

Adapting to a changing environment: Simulating the effects of noise on animal sonification

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

Abstract

Adaptation is an important capability in a fast-changing world. What factors allow an animal population to adapt to external changes in their environments? What effects do those changes have on the animal populations that do adapt? This paper explores these questions in the context of intraspecies communication in a noisy soundscape. Using a simulated soundscape and populations generated using Neuroevolution of Augmenting Topologies (NEAT), the same scenario is played through many times to understand the range of possible outcomes given an initial population and a set of noise conditions. While noise is found to have minimal effect on the best possible scenario, it affects how often that scenario is reached. The onset of noise is also found to impact the complexity of the evolved neural networks.

Introduction

Animal populations have long had to contend with and adapt to changing environments, but this evolutionary imperative has been accelerated in recent decades and centuries by human land-use patterns, technological innovations, and climatic impacts. Many studies have tried to understand how and how well animals respond to environmental changes in the real world (e.g., Radchuk et al. 2019). Simulations can add to this understanding by offering the opportunity to replay scenarios and understand the range of possible outcomes given the same initial conditions.

These types of simulations have been used to study a range of adaptive features in virtual populations. Wagenaar and Adami (2004) tested the importance of evolutionary history in the composition of a population and found that populations were often able to adapt to new environments quickly, but evolutionary history remains an important factor in their overall fitness. Kashtan et al. (2007) demonstrated that populations can evolve more rapidly towards a goal if the environment they are in is shifted constantly.

This study is interested in particular in sonic adaptation: how do simulated populations adapt their virtual sonic output in the presence of noise? It builds on a previous study (Kadish et al., 2019) that examined the dynamic interplay between two populations attempting to communicate in a simulated soundscape.

The simulations address 4 key questions: How does noise shift the communication patterns of a population? How does noise impact the fitness of species attempting to communicate in a soundscape? How does prior population diversity and neural network complexity impact the species' response to the onset of noise? How do those factors change in response to noise?

Background

The Anthropocene is the proposed geologic epoch in which humans have become the dominant force on the planet at a geologic scale, impacting everything from its climate to its geology and ecology (Crutzen, 2002; Tsing, 2016). One defining feature of the Anthropocene is the pace of change. In past geologic epochs, ecosystems changed slowly over time and — rarely — in a flash as a volcano erupted, a forest fire burned an area to the ground, or a landslide reshaped a mountain. Now, ecosystems change constantly and rapidly as the climate shifts, pathogens spread at the pace of modern travel, and new neighbourhoods are built on the outskirts of cities (Otto, 2018).

In this context, animals are forced to be ever more adaptable to rapidly-changing habitats. One domain in particular that is changing rapidly in the Anthropocene is the sonic domain. Soundscape ecologists refer to human-produced sounds as anthrophony (Krause and Gage, 2003; Pijanowski et al., 2011) and posit

that they compete with biologically-produced sounds (biophony) for bandwidth in the finite domain of a soundscape’s acoustic spectrum.

A growing body of research in soundscape studies and bioacoustics from the past 2 decades (Blickley and Patricelli, 2010; Shannon et al., 2016; Rosa and Koper, 2018; Raboin and Elias, 2019) has demonstrated the negative impacts of anthropogenic noise on many animal populations, but also the adaptability of some populations and their ability to shift their communication strategies in noisy environments. Animals rely on sonic information to find mates, identify and warn each other of approaching predators, and claim territory. Any additional sounds in the soundscape have the potential to disturb these information pathways, but anthropogenic sounds are often particularly disruptive because they can grow rapidly and remain a force in the ecosystem for extended periods.

Animal populations have been shown to have a range of responses to the onset of noise in their ecosystem. Some, like the spadefoot toad population in a harsh California desert overcome by the sounds of military jets flying overhead, simply go into rapid decline (Krause, 2008). In these cases, noise can stress animals (Tennessen et al., 2014), causing them to spend more time and energy looking for predators, or interfere with their mating rituals and their ability to reproduce (Shannon et al., 2016).

Others, however, have shown remarkable abilities to shift their vocalizations in response to environmental noise, with varying degrees of success. Brumm (2004) found that nightingales in noisier environments sang louder, demonstrating an animal example of the Lombard effect in which people reflexively speak louder in noisy environments. Yet other animals didn’t adjust the volume of their calls at all in the presence of noise, instead modulating only the call duration, rate, timing, and frequencies (Slabbekoorn and den Boer-Visser, 2006; Fuller et al., 2007; Love and Bee, 2010).

