

Robophony: A new voice in the soundscape

David Kadish
IT University of Copenhagen
itu.dk
davk@itu.dk

Soundscape studies typically distinguish between three sources of sound: biophony, sounds produced by animals, plants, and other biological entities; geophony, sounds produced by non-living features such as water or wind; and anthrophony, sounds produced by humans and human technology. Recent developments in the study of hybrid ecosystems and ecological robotics challenge these categorizations. A series of four soundscape interventions are considered, which lead to the proposal of a new category: robophony. These interventions — robots and autonomous digital agents — operate in feedback loops with the existing bio-, geo-, and anthrophony in the ecosystem. The properties that emerge from these cases: site specificity, hybrid sourcing, and layered temporality form the basis of this new category of sound in the soundscape.

robophony, soundscape ecology, robots, NEAT

1. INTRODUCTION

The categorization of sound in a soundscape is necessarily imperfect. Sounds are fluid, overlapping and — in the age of digital reproduction — often removed from their sources and therefore difficult to group. And yet, the categorization of sound is an important part of the analysis of a soundscape. It helps researchers to understand how a sound enters a soundscape, what role it plays in structuring the sonic environment, and how it relates to other parts of the soundscape (Pijanowski, Villanueva-Rivera, et al., 2011). A set of thoughtfully constructed categorizations can help researchers and practitioners identify patterns and discuss the features of a soundscape in a meaningful way.

Schafer (1977) presented the first taxonomy of sound objects in his 1977 text “The tuning of the world”. The labels were based on the catalogue headings in a collection of literary descriptions of sounds gleaned from written documents that his team had gathered. The headings were compiled as entries were added to the catalogue, so that the taxonomy is more the result of a generative process than a deliberate organization. They capture the whole set of catalogued sounds into a three-layer hierarchy, with the top layer dividing sound objects into the supercategories of natural sounds, human sounds, sounds and society, mechanical sounds, quiet and silence, and sounds as indicators.

The field of soundscape ecology, which developed in part out of Schafer’s work, uses a different system

of categorizing sound sources. Bernard Krause, one of the pioneers of the field, used a tripartite division of biophony, geophony, and anthrophony to group sounds in a 2003 technical report on soundscape as an indicator of ecosystem health with Stuart Gage. The categories of geophony and anthrophony was split further in Krause’s “Anatomy of the soundscape: Evolving perspectives” (2008). Here, he used electromechanical, physiological, controlled, and incidental as the four subcategories of anthrophony, covering anthropogenic sounds from repetitive mechanical from whirring motors to the crunching of leaves underfoot as a person walks through the woods, while the subcategories of wind, water, weather, and geophysical forces subdivide geophonic sounds into more precise groups.

These categories are operationalized in different ways by Krause and his contemporaries. They are deployed in the production of a framework for understanding the dynamic relationships within a soundscape and as way of understanding which types of sounds dominate over different temporal and spatial configurations (Pijanowski, Farina, et al., 2011). They also help ecologists to measure overall ecosystem health (Pijanowski, Villanueva-Rivera, et al., 2011), understand the effects of different sounds of aspects of animal life (Pijanowski, Villanueva-Rivera, et al., 2011), and establish the relational dynamics between the different categories of sound (Gasc et al., 2017).

The model of these relationships in Pijanowski, Farina, et al. (2011) gives specific attention to the feedbacks between the whole soundscape and these categories of sound. This describes the way in which, for example, anthropogenic additions to the soundscape sounds can mask bird calls, prompting them to shift, thus prompting a secondary soundscape change.

One thing missing from this model, however, is a way of accounting for the pace of these feedback mechanisms. If one wants to understand the dynamics of a soundscape — which is a key goal of soundscape ecology — it is important to be able to discuss the speed at which different types of sound shift in response to other shifts in the soundscape.

It also places little emphasis on the biosemiotics of the sounds that are entering the soundscape. Under the current model, the sound of a car engine and the playback of recorded soundscapes over a loudspeaker are both treated as anthrophony, but the two sounds would likely be perceived in wildly different ways by inhabitants of an ecosystem, human or nonhuman.

This does not require a wholesale re-imagining of the system currently in use. Bio-, geo-, and anthrophonic sounds are largely internally consistent in terms of their feedback and response rates. However, a new set of actors is emerging that requires its own category in this formulation.

Learning robotic systems and autonomous agents, particularly those designed to engage directly with the sonic ecosystem, have the ability to change their behaviour in a manner similar to the ways that animals shift their sonic outputs in response to environmental changes. However, robots operate on sped-up timescales, shifting behaviours in ways that are unencumbered by established group dynamics and highly-specified morphologies. This paper proposes a new category of ecological sound, *robophony*, to capture this emerging class of sound objects.

