in the naked mole-rat. *Science*, 371(6528):503–507, 2021.

[35] Mariana Lenharo. Can AI help us talk to dolphins? The race is now on, May 2025. News article.

[36] Yossi Yovel and Oded Rechavi. Ai and the doctor dolittle challenge. *Current Biology*, 33(15):R783–R787, 2023.

[37] Jacob Andreas, Gašper Beguš, Michael M. Bronstein, Roee Diamant, Denley Delaney, Shane Gero, Shafi Goldwasser, David F. Gruber, Sarah de Haas, Peter Malkin, Nikolay Pavlov, Roger Payne, Giovanni Petri, Daniela Rus, Pratyusha Sharma, Dan Tchernov, Pernille Tønnesen, Antonio Torralba, Daniel Vogt, and Robert J. Wood. Toward understanding the communication in sperm whales. *iScience*, 25(6):104393, 2022.