As with many ecological phenomena, the precise interplay between anthropogenic noise and animal responses can be difficult to study. External factors are difficult to control and replications of controlled experiments are costly and time-consuming. Evolution is historically contingent and often dependent on small variations of initial conditions and the occurrence of rare events (Blount et al., 2018). Computer simulations

and computational evolution allow researchers to “replay the tape” (Gould, 1989; Sepkoski, 2016), observe the patterns and trends that emerge from multiple simulations of the same scenario.

This study builds on prior work in simulating the evolution of animal communications in virtual soundscapes. In Kadish et al. (2019), Neuroevolution — the computational evolution of neural networks — was used to evolve the communication of populations of different species in a simulated soundscape to test the formation of acoustic niches. Here, a single species is evolved under a series of noisy conditions to better understand the impact of noise on these processes.

Evolving Neural Networks

The experiments in this paper use neural networks that create and process the sounds. The neural networks are evolved by the popular Neuroevolution of Augmenting Topologies (NEAT) algorithm (Stanley and Miikkulainen, 2002a), which can optimize both the weights and the neural networks’ architecture. Crucially for this study, NEAT divides the population into subspecies¹ based topological similarity. That way, individuals compete primarily with close relatives instead of with the population at large. The distance δ between two network encodings can be measured as a linear combination of the number of excess (E) and disjoint (D) genes, as well as the average weight differences of matching genes (\bar{W}): $\delta = E + D + \bar{W}$. If a genome’s distance to a representative member of the species is less than δ , a compatibility threshold, the genome is placed into this species. For a more detailed description of NEAT, we refer the interested reader to Stanley and Miikkulainen (2002b).

Ecosystem

The experiments in this paper make use of an ALife soundscape ecology platform originally developed for Kadish et al. (2019) and described in detail therein². The soundscape is a 9-band (or channel) space in which messages are passed between senders and receivers. Senders and receivers are independent NEAT populations that attempt to encode (the senders)

¹Typically, these are referred to as species within NEAT, but the idea of species in this simulation is distinct from NEAT’s conception of speciation. NEAT species are referred to as *subspecies* in this paper to avoid confusion.

²All of the code used in these experiments will be made available at <https://github.com/dkadish>.

and decode (the receivers) a message passed through the soundscape. They are organized into species consisting of a sender population and a receiver population, that try to communicate with each other.

In this experiment, these populations have to contend with noise on some of the channels. To communicate, senders — using their evolved neural networks — convert a 3-bit message into a “sound” that has values between 0.0 and 1.0 on each of the 9 channels. Messages are passed one-at-a-time through the soundscape and “noisy” channels have 1.0 added to the messages before the receivers can decode them.

Fitness Calculations

In every generation, each sender sends the 7 non-zero 3-bit messages to every receiver in the population, for a total of 350 messages. The fitness for both senders and receivers is based on how well the receivers are able to decode those messages. The total fitness for an individual is the sum of the scores for each message.

The base message score reflects how closely each bit of the message decoded by a receiver matches the original bit. Each bit decoding is scored between 0 and 1 and the product of those scores is multiplied by 3 for a maximum possible bit score of 3 for each message. This bit score is adjusted by a multiplication factor which is a function of the number of correctly identified bits. The details of these calculations can be found in Kadish et al. (2019).

From here, the fitness calculations for the senders and receivers diverge slightly. Receivers are rewarded for correctly identifying the messages from members of their own species³. Senders, on the other hand, are penalized for the total volume of their messages.

This volume penalty is new for this experiment and was not employed in Kadish et al. (2019). It was added to reflect the costs of producing loud sounds that are born by living populations in the wild. There is an energetic cost to producing loud sound. Sometimes is it simply shifting behaviour to vocalize louder, but other times louder vocalizations require the evolution of physical features. To varying degrees, both of these shifts often require extra energy, which means increased feeding pressure on the population or less energy available for other uses. Additionally, the

³This reward is unimportant for these experiments as there is only a single species in the ecosystem, but it helps to explain why the receiver fitnesses are higher than those of the senders.

louder sounds themselves increase the potential for both predators and prey to overhear the calling. The result may be a higher risk of predation or greater difficulty in finding food as prey scatter when they hear the calls (Brumm, 2004).

The volume penalty for each message was calculated as the sum of the sender’s output on each of the 9 channels, multiplied by a volume penalty factor which was set at 0.2 experimentally⁴.

Experiments

In order to test the 4 questions posed in the introduction, we conducted a series of 5 experiments. The first 2 experiments (1-2) test the evolution of populations from a randomly generated initial population, under constant noise conditions. The remaining 3 experiments (3a-c) examine evolutionary response to the onset of noise in a previously noise-free environment. Each experiment is run 50 times to gather information about the patterns that emerge from different environmental and initial conditions.

