Genetically Sonified Organisms: Environmental Listening/Sounding Agents

Doug Van Nort
School of the Arts, Media, Performance & Design
York University, Toronto

Abstract

This paper presents the Genetically Sonified Organisms (GSOs): environmental listening and sounding agents that evolve over long periods of time in call and response with their given acoustic environment. The conceptual and ethical implications of employing these sustainable, solar-powered devices within an established natural acoustic ecology are discussed, as well as a report on their behavior in a long-term public land-art context. The GSO design is explained, with behaviour that includes a welcomed discrepancy that was discovered between their theoretical evolutionary functioning and unexpected divergent behaviors that emerged in the field, through their relationship to solar energy among other environmental conditions.

Introduction

The work presented in this paper takes as its context for exploration the acoustic ecology (Wrightson 2000) of a given site, rather than a musical concert hall, theatre of gallery context. The project manifested as a piece of environmental sound art, taking the form of an installation within a group sound art show at the Fieldwork site (Fieldwork 2018) in rural Ontario, dedicated to land art explorations. The project is concerned with creating a piece that is in dialogue with its acoustic environment and is designed so as to have awareness of its own place within this through adaptive traits, and as such can be considered as a piece of ecological sound art as articulated by (Gilmurray 2017). This research shares the spirit and intention of the ecoacoustic approach of (Burtner 2011) along two important dimensions:

(i) The goal here is to move away from a soundscape approach based on a sampling of the natural environment and relocation of this into an electroacoustic studio context, and instead invites the public to listen to the world as a primary source of aesthetic reflection, becoming active participants in the process.

(ii) To move further away from anthropocentrism and human mediation by regarding humans as one type of actor within a complex ecosystem of sound-making agents including mammals, insects, birds and machines.

Figure 1: GSO immersed in a dense natural environment

However unlike the ecoacoustic approach, the work here does not introduce technology as a form of prosthesis that mediates and amplifies human-environment interactions, but rather as a material and computational condition that is introduced into an existing eco-system as another agent amongst a complex web of life. It was this motivation that led me to take a metacreative approach to the construction of the work. I will describe here a bit more about the project motivations and context, its design, and finish with reflections that result from the realization of the work over a period of months, with focus on the environment-machine agent interactions that unfolded.

GSOs

This project is called Genetically Sonified Organisms, or GSOs. It was commissioned for a land-based sound art show, and ran from mid-May until mid-November of 2017. The name is a conscious reference to the phenomenon of genetically modified organisms (GMOs) in our food supply, as a means of highlighting the non-neutral act of placing a set of human-made sound-making objects into an existing acoustic ecology, and the potential for disruption as is present in any site-specific work of sound art. It also raises this reference as a point of distinction: I would argue that GMOs have a narrow motive of human economic efficiency and financial optimality, while this project explicitly focuses on an expansive and dialogic relationship with the non-human environment that does not privilege human financial gain as its primary motive.

In this project the medium of focus is the acoustic environment, and this dialogue is considered by constructing a context that encourages the emergence of sonic patterning that results from the interplay of various agents in a natural environment, including humans and non-humans, both computational and biological. By introducing computational agents into the environment, the project conceptually is focused on shifting away from prevalent views on artificial intelligence that tend towards techno-centrism by raising AI systems to human and supra-human levels of expected behaviour, instead shifting the conversation towards a view based on listening to the living environment itself as a source of non-human intelligence. Computational agency is then framed, as an insect or bird, to be one small part of this larger intelligence.

The physical component of the work resides in the creation of a set of 'creatures' that are designed to communicate with one another and with the large environment via the medium of sound. The means of communication begins as a simple call/response from a pre-selected set of simple tones, chosen for their aesthetic effect but also for their functional qualities: residing within the range of human hearing and sufficiently distinct from any animal or human generated sound present at the site. Each creature is designed to respond to sounds that are similar to their known vocabulary, while adapting their call based on the difference found between their own lexicon of calls and those that they repeatedly hear in the environment. As an aesthetic and compositional work, the interest with the project is in creating a piece that is pointillistic and filled with fairly discrete and engaging tones/noises/etc. More importantly, as a work that is aimed at being in conversation with the site, the conceptual interest is in discovering the ways in which the work both influences and is imprinted upon by the acoustic ecology of the site, including birds, insects, various mammals and the humans who pass through. One of these computational creatures, situated at one location of the project site, can be seen in figure 1. Three specific questions that the project raised at its outset were as follows:

(i) Will the set of computational creatures evolve as a collective?

