

An artificial life approach to studying niche differentiation in soundscape ecology

David Kadish and Sebastian Risi and Laura Beloff
IT University of Copenhagen, Copenhagen, Denmark
davk@itu.dk

Abstract

Artificial life simulations are an important tool in the study of ecological phenomena that can be difficult to examine directly in natural environments. Recent work has established the soundscape as an ecologically important resource and it has been proposed that the differentiation of animal vocalizations within a soundscape is driven by the imperative of intraspecies communication. The experiments in this paper test that hypothesis in a simulated soundscape in order to verify the feasibility of intraspecies communication as a driver of acoustic niche differentiation. The impact of intraspecies communication is found to be a significant factor in the division of a soundscape's frequency spectrum when compared to simulations where the need to identify signals from conspecifics does not drive the evolution of signalling. The method of simulating the effects of interspecies interactions on the soundscape is positioned as a tool for developing artificial life agents that can inhabit and interact with physical ecosystems and soundscapes.

Introduction

Artificial life experiments have become important tools for exploring biological phenomena. In particular, they have allowed researchers to study the relationships between evolutionary processes and ecological theories (Aguilar et al., 2014), like the emergence of interspecies relationships like mutualism and parasitism (Watson et al., 2000).

One area of ecology that has received little attention thus far from artificial life (ALife) studies is soundscape ecology. The field of soundscape ecology has been formalized by researchers over the past decade (Pijanowski et al., 2011a), building on earlier conceptions of the soundscape (Schafer, 1977). One of its foundational theories is the acoustic niche hypothesis (ANH) (Krause, 1987), which applies the concept of ecological niches — the distribution of resources that are used by a species in an ecosystem (Pocheville, 2015) — to the soundscape.

The manuscript is slightly amended from the published version to correct an error in Figure 3a. The figure in the published version plots the messages from a different run of the simulation than the one shown in Figure 3b. This in no way changes the results of the study.

This experiment tests the proposed mechanisms for the formation of these niches in a virtual soundscape in order to understand how species change vocalizations in response to one another. It models the behaviour of two species in a virtual ecosystem and tracks how their calls shift through the audio spectrum in response to different evolutionary pressures. Through the experiment, evolutionary pressure to communicate within a species is found to play a significant role in the formation of acoustic niches.

In examining the emergence of communication between artificially evolved species, this study draws from a body (Arita and Koyama, 1998; Wagner, 2000; Sasahara and Ikegami, 2007) of ALife-based studies of communication including the work of Floreano et al. (2007) in emergent communication between robotic agents. However, it is distinct from these previous studies in its focus on the effect of the emergent communication on the ecological phenomena of niche differentiation.

The main contributions of this study are the development of a simplified model of a soundscape for the purpose of rapid experimentation and in-depth analysis of population-soundscape dynamics, and the demonstration of the ANH on this model.

Background

In the physical world, the concept of soundscape — the collection of the acoustic features of a landscape — has roots and influences in a diverse array of academic fields (Lyonblum, 2017). It grew initially out of the arts and cultural studies work of Westerkamp (1974), Schafer (1977), and Truax (1978), but has since expanded into the sciences. In the field of ecology, the soundscape is considered an important ecological resource and its composition is thought to indicate the diversity and stability of the ecosystem (Pijanowski et al., 2011b). Though the field of soundscape ecology was only proposed relatively recently (Pijanowski et al., 2011b), the application of ecological principles to the study of soundscape has a longer history. Notably, the concept of ecological niches was first introduced in the context of sonic resources by Krause

(1987) as the acoustic niche hypothesis.

Acoustic niche hypothesis (ANH)

The acoustic niche hypothesis expands the concept of ecological niches to the spectro-temporal plane of the soundscape. Krause proposed that, in the same way that niche differentiation leads to species making use of the range of physical resources available in an ecosystem, species tend to differentiate their use of an ecosystem's sonic resources. This differentiation, according to Krause, occurs spectrally in the sonic frequencies that animals use for vocalization and temporally in the time-based patterns of their sounds. The theory holds that older, more mature ecosystems should show a greater degree of differentiation between the auditory niches that long-established species occupy.

The ANH describes the result of acoustic differentiation, but Endler (1992) proposed the primary mechanism for this evolutionary driver: sexual selection based on a mate's ability to hear a call and the ability to maintain territory. In this formulation, vocalizations and auditory receptors have co-evolved to maximize the reception of signals from members of one's own species (conspecifics), while minimizing interference from members of other species (heterospecifics).

This type of spectral differentiation has been observed numerous times in the wild: in the calls of certain species of frogs (Feng and Schul, 2007); in the buzzing of cicadas (Sueur, 2002); and in the overall division of a soundscape in Borneo between a series of birds, gibbons, and accompanying insects (Krause, 2008). However, it has proved difficult to experimentally probe the formation and division of spectral niches, due to the lengthy timescales that would be required to allow evolutionary processes to progress (Miller, 1995) and the complexity of the systems and soundscapes that are encountered "in the wild" (Wheeler et al., 2002).