What follows is a summary of the history of the contemporary categorizations of soundscape and their place in a dynamics-based conception of these categories. Examples are brought to motivate the creation of the category of robophony, going into detail about two cases which represent the primary motivators. Finally, the category of robophony and the issues that remain with its instantiation are discussed.

2. BACKGROUND

All categories represent an imperfect flattening of an idea space. At their best, however, categorizations and taxonomies of objects and phenomena can bring clarity and new analytical perspectives to a field of study. Since the formal definition of the soundscape in Schafer's *The tuning of the world* (1977), two main taxonomies have been used to group sounds. Schafer's own categorizations were presented in Chapter 9 of his book, while those used by Krause were developed as part of an understanding of soundscape ecology beginning in the late 1990s.

2.1. Schafer's Taxonomy

Schafer dedicates an entire chapter of his *The tuning of the world* to the subject of "classification". He notes that — depending on one's perspective — sounds might be categorized according to their acoustic, psychoacoustic, semiotic-semantic, or emotional-affective qualities. He first describes a system of classification according to physical characteristics — the duration, frequency, dynamics, internal fluctuations, mass, and grain — before noting that these describe isolated sound events, thus limiting its utility in the study of soundscape ecology.

Most pertinent to this discussion, however, is Schafer's subsequent description of a taxonomy based on the semiotic and semantic content of a sound. This taxonomy emerged from a process of cataloging written descriptions of sound. The World Soundscape Project (WSP) team had engaged in a lengthy process of collecting what Schafer calls earwitness accounts and sorting them into a catalogue. Schafer's taxonomy is drawn directly from the categories and subcategories that emerged from this study.

The taxonomy consists of six primary categories: *natural sounds*, *human sounds*, *sounds and society*, *mechanical sounds*, *quiet and silence*, and *sounds as indicators*. These categories — with the exception of *quiet and silence* — are broken into secondary and often tertiary subcategories. In total, 46 secondary subcategories and an extensive set of tertiary subcategories serve to categorize every sound that the WSP catalogued from written sources.

Because the taxonomy is derived from the cataloging process, it doesn't have a pre-conceived organizing principle. It is, however, biased by the types and origins of literary sources that are chosen as some of its subcategories reveal. There is, for example, a *town soundscapes* subcategory of *sounds and*

society which features a single tertiary subcategory: *Britian and Europe, etc.*

It is also important to note that the categories are based on a semiotic approach to sound. Specifically, the approach is anthroposemiotic (Emmeche, 2007), prioritizing the human understanding of meaning from sound, as evidenced by subcategories such as *eating* (Schafer, 1977). This approach is useful for creating categorizations of sound for humans, but is less relevant to the categorization of sounds in soundscape ecology. In ecological terms, a biosemiotic approach that accounts for the interpretation of sonic signs by nonhumans in the ecosystem is crucial. A human semiotic category does not necessarily have any meaning to other ecosystem inhabitants. The study of soundscapes in an ecological context necessitates a new taxonomy.

2.2. The Krause-Gage Taxonomy

The field of soundscape ecology was not formally defined until a special issue of the journal *Landscape Ecology* in 2011, but the taxonomic framework used by soundscape ecologists has roots in much older work. From their introduction to the field, (Pijanowski, Farina, et al., 2011) use a tripartite system for grouping sounds by origin: geological, biological, and anthropogenic. These terms have roots in a 1997 book by Krause and a 2003 report by Krause and Gage.

Unlike Schafer, Krause and Gage are deliberate in the purpose behind the construction of these categories in their 2003 report. In an analysis of the relationship between acoustic indicators and ecosystem states, they point out that the “classification will assist in identifying the introduced elements that may cause stress or change not otherwise noticed by traditional visual evaluation”.

The notion of what is “introduced” provides a clear link to work in landscape ecology and conservation such as (Wiens, 2009). It sets up the later development of soundscape conservation (Dumyahn and Pijanowski, 2011), in which the authors discuss the need to preserve “natural sounds”.

For their part, Krause and Gage are clear about which sounds they consider to be natural: “In most environments today, soundscape signatures are comprised of two natural components, biophony and geophony, and a probable human component that includes the third, anthropony” (2003).

This characterization immediately raises some significant conservation questions. Are the biophonic sounds produced by introduced species themselves not considered introduced? Are they “natural”

in this analysis? What about the reproduction of biophonic sounds through loudspeakers and playback devices? If a recorded bird song is played over a speaker within their native range, is that considered biophonic or anthroponic? Is it something introduced or is it a part of the natural soundscape?

2.3. Hybrid Ecosystems

These are important underlying questions in soundscape ecology. There are objects and sounds that blur these lines between the categories in the Krause-Gage taxonomy. They complicate the identification of “introduced” and “natural” elements in the soundscape. These works challenge the understandings of soundscape and landscape conservation that underpin the basic sonic taxonomy of soundscape ecology.