(ii) Will they diverge and specialize uniquely to their local neighbours, or will the entire set adapt collectively towards some set of sounding agents that are present within the environment?

(iii) Will there be periods of adaption and change that vary clearly with the transitions of the seasons?

Before engaging these questions, I will describe the GSO design in more detail.

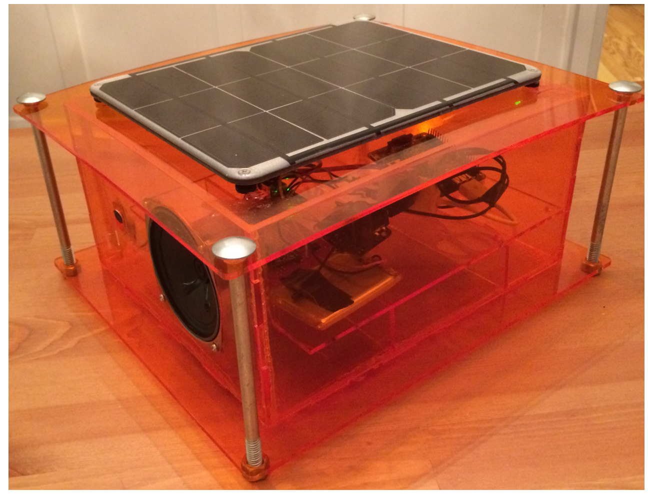

Figure 2: Protoyping GSO Enclosure and Circuitry. Speaker and Microphone on left-facing edge, solar panel on top.

GSO Design

The design of the GSOs encompasses three areas: material, sonic and behavioural components.

GSO: Materials The design of the GSO enclosure as well as its circuity began from an eco-ethical perspective: I felt that the devices must be weather proof in a fashion that does not contaminate the site, and they must be solar powered. While the intention was to point "outward" at the acoustic environment and away from the GSO as an art-object, there was no attempt to hide them at the site. As such, consideration was given to their form in terms of the intersection of their sonic functionality as well as in being an aesthetic object and reference for a biological creature. This resulted in the design seen in figure 2: the solar panels sitting on top of the device (appearing almost like a beetle's shell), and a microphone and speaker sitting at the "face" of the creature, as can be seen on the left-facing edge. This design allowed for easy access to the circuitry, and silcon was used for waterproofing. The enclosure was constructed of laser-cut acrylic, with a mixture of orange and blue creatures, these colors being chosen in order to stand out rather than blend in to the natural environment with the hopes of not confusing the wildlife present. The electronics employed a Raspberry Pi 3 for computation, Adafruit circuitry for audio and power management, and an Arduino Lillypad for power cycling and sleep/wake scheduling.

GSO: Sound Design The sound design also began from an ethical point of view. As with the enclosure, the question was the extent to which the sound should blend in or stand apart from the existing sonic environment. Knowledge of this sonic environment was gained through visits and field recording trips to the site, as well as discussion with the curator who has extensive experience at the location. My intuition was to create a sonic language that stood apart from the existing soundscape, in order to not disrupt mating and other acoustic communication rituals. This was

