THE SENSOR CITY INITIATIVE: COGNITIVE SENSORS FOR SOUNDSCAPE TRANSFORMATIONS

Daniel, STEELE[1,2], Dirkjan, KRIJNDERS[2], Catherine, GUASTAVINO[1]

[1]Center for Interdisciplinary Research in Music Media and Technology (CIRMMT) and School of Information Studies, McGill University, 3661 Peel St., H3A 1X1, Montreal, Quebec, Canada

daniel.steele@mail.mcgill.ca, catherine.guastavino@mcgill.ca

[2]INCAS3, P.O. Box 797, 9400 AT Assen, The Netherlands

dirkjankrijnders@incas3.eu

Abstract

The authors introduce a novel urban measurement system. Sensor City, located in the Netherlands, is a large network of outdoor sensor nodes with high-bandwidth communication to a central database and GIS. The network expands the ability to systematically investigate human perception and evaluation in real urban settings and it is the first known cognitive sensor network available to soundscape researchers. The emphasis for this paper is on soundscape evaluation, or the way the acoustic environment is perceived. The authors outline the challenges posed by a cognitive sensor system and how it can approach a more human-like understanding of soundscape, such as through the detection of meaningful events, deciding what data to record, and in-field sensor placement. The authors then explore the potential benefits to the host city in a number of domains and highlight upcoming research directions that rely on this technology, such as the automatic judgment of appropriateness in an urban setting.

Keywords: cognitive sensor network, soundscape, urban design, GIS Ostrava 2013

INTRODUCTION

Sensor City

In 2009 the Dutch Province of Drenthe and the Municipality of Assen teamed up with local research institutes to create a city-wide sensor network named Sensor City. The project offers a novel approach to sensor research in both the size of the installation and the scope of its objectives. Sensor City primarily creates a robust organization of technological and knowledge-based resources on which subsequent projects can be built. The network consists of 150 measurement nodes connected to a central hub via a private, dedicated fibre cable. Each node has a dedicated power source and a secure cabinet for storing processing hardware. This arrangement, redundant to city services, allows the network to be run at a low cost, with unlimited bandwidth, and without interference from or dependence on third parties. The network has also been designed to be robust enough to feasibly run far into the future. Most importantly, the capability of acquiring information from all of these nodes simultaneously enables researchers to instil the network with more cognitive-like powers.

Human Evaluation

The Sensor City project aims to balance the rigorous quantitative evaluation of the daily sonic environment with the more subjective aspects of human perception. Measurements are taken through many types of sensors simultaneously, including microphones that can take full sound recordings from which level and other content can be derived, temperature, light, pollution, and precipitation sensors, and magnetometers for measuring vehicle movements. Researchers expect to be able to combine physical measurements obtained by the network with research into human perception, evaluation and, eventually, behavior to provide better insight about urban conditions. At the project start these classifications will be coarse, but will be refined in subsequent iterations of the technology informed by both machine learning and human testing. The physical environment will be captured by combining public GIS with information about land use, points of interest, zoning plans, etc. Figure 1 provides a plan view of Assen with an indication of existing (represented by pink, blue, and green circles) and expected node placements (in red). Finally, observations, questionnaire studies, and laboratory experiments will provide insights into situated perception and cognition in everyday life context. These data will then inform the next iteration of automated classification algorithms to detect specific events and discard less meaningful but abundant data. Eventually the sensor system, with context information provided by the GIS will be used to predict the perception and evaluation of outdoor spaces for different user groups. We call such a sensor system a cognitive sensor system in that it bridges the gap between physical measurements and automated classification on one hand and perceptions and meanings attributed by people on the other hand. According to Shenai and Mukhopadhyay (2008), "A cognitive system is one that can perceive the environment and adapts to it, can make intelligent decisions based on its knowledge that effect changes in the environment, can self-manage, and self-heal."