ALife approaches

Where ecological phenomena have been difficult to experimentally investigate, researchers have proposed that ALife approaches can be a mode of inquiry that allows for the manipulation of particular conditions and the rapid collection of large quantities of data about a simulated ecological system (Miller, 1995). In 2018, Eldridge and Kiefer proposed synthetic acoustic ecology (SAC) as a toolset for exploring questions in the field of soundscape ecology using ALife methods in virtual ecosystems. Their study examined one of the assertions of ANH (Krause, 1987) — that one can identify the maturity of an ecosystem by examining its acoustic signature. Using a multi-agent system model, they demonstrated that patterns emerge in two common acoustic indices that indicate the stability of a model ecosystem.

Niche differentiation mechanisms

The study in this paper uses a virtual soundscape to test hypotheses in soundscape ecology, building on the work of Eldridge and Kiefer (2018). While Eldridge and Kiefer's study focused on the verification of acoustic biodiversity metrics, this study examines the mechanisms that breed interspecific diversity and intraspecific convergence in the vocalizations of communities in a soundscape. In particular, it is designed to test Endler's hypothesis (1992) that the ability to identify vocalizations from members of the same species drives acoustic niche differentiation.

The acoustic niche hypothesis posits that soundscapes niches are differentiated on both spectral and temporal levels, so that species ensure that their calls are isolated in both frequency and time. In order to simplify the modelling and analysis and to allow for a deeper examination of the effects of differentiation, this study focuses only on the spectral component of this differentiation.

Approach

The experimental setup for testing the drivers of acoustic niche differentiation consists of a set of evolving populations and a soundscape that they communicate within. The experiment tests two hypotheses: the alternative hypothesis (H_1), that acoustic niche spectral differentiation is driven a need to identify signals from potential mates or territorial rivals of the same species; and the null hypothesis (H_0) that spectral differentiation in acoustic niches is not driven by the need to distinguish the species of the signaller.

In order to facilitate rapid experimentation and ease the analysis of the emergent signalling systems, the experiments use a simplified, discretized model of a soundscape instead of a full-spectrum, temporally-varying acoustic space. Sounds are modelled as 9-bit vectors that represent the use of 9 available frequency bands in an instantaneous signal. These simplifications allow the repetition of the experiments many times with a large number of generations and individuals, such that results reflect general trends in the dynamics of these systems rather than the peculiarities of any single simulation. The entire system is illustrated in Figure 1 and described in detail in the sections below. Lettering in brackets refers diagram labels in Figure 1.

Populations

In soundscape ecology in the physical world, the communicative process is often assessed in two parts: sender and receiver. Every individual, of course, is both sender and receiver, but the processes experience different evolutionary pressures; "[n]atural selection favors signals that elicit a response in the receiver that increases or maintains the fitness of the sender" (Endler, 1992). The same is true in reverse, such that the sender and receiver of a particular species evolve alongside one another, but with slightly different driving forces.