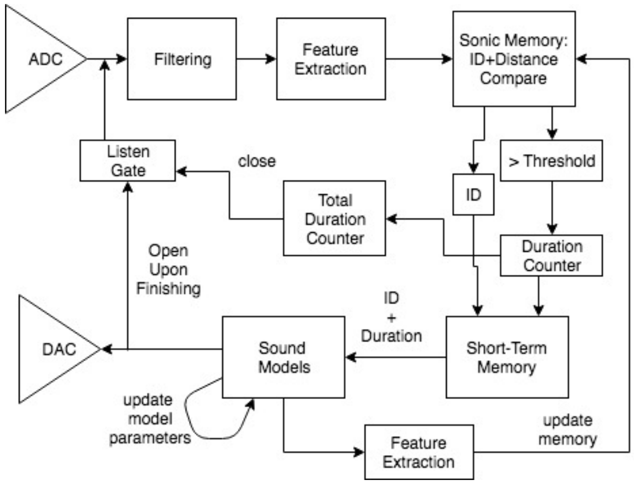

Figure 3: System Diagram of the GSO Listening/Learning Process

further reinforced through discussions with a colleague who specializes in bioacoustics research. At the same time, the project walks a fine line in that the goal is indeed for the GSOs to engage in sonic interaction with biological agents, to adapt to their environment and form a dialogue with these creatures through mimicry. For this reason, I decided to utilize physical modelling synthesis techniques that were based upon animals that may be present at the site. These methods were created in Pure Data, and drew from the synthesis techniques presented in (Farnell 2010) that the author refers to as a practical modeling of "lifeforms". Each GSO began with a lexicon of twenty identical tones, with each drawing from models of wildlife that may be present at the site: crickets, cicadas (friction models), frogs, flies, bees (source-filter models), birds (avian syrinx model), and large mammals (vocal tract models). The model choices were thus not made for aesthetic reasons, and in fact this was a compromise in that some tones would likely not have been chosen were this a studio composition. Instead, they were designed for their initial abstractness from the soundscape of the site, yet were created for their potential to converge over long periods of time to very similar sonic output as certain wildlife that was known to inhabit the site.

GSO: Behaviour Design As a central focus of this current discussion on metacreation, I will devote the following section to this description.

GSO Agents in the Lab

The behaviour of the GSOs is built around the sound models employed, the affordances of their physical construction and circuity, and most importantly the desire for an adaptive, dialogic relationship with the site. In searching for intelligence within the larger sound field of the site, the project follows on the pioneering work of David Dunn (Dunn and Van Peer 1999). Notably his piece Mimus Polyglottos, in which he introduced square waves to a group of mockingbirds in order to observe their mimetic abilities. Much like the superb lyrebird of Australia, the mockingbird species is an exceptional mimic who is capable of reproducing various species as well as complex and "unnatural" human-made machines such as chainsaws or cell phones. While mating and defense of territory appear to be clear functional reasons for this behaviour, the exact range of explanations remains an open area of research, with debates on the extent to which the behaviour results from learning or from evolutionary convergence (Kelley et al. 2008). The fact that mimids themselves integrate both machine and biological sources points to sonic mimicry as a well-articulated entry point into examining emergent intelligences in environments that integrate both biological and technological agents. Providing the GSOs with sound models based on existing animals affords them the ability to *potentially* reach the level of mimid over time.

Beyond the specific case of mimids, vocal convergence is a common trait across various species. With the GSOs I focused on a basic one-on-one vocal adaptation following a paradigm of listening for a relevant call, and attempting to match this with a response. While this certainly manifests within species, there are many cases of cross-species convergences related to bonding and the specifics of the environment (Tyack 2008). This justifies the interaction type, and at least opens up the potential for GSO-animal convergences based on positive traits such as bonding, which would be one desirable outcome of the project.

Following this behaviour model, and working with the constraints of Pure Data on a Raspberry Pi 3, the GSO behaviour model focused on the following steps:

(i) Listen (silently) to the environment for sounds (analyzing features) until a sound that is close to their lexicon of calls is heard.

(ii) Make note of which sound this is close to, and the duration in short-term memory (30 seconds max call length).

(iii) Continue making note of all sounds that are close to the lexicon until there are none for a duration of 10 seconds.

(iv) Stop listening, and "reproduce" this sound or sequence of sounds by making the call that most closely matches this, while listening internally (analyzing features) to their own sound.

(v) Update the long-term memory of calls, accessed by the distance search function.

(vi) Increment the list of sound synthesis parameters for each model that was invoked, so that the next output moves in the "direction" (relative to parameter space) of the sounds most recently heard.

