

Collaboration and Consensus in Listening

ANNA BARNEY AND SALOMÉ VOEGELIN

ABSTRACT

This article adapts a conversation on a network project, *Listening across Disciplines*, which brought together artists, musicians, scientists, technologists and social scientists to discuss the use, value and application of listening as a shared methodology of inquiry and communication. The discussion focuses on one of the key issues emerging from this network: the question of consensus and collaboration in the development of a shared listening methodology.

This conversation brings to public consideration one of the key discussion points of the AHRC (Arts and Humanities Research Council, U.K.)-funded network project *Listening across Disciplines* <www.listeningacrossdisciplines.net>: the issue of consensus and legitimacy about the listened to and the heard, which touches the core of a cross-disciplinary engagement in listening.

Anna Barney, professor of biomedical acoustic engineering at the University of Southampton, and Salomé Voegelin, Reader in Sound Arts at the London College of Communication, University of the Arts London, led the project, which brought together artists, musicians, scientists, technologists and social scientists, as well as scholars and practitioners from the humanities, to discuss the use, value and application of listening as a methodology of inquiry and communication. The network investigated the aesthetic, analytical, data gathering and diagnostic function of listening and debated the legitimacy and evaluation of the heard in different disciplinary contexts. It asked: How might we be able to work collaboratively on the production of sonic knowledge and its application? What faith have we in reliability and consistency of sound, what scientific gains can we make if its unseen terrain achieves legitimacy and what aesthetic and conceptual possibilities then become accessible and thinkable?

Anna Barney (biomedical acoustic engineer), Institute of Sound and Vibration Research (ISVR), University of Southampton, Highfield, Southampton, SO17 1BJ, U.K. Email: <ab3@soton.ac.uk>. ORCID: 0000-0002-6034-1478.

Salomé Voegelin (artist, writer), Creative Research into Sound Arts Practice (CRISAP), London College of Communication, University of the Arts London, Elephant and Castle, SE1 6SB, U.K. Email: <s.voegelin@lcc.arts.ac.uk>. ORCID: 0000-0002-8970-3997.

See <mitpressjournals.org/toc/lmj/-/28> for supplemental files associated with this issue.

To bring into discourse a key finding of the network, this article takes the form of a conversation between the two lead researchers, conducted on their way home from presenting their ideas on listening as a democratic competence at the 4th Council of Europe Platform Exchange on Culture and Digitisation in Karlsruhe in October 2017. This conversation focuses on the status and importance of consensus in science and art, where listening and sound, as invisible modalities, highlight the fragility and perhaps the impossibility of agreement. The issue of establishing consensus seemed to be at the core of our cross-disciplinary network discussions: more specifically, the need on the side of the scientist to find consensus and the desire on the side of the artist to disrupt and question its possibility.

Anna Barney: I think what came out of the network meetings was the idea that, in science, the goal is to reach a consensus. When you use the visual language of science—graphs and so on—there's an agreed grammar of the visual: that time goes from left to right, the thing you're measuring goes on the vertical axis and so on. That allows us to share information through the visual that's not just the pictures of things but the trend in things, the way things are behaving. So when my colleague and I look at a graph, we should see the same things because there's this consensus. A lot of what we have to do is show our data in prescribed ways in order to ensure we're talking the same language. From the network discussions, it seemed that in art, other people taking the same meaning from what was happening was the least of your concerns.

Salomé Voegelin: Very often, in art it's about value rather than meaning. It's about whether a work is considered to be a strong artwork, how it affects an aesthetic understanding and how it's evaluated in discourse. But it's also about ambiguity, actually not understanding it in that sort of scientific way but being drawn into it, into its inarticulate and unresolved material and speculative presence and what it might produce from there. So there's a consensus that it is strong as a work of art, but there's not consensus on what it may mean. There isn't a resolution.

AB: When I draw a graph, I want to evoke in the other person the same thing it evokes in me, in terms of an understanding of the data—it goes up here, it goes down here, there's a peak, a trough—I want them to see that and come to the same conclusion about what the data's doing as I have.