Experiments 1 and 2 begin with randomly generated populations of senders and receivers and evolve them in a constant noise environment for 300 generations. Experiment 1 establishes a baseline for the patterns of communication that are evolved in a noise-free environment. Experiment 2 is conducted in an environment with noise on 3 channels (0, 1, and 2). A comparison of the evolved patterns of communication in these experiments helps to address questions about the fitness impact of noise and the ways that it shifts communication strategies. Experiment 1 also provides a set of evolved populations that are used in Experiments 3a-c.

Experiments 3a-c deal with the response of an evolved population to the sudden introduction of noise in a previously noise-free environment. 3 of the 50 populations from the final generation of Experiment 1 are evolved for a further 300 generations as noise is added to channels 0-2. The 3 populations were selected for their similar final fitness in Experiment 1 and their different levels of use of channels 0-2 — the channels which become noisy in these experiments. The population in Experiment 3a uses none of the

⁴The maximum volume penalty for each message then is 1.8, while the maximum pre-penalty score is 5.148. The penalty used here is enough to exert evolutionary pressure, but not to overshadow the impact of producing a decodable message.

Table 1: Parameters for each experiment

No	Runs	Generations ¹	Noise	Onset
1	50	300	-	-
2	50	300	0, 1, 2	0
3a-c	50	600	0, 1, 2	300

Table 2: NEAT Parameters

Parameter	Value
Population	50
Node add probability	0.2
Link add probability	0.5
Weight mutation power	0.5
Stop condition	300 Generations

channels which become noisy, providing a baseline for how the population continues to evolve relatively undisturbed. The populations in 3b and 3c use 1 and 2 of the noisy channels respectively. Together, these experiments help to illuminate how different levels of interference impact a population. They also allow for the comparison of vocalization impacts and populations with different base levels of diversity and neural network complexity.

Experimental Parameters

The same NEAT parameters were used in all 5 experiments and a summary of the most important ones are found in Table 2. Other experimental parameters — such as the noise conditions — are found in Table 1.

Results

Constant Noise Environment: Experiments 1 & 2

The average spectra for all 50 runs of Experiments 1 and 2 are shown in Figure 1. Though the vocalizations were spread over all 9 channels in the first generations of both experiments, populations in Experiment 2 were less likely to end up using the noisy channels (0-2) than the noise-free ones (3-8). The effect of noisy channels on randomly generated populations is, as expected, to drive them towards noise-free channels.

The fitness of the senders and receivers in Experiments 1 and 2 can be seen in Figure 2. Though the effect is not large, the average fitnesses for both senders and receivers of Experiment 2 are consistently

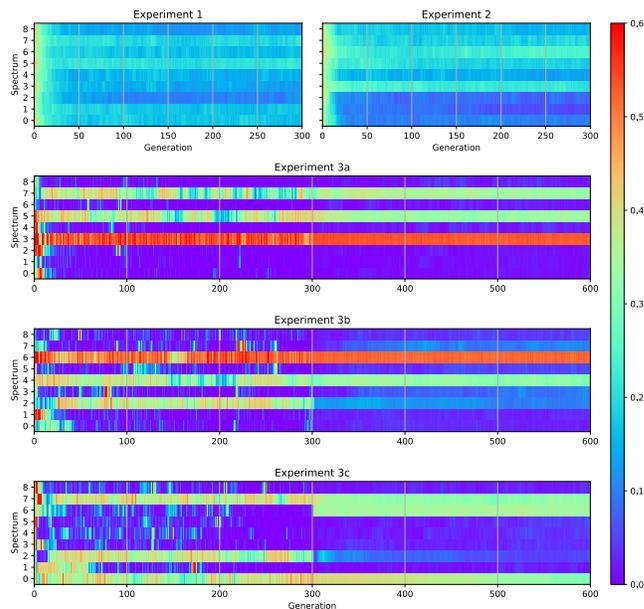


Figure 1: Average message spectrum over 50 runs. Experiment 1 is run without noise and the messages are spread over the 9 channels. Experiment 2 has constant noise on channels 0, 1, and 2, resulting in those channels being less heavily utilized on average. Experiments 3a-c are individual, noise-free runs for the first 300 generations; noise is then added to channels 0-2, and generations 300-600 are simulated 50 times, with the plots showing the averages of those simulations.

lower than those from Experiment 1. The effect is visible from a different perspective in Figure 3, where the distribution of individual fitnesses from the final 150 generations of the 50 runs is shown in a violin plot.