Ecologists have already begun to acknowledge these new ecological coalescences. Hobbs, Higgs, Hall, et al. (2014) propose that, instead of describing them as natural or anthropogenic, ecosystems can be described as *historical*, *hybrid*, and *novel* to varying degrees. This reflects the understanding that few ecosystems — including vast tracts of forest that are often regarded as pristine, like the Amazon (Roosevelt, 2013) — are free of human influence and that the state of an ecosystem is in constant flux (Hobbs, Higgs, and Harris, 2009). The notion of a historical ecosystem considers whether an ecosystem is currently situated within a dynamic range that correlates with the past or whether human activity has fundamentally shifted the ecosystem’s composition.

An ecosystem where, for example, a particular species has gone extinct or a new species has been introduced is perhaps no longer historical, but could retain many of its historical dynamics with some human intervention. Then, it is considered hybrid: not exactly within its former range, but retaining much of its former character or significant properties. Novel ecosystems have undergone a wholesale change in composition — perhaps they have lost an irreplaceable keystone species or it is the site of a rehabilitated mine — so it may not be possible to reconstitute its former makeup.

Conservation priorities might be set differently for each type of system such that maintenance of a historical state may be the priority in a historical ecosystem, hybrid and novel ecosystems can be managed to conserve a general ecosystem function or a dynamic biodiversity. In the context of hybrid and novel ecosystems, anthropogenic sound — like other anthropogenic interventions — might play a role in supporting or reconstructing the ecosystem and its

soundscape. These types of interventions would be complex and fraught, but the anthropogenic sounds that could be used might not have the kinds of mostly negative connotations that Krause tends to associate with the anthrophonic sounds of car engines and airplane overflights.

2.4. Ecobots

What could produce the kinds of sonic interventions that might contribute to the soundscapes of hybrid and novel ecosystems? The introduction of a new biological species — either intentionally or unintentionally — would likely make an impact on the soundscape. But this discussion focuses on another potential sonic actor: robots.

In their (2017) paper, Wynsberghe and Donhauser discuss a category of robots that act in ecological manner which they call ecobots. These are robots whose design purpose is to perform an ecological function, such as helping to manage the proliferation of an overpopulated species or consuming toxins from a contaminated ecosystem.

The concept of ecobots is important here because they stand apart from what Wynsberghe and Donhauser (2017) call robots-in-ecology. These are robots that happen to be operating in ecosystems, but that aren't necessarily responsive to the ecosystem or performing an ecological role. Any sounds that these robots-in-ecology produce, such as the sound of the rotor of a drone that is hovering over an ecosystem, would fall into the same anthrophonic category as that of a ATV engine.

Ecobots, however, have a different relationship with the ecosystem. Their actions are responsive to changing conditions in the ecosystem. They become part of the ecosystem's flows and feedbacks in a manner that is in some ways similar to that of an animal. Their contributions to the soundscape are not quite the same as those of a robot-in-ecology. Their sounds carry different meaning.

3. ROBOTS IN THE SOUNDSCAPE

This section explores these provocations through a series of four case studies, the last of which is a work in progress by the author. Each of the cases discussed here has a digital component operating with some degree of autonomy, producing sound in an outdoor ecosystem and could be considered to be ecobots.

3.1. David Dunn

David Dunn is a composer and researcher whose work has often involved computer-mediated sonic

engagements with ecosystems. Two works are of particular interest here: *Sonic Mirror* from 1986-1987 and *Autonomous Systems* from 2003-2005.

Both works are attempts to engage soundscapes, using computers as part of a sonic feedback mechanism. Their roots can be traced to the rise of cybernetics and the availability of portable computing technologies, and they are steeped in cybernetic concepts such as feedback, complexity, and emergence.

Sonic mirror can be understood as an attempt to insert computational node within the larger ecosystem. "The original concept was conceived as a stationary cybernetic sound sculpture capable of processing acoustic data within an outdoor environment. Eventually the sculpture might function as an autonomous system structurally coupled to its surrounding environment in a manner that might allow for 'learning' between components" (Dunn, 2002).

Autonomous Systems is something of a follow-up to *Sonic Mirror*. In this work, the soundscape is recorded, processed, and replayed in a repeating, cyclical process. As the work records and replays its own shifting of the soundscape, the animals in the ecosystem engage it and shift their behaviour, and thus the behaviour of the machine as well.

These works are interesting as early examples of digital sonic systems operating in a feedback loop with a whole ecosystem. They bring a cybernetic approach to interventions in the soundscape of an ecosystem and, in doing so, create systems that are simultaneously reliant on human design and invention and bio- and geophonies. The sounds that they makes are not neatly human or non-human, digital or analog. The are ontologically uncertain.