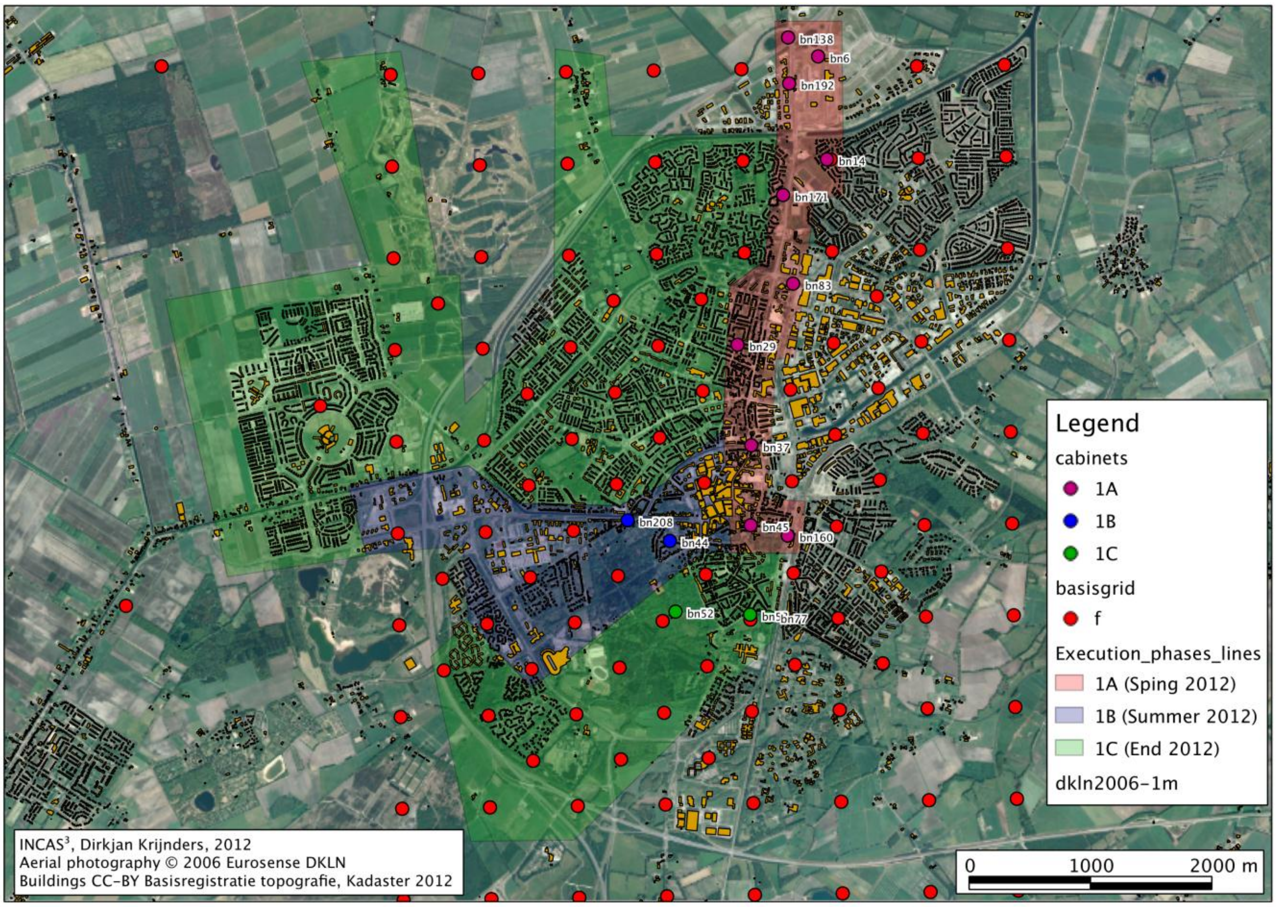

Fig. 1. Plan view of Assen, the Netherlands with Sensor City superimposed

Sensor City further provides a unique opportunity to study outdoor urban and peri-urban environments systematically and reproducibly. Sensor City currently supports two major projects: one focusing on urban mobility, the other on urban sound. This paper focuses on the sound project.

Sensor City Sound

The sonic environment, or the auditory component of the urban environment, is measured by recognizing and classifying sound events. When a human perceives the sonic environment, it is considered to be the soundscape. With the aid of the cognitive sensor technology, we aim to move beyond simply recognizing events, but to actually process them as well.