This representation shows that the fitnesses in Experiment 2 are lower not because the most fit runs did not achieve the same level of fitness, but because Experiment 2 produced a thicker “tail” of runs in which populations were unable to find the most fit communication strategies. In other words, it was still possible to find very fit strategies under noisy conditions, but populations were less likely to do so than under noise-free conditions.

Noise Onset: Experiments 3a-c

The average spectra for all 50 runs of Experiments 3a-c are visualized in Figure 1. For each of these experiments, the first 300 generations are from a single, noise-free simulation. The final 300 generations of the



Figure 2: Fitness from all experiments. Fitness levels from generations 300-599 of Experiments 3a-c have been transposed to generations 0-299 so that they can be compared to the fitness responses of Experiments 1-2 after the onset of noise.

spectral images in Figure 1 are the average spectra over the 50 runs of the experiment.

Communication patterns It is clear from Figure 1 that the onset of noise impacts the communication of the populations in the three runs quite differently. The populations in 3a show little shift in their communication pattern, as is expected given their minimal use of the 3 noisy channels. There are a wide range of responses from the populations in Experiment 3b. Many rapidly stop signalling on channel 2 as it becomes noisy, though some runs retain the use of that channel despite the noise. Some of the runs eventually begin to use channels 3 or 7 as a replacement. In Experiment 3c, the populations of most runs quickly switch from using channel 2 to the noise-free channel 6, yet they fail to stop using channel 0 which is also noisy.

The effect of noise on communication patterns here is complicated. One noisy channel is rapidly abandoned (3c, #2), another is vacated by a portion of the simulated populations but retained by others (3b, #2), while a third is still used by nearly all of the simulated populations after the addition of noise (3c, #0). The simulation trajectories are also non-uniform, with little variation among the 50 runs of Experiments 3a and 3c but a wide diversity of responses to the onset of noise in Experiment 3b.

Fitness The average fitnesses of these 3 experiments are plotted in Figure 2. The plots show the fitness levels of the last few generations of the initial noise-free runs as well as 300 generations after the onset of noise. All 3 initial populations were highly fit before the onset of

noise and the populations in Experiment 3a retain that high fitness level throughout the experiment, basically unaffected by noise on unused channels.

Experiments 3b and 3c experience a precipitous drop in fitness at the onset of noise, with 3c dropping further than 3b as 2 of its 3 in-use channels become noisy. However, the average fitness of the populations in Experiment 3c recovers more quickly than those in 3b, before being overtaken again later in the simulation. Neither reaches the pre-noise fitness of the initial population and by the end of the experiment they remain below the steady-state average fitnesses of Experiments 1 and 2.

Figure 3 shows that the most fit runs from Experiments 3b and 3c achieved similar fitness levels to the other experiments. However, the distribution of fitnesses — or how likely any individual run was to achieve a high level of fitness — was poor by comparison. The violins for 3b reflect the diversity of evolutionary paths taken by the 50 runs of that experiment. The upper bulge is created by a cluster of runs that found 3 channels to communicate on, while the lower bulge is formed by runs that dropped the use of channel 2, but were unable to find a replacement. In contrast, the violins for 3c are more conical. There is a cluster of higher-fitness populations that tapers into a long tail of poorer performers. Due to the continued use of a noisy channel by the large bulk of the populations, even the higher-fitness cluster sits below the high-fitness clusters of the other 4 experiments.

Diversity and Complexity The diversity of the populations and size and connectivity of those populations' neural networks from Experiments 3a-c are shown in Figures 4 and 5. Two of this study's questions are related to these data: how do diversity and connectivity impact the populations' responses to noise and how does the onset of noise impact diversity and connectivity?

The plots in Figures 4 and 5 show that there are large differences between the sender populations of Experiment 3b and those in Experiments 3a and c, but that the receiver populations have similar levels of diversity and complexity across the 3 experiments.