SV: In art, at least in post- or post-postmodern works, I don't think that's the case. In modernism you still had the idea that there was an agreed-on and singular interpretation that told you as a viewer how to view this painting or sculpture and therefore how to read it and understand it. But this understanding is not only held within the work, it is also held within who the artist is, her other works and the context of the work: whether it's in the Tate or it's in an independent gallery or whether it's in somebody's living room. All these aspects make the experience and reading of the same work entirely different. I guess the graph, the scientific graph, is the same whether you look at it in your bathroom or in your laboratory.

AB: It should be independent of who produced it, but it might not be independent of where it's published because you have the reassurance of a good journal having done a proper review process and somebody having looked at that graph, so when you say, "See, this data increases, decreases, changes in this way," they will have agreed before it's gone to publication. But if we want to achieve this in sound, then for sound to work for science in a similar way, something like sonification has to have an agreed grammar.

SV: What you just said in relation to the context does play a big role in art as well, if it's reviewed in a respected journal or shown in a well-known gallery or museum. I'm more inclined to follow the invitation of the work and engage in its expression. In a sense, I'm going to trust it to be a work of art and to be potentially a good piece of art that I give my time to. I don't have to like the work—art is rarely about liking—but I have to want to engage in it. The assessment in the journal could still be wrong, and it could be verifiably wrong—somebody could say, "No, I'm going to write a counter-review"—but there is a trust that if something is reviewed in *frieze* magazine, say, you have a sense of its worth, and think: I'm going to take this seriously.

AB: In science, if it's in a good journal, you can't be sure if it's a true description of the world, but you can probably be sure that it's self-consistent and that it doesn't contradict any of the major known theories that have already been accepted.

SV: But it isn't future-proofed, so in the future it could be still contradicted.

AB: Yes, and it could certainly be a simplified representation compared to what you might come up with in the future.

SV: Until now we've been talking generally about art and science, but the aim of *Listening across Disciplines* is to really cross them over and say: If we take an artistic sensibility of

sound into a scientific knowledge frame and insist the listened to has to be as trustworthy, as reliable and as repeatable as visualizations, we need a consensus, a shared frame of interpretation and reference.

It was interesting when I talked to Steven Goldfarb at CERN in Geneva to discover that they use a lot of visualization mechanisms that do interesting movements, or visualization graphs that show events, and all sorts of other visual displays of what goes on invisibly in the Large Hadron Collider. To me it appeared as if they were looking at graphic scores.

AB: I think that's interesting. They are probably looking at Feynman diagrams—certainly something of that sort; these were only relatively recently invented by Richard Feynman, probably in the late '40s. They were something Feynman, who was a very eminent physicist in the early subatomic particle era, invented for himself to visualize how these collisions were working; before that, those diagrams didn't exist, so in about 70 years' time they've become an absolutely standard tool and the consensus is there.

SV: So it is not that our ability to read visual information is innate. We learn it at school or at university: We want to catch a train, so we learn how to read a train timetable; we learn how to read a *Financial Times* graph. We have to learn how to read data in that way.

AB: I suppose we choose visualization methods that somehow give us some intuition into the data we're looking at. So it's hard to say how much that language is innate and how much the language you speak is a function of how your brain's structured. You probably pick diagrams that seem to speak in some way to your intuition about the data or where, as a person drawing the diagram, you feel the construction of the diagram is in accord with your understanding of what's happening.

SV: Which nevertheless seems quite cultural, and so if we had a listening education, at university or even at school level already, we could listen to sounds, not just for their semantic information—this is the sound of the train—but in a much more discrete and autonomous way. To hear and interpret frequency changes, repetitions, intensities, etc. could be used scientifically maybe just as accurately as a visualization.

Therefore, it's about developing audification and sonification, charts of sound that could become as useful as visual ones, and then having the commensurate training and the experience to understand what is heard—not as the sound of a thing, but as sound generating insight into a process. Nobody says, "I wouldn't know how to look, therefore I can't use these visualization devices," but we're quite happy to say, "I wouldn't know how to listen."