3.2. Ian Ingram, The Woodiest (2010)

Artist Ian Ingram's *The Woodiest* (2010)¹ is an "auto-erotic, hermaphroditic, all-in-one woodpecker love-bot" according to the accompanying artist statement. It is inspired by the mating ritual of the Pileated Woodpecker, which consists of some initial drumming by an individual — a form of long-distance communication — followed by drum-tapping by the pair, once the second bird has been attracted to the site (Kilham, 1979).

Ingram (2010) notes that this appears to be akin to foreplay on the part of the birds; as such, *The Woodiest* is a robotic system that can engage in the entire act on its own. The robot listens for territorial drumming by a biological

¹The Woodiest documentation: <https://vimeo.com/16213036>

woodpecker and, upon hearing this drumming, the male artificial woodpecker subsystem responds with its own territorial drum. The female subsystem then responds with a courtship drum-tap, completing the encounter from a sonic perspective. The robot simultaneously engages under-explored areas of robotics for animal pleasure as well as robotic self-pleasure.

In video documentation of the work, the robot's drumming sounds quite similar to that of the woodpecker that it responds to, with the addition of the whine of the servo motor before and after the drumming and the layering of the vibration motor sounds on top of the rapping of the artificial beak on the hollow tree. The robot is notably unable to produce the higher-frequency vocalizations that the woodpecker demonstrates.

These differences mark the robot as an imposter to the human listener. But it is unclear if the woodpeckers — or other species in the ecosystem — experience a semiotic difference between *The Woodiest* and a mating pair of biological woodpeckers. This raises a question of audience for the taxonomic system under examination. Whose perception of the sound is important? In some sense, the sound is anthroponic. Ingram designed and built *The Woodiest* from human mined, fabricated, and assembled components. He decided the parameters within which *The Woodiest* would make noise and also chose to install *The Woodiest* on a particular tree in a particular forest at a particular time. And yet, *The Woodiest* makes no sound unless a woodpecker triggers it. The sound is necessarily collaborative. It is both bio- and anthro- and also somehow neither.

3.3. Richard Vitols, *Woodpecker* (2016)



Figure 1: One woodpecker from Vitols' *The Woodpecker* (2016). Used with permission of Richard Vitols.

Ingram's are not the only artificial woodpeckers to be found outdoors. The *Woodpecker*² is a series of

²Woodpecker documentation: <https://vimeo.com/180702278>



Figure 2: A broken woodpecker from *The Woodpecker* (2016). Used with permission of Richard Vitols.

30 artificial woodpeckers, installed in a forest near Dusseldorf in 2016 by artist Richard Vitols. The titular woodpeckers are, according to the artist, an attempt to rescue the local trees from a coming infestation of insects.

The inspiration for the work emerges from the intersection of two phenomena: the first is the ability of insects to respond to the sound of woodpeckers and the growing detrimental effect of insect infestations on the global forest canopy; the second is the projected decline of bird populations. This led Vitols to wonder if a robotic woodpecker impersonator might be able to help the trees survive an insect onslaught.

The woodpeckers themselves are simple robots, constructed from a basic microcontroller, a solar panel and battery pack, and a solenoid (linear) motor that hits the tree to produce the woodpecker sound. They were installed in the forest for four weeks, knocking on trees to produce pecking sounds when their batteries were charged. They were examined once per week and some that were broken — such as the one in Figure 2 — were removed and retired.

The solenoid motors don't produce the rapid beats of the vibration motors in Ingram's work. But these artificial woodpeckers aren't attempting to perform a woodpecker mating call. They are simulating the sounds of a woodpecker foraging and the sequential hammering they produce is difficult to distinguish from an actual woodpecker.

Woodpecker is unique, in part, due to its materiality. Few roboticists discuss their work in terms of breakdown and decay, yet Vitols takes explicit note of the robots that have broken or been destroyed (see Figure 2). This points to two different temporal cycles in digital producers of sound. There is the working lifecycle and the material lifecycle. The working lifecycle can be as short as a few days or weeks in

this case, but the copper and silicon that form the electrical circuits that power the robot won't decay for many biological generations.

At the same time, these digital actors are unable to self-repair or reproduce. They have no way of continuing their existence beyond their individual bodies. Without human intervention, their direct temporal mark on the soundscape is fixed in a manner quite distinct from the biological and geological sources of sound.

3.4. The rowdy krause (2019-)

The rowdy krause is a work-in-progress to evolve a voice for a new endemic species within an ecosystem. The robot's vocalization is based on Krause's (1987) acoustic niche hypothesis (ANH), which contends that animals differentiate their class across the frequency spectrum in a manner analogous to the differentiation of resource use in the traditional understanding of ecological niches.

The rowdy krause begins by listening to the existing soundscape. It listens in particular for interesting sounds that rise above the background din of the soundscape. Within those sound events, it pays attention to which parts of the frequency spectrum are most heavily used and which are left open for a new species to occupy.