The ability of this research to transform the landscape through better evaluation is palpable, but there are some issues to be worked out. First, it is important to understand that soundscape interventions can be made within and between large, more established disciplines (Steele, Luka & Guastavino, 2012). Those working on this project must understand the impacts of this research across multiple domains and take into account the particular needs of multiple stakeholders. Second, we must keep in mind that since soundscape is the human understanding of the sonic environment, our findings can be extremely individually or culturally specific.

COGNITIVE SENSORS AND SOUDSCAPE

What is soundscape?

The concept of soundscape has garnered increasing research attention over the last decade in different fields of practice (e.g. urban design, wilderness management, noise control, transport, psychoacoustics) involving researchers from diverse disciplinary backgrounds (e.g. acoustics, music, architecture, ecology, psychology, sociology, geography, urban planning). Various definitions and synonyms have been proposed (see Botteldooren et al., 2011; Niessen, Cance, & Dubois, 2010; Brown, Kang & Gjestland, 2011); but nearly all uses of the term soundscape have in common the emphasis on the way the acoustic environment is perceived and understood by the individual, by a group, or by a society (Botteldooren et al., 2011; Brown, Kang & Gjestland, 2011; Truax, 1999).

Central to soundscape research is the recent paradigm shift from quantitative analytic approaches (e.g. psychophysics) to more qualitative cognitive approaches focusing on meanings attributed to soundscapes in relation to human activities (e.g. Dubois, Guastavino, & Raimbault, 2006; Guastavino, 2007). Indeed, there is converging evidence that soundscape cannot be assessed/measured exhaustively in terms of acoustic measurements as humans typically evaluate the meaning of the activity, source and/or agents producing sound when evaluation soundscapes. These judgments rely on cultural values attributed to the different types of activities and not on inherent properties of the sounds produced

Soundscape and urban design

The field of urban design emerged as a response to the tendency of urban planners to design in plan view and ignore the city user. Urban designers now preoccupy themselves with designing experience at the human scale, but they lack common metrics for evaluating good design. Introducing humans into the design picture motivates a change in the purpose of design – not all people are satisfied by the same product or treatment – so diversity and individuality must be taken into account. This distinction is actually built into the definition of soundscape. The acoustic environment is the sum of all of the sounds in a given space, or in other words, the sounds that transpire independent of the listener. The soundscape is the way an individual (or society) perceives and understands the acoustic environment. But why involve urban designers in the soundscape debate? In fact, it has been demonstrated that the best outcomes for urban soundscapes appear at the intersection of acoustic and urban design (Steele, Luka, Guastavino, 2012). Yet very little work has ever been done to understand how exactly specific urban morphologies lead to specific soundscape outcomes, and Sensor City offers that opportunity.

Urban sound information has traditionally been reported with only one property of the sound being measured - its sound pressure level (SPL). The data is then superimposed on a noise map that greatly oversimplifies the context and source of each of those sounds. Noise maps still maintain value, however. Noise annoyance grows steadily with loudness (Steele & Chon, 2007) so noise maps are useful for predicting residential annoyance. However, measurement of sound levels alone fails to take into account the context and meaning of the sound, the presence of multiple sources, and the interaction of other modalities (especially vision). GIS data, as it is being used by Sensor City, can help to shed some light on these various issues by telling us, for example, the particulars of the built form or zoning considerations. GIS data for the Netherlands is extremely precise, complete, and recent, and Sensor City researchers are already aiming at systems designs that can take this advantage into account.

Soundscape researchers, in response to these criticisms of noise maps, are looking for better ways to represent soundscape from the plan view. A thorough soundscape map would likely involve multiple layers to deal with the multifaceted features affecting soundscape judgments. Sound recognition is widely considered to be an important first step in the process of updating urban sound maps, since identification is key to understanding context. Krijnders et al. (2010), conceive a sound recognition system driven by signal processing and, if available, the known context of the signal. This entails that the incoming sound spectrum is separated in tonal parts, pulse-like parts and broadband parts, for which different features are extracted. Standard machine learning algorithms then classify the parts to arrive at a final, more cognitive recognition. Raimbault and Dubois (2005) go further and propose charting soundscape impacts, with different categories of soundscape clearly identified, such as transportation, people presence, and natural sounds. They suggest that such a system will improve spatial analyses and facilitate visual communication. Visual communication is crucial for representing soundscapes because change-makers in this domain, such as architects and urban designer, are trained and operate almost exclusively in the visual modality (Steele, Luka, & Guastavino, 2012).