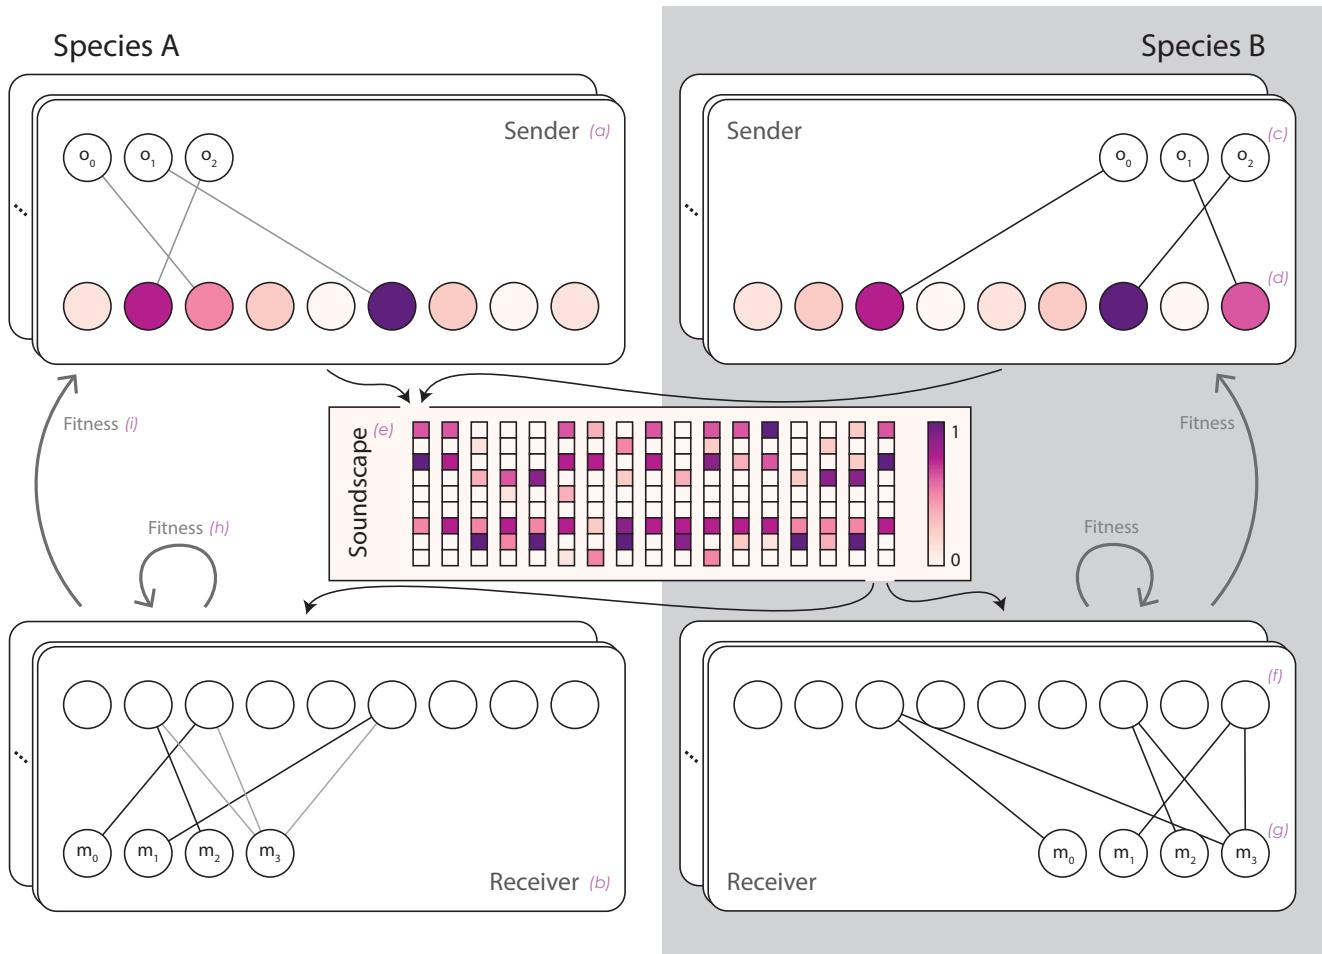


Figure 1: The experimental setup. Senders (a) encode a 3-bit message ($o_0..2$) into the 9-band soundscape (e) using a neural network with 3 inputs (c) and 9 outputs (d). Receivers (b) “hear” encoded messages from all species’ senders and predict the original message ($m_0..2$) and whether the message originates from a conspecific (m_3) using the 9 inputs (f) and 4 outputs (g) of their neural networks. Sender fitness (i) depends on how well conspecific receivers identify their species and decode their messages. Receiver fitness (h) depends on how well they identify the species of all senders and how well they decode messages from conspecifics.

The populations in this experiment are modelled as artificial neural networks, which are optimized with the neuroevolution of augmenting topologies (NEAT) algorithm (Stanley and Miikkulainen, 2002). NEAT models individual phenotypes as neural networks with a fixed number of inputs and outputs and an evolvable internal structure and connectivity. This allows the population to begin with simple neural structures and to evolve complexity as necessary to achieve the task.

For the experiment presented here, each species actually consists of two NEAT populations: a population of senders (a in Figure 1) and a population of receivers (b in Figure 1). The populations communicate over a simulated soundscape (e) that consists of 9-bit vectors, interpreted as acoustic frequency bands which can be used by senders to transmit messages. The senders encode 3-bit messages ($[o_0, o_1, o_2]$) into a representation in the 9 frequency bands using their 3-input (c), 9-output (d) neural network structures. The structure of 2 species encoding 3-bit messages into a 9-band soundscape allows for the development of relatively complex messaging while allowing the soundscape to remain undersaturated as each species could theoretically communicate in only 3 of the 9 bands. The frequency bands form the inputs to the 9-input (f), 4-output (g) receiver neural networks. The first 3 outputs ($[m_0, m_1, m_2]$) of the receiver network are its estimation of the original message and the final output represents the receiver's prediction of whether the message comes from a conspecific ($m_3 \geq 0.5$) or from a heterospecific individual ($m_3 < 0.5$).

The soundscape (e) is shared among species but messages are received serially in order to decouple timing effects; therefore, each receiver “hears” messages from the senders of all of the present species, but receives them one at a time. Additionally, any spatial arrangement of the individuals is not considered as part of this experiment, so each receiver “hears” the signals from every sender at the same “volume” with no attenuation due to a distance or set of obstacles between them.