The system diagram outlining this process is shown in figure 3. A given GSO continually listens to sound that has been bandpass filtered, analysing a set of eight averaged sound features: fundamental frequency (using fiddle), spectral centroid, zero crossing rate, and the first five MFCC values. This vector is compared to a table of values that represent the current lexicon of calls, and a distance value as well as the closest ID are output. This comparison utilizes the TimbreID toolkit (Brent 2010). Each time the distance is below a given threshold for 1-30 seconds, the duration and model ID are held in short-term memory. If either 10 seconds has passed with unrecognizable sound sources, or if the GSO hears 30 seconds worth of recognizable sounds, the system stops listening and proceeds to reproduce the sequence of calls. The ID and durations are fed to the bank of sound synthesis models, and the call is output. This output is analyzed internally, and these analysis values are added to a larger long-term sonic memory that contains many more members than the current lexicon of 20. A distance comparison of the input to this larger database of analysis values is computed for all data points that correspond to the currently-chosen model ID. This long-term sonic memory pairs analysis values with the model synthesis parameters that produced them, allowing for the synthesis parameters that are closest to the currently-heard call to be known. This allows for the chosen sound model to be updated slightly in the direction of the given call. This process is depicted in figure 4. While a GSO could theoretically produce an output that is quite close to the heard call right away, a design decision for the project was to update the model synthesis parameters in a very incremental fashion. In particular, the increment was calculated such that if the same sound were heard 10 times per day, the model would (theoretically) converge to the point of matching this sound source over the course of five months, which represented the initial duration of the exhibition. This incremental update is a key part of the evolutionary aspect of the system: the large sonic memory is regarded as the GSOs cognitive understanding that links the heard sound to all sounds it has created an understanding of previously (i.e. in practice, mapped analysis to synthesis parameters), while the running memory of the current 20 calls represents its "embodied" knowledge of what sound it can currently produce. In this sense, adapting sound model parameters is akin to learning the articulatory mechanisms required to produce a given sound, and thus is in keeping with a long-form evolutionary convergence. Of course, a linear increment of model parameters might eventually lead to outputs that occupy a drastically different part of analysis space (and relatedly, a sound that is heard as drastically different). This was again by design, as such a diversity of calls opens up the possibility of attracting a wider diversity of responses from other inhabitants of the pond environment, and falling into points of stability in parameter space would thus suggest repeated interactions with the environment.

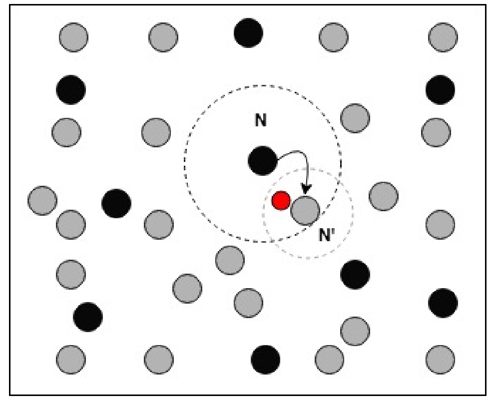

Figure 4: GSO Listening/Adapting, reduced here to 2-D for ease of visualization: Input sound analysis vector (red) is compared to active lexicon of calls (black), and found to be closest to value N. It is then compared to a much larger database of all known analysis values related to the same model (grey). It is found that N' is the closest sound overall, and so the synthesis parameters for the model corresponding to N are incrementally updated so that they are closer to the synthesis parameters that produced N'.

GSO Agents in the Wild

The previous section addresses the reasoning for choosing the behaviour model, and the design of a single GSO. However the true interest was in how these creatures, as a collective, would dynamically interact with each other and the soundscape of the project site. This chosen site is a pond located at the back of the Fieldwork property, roughly 50 feet in diameter with levels varying drastically across the seasons. A particularly wet moment for one of the placements can be seen in figure 5, which was taken during the install process. In fact, figure 1 and figure 5 are the same location, the former picture taken in September and the latter in May. This site was chosen as the pond is known to be home to a very acoustically rich and diverse set of creatures: beavers, bullfrogs, peepers, ravens, geese, songbirds, crows, occasional bears and other mammals, crickets, and cicadas are among the local inhabitants and visitors that one might hear. I engaged in a series of call and response vocal exercises during the spring (recording these), which informed my understanding on how loud (and sensitive) the GSOs needed to be. The goal was for these creatures to be placed across the pond, and have the ability to hear one another. This effectively determined loudness, microphone sensitivity and power consumption requirements. While considering ground stability, unobtrusiveness to visitors, and sunlight access, five locations were identified around the pound. The GSO creatures elicited much discussion and surprise during the day of installation, demonstrating clear moments of conversation between one another as they had done during the development process. Over the course of the exhibit, they would prove to present a more complex and unexpected pattern of behaviour.