In order to encourage the emergence of a complex, but biologically-plausible voice, the robot uses a mammalian vocal tract simulator to produce sound. It is able to control the voicebox, throat, tongue, lip and nasal cavity of a virtual vocal tract called Pink Trombone³. This creates a context for the robot to produce sounds that could be biological in origin.

The vocal tract is controlled by a neural network that is evolved using neuroevolution of augmenting topologies (NEAT) (Stanley and Miikkulainen, 2002). NEAT allows for the evolution of both the structure and the weighting of the neural network, so that it can add complexity as needed. The neural network takes a representation of the last sound that it made as input — so that it can “hear” itself — and outputs the next shape of the trachea, epiglottis, tongue, and lips, as well parameters to control the nose and voicebox according to the current network structure.

Once the neural network-controlled vocal tract produces a vocalization, the NEAT evolutionary process assesses how well this sound fits into the unused frequencies in the soundscape. Over time this process is repeated for many different evolved voices. Neural networks whose voices better match the predictions of the ANH are selected more often

³Neil Thapen's Pink Trombone: <https://dood.al/pinktrombone/>

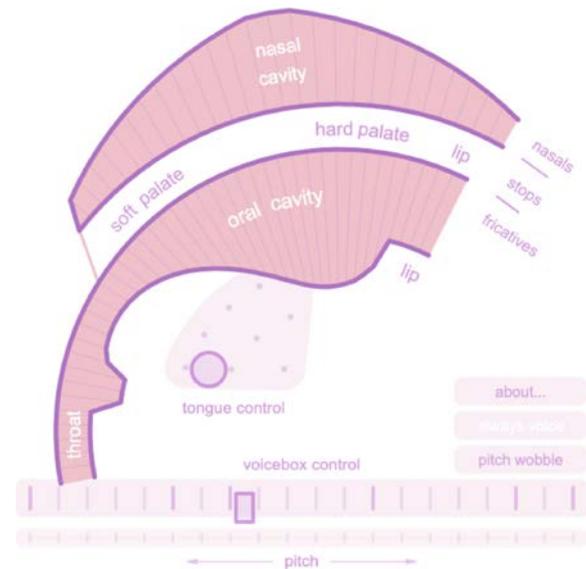


Figure 3: A screen capture of the Pink Trombone vocal tract simulator. Used with permission of Neil Thapen.

for reproduction in the evolutionary process and are more likely to pass their traits to the next generation of neural networks. This leads to a series of vocalizations that make increasingly better use of the least used frequency bands in the soundscape⁴.

The rowdy krause continues to listen to the soundscape as it shifts in response to changing seasons as well as the presence of the new robotic voice. It continually incorporates what it hears into its understanding of the sonic structure of its ecosystem, so that the vocalizations that it produces are always shifting in response to changes in the soundscape.

Currently, the rowdy krause has been tested in the lab using recorded soundscapes to drive the evolutionary process. However, field trials are planned for a peri-urban forest garden ecosystem, as well as a set of urban garden environments.

The rowdy krause represents another dimension in the discussion of robophony. Unlike Dunn's work, it is not a reprojection of past sound events back into the soundscape. It is also not an analog percussive sound like the two artificial woodpeckers. This voice is entirely synthesized, and yet it maintains some of the characteristics of biophony. Its generative frameworks — the vocal tract, neural network, and evolutionary process — reference biological objects and phenomena. It is also explicitly responsive to the soundscape in a way that mimics the theoretical response of a biological species.

⁴A demonstration of this process is available online at <https://vimeo.com/359044847>

The rowdy krause is clearly not biological. It has a human builder and programmer and can be installed or removed on a whim. Its actions are bounded by human decisions about the constraints of the vocalization system and the neural network that controls it and yet it is relatively free from human control within those bounds.

Because its calls are new to the ecosystem, the biosemiotic significance of them is unclear. It is possible that the sounds it produces, in the perception of some of the inhabitants of an ecosystem, is close enough to that of a potential predator, mate, or prey to carry some meaning. But the human semiotics of the rowdy krause are equally unclear. Is it considered noise on par with that of an engine? Or is it more akin to bringing a dog into an ecosystem?

4. ROBOPHONY

The four cases discussed in the previous section are distinct from one another, but each points to a type of hybridity that exists somewhere between the categories of anthrophony and biophony. Sonic Mirror and Autonomous Systems introduce the notion of a dynamic soundscape component that uses a system of feedback between anthropogenic and bio-geological actors to produce a dynamic sonic component that is hybrid in origin. The Woodiest's sonic emanations raise questions about the biosemiotics of robotically-produced sound. Woodpecker adds to this the temporal dynamics of a soundmaking technical object and distinguishes these robots further from their biological analogs in their inability to heal and reproduce and their sensitivity to the element and predation. Finally, the rowdy krause prompts the review of questions about the origins of biologically-informed, computationally generated sound and the semiotics and biosemiotics of an introduced digital species.