Fitness

In a communicative process, the evolutionary pressure on senders and receivers is related but differs in some crucial aspects. The fitness functions used in this experiment reflect these differences. Since communication for mating and territorial maintenance is hypothesized to drive acoustic differentiation (Endler, 1992), the sender is indifferent to how its messages are interpreted by receivers from other species. The receiver, however, processes all messages regardless of their origin; it has to learn how to differentiate messages from conspecifics from those of heterospecifics.

Following this reasoning, the fitness of the sender (i) is formulated to reflect how well its message is understood — or correctly decoded — by the receivers of its own species; it does not depend on how the receivers of the another species

process its messages. The fitness of the receiver (h) reflects both how well it is able to distinguish the species of the sender *as well* as whether it is able to correctly decode the message.

The ability of a receiver to perform these two tasks — identifying messages from conspecifics and decoding messages — is formulated into components of the the fitness function as f_s (species identification fitness, Equation 1) and f_d (message decoding fitness, Equation 2). m is the decoded message where the first three components ($m_{0..2}$) are message as decoded by a receiver. The fourth value output by the receiver (m_3) determines whether the receiver has identified this message as coming from a conspecific ($m_3 > 0.5$) or from a member of another species. The original message is a three-bit string represented by o_i .

$$f_s(m) = \begin{cases} f_{adj}(1 - |1 - m_3|) & \text{if same species} \\ f_{adj}(1 - |0 - m_3|) & \text{if different species} \end{cases} \quad (1)$$

$$f_d(m) = 3 * \prod_{i=0}^2 f_{adj}(1 - |o_i - m_i|) \quad (2)$$

To achieve the desired fitness formulations, these equations are applied in different ways for senders and receivers by adjusting the enabling/disabling coefficients e_s and e_d in Equation 4. For each message produced, a sender's fitness is based on the interpretation of the message by all receivers *from its own species*. Equation 4 is applied for each receiver from the sender's species with $e_s = 1$. The value of e_d depends on whether the species is identified incorrectly ($e_d = 0$) or correctly ($e_d = 1$). IfR the species is incorrectly identified, then the interpretation of the message is of no consequence, which is why the fitness of the message decoding is ignored.

Receivers “hear” messages from the senders from both species and their ability to identify and ignore messages that are not from their species is an important component of their fitness. For each message that a receiver “hears”, f_s is calculated as part of its fitness ($e_s = 1$). If the receiver correctly identifies that a message originated from a member of its own species, it receives an additional score for decoding the bits of the original message (f_d) and a bonus multiplier (f_b) for correctly identifying multiple bits ($e_d = 1$), as described in Equation 4.

$$f_b(e_s, N) = \begin{cases} 1.0 & \text{if } e_s = 0 \\ \prod_{i=0}^N \left(\frac{i}{10} + 1\right) & \text{if } e_s = 1 \end{cases} \quad (3)$$

$$f_t = (e_s f_s(m) + e_d f_d(m)) * f_b(e_s, N) \quad (4)$$

One detail that requires some explanation is the adjustment function (f_{adj}) applied to the fitness equations for species identification (f_s) and message decoding (f_d).

The results that these equations evaluate are treated as binary in the operation of the system but the receivers produce output as decimal numbers between 0 and 1. If the receiver outputs $m_3 = 0.6$ for a message from a member of its own species, the consequence is no different from $m_3 = 1.0$ — the receiver has correctly decided that the message should not be ignored. However, an application of Equation 1 without f_{adj} would result in quite different fitnesses for the two outputs. Equation 5 creates a sharp rise in the fitness, centred around a value of 0.5 without producing a discontinuity, which was found to create an effective fitness landscape for the evolutionary process.

$$f_{adj}(x) = \frac{1}{2} (\tanh(8.0 * (x - 0.5)) + 1) \quad (5)$$

Null model and hypothesis

The model used to test the null hypothesis (H_0) uses a modified formulation of the fitness functions. The null hypothesis is that the need to identify messages from members of the same species *does not* play a role in niche differentiation. Therefore receivers are assumed to be able to know *a priori* which messages come from senders of their own species and no fitness is assigned for the task of species identification in this null model.

In the null version of the model, this results in the receivers only processing messages from members of their own species and ignoring messages from the other species. Senders and receivers are evaluated with the fitness function in Equation 4 with $e_s = 0$.

Results

We ran simulations of our ecosystem with senders and receivers for two species. Each population consisted of 50 individuals and the simulation was run for 300 generations. The results discussed here are averages and standard deviations from 20 independent simulations. Additionally, the results from a representative example simulation are highlighted in figures and throughout this section in order to discuss specific features of an individual simulation.