AB: You presumably could construct an aural grammar that did the same work for you. You would need to teach people what they were listening for and you would need to train their ability to discriminate aurally, as you train a musician to be able to pitch a musical interval.

SV: Yes, you still have to have at least a linguistic equivalent, a schematic, of “this is this” and “this means that.”

AB: What I would hope to avoid is simply mimicking the visual by changing some aspect of the sound. We wouldn't have achieved much if all we did was increase the pitch as the graph went up, decrease the pitch as the graph went down—that would just seem to me to be another way to access the same.

SV: I think that represents, for me, exactly the problem with sonification. Visualization goes back to the thing—what's there, what's happening in the scientific experiment—and then visualizes that event from there. Whereas sonification very often happens from already translated data sets; therefore, it is already engaging in an interpretation. I would be interested in starting with “what's there, what's happening” in sound also, and for this we need to start the data gathering in sound, from which we can then develop systematized interpretations and “sonic graphs.” In your case, where you listen to lungs, we can actually use that moment of listening to develop a sonification understood as a repeatable representation of the heard, instead of relying on visual graphs that tell us that this lung is diseased. The aim would be to find a trust in listening to the lung, to the “what is there” of the sonic, to establish a consensus of how to represent and read that listening.

AB: Listening to the lung directly—I think you would find that the clinicians have a trust in that; they hear the sounds that I try to detect automatically and they're confident they are hearing wheezing or breath sounds. What they're not confident of is that they can get a reliable diagnosis from that sound, because sometimes similar but different diseases produce the same sound. When we do a visual analysis of that, we can use something called a time-expanded waveform: We plot it and then we make it much bigger and stretch it out in time. What we don't usually do then is play it back in that slow time, but that's something that you could think about—not just looking for the features that you see in a visual expansion, but playing it more slowly, more quickly to get a sense of quality of the sound.

SV: I think that's what I mean and hope for. It would be fantastic to have a listening education, so people would feel as literate and as confident with sound as with the visual and could make and “read” sonic/sonification data and find consensus on its meaning. Because that would obviously be like listening in the scientific arena. But I also think I'm interested in bringing artistic ambiguity and doubt into the verifiable, bringing the sensibility and concepts of the invisible and the contingent to scientific reliability. I mean, my best example in what remains inexplicable when you insist on measurements is this: When you have noisy neighbors and you complain about them, and then the noise pollution people come and they measure that sound in decibels and say, “No, this isn't loud enough for us to register a complaint, this isn't annoying because my reading of it, in decibels,” meaning it's not an-

noying. But you know it is, because to you there is something inexplicable about what's happening, why you can't live with it—but that doesn't seem to count.

It's not only the sound as a measurable entity, it's the experience of the sound that falsifies the reading and brings your body into its measurement. It brings your body in and that means that in a certain sense it becomes multisensorial and it also becomes about you, about your subjectivity. And I suppose that's where the consensus becomes very difficult, because if you love that sound, you think: This is wonderful, I like this, I can stay here a while; but if you don't like it, the very same sound becomes something very different, because your attitude and your sense of it is different. And obviously, scientifically it's not very valuable if you listen to patients and you don't like them, or you don't feel comfortable that morning, or the room you're listening in is different and annoys you, it's cold and you hear something else; that sort of enters the processes of interpretation and rattles the consensus. But I'm sure this subjective and contingent awkwardness is still there in science, is it not?

AB: Well, of course there's an advantage to recording the sound, which is you can take it away and listen again somewhere else. But the traditional auscultation isn't recorded, listeners just get what they get on the day; if they're not feeling well, if they're in a noisy environment, if they're stressed, then it will be a different thing. Something that's interesting is what you could add to the measurement of sound. I'm already recording sound, so I could listen to it and try to analyze it from just listening. In fact, that's what the clinicians do, they don't record but just try to analyze from listening. But something we happily do for the visual but not the aural in science is to rearrange the information, to look from a different perspective. Maybe similarly you could put all the loud sounds first and all the quiet ones later or something; that way you could reconstruct the signal differently.