One of the major differences is in the pre-noise sender diversity levels. Experiment 3b had 3 subspecies at generation 300, while the other experiments had only 1. While most runs of 3b

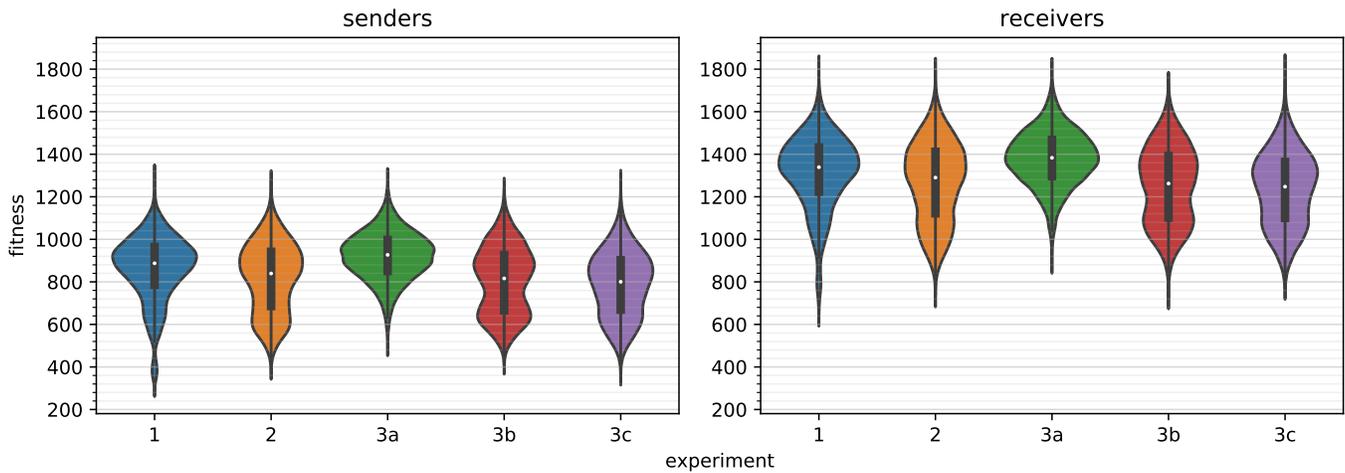


Figure 3: Violin plots — enhanced box plots that use kernel density estimates to show the placement of individual samples — showing how the fitness of individuals in the last half of each experiment is distributed. The experiment average is shown as the white dot in the centre of each violin and the boxes show the quartiles. Wider plots (3a) indicate that individuals’ fitnesses were tightly clustered, while longer plots show a greater distribution of fitness levels within an experiment. The experiments all have a similar upper limit, demonstrating that the noise doesn’t impact the possibility of populations finding highly fit communication strategies. However, the range of distributions and average fitnesses shows that the noise impacts their likelihood of doing so in practice. For example, Experiment 2 has a wider tail than Experiment 1, as noise hampered the communication strategies of some populations.

experienced a precipitous drop in diversity shortly after the onset of noise, they also tended to rebound relatively rapidly to a higher level of diversity than the other experiments. After about 100 generations, the average sender diversity of all 3 experiments flattened out somewhat to relatively steady values between about 1.5 and 2.5 subspecies.

The receivers had a larger range of diversity at generation 300, but the range quickly narrowed after the onset of noise. All three experiments experienced an immediate drop in the average species count, though the populations in Experiment 3b followed this with a subsequent spike.

For the remainder of the first 100 generations after the onset of noise, each experiment behaved quite differently: 3a followed its drop with a gradual rise; 3b’s diversity fell again after its spike and then formed a near-parabolic arc; and the diversity of 3c’s populations jumped back up to a plateau after about 20 generations. From generation 400 and on, the receiver diversity in all 3 experiments remained stable.

Another large difference between the experiments at generation 300 is the average complexity of the senders’ neural networks. Experiment 3b enters the

noisy phase of the experiment with a lower average sender node and connection count than Experiments 3a and 3c.

The 3 experiments progress differently for the sender populations: experiment 3a sees little change in the average node count and a small gradual drop in connections; 3b has an immediate drop in both node and connection count and then remains relatively stable for the duration of the simulation; and 3c has no immediate change, followed by a sharp rise in both factors after about 25 generations, coinciding with the sharp reduction in species count.

The receivers from all 3 experiments, however, show rather similar trajectories. All generally experience steady increases in the number of nodes and connections over the course of the 300 noisy generations. Experiment 3b trends slightly higher than 3c, which trends slightly above 3a, but the paths are similar. There is an exception shortly after generation 300 in the connection count, where 3b experiences an immediate steep drop from the number of pre-noise connections and 3a experiences a similar drop, offset by about 20 generations.