For each simulation, we generated spectrograms that mirror the type of chart that is often presented in studies of soundscapes (Krause, 1987; Pijanowski et al., 2011a), except that the x -axis of these plots represents generations instead of real-time auditory signals. These diagrams, such as the one seen in Figure 2, show how the two species's use of the frequency bands shifts from generation to generation. The initial populations's encoded messages are randomly distributed across the 9 frequency bands, but the signals converge over the course of the first 50 to 100 generations into a subset of bands used primarily by one species. In this example, after an initial series of about 100 generations, both species show consistent use of 3 bands — 0, 2, and 4 for *Species A* and 1, 5, and 7 for *Species B* — for the remainder of the simulation. *Species A* develops and then eventually

scales down the use of band 3 and band 8, but *Species B*'s use of 1, 5, and 7 remains remarkably stable through most of the latter 200 generations.

The spectral plots provide a useful visual representation of the divergent signals, but the actual level of separation can be quantified further and visualized in another manner. Figure 3a shows a mapping of the high-dimensional messages to two-dimensions using t-distributed Stochastic Neighbour Embedding (t-SNE), plotted for particular generations of interest. The encoded messages generated by senders from the two species can be seen to rapidly separate into clusters from an initial state of near-total overlap. This can be further examined in the plot below the cluster maps (Figure 3b) which shows the silhouette score for the clusters over the course of generations. The silhouette score is used in the evaluation of clustering algorithms and is a measure of the density of clusters (Rousseeuw, 1987), where a score of 0 indicates overlapping data and a score of 1 indicates dense and well-separated clusters. The rapid rise of the silhouette score here indicates the splitting of the spectrum audio spectrum between the senders in relatively few generations.

The plot shows the average and standard deviations of the silhouette scores from the series of 20 trials of H_1 (dark grey) alongside the silhouette score from the specific run from which the clusters in the plot above were derived (pink). In addition, it shows the average and standard deviation of silhouette scores from 20 trials of the null model H_0 (light grey). A test of the hypotheses using Welch's t-test — because the variance of the samples cannot be assumed to be equal — reveals that the difference between the two models is significant after generation 4 ($P < 0.01$), with an average P-value of 15×10^{-5} for latter 295 generations.

While the null model does produce a level of clustering of the species' messages, this is to be expected as a result of the selection of frequency bands on which to communicate. However, in the null model, this selection is not competitively driven by the presence of the other species. In H_1 , the receivers of the two species drive their senders towards diverging frequency bands as their fitness increases with their ability to identify messages from their own species and reject those from the other.

We also examined the actual performance of the species with regard to their ability to recognize and decode messages from their conspecifics. Figure 4 shows the scores of the senders and receivers from a species over the course of 300 generations. On average, the proportion of messages that are correctly identified as being from members of the same or other species (red) rises sharply in the first generations before steadyng near 80%. The proportions of bits that are correctly decoded and messages that are fully decoded correctly are slower to rise, but continue to do so throughout most of the evolutionary process.

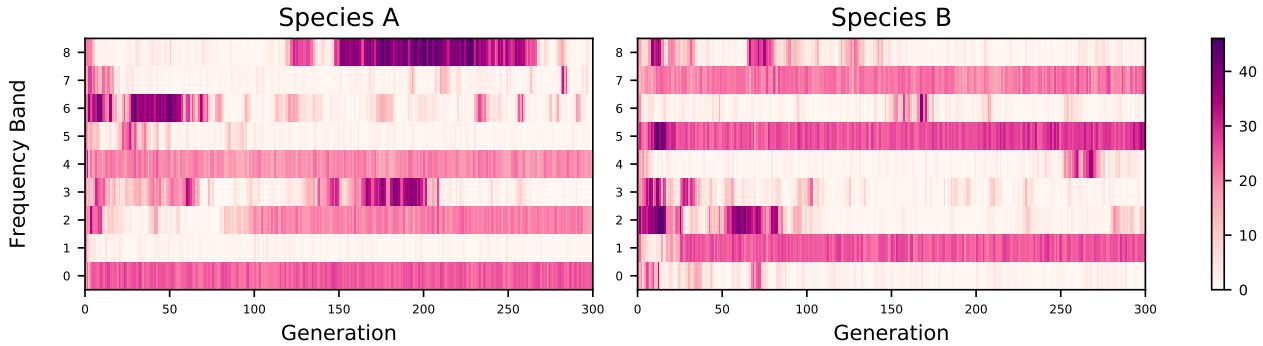


Figure 2: The use of the 9 frequency bands by the population of senders in the example simulation of H_1 . The graph shows the changing use of the frequency bands over 300 generations. In the first generations, both species’s signals are spread across the 9 bands such that the signals from the two species overlap. These signals converge rapidly to a smaller subset of the available bands. By about the 50th generation, there is little overlap between the two species — *A* uses mainly bands 0, 4, and 6 while *B*’s signals are concentrated on 1, 5, and 7 — though there is some use of bands 2 and 3 by both species. In this example, *Species B* uses band 2 heavily but intermittently until just before the 100th generation, when it ceases almost all activity on the channel and *Species A* begins to make consistent use of it for the remainder of the simulation. By the generation 300, both species have converged to the near-exclusive use of 3 channels: 0, 2, and 4 for *Species A* and 1, 5, and 7 for *Species B*.