SV: So saying that, there is almost an admission that even within the medical listening data there is a potential for “composition,” a subjective intent and manipulation, and that depending on what order things are in, I will hear them differently?

AB: Yes, definitely.

SV: That maybe is another reason why listening isn't scientifically trusted, because it has that potential to become ambiguous, changeable. I assume if I change the order of my recordings of the very same lungs, I might interpret something differently.

AB: Yes, except that I don't think anybody does change the order for scientific listening; you just record and then you think of it as a realistic thing, so you play it back. You're always trying to get as close as possible to the real sound, in some sense. You try to have the best possible stethoscope and recording equipment, the highest sampling rate, so that

what you get later is exactly what you got at the time. In fact, we might be able to manipulate the recording to give more information.

SV: This is such an interesting point, because in a way we're talking about field recording. Your field recording might happen on the body rather than in the geographical or urban field but it seems to still be a field recording, which could potentially share a history with phonography. For the Naturalists at the beginning of the twentieth century, they felt they must have the best equipment, because there is an authentic world out there that we can capture. I think that within art this has been critiqued and developed. Contemporary field recording tries to acknowledge that if I go in the field, I'm also in that field as a recordist, and my presence will change that field. I make decisions as to what microphone to use, as to where and how I walk, where I go and, implicitly or explicitly, I'm in the recording with my body and all my decisions. And of course, because it's not really there as an authentic recording but rather there is an acknowledgment that there is a composition going on, there is a recomposing of the authentic world going on, and maybe there is no authentic world. So there can't be a consensus of how it is experienced. Is, then, listening in science just ignoring that component and thinking, well, if I have the best of what is currently available and considered the best, I will reach the real world?

AB: The clinician as listener hears sound generated in and transmitted through the body of the patient, so the clinician's desire might be to get the most authentic, high-fidelity reproduction of what they heard at the time. The scientist, in some senses, wants to know more how that sound is produced, and because of that, they're trying to get rid of the body of the patient and to go back to the sound source. If you take a crackle at one of the bronchioles, it's probably like just a pop inside the lung, but because that sound then travels through the lung and other tissues, what you hear on the surface is a sort of smeared version of that pop—and that's what the clinicians try to reproduce. But the scientists try to reproduce the original pop by taking out the ribs and the lung tissue, the airspace, the fat of the person, to get back to the mechanism that produced it. So there are already two different "authentic" sounds we might be aiming for.

SV: But if we wanted to establish a consensus of listening, would we have to start writing in scientific journals about these bodies and about what is being bracketed out in order to get the authentic sound?

AB: I think you would have to be very clear in any sort of sonification process what sort of information you are trying to elicit. Because the point about recording the data is then to extract not a truth so much as a set of useful information, in scientific terms. For the clinician, the sound on the surface that they measure with the stethoscope is what they're trying to extract, and from that they're trying to take the information about whether sound occurs in the inspiration or the ex-

piration—what kind of sounds, normal or abnormal? That's one set of information; the scientist is more trying to extract the original sound, perhaps. The clinician wants a diagnosis of the disease; the scientist might want to understand and model the sound-producing mechanism.

SV: And does the scientist always already know what she's listening for?

AB: No, because it's not something we've ever heard; you can't get inside somebody's lung and listen to it.

SV: So there is an experiment . . .

AB: There's a model.

SV: There is a model, but is there the time and the leisure to just say, "Well, I'm just going to do field work on lungs and see what I can hear?" Is there the leisure and the kind of attitude for that, like now I think, maybe if we just give it a go, we might hear new stuff we never heard before?

AB: It's, relatively speaking, easy to develop a research project where you take people who've either got a disease or been sent for a treatment and to measure them as they arrive. It's very difficult to find an appetite for healthy-study funding to know what normal lungs sound like. It's just harder to find people who are interested to support that kind of research.