The hybrid nature of these objects points to the need for a new category of in the taxonomy of sounds in soundscape ecology: robophony. Broadly, robophony is the set of sounds produced by ecological robots in a soundscape. Informed by the cases in the previous section, it encompasses the following distinguishing features: site specificity, hybrid sourcing, layered temporalities.

4.1. Site Specificity

The notion of site specificity is borrowed from mid-21st century art and is often — though not always — invoked in the context of monumental land art works such as Christo and Jeanne-Claude's *Surrounded Islands* (1983) or the many ecological interventions of Helen and Newton Harrison. Here, site specificity

refers not to the physical object itself, but to the particular sounds that it produces.

In fact, the sound-producing objects themselves in the four cases are for the most part not specific to a particular site. Dunn's computer equipment for Sonic Mirror and Autonomous Systems could be set up almost anywhere, as could the physical body of the rowdy krause. The two woodpecker-based works are designed to sit on a tree, but they were not built for the particular trees — or even the particular species of trees — that they ended up on. They feature adjustable straps so that they can be mounted on a tree of the artist's choosing after the fact.

However, the sonic output of all of these works is a result of the particular configuration of their surroundings. Dunn and Kadish's work are most explicit in this regard. Dunn's work uses the live soundscape as the raw material for the computational modulation and reprojection. The soundscape is part of the cybernetic system that produces new sound. It does not exist without the specificity of its place. In the rowdy krause, the existing soundscape is an extra step removed from the eventual sound that is produced, but it is no less a part of the work. The rowdy krause's voice is evolved to fit the ecosystem's soundscape and is a result of the combination of existing biophony, geophony, and anthrophony of the site.

These works could inhabit another site, but they would at that point be new works. Their voices would be different, would evolve differently, and their robophonic contribution to the soundscape would be change.

This holds to a lesser, but not inconsequential degree for The Woodiest and Woodpecker. The Woodiest's sonic projections are a product of its sight insofar as it exists in dialogue with resident woodpeckers. Its pecking sequence is activated only in response to another woodpecker's call. As such, the particularities of its actions can be understood as the response to a particular community.

Woodpecker is not responsive to external sounds, but its own sound — like that of The Woodiest — is not entirely of itself. The source of sound in both of these works is the percussive action of the robot meeting the tree. In fact, the main source of sound in both of these robots is the vibration of the tree or branch. Its materiality, its age, texture, and solidity, all determine the resulting sound. These sounds are specific: to the tree, to the branch, to the meeting of biological matter and technological artefact.

As devices built by humans, it is tempting to categorize the sounds that these works produce

as anthrophony. They would likely be understood as electromechanical or controlled sounds under Krause's taxonomy (1987), both subcategories of anthrophony. But site specificity is a trait more often associated with biophony. In fact, Krause's acoustic niche hypothesis (1987) along with another foundational theory of soundscape ecology — the acoustic adaptation hypothesis which concerns the adaptation of animal calls to a particular physical environment (Morton, 1975) — explicitly frame biophony as site specific. Neither of these categories are quite sufficient to capture the acoustic impact of these new ecological actors. Robophony is therefore positioned as originating from devices constructed by humans, but with sound that it embedded in a local context and specific to a place.

4.2. Hybrid sourcing

The Krause-Gage taxonomy is source-based — it categorizes sounds according to the object that produces the sound. However, the source is not always clear. When a recording of a bird call is played back into an ecosystem, is the source biological or technological and therefore anthropogenic?

This question implicitly refers to a phenomenon that Schafer (1977) terms *schizophonia*. The sound is separated from its source, creating a disconnect from its original location and meaning. If a bird produces an alarm call, it signifies that there is danger for themselves, other members of their species, and possibly for others as well. When that call is recorded and played back, the sign has — in biosemiotic terms — lost its object (Emmeche, 2007).

This describes the situation of *The Woodiest* and *Woodpecker*. Though neither work features the playback of recorded sound, both sign the presence of a woodpecker through sound without the actual presence of that woodpecker as an object. *The Woodiest* signals a biological woodpecker that a mating pair is present and that this is their territory, but were the bird to claim the territory for themselves, there would be no consequence normally associated with infringing on a fellow woodpecker's territorial claim. *Woodpecker* similarly alerts insects to the presence of a predatory woodpecker without the possibility of capture and consumption.