SOUNDSCAPE AND GIS – SYSTEMATIC MULTIMODAL SOUNDSCAPE OBSERVATION

The Sensor City network represents a new approach for soundscape research in a number of ways. First, the multimodal interactions of soundscape will be considered. Because the sensor nodes will be collecting data not only about sound, but also weather, time, and special circumstances, it will be possible to systematically investigate the relative contribution of the audio content of soundscapes and other sensory inputs to the understanding of soundscapes *in situ*. The idea that soundscape is influenced by non-auditory factors is well-established (Yong Jeon, Jik Lee, Young Hong & Cabrera, 2011); Second, Sensor City enables soundscape evaluation on a large scale – the scale of the whole city, including residential and work areas, as well as restorative and other recreational landscapes. Third, Sensor City evaluates soundscapes in time, extending previous studies by offering longitudinal capabilities, tracking the evolution of time at various scales (i.e. hourly, weekly, seasonally, generationally). At the same time, the introduction of the temporal aspect of soundscape poses interesting challenges for representation with GIS. How does one represent something that exhibits strong patterns (such as the Monday to Friday workday and associated traffic movements) but that has subtle changes over time? In fact, city mobility and built form are inherently intertwined, so it is worth looking at an example of how Sensor City might learn from the particulars of urban movement.

Considering the heavy influence of automotive traffic on the perception of soundscape (Guastavino, 2006), it follows naturally that mode share, or the distribution of the population using certain modes of transportation for daily activities, can have a large effect on soundscape. Since the Netherlands is known for having a remarkably high bicycle mode share (in Assen, nearly half of the population routinely rides their bicycles), the interactions between mode choices can provide for interesting results. The Sensor City network could automatically determine the proportion of bicycle to automobile (and pedestrian) traffic at any given point through magnetometer and microphone data. This information would be hard to obtain with traditional counting techniques because it requires all-day monitoring, and tracking the less predictable movements of people on foot and bicycle. Instead, the combination of microphone and other data could provide permanent long-term monitoring to account for different morphologies all over the city at the same time, at different times of day, at different weather conditions, and during special events. This data could be useful for helping the city to determine which areas of the city officials should concentrate on if they wish to alter the mode share in a particular direction, such as having more people ride bicycles to a popular racing venue on the outskirts of the city or music festivals. Assen could take data at various points in the city and find out if people from a particular neighborhood are taking bicycles or cars to an event and perhaps provide better incentives for bicycling to the event if they wanted to save resources by not having people park their cars at the event. At the same time, continuous data would be available to help the city determine if their installation or educational program provided the intended outcome.

But here a critical problem emerges: how does one select sensor locations to achieve “meaningful” data? It is paradoxical to both choose interesting locations and remove systematic bias in site selection. Do we aim for activity centers where things might happen, such as a park bench where someone is likely to sit and read a book? Or do we aim to achieve a fair representation across the municipality’s different morphologies to maximize our understanding of design patterns? Or finally, do we select sites at random to equalize the potential for interesting findings across the municipality, such as dead spaces? Sometimes, unexceptional spaces can be ignored despite the important lessons and feedback that can be obtained from them. Figure 2 shows an unremarkable tunnel for bicyclists crossing underneath a busy road, which could help us to reproduce and validate (ecologically) Korte and Grant's (1980) finding that people move more quickly through noisy spaces. Should this space be ignored for not being particularly positive or is its commonplace design conveying meaningful information about how humans use built form?