SV: As a field recordist you can just say, "I'm going to go to Cardiff," or wherever you'd want to go, "and I'm going to listen to these areas, and they're there for me to listen to; I don't need their permission, I need to have equipment, I need to have time, I have to have money to fund myself, but I can go and listen and see what I hear."

AB: In order to do scientific research on people you have to have an ethical approval for it, and one aspect of that is to ensure that you don't take a set of people and experiment on them for no good reason. You have to show there might be a benefit, and it's easy to show that if you measure people with a disease, you could point to a better outcome—well, perhaps it's not easy to show, but if you put forward a clear benefit to the ethics panel it may be shown that the subjects will be inconvenienced by a slightly longer appointment, but that may be balanced by, for example, quicker diagnosis, better treatment at reduced cost.

SV: So you always have to already have an aim.

AB: Yes.

SV: This I think is one of the fundamental differences. Art is often purposefully aimless and celebrates the nonfunctional and even the unreasonable.

AB: I think science is never meant to be purposeless.

SV: So I suppose talking about field recording as a field recording of the body implies making compositions, really divorcing yourself from the demand of authenticity, because you are creating a new world, a new experience in sound. This is a bit like when you do film sound; when you want to represent angry door slamming, what you hear is most often not a door being slammed but a sound that is constructed to give the impression of a door being slammed angrily. Maybe it is just a plank of wood being smashed, or even less-related Foley sound. But due to the context of the narrative and the visual, we believe it. So we're creating a consensus that a door is being slammed, not through recording the authentic, but through artifice; we're creating rather than recording worlds. And so it's almost like agreeing on a possible world rather than on a real world.

AB: I think that's very different for the scientist trying to describe the real world; the contribution of sound art to a scientific study could be in the choice of sonority. It would be in how you put together the sound to get the information, whether that was reconstructing in some way an existing sound recording or adding sound to something that was not normally sonic. I think the artistic sensibility is what would allow you to go beyond simply reproducing the visual through sound.

SV: Every clinician will have their own ears and bring with them that fallibility of the body that theory tries to get rid of, or at least theorize into an abstracted entity. For me it is here and on that fragile body that consensus has to be found, maybe as something passing, as something contingent, as something that might remain doubtful but is nevertheless attempted.

AB: Well, certainly in scientific and technical research, the idea that the clinicians all differ and that some of them may have impaired hearing or take a subjective view of how things sound is something we try to remove, and the argument put forward is: We should do the automatic recording and automatic analysis, which takes out that subjectivity of the clinician—get back to the true sound, in some sense, even when that sound's recorded through a variable body (fatter, thinner, with comorbidity), which affects the way the sound is transmitted. We want to take out the body of the clinician in the automatic recognition; we can't take out the body of the patient, but I think we would, if we could, just have the pure lung with its pure sound creation.

SV: But then surely, especially if you miss out other morbidity signifiers, you could actually hear all sorts of things about this body? As an artist I'd find it really fascinating to research all the stuff we can hear in the body without an aim, as a field recording that doesn't already know what it's looking for.

AB: It's very difficult to fund scientific healthy cohort studies.

SV: And so that listening is already predetermined by illness rather than by just listening to what something sounds like . . .

AB: Yes.

SV: . . . or what sonic data—what sonic information—could give us?

AB: I think medical students in training listen to each other and generally hear healthy lungs, but once you're looking at scientific research you're nearly always looking at people who are there in the hospital because they've been sent, through some concern or other, to a clinic.

SV: Which is also very interesting. This need for consensus on the result, that's what it sounds like to me, requires consensus in the reading, and then, given that the result is a diagnosis, a consensus can be had about how something could be treated. Whereas in art there is an experience being enabled, or that one is invited into, that is not verifiably the same for each audience member. If we go to the concert or we go into an exhibition, the demand is not that we agree with each other, necessarily, on what is there, but we have to agree that something happens, and that something strong happens.