Discussion: Chaos and the Emergent Mind of the Pond

Some readers will recognize this section heading as a title of a composition by David Dunn. I invoke it here not as an explicit reference to the content of this work, but rather as a poetic and accurate description of the GSOs during their tenure at the Fieldwork pond. In invoking "mind" here I am course reflecting upon the nature of agency that was or was not present. *Where* this was located is not as clear in more traditional musical metacreative contexts, which typically involves one or more performer(s) interacting with one more more system(s). Certainly these paradigms from

Figure 5: GSO during installation process.

the musical context differ from the GSO project, paralleling the abundance of discussions around music vs. sound art. That said, composition in this context can be seen as as the basic set of behaviour rules for the system. Looking at GSO behaviour through the lens of composition, providing simple call/response rules towards an emergent larger form can then be seen as similar to Pauline Oliveros's Sonic Meditations (Oliveros 1974). Making the statement that this properly resides within the realm of composition imports the same problems and prospects as the sonic meditations, or other work (such as Dunn) that has looked "outside" to the larger world in order to find musical structure. Certainly, in the GSO project context any sense of emergence that comes from adaptive or evolutionary rules differs from the more common views on evolutionary algorithms found within musical metacreation such as (Miranda and Al Biles 2007) or (Biles 2013), whose focus remain on simulations of evolutionary behaviour within a closed computational environment, possibly with a set of influencing channels of musical input that determines fitness (traditionally found in systems focused on machine improvisation, as opposed to composition). At the same time, in the search for locating agency within the GSO project we can productively look to discussions around autonomy that arise within the MUME literature.

Notably, in (Eigenfeldt et al. 2013) the authors outline a hierarchy of seven distinct layers of autonomy applicable to metacreative systems, each having different qualities, ranging from independence up to volition. Regarding a single GSO creature we would be hard pressed to advance this far along with hierarcy: certainly a given GSO has independence (level 1) as well composed (level 2) and generative reactions to input (level 3), simply by virtue of its feedback-adaption routine. Moving beyond this level is when the taxonomy breaks down for this particular context. Having experienced many systems within the the MUME community, and even compared to my previous work in the field (Van Nort, Oliveros, and Braasch 2013), a GSO creature and its Pure Data algorithms running on a low-powered Raspberry Pi certainly does not possess the same level of computational complexity, and can not be considered to be proactive (level 4 of the taxonomy). However I am taking the full GSO collective, situated in its intended environmental context of the pond, having a specific material conditions of construction, as the non-human agency for consideration here. In "zooming out" to this level, we must consider the connection to place but also the material idiosyncrasies that ultimately defined the Genetically Sonified Organisms. In particular these creatures were built with a (staggered) sleep/wake cycle wherein they would be active for 1-2 hours, then sleep for 1-2 hours in order to recharge their battery. They were thus diurnal by design, and the Arduino and timer which regulated this were powered by a separate watch battery. However, through an oddity of the circuitry implementation, these regulating components were directly correlated with solar levels. The result was that when the main battery and the solar levels were both low, the timer would freeze and thus shift out of phase with "clock time". The result was that creatures became noctural over time, pushing back against the scheduled hours for public experience of the work. Taking this as part of their perceived agency, these creatures indeed exhibited proactivity (level 4) and adaptability (level 5) through their individual evolutions of response to both sun as well as acoustic context. They were received as obstinate at times by myself in trying to document the clear inter-GSO interactions that were exhibited in the lab context, as well as by visitors who wanted behaviours to conform to their visit times at a location that was far from a city centre. They were received as surprising in a variety of situations, for example when a clear call-response was perceived between GSO and bullfrog during the spring, between a raven in the summer, and with rainfall in the Fall. Meanwhile, their large-scale adaptability across seasons presented its own unique trace of this larger intelligence. The GSOs recorded each increment of their output over minutes, days and months and saved this to text. In examining this, different sides of the pond indeed exhibited similar tendencies, with outputs converging closer to bird-like sounds on one side, and insect-like sounds on another. Each unit maintained a semblance of uniqueness, in contrast to their initial uniform lexicon of tones.