Discussion

The results presented in the previous section demonstrate that it is possible to drive spectral differentiation in the acoustic signature of an agent through an impetus to communicate with other members of the same species. An analysis of the distance between intraspecies messages and interspecies messages shows a significant difference between the test of the alternative hypothesis (H_1) and the null hypothesis (H_0), as seen in Figure 3b. Moreover, Figure 2 provides a visual reference for the division of the spectrum in a selected simulation of the alternative hypothesis (H_1). The spectrum has been split between the two species after the first 100 generations, such that *Species A* primarily makes use of bands 0, 2, and 4 while *Species B* relies on bands 1, 5, and 7. It is interesting to note that, in the first 100 generations, band 2 is used mainly by *Species B*, however this changes around generation 90 as *Species A* begins to use the band regularly. Once *Species A* establishes regular use of the band, *Species B* never returns to it with any stability for the remainder of the simulation.

In models of the null hypothesis (H_0), the two species occasionally achieve a level of differentiation of their messages, however this occurs only by chance. In both models, species tend to converge to the primary use of roughly 3 of the 9 available channels for communication. Three channels is the fewest that can be used to encode the three-bit message and it is often the easiest solution for the evolving neural networks to find. However, in the null case, the channel selection is not driven by competition between the species, only by cooperation within a species. This lack of competition often leads to overlapping channel selections, which in turn, is responsible for the lower silhouette scores

for the null models (Figure 3b).

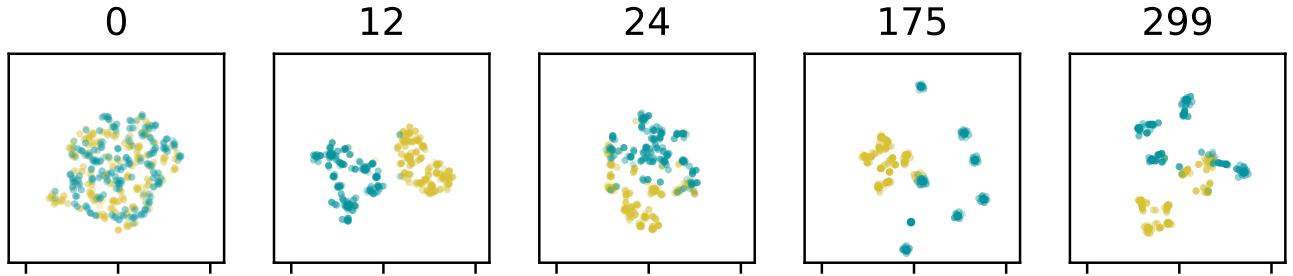
While these results cannot be taken as confirmation of the proposed mechanism of the ANH, they demonstrate that the mechanism is plausible. The drive to produce signals that are identifiable and understandable to members of one’s own species within the finite resource that is a soundscape results in the formation of acoustic niches for vocalizing species.

This study also demonstrates the efficacy of a highly simplified model in demonstrating the plausibility of a particular mechanism for the formation of patterns within a soundscape. It compliments the work of (Eldridge and Kiefer, 2018), which explores the way that common acoustic indices respond to changing populations and signals, and presents another application for a synthetic acoustic ecology. Together with other types of computational studies of soundscapes (Eldridge and Kiefer, 2018), this paper lays the foundation for a method of rapidly interrogating evolutionary acoustic processes. In addition to providing insight into ecological studies, research in this area can also be used to inform the development and analysis of evolutionary acoustic agents live “in the wild” and interact with biological ecosystems.