AB: That's a significant difference. When you do scientific studies, supposing you're listening with some kind of automatic system you've developed, and you want to compare that to a clinician listening, you would go through a process called consenting with clinicians. You get, say, two clinicians to listen to your recordings and say yes, they hear this sound or no, they don't. In the cases where one clinician says yes and the other says no, you make a research decision about whether you're going to grade those as *no* cases or *yes* cases, so that when you ground-truth for your automatic system, it agrees with the consensus data from the clinicians. But where the clinicians disagree you force a consensus by saying, every time one says yes, it's *yes*, or every time only one says yes, it's a *no*, and expecting your automatic system to do the same job despite this slightly artificial consensus where you've actually got a *maybe*.

SV: And so it's about sound being ephemeral, invisible and uncertain but nevertheless legitimate as a holder of knowledge that could be accepted. There could be an appreciation that there is uncertain yet legitimate insight "in there" that's worth listening out for—even if that means we need to change education, we need to add something to education, we need to broaden our research tools and what we do with them and maybe change our starting point.

AB: I think there is a potential to use it in order to achieve consensus in the scientific sense, even when consensus in the scientific sense might be flawed so it isn't truly a consensus. But there does need to be some sort of agreed grammar in the same way that there is an agreed grammar in the visual that allows you to get to that point.

SV: If you talk about sound work, we can talk about it in musical language, which has some sort of credibility and a

semantic system of its own, and we can get somewhere and have a consensus, or we can talk about it in terms of visual arts where it finds a consensus—not in terms of what it does, but whether it is a valuable piece of artwork, whether it's a good piece of artwork—but these remain in effect visual languages that cannot truly reach the sonority, the stuff that sound does by itself for which there aren't words. So it seems that at least one consensus is that sound does something, it offers something, but because we might not have a consensual vocabulary or a sharable vocabulary, a common terminology, we can't necessarily make it count, it can't necessarily participate as an equal partner in conversation.

AB: But another area where we may have consensus is that, in clinical work, the sound is an indicator to something else; if you listen to a lung and you hear some particular sounds, it is an indicator that you should send that patient to have a more thorough investigation. It's a pointer to another process. In some senses, I think that's what you're saying about sound as art as well that's not dissimilar, that it's not an end in itself, it elicits some other action or feeling or sense or process.

SV: Yes, very often sound doesn't necessarily invite only sound and listening; its ephemerality actually invites a more multisensory engagement. Its lack of articulation in many ways heightens the awareness of what else I do not engage in my perception of the real. Not having a language to fall back on I have to sort of stumble a bit, almost like I'm deprived of something, of clarity and sense; therefore I have to experiment and expand my perception to explore in what other ways I could make sense of things, which I think is partly what's exciting and is where sound crosses over into other research areas and can motivate curiosity: where, all of a sudden, you're not quite sure of the way, how you're supposed to read this design, this body, or whatever it is, and then you have to start to engage in another way.

AB: So that does have quite strong parallels, I think.

Manuscript received 27 February 2018.

ANNA BARNEY is professor of biomedical acoustic engineering at the University of Southampton, U.K. Her work considers application of signal processing techniques to sound recordings in the areas of medicine and health science to enhance diagnosis and monitoring of disease progression. She has a particular interest in sound developed within the lungs and the respiratory tract, both speech sounds and the sounds of breathing in healthy and in diseased states.

SALOMÉ VOEGELIN is a Reader in Sound Arts at the London College of Communication, University of the Arts London. She is an artist and writer engaged in listening as a sociopolitical practice of sound. Her work and writing deal with sound, the

world sound makes: its aesthetic, social and political realities that are hidden by the persuasiveness of a visual point of view. She is the author of *Listening to Noise and Silence* and *Sonic Possible Worlds* (Bloomsbury, 2010 and 2014); the third book in the series, *The Political Possibility of Sound*, will be published at the end of 2018.

Together Barney and Voegelin lead the AHRC-funded Listening across Disciplines network: a pan-European network that seeks to establish listening as a cross-disciplinary methodology and studies the ways in which sound can generate new knowledge and create innovative modes of knowledge production <www.listeningacrossdisciplines.net>.