Conclusion and Future Work

It would be unwise and inaccurate for me to suggest that discussions of agency and autonomy in the Genetically Sonified Organisms project fits neatly in line with the main stream of the rich discussions emerging form the MUME community, largely focused on compositional or improvisational agents for musical performance. At the same time, just as contemporary musical discourse was opened up by looking outside of the concert hall through the works of Cage, Oliveros, Dunn and many others, I think that considering metacreative systems that are taken up as part of a larger ecological framework such as the GSOs can help to enliven new discussions around the complex networks of agency that arise when biological and computational agents, both human and non-human, begin to inform one another as actors within a mutual ecosystem. This not only parallels discussions emerging within interactive system design (Keller and Lazzarini 2017) and in larger critical thought (Latour 2014), but is in keeping with an "ethics of engagement", touching base with non-anthropocenic perspectives within the field.

Acknowledgements

I would like to thank Kieran Maraj for his excellent work in translating my abstract enclosure and circuitry designs into concrete and robust real-world objects. I also would like to thank Gayle Young for her initial connection to the wonderful Fieldwork collective, Susie Osler for her curatorial eye and overall support, and Cam Gray, Jerrard Smith, Diana Smith and Sheila McDonald for their assistance and overall openness of spirt throughout the installation and maintenance of the project. Development of this work benefitted from the DisPersion Lab environment, supported by the Canadian Foundation for Innovation and the Canada Research Chairs program.

References

Biles, J. A. 2013. Straight-ahead jazz with genjam: A quick demonstration. In *MUME 2013 Workshop*.

Brent, W. 2010. A timbre analysis and classification toolkit for pure data. In *ICMC*.

Burtner, M. 2011. Ecosono: Adventures in interactive ecoacoustics in the world. *Organised Sound* 16(3):234–244.

Dunn, D., and Van Peer, R. 1999. Music, language and environment. *Leonardo Music Journal* 63–67.

Eigenfeldt, A.; Bown, O.; Pasquier, P.; and Martin, A. 2013. Towards a taxonomy of musical metacreation: Reflections on the first musical metacreation weekend. In *Proceedings of the Artificial Intelligence and Interactive Digital Entertainment (AIIDE?13) Conference, Boston*.

Farnell, A. 2010. *Designing sound*. MIT Press.

Fieldwork. 2018. http://www.fieldworkproject.com. "[Online; accessed 25-May-2018]".

Gilmurray, J. 2017. Ecological sound art: Steps towards a new field. *Organised Sound* 22(1):32–41.

Keller, D., and Lazzarini, V. 2017. Ecologically grounded creative practices in ubiquitous music. *Organised Sound* 22(1):61–72.

Kelley, L. A.; Coe, R. L.; Madden, J. R.; and Healy, S. D. 2008. Vocal mimicry in songbirds. *Animal behaviour* 76(3):521–528.

Latour, B. 2014. Agency at the time of the anthropocene. *New literary history* 45(1):1–18.

Miranda, E. R., and Al Biles, J. 2007. *Evolutionary computer music*. Springer.

Oliveros, P. 1974. *Sonic meditations*. Smith Publications Baltimore, MD.

Tyack, P. L. 2008. Convergence of calls as animals form social bonds, active compensation for noisy communication channels, and the evolution of vocal learning in mammals. *Journal of Comparative Psychology* 122(3):319–331.

Van Nort, D.; Oliveros, P.; and Braasch, J. 2013. Electro/acoustic improvisation and deeply listening machines. *Journal of New Music Research* 42(4):303–324.

Wrightson, K. 2000. An introduction to acoustic ecology. *Soundscape: The journal of acoustic ecology* 1(1):10–13.