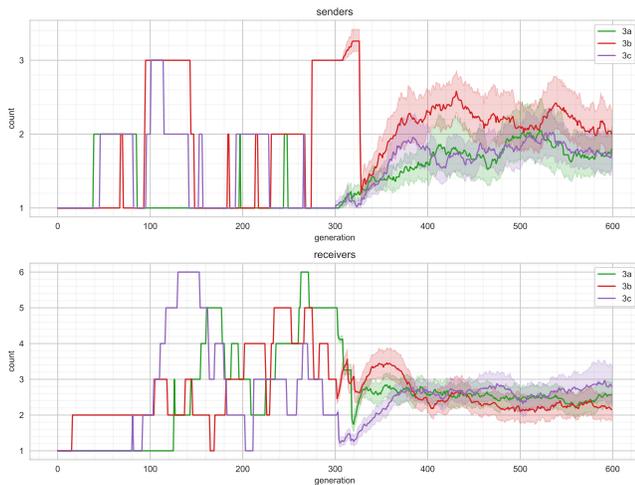


Figure 4: Number of subspecies (NEAT speciation) at each generation. The NEAT species count is directly related to population diversity.

Discussion

Communication Strategies

The first question in the introduction focuses on communication patterns. How does the presence or onset of noise shift those patterns? The expectation was that the presence of noise on some channels would drive populations away from those bands and onto noise-free bands. All five experiments showed a tendency for that to occur, but not for all runs of the experiment. Some runs dropped noisy channels without ever finding a replacement, others continued using the channel despite the noise.

When noise was added to an existing population, some channels were more easily abandoned than others. Many — though not all — of the populations of Experiments 3b and 3c stopped using the noisy channel 2, but most of the populations in 3c continued to use channel 0 despite the presence of noise. What caused this difference in response to noise?

The design of these experiments doesn't allow for a definitive answer, but it seems likely that the degree of use of the channel contributes in some way. In Experiment 3c, channel 0 is used to a greater degree than channel 2 at generation 300 — in fact, the average signal channel 2 appears to be declining even before the onset of noise. Perhaps because of this, the presence of noise on channel 2 had a greater impact on the decoding of messages than the noise on channel 0, causing populations to abandon it more rapidly.

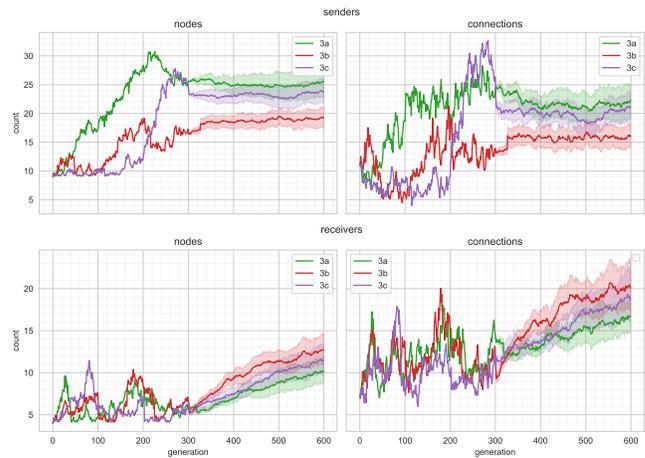


Figure 5: Average number of neural network nodes and connections in Experiments 3a-c.

There is also a difference in how channels were chosen to replace an unused channel. While the populations from Experiment 3b selected a range of open channels over many generations, the majority of 3c's populations almost immediately picked up channel 6. Why was the switch so rapid and consistent in one population and slow and varied in another? The base population in Experiment 3c had sent signals on channel 6 a few times earlier in its evolutionary history. Perhaps the changes required to fully move to that channel required few genetic changes and so it was easy for most populations to do so. Is Experiment 3b simply a good example of the contingency of evolution (Blount et al., 2018; Ogbunugafor and Eppstein, 2019), or is there something specific about the populations' structures that allowed them to find such varied communication strategies?

A future experiment that focused on selecting initial populations with a range of channel usages and evolutionary histories could help to disentangle some of these factors. Such a study should also attempt to shift the noise conditions instead of simply the initial populations to test the ability of the populations to switch to different mixes of open channels.

Fitness

The impact of noise on population fitness is relatively clear. Noise is most detrimental when it appears on parts of the spectrum that are already being used for communication. Its effect on the maximum fitness potential is minimal — meaning that it is possible to

find a very fit communication strategy, so long as there are noise-free channels available — but its negative effect on the average fitness over 50 runs and therefore the likelihood that any one population will achieve a high fitness level is observable.

It is important to note, however, that this result is only in isolated simulation. In real ecosystems, fitness is not a number, but a likelihood of survival and reproduction. There are multiple species cooperating or competing for resources, which means that events have secondary and tertiary effects. While the populations in 3a are relatively unaffected by noise in simulation, in the real world this may not be the case. They might benefit from the opening of the rest of the acoustic spectrum as populations — such as those in 3b and c — decline due to noise-induced stress, or their loss might deplete a source of food or cooperation causing the populations in 3a to suffer secondary stresses.