The origin and biosemiotics of the other two cases are perhaps even more confounding. In *Sound Mirror*, *Autonomous Systems*, and the rowdy krause, the existing soundscape in an ecosystem acts as the base material for the sound that is produced. Dunn's work processes that soundscape and reprojects the result back into the ecosystem. The sound is partially bio-, geo-, and anthrophonic in origin, depending

on the composition of the soundscape that the works record during their operation. But there are another set of anthropogenic forces that act on it: the hardware system of computer, microphone and speaker; and the software system of algorithms that modulate and process the incoming sound. It cannot be said to be purely anthrophonic, nor is it bio- or geophonic.

The rowdy krause evolves something to actively differentiate itself from what it hear — but the material from which it differentiates itself is likely a mix of biophony, geophony, and anthrophony, depending on the ecosystem that it inhabits. It continues to evolve its voice in relation to the soundscape as the soundscape shifts in response to its presence. This feedback is most likely to occur in the domains of bio- and anthrophony as geological sounds are not likely to change in the short term in response to the rowdy krause. This means that the call of the rowdy krause is also hybrid in origin. It may be primarily generated by the anthropogenic computational system, but it is a result of all three of the Krause-Gage phonic taxa.

For all of these works, the framework for their sounds are set by their human designers. But they would not exist or would exist differently without the bio- and geophonic components of the soundscapes that they inhabit. Their sources are distributed between biological, geological, and anthropological sources and they therefore fit neatly into none of those categories. This points to the position of robophony as a set of sounds with hybrid sources that fail to fit neatly into the three other categories.

4.3. Layered temporalities

The hybrid sourcing of robophonic sound also gives way to a complex network of temporalities. It is important to understand the dynamic timescales of each of the sources of sound in a soundscape, because these timescales impact the rate of feedback between sources.

Patterns of geophony tend to shift on geological timescales — though anthropogenic activities are accelerating these as well — meaning that the sound of rain, for example, is not immediately responsive to most biological changes in the landscape. Biophony and anthrophony operate on a range of timescales. Evolutionary changes operate on relatively slow cycles that are depending on the lifespan and mating rates of the species in question. But behavioural changes can occur much more rapidly.

Computational timescales can be sped up even further, calculations taking place at the microsecond scale, and rendered without a perceivable delay

for many biological audiences. The material flows of computational technologies, however, tend to operate on more geological timescales. Biomass breaks down quickly, but silicon and copper remain in the landscape for many biological generations.

Sound Mirror and Autonomous Systems perform this kind of near-instantaneous computational temporality. But their performances also induce rapid — yet still comparatively slow — engagement from nearby biological actors. The performance mixes the instantaneous computational timeframe with the behavioural timeframe of biological entities.

The rowdy krause does not directly incorporate existing sounds into its own vocalization, but it does perform computational evolution at a pace far faster than biological evolution is able to act. In its offline version, it can simulate multiple generations — birth, life, mating, and death — of about 20 individuals in under 15 minutes. The version that will eventually inhabit an outdoor ecosystem will evolve more slowly than this, but still at a pace unavailable to biologically evolving species.

Woodpecker serves as a reminder that although robotic material flows are geologically slow, their functional flows can still be quite rapid. Figure 2 shows a robotic woodpecker that was destroyed by the elements after less than four weeks in operation. Robots like these are sensitive to moisture, cold, and animal attacks, meaning that although the materials may persist for many generations, they may not contribute to the soundscape for very long.

4.4. Future robophonies

The four cases that are presented above together make a case for the addition of robophony to the taxonomy of sound in the soundscape. But there are plenty of examples of technologies that do not quite qualify as robophony at the moment, but could in the near future.

Many of these are sold as sonic pest control devices. They play high pitched sounds to deter bats, mice, and mosquitos. One plays tones in the 400-1000Hz range in an attempt to drive moles away. Currently, these devices typically lack a sensory system to determine their effect on the ecosystem and whether they are “working” or when their target is nearby.

However, it is not unreasonable to imagine these things getting smarter. As the harmful effects of indiscriminate sonic output are better understood one could imagine that instead of constantly playing a static tone to deter rodents, they might begin to play varying but relevant sounds only when they detect a rodent. This type of feedback between the

rodent and the operation of the sonic agent could qualify it as robophony.

5. CONCLUSIONS

Taxonomies are inherently flawed, but they inform our understanding of a system and its component parts. The categories of sound in a soundscape have served soundscape ecology well thus far, enabling discussion about the composition of a soundscape and the impacts of human activity on the nonhuman inhabitants of an ecosystem.

But the cases presented here demonstrate that these categories are insufficient to capture the dynamics of a growing number of biologically interactive technological agents. They are human-made, but not human-driven. They are temporally complex, and they produce sound specific to their adopted habitat.

These agents, entities, actors in the soundscape require a new descriptor, one that captures these properties and, in doing so allows for the consideration of the ethics, biosemiotics, and new feedback loops that accompany their presence. This descriptor is *robophony*.