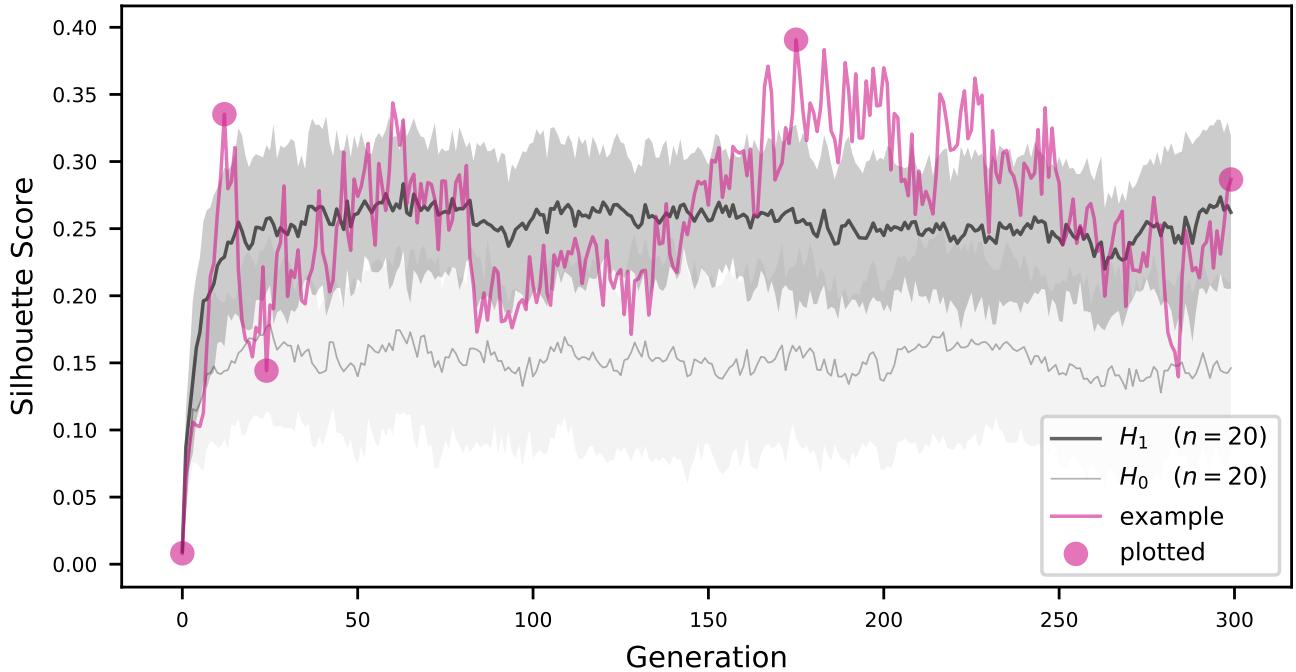
Conclusion

Though the experiment presented here is based on a highly simplified model of a physical ecosystem, it demonstrates that it is possible to rapidly and repeatedly test some of the basic principles of soundscape ecology. As predicted, the experiment was able to demonstrate the important role of intraspecies communication in the partitioning of the acoustic resources of an ecosystem.

This has important implications for the development of



(a) Cluster diagrams of messages in selected generations, mapped to 2 dimensions using t-SNE. Each point is a message generated by a sender in a single generation (labeled above the plot) of the simulation. The selected generations are marked on the silhouette score plot below with pink circles. The different colours represent messages originating from members of the two different species. The messages in generation 0 are scattered randomly from both species as the initial neural network connections for the senders are randomly generated. The messages rapidly converge to two clusters by generation 12. However, these clusters are still evenly spaced internally, as the initial selection pressure is mainly to differentiate messages between the two species. In later plots, for example in those from generations 175 and 299, smaller clusters form within the messages from a single species as the senders from each species converge on representations for particular bits and messages. This clustering drives the increasing bit and total scores in Figure 4.



(b) Silhouette score of the encoded messages, grouped by species, over the course of 300 generations. Scores reflect the validity of the message clusters when grouped by species, averaged over 20 runs of the simulation, and plotted with the standard deviation in the background. An example of an individual run is also plotted (pink) and the generations of that run that are plotted in the cluster diagram above are noted. The difference between the alternative hypothesis (H_1) and the null hypothesis (H_0) is significant ($P < 0.01$) after generation 4. The average P-value after generation 4 is 15×10^{-5} .

Figure 3: Cluster validity scores over 20 runs of the simulation. Message clusters are shown above for selected generations of an example run.

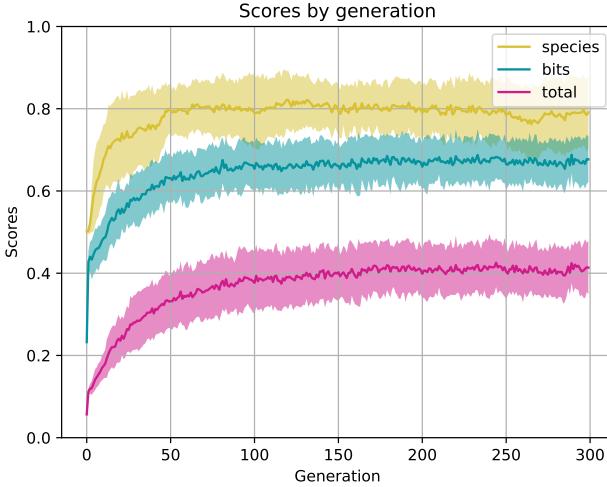


Figure 4: Performance of a species, showing the receivers' ability identify the species (yellow), and the rate at which sender-receiver pairs were able to correctly identify individual message bits (blue) and the entire message (pink).

hardware-based ALife agents for the production of sound in a physical, hybrid ecosystem. It suggests that, if one of the goals of that agent is to identify a niche for itself in the soundscape, it is important to co-evolve the auditory production with auditory perception to drive the vocalizations into an empty portion of the spectrum.