Diversity and Complexity

Two of the study's initial questions relate to the population diversity and neural network complexities: What effect do these factors have on a population's response to noise and how are these factors affected by the onset of noise?

From the results, it seems likely that sender diversity in the population in Experiment 3b contributed to the diverse range of responses to the onset of noise. In most runs of that experiment, 2 subspecies died off rapidly approximately 25 generations after the onset of noise. That means that the fitnesses of those subspecies were no longer improving, but the fittest subspecies remained to seed the sharp increase in diversity that the populations experienced for the next 75 generations.

The role of neural network complexity is more difficult to discern. For the receivers of all 3 experiments, it seems as though a gradual increase in receiver network complexity conferred some benefit in a noisy environment as the complexity grew steadily from the onset of noise, but not beforehand. But the exact nature of this benefit is unclear.

In the case of the senders, the network complexity didn't change much after the onset of noise. Experiment 3c's sharp drop can be understood as an elimination of the nodes which were contributing to signalling on channel 2, which was abandoned shortly after generation 300. And the jump in complexity in 3b

signals that the species that were eliminated were those on the lower end of the complexity range, leaving only higher complexity individuals behind.

In fitness terms, the drop in complexity in 3c appears to have caused a spike in fitness, but only from its very lowest point immediately after the onset of noise. The drop in complexity, then, fixed a specific, short-term problem, but the remaining rise in fitness seems to have occurred with little further shift in complexity. The spike in complexity in Experiment 3b and the gradual drop in 3a have no notable effect on fitness either. Higher complexity seems to have had range of effects on fitness and communication in these experiments.

It is difficult, then, to say how the complexity of the populations' neural network impacts its reaction to noise in general. The effect of noise onset on network complexity is somewhat more clear: senders seem to undergo some changes in the first set of generations after the onset of noise, often as sharp adjustments upwards or downwards, and then remain relatively constant; receivers, on the other hand, tend to experience a gradual complexity gain.

Conclusions and Future Work

Due to the somewhat exploratory nature of these experiments, some of the conclusions have been limited to the specific trials that have been conducted here. Further experimentation could answer specific questions that have arisen in this work in a manner that is more generalizable.

For example, a study on the introduction of noise on different channels that are being used to varying degrees could illuminate whether the intensity of channel use at the point when noise is introduced is a factor in the ability of a population to abandon it. Or a study of the responses to noise of populations at different levels of diversity and complexity could come to more concrete conclusions about whether the pre-noise conditions impact the diversity of responses.

This study lays the foundation for such future work, but it also reveals some basic properties of communication under noisy conditions: Populations tend to move away from channels with noise. The presence of noise on channels in use degrades the ability of populations to reach the highest levels of fitness. And the onset of noise tends to halt the growth of sender neural network complexity and stimulate the steady growth of receiver neural network complexity.