6. REFERENCES

- Dumyahn, Sarah L. and Bryan C. Pijanowski (Nov. 2011). “Soundscape conservation”. In: *Landscape Ecology* 26.9, pp. 1327–1344. ISSN: 0921-2973. DOI: 10.1007/s10980-011-9635-x.
- Dunn, David (2002). *Four Electroacoustic Compositions*.
- Emmeche, Claus (Dec. 2007). “A biosemiotic note on organisms, animals, machines, cyborgs, and the quasi-autonomy of robots”. In: *Pragmatics & Cognition* 15.3, pp. 455–483. ISSN: 0929-0907. DOI: 10.1075/pc.15.3.06emm.
- Gasc, Amandine et al. (Jan. 2017). “Future directions for soundscape ecology: The importance of ornithological contributions”. In: *The Auk* 134.1, pp. 215–228. ISSN: 0004-8038. DOI: 10.1642/AUK-16-124.1.
- Hobbs, Richard J., Eric Higgs, Carol M. Hall, et al. (2014). “Managing the whole landscape: Historical, hybrid, and novel ecosystems”. In: *Frontiers in Ecology and the Environment* 12.10, pp. 557–564. ISSN: 15409309. DOI: 10.1890/130300.

- Hobbs, Richard J., Eric Higgs, and James A. Harris (Nov. 2009). "Novel ecosystems: implications for conservation and restoration". In: *Trends in Ecology & Evolution* 24.11, pp. 599–605. ISSN: 01695347. DOI: 10.1016/j.tree.2009.05.012.
- Ingram, Ian (2010). *The Woodiest*. Tech. rep. www.ianingram.org.
- Kilham, Lawrence (July 1979). "Courtship and the Pair-Bond of Pileated Woodpeckers". In: *The Auk* 96.3, pp. 587–594. DOI: 10.1093/auk/96.3.587.
- Krause, Bernard L. (1997). *Into a wild sanctuary : a life in music & natural sound*. Berkeley Calif: Heyday Books, p. 200. ISBN: 9781890771119.
- Krause, Bernie L. (1987). "Bioacoustics, Habitat Ambience in Ecological Balance". In: *Whole Earth Review* 57, pp. 14–18. ISSN: 0749-5056.
- (2008). "Anatomy of the soundscape: Evolving perspectives". In: *AES: Journal of the Audio Engineering Society* 56.1-2, pp. 73–80. ISSN: 15494950.
- Krause, Bernie L. and Stuart Gage (2003). *Testing Biophony as an Indicator of Habitat Fitness and Dynamics*. Tech. rep. Glen Ellen, CA: Wild Sanctuary Inc, p. 18.
- Morton, Eugene S. (Jan. 1975). "Ecological Sources of Selection on Avian Sounds". In: *The American Naturalist* 109.965, pp. 17–34. ISSN: 0003-0147. DOI: 10.1086/282971.
- Pijanowski, Bryan C., Almo Farina, et al. (Nov. 2011). "What is soundscape ecology? An introduction and overview of an emerging new science". In: *Landscape Ecology* 26.9, pp. 1213–1232. ISSN: 0921-2973. DOI: 10.1007/s10980-011-9600-8.
- Pijanowski, Bryan C., Luis J. Villanueva-Rivera, et al. (Mar. 2011). "Soundscape Ecology: The Science of Sound in the Landscape". In: *BioScience* 61.3, pp. 203–216. ISSN: 1525-3244. DOI: 10.1525/bio.2011.61.3.6.
- Roosevelt, A.C. (Dec. 2013). "The Amazon and the Anthropocene: 13,000 years of human influence in a tropical rainforest". In: *Anthropocene* 4, pp. 69– 87. ISSN: 22133054. DOI: 10.1016/j.ancene.2014.05.001.
- Schafer, R. Murray (1977). *The tuning of the world*. A.A. Knopf, p. 301. ISBN: 0394409663.
- Stanley, Kenneth O. and Risto Miikkulainen (June 2002). "Evolving Neural Networks through Augmenting Topologies". In: *Evolutionary Computation* 10.2, pp. 99–127. ISSN: 1063-6560. DOI: 10.1162/106365602320169811.
- Vitols, Rihards (2016). *Woodpecker*. Tech. rep. vitols.xyz.
- Vitols, Rihards and Regine Debatty (2016). "The Woodpecker: Could fake birds save our forests?" In: *We Make Money Not Art* August.
- Wiens, John A. (Oct. 2009). "Landscape ecology as a foundation for sustainable conservation". In: *Landscape Ecology* 24.8, pp. 1053–1065. ISSN: 0921-2973. DOI: 10.1007/s10980-008-9284-x.
- Wynsberghe, Aimee van and Justin Donhauser (Oct. 2017). "The Dawning of the Ethics of Environmental Robots". In: *Science and Engineering Ethics*. ISSN: 1353- 3452. DOI: 10.1007/s11948-017-9990-3.