In a broader sense, this experiment sets out the foundation for a method of testing ideas for hardware-based agents in software simulations to understand the possible dynamics once they are released in the field. It grounds the inquiry into a complex phenomenon with a concrete example that solidly demonstrates the theoretical basis for a physical experiment through repetition and statistical analysis on a scale that is difficult to achieve in the field. And it demonstrates the feasibility of a key theory in soundscape ecology.

Acknowledgements

Thanks to the REAL Lab, especially Kasper Støy, Rosemary Lee, Djordje Grbic, Miguel Gonzalez Duque, Mads Johansen Lassen, and Niels Justeen for their suggestions.

References

- Aguilar, W., Santamaría-Bonfil, G., Froese, T., and Gershenson, C. (2014). The Past, Present, and Future of Artificial Life. *Frontiers in Robotics and AI*, 1:8.
- Arita, T. and Koyama, Y. (1998). Evolution of Linguistic Diversity in a Simple Communication System. *Artificial Life*, 4(1):109–124.
- Eldridge, A. and Kiefer, C. (2018). Toward a Synthetic Acoustic Ecology: Sonically Situated, Evolutionary Agent Based Models of the Acoustic Niche Hypothesis. In *The 2018 Conference on Artificial Life*, pages 296–303, Cambridge, MA. MIT Press.
- Endler, J. A. (1992). Signals, Signal Conditions, and the Direction of Evolution. *The American Naturalist*, 139:S125–S153.
- Feng, A. S. and Schul, J. (2007). Sound processing in real-world environments. *Springer Handbook of Auditory Research*, 28(Bregman 1990):323–350.
- Floreano, D., Mitri, S., Magnenat, S., and Keller, L. (2007). Evolutionary Conditions for the Emergence of Communication in Robots. *Current Biology*, 17(6):514–519.
- Krause, B. L. (1987). Bioacoustics, Habitat Ambience in Ecological Balance. *Whole Earth Review*, (57):14 – 18.
- Krause, B. L. (2008). Anatomy of the soundscape: Evolving perspectives. *AES: Journal of the Audio Engineering Society*, 56(1-2):73–80.
- Lyonblum, E. Z. S. (2017). *An Omnivorous Ear: The Creative Practice of Field Recording*. Doctor of philosophy (phd), University of Cambridge.
- Miller, G. (1995). Artificial life as theoretical biology. How to do real science with computer simulation. *Cognitive Science Research Paper*, 378(5):1–33.
- Pijanowski, B. C., Farina, A., Gage, S. H., Dumyahn, S. L., and Krause, B. L. (2011a). What is soundscape ecology? An introduction and overview of an emerging new science. *Landscape Ecology*, 26(9):1213–1232.
- Pijanowski, B. C., Villanueva-Rivera, L. J., Dumyahn, S. L., Farina, A., Krause, B. L., Napoletano, B. M., Gage, S. H., and Pieretti, N. (2011b). Soundscape Ecology: The Science of Sound in the Landscape. *BioScience*, 61(3):203–216.
- Pocheville, A. (2015). The Ecological Niche: History and Recent Controversies. In *Handbook of Evolutionary Thinking in the Sciences*, pages 547–586. Springer Netherlands, Dordrecht.
- Rousseeuw, P. J. (1987). Silhouettes: A graphical aid to the interpretation and validation of cluster analysis. *Journal of Computational and Applied Mathematics*, 20:53–65.
- Sasahara, K. and Ikegami, T. (2007). Evolution of Birdsong Syntax by Interjection Communication. *Artificial Life*, 13(3):259–277.
- Schafer, R. M. (1977). *The tuning of the world*. A.A. Knopf.
- Stanley, K. O. and Miikkulainen, R. (2002). Evolving Neural Networks through Augmenting Topologies. *Evolutionary Computation*, 10(2):99–127.
- Sueur, J. (2002). Cicada acoustic communication: potential sound partitioning in a multispecies community from Mexico (Hemiptera: Cicadomorpha: Cicadidae). *Biological Journal of the Linnean Society*, 75(3):379–394.
- Truax, B. (1978). *The World Soundscape Project's Handbook for acoustic ecology*. Music of the environment series. A.R.C. Publications, 1 edition.
- Wagner, K. (2000). Cooperative Strategies and the Evolution of Communication. *Artificial Life*, 6(2):149–179.

Watson, R. A., Reil, T., and Pollack, J. B. (2000). Mutualism, Parasitism, and Evolutionary Adaptation. In Bedau, M., editor, *Artificial life VII : proceedings of the seventh International Conference on Artificial Life*, page 564. MIT Press.

Westerkamp, H. (1974). Soundwalking. *Sound Heritage*, 3(4).

Wheeler, M., Bullock, S., Paolo, E. D., Noble, J., Bedau, M., Husbands, P., Kirby, S., and Seth, A. (2002). The View From Elsewhere: Perspectives on ALife Modeling. *Artificial Life*, 8(1):87–100.