References

- Blickley, J. and Patricelli, G. (2010). Impacts of Anthropogenic Noise on Wildlife: Research Priorities for the Development of Standards and Mitigation. *Journal of International Wildlife Law & Policy*, 13(4):274–292.
- Blount, Z. D., Lenski, R. E., and Losos, J. B. (2018). Contingency and determinism in evolution: Replaying life's tape. *Science*, 362(6415):eaam5979.
- Brumm, H. (2004). The impact of environmental noise on song amplitude in a territorial bird. *Journal of Animal Ecology*, 73(3):434–440.
- Crutzen, P. J. (2002). Geology of mankind. *Nature*, 415(6867):23–23.
- Fuller, R. A., Warren, P. H., and Gaston, K. J. (2007). Daytime noise predicts nocturnal singing in urban robins. *Biology Letters*, 3(4):368–370.
- Gould, S. J. (1989). *Wonderful life : the Burgess Shale and the nature of history*. W.W. Norton, New York, 1st ed. edition.
- Kadish, D., Risi, S., and Beloff, L. (2019). An artificial life approach to studying niche differentiation in soundscape ecology. In *The 2019 Conference on Artificial Life*, pages 52–59, Cambridge, MA. MIT Press.
- Kashtan, N., Noor, E., and Alon, U. (2007). Varying environments can speed up evolution. *Proceedings of the National Academy of Sciences*, 104(34):13711–13716.
- Krause, B. L. (2008). Anatomy of the soundscape: Evolving perspectives. *AES: Journal of the Audio Engineering Society*, 56(1-2):73–80.
- Krause, B. L. and Gage, S. (2003). Testing Biophony as an Indicator of Habitat Fitness and Dynamics. Technical report, Wild Sanctuary Inc, Clen Ellen, CA.
- Love, E. K. and Bee, M. A. (2010). An experimental test of noise-dependent voice amplitude regulation in Cope's grey treefrog, *Hyla chrysoscelis*. *Animal Behaviour*, 80(3):509–515.
- Ogbunugafor, C. B. and Eppstein, M. J. (2019). Genetic Background Modifies the Topography of a Fitness Landscape, Influencing the Dynamics of Adaptive Evolution. *IEEE Access*, 7:113675–113683.
- Otto, S. P. (2018). Adaptation, speciation and extinction in the Anthropocene. *Proceedings of the Royal Society B: Biological Sciences*, 285(1891):20182047.
- Pijanowski, B. C., Farina, A., Gage, S. H., Dumyah, S. L., and Krause, B. L. (2011). What is soundscape ecology? An introduction and overview of an emerging new science. *Landscape Ecology*, 26(9):1213–1232.
- Raboin, M. and Elias, D. O. (2019). Anthropogenic noise and the bioacoustics of terrestrial invertebrates. *The Journal of Experimental Biology*, 222(12):jeb178749.
- Radchuk, V., Reed, T., Teplitsky, C., van de Pol, M., Charmantier, A., Hassall, C., Adamik, P., Adriaensen, F., Ahola, M. P., Arcese, P., Miguel Avilés, J., Balbontin, J., Berg, K. S., Borrás, A., Burthe, S., Clobert, J., Dehnhard, N., de Lope, F., Dhondt, A. A., Dingemanse, N. J., Doi, H., Eeva, T., Fickel, J., Filella, I., Fossøy, F., Goodenough, A. E., Hall, S. J. G., Hansson, B., Harris, M., Hasselquist, D., Hickler, T., Joshi, J., Kharouba, H., Martínez, J. G., Mihoub, J.-B., Mills, J. A., Molina-Morales, M., Moksnes, A., Ozgul, A., Parejo, D., Pilard, P., Poisbleau, M., Rousset, F., Rödel, M.-O., Scott, D., Senar, J. C., Stefanescu, C., Stokke, B. G., Kusano, T., Tarka, M., Tarwater, C. E., Thonicke, K., Thorley, J., Wilting, A., Tryjanowski, P., Merilä, J., Sheldon, B. C., Pape Møller, A., Matthysen, E., Janzen, F., Dobson, F. S., Visser, M. E., Beissinger, S. R., Courtiol, A., and Kramer-Schadt, S. (2019). Adaptive responses of animals to climate change are most likely insufficient. *Nature Communications*, 10(1):3109.
- Rosa, P. and Koper, N. (2018). Integrating multiple disciplines to understand effects of anthropogenic noise on animal communication. *Ecosphere*, 9(2):e02127.
- Sepkoski, D. (2016). “Replaying Life's Tape”: Simulations, metaphors, and historicity in Stephen Jay Gould's view of life. *Studies in History and Philosophy of Science Part C :Studies in History and Philosophy of Biological and Biomedical Sciences*, 58:73–81.
- Shannon, G., McKenna, M. F., Angeloni, L. M., Crooks, K. R., Fristrup, K. M., Brown, E., Warner, K. A., Nelson, M. D., White, C., Briggs, J., McFarland, S., and Wittemyer, G. (2016). A synthesis of two decades of research documenting the effects of noise on wildlife. *Biological Reviews*, 91(4):982–1005.
- Slabbekoorn, H. and den Boer-Visser, A. (2006). Cities Change the Songs of Birds. *Current Biology*, 16(23):2326–2331.
- Stanley, K. O. and Miikkulainen, R. (2002a). Evolving Neural Networks through Augmenting Topologies. *Evolutionary Computation*, 10(2):99–127.
- Stanley, K. O. and Miikkulainen, R. (2002b). Evolving Neural Networks through Augmenting Topologies. *Evolutionary Computation*, 10(2):99–127.
- Tennessen, J. B., Parks, S. E., and Langkilde, T. (2014). Traffic noise causes physiological stress and impairs breeding migration behaviour in frogs. *Conservation Physiology*, 2(1):cou032–cou032.
- Tsing, A. (2016). Earth Stalked by Man. *The Cambridge Journal of Anthropology*, 34(1):2–16.
- Wagenaar, D. A. and Adami, C. (2004). Influence of Chance, History, and Adaptation on Digital Evolution. *Artificial Life*, 10(2):181–190.