Fig. 2. A tunnel for people on bicycles passing underneath a busy road. Assen, the Netherlands

WORKING WITH THE COGNITIVE ELEMENT

Previous researchers have taken on the problem of meaning in soundscape research, asserting that laboratory testing takes human subjects outside of the context of the environment where the sound is experienced. To address this problem, the concept of ecological validity has been extended to soundscape reproduction (Guastavino, Katz, Polack, Levitin, & Dubois, 2005). In the case of sensor city data, combined with intended questionnaires and other data collection techniques, the concern of ecological validity can be mitigated. As well, the social considerations, such as the fairness aspects of noise exposure (Maris, 2007) can be dealt with systematically.

Since soundscape is the culmination of perception and understanding of the acoustic environment, it is important to understand that there are elements of soundscape cannot be measured directly by a set of one-dimensional sensors. It will be important to also determine which of those real-world elements cannot be determined by a proxy measurement either – we may be able to predict annoyance but not overall well-being or aesthetic merits. These elements will encompass many domains, so it is important to think in a multidisciplinary way to understand the limits of Sensor City. For example, the task of measuring traffic noise will be quite easy. Further, the effect of traffic noise on soundscape perception is fairly well understood. However, soundscape is generally improved when residents of a certain location know their neighbors. There is no direct way to measure familiarity with neighbors with a sensor, but it is perhaps conceivable to measure it indirectly. Sensors could detect a spontaneous interaction in a residential setting and check for low aggression in the voices. In the event of high aggression or no verbal communication, we might infer that neighbors are not familiar with each other in the vicinity of this sensor set, whereas low aggression and extended outdoor conversations might indicate the opposite. Lastly, it has been demonstrated that soundscape can be considered good when residents perceive they have a safe and short walking journey to a nearby park (Gidloff-Gunnarsson & Ohrstrom, 2007). While it's true that it is quite easy to measure the physical distance between a residence and a park, it is impossible to understand the intricacies of an individual's perception about whether it is a good or a bad walk to that park (i.e. someone does not like a particular restaurant passed on the walk). In general, elements from the "understood" portion rather than the "perceived" portion of the soundscape definition will be those that are hardest to automate, thus the cognitive sensor movement.

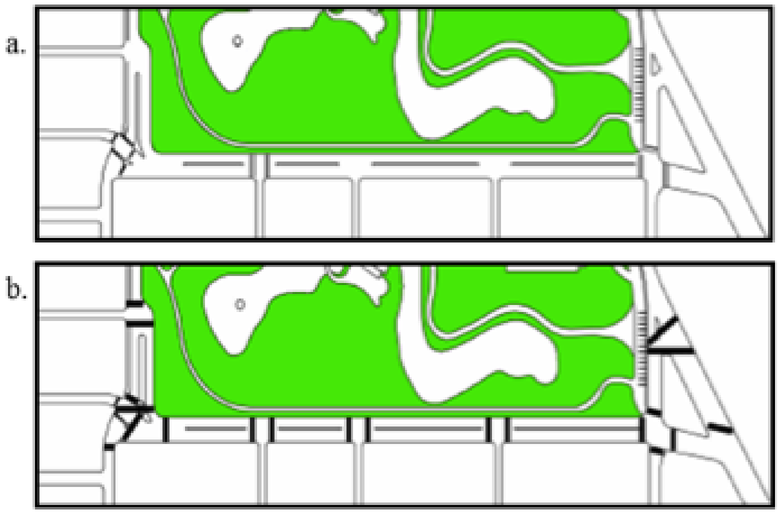

Fig. 3. An urban park surrounded by dense residential units. a) In its current form, the park has poor pedestrian access, b) an intervention is proposed that installs new crosswalks (black) at every intersection. The intervention will improve soundscapes for nearby residents inside and outside of their homes. Parc La Fontaine, Montreal, QC, Canada. (Image credit: Kris Steele).

Possible outcomes from the data

An upcoming movement in urban design will be design validation, or tests that demonstrate a particular design operates as expected. Surprisingly, a culture of evaluation of the performance of the built environment by its designers is not standard practice. Evidence-based architectural design has been largely limited to designs for healthcare. Urban designers need to embrace the many calls that have been made for post-occupancy evaluations and learn how to capitalize on poor results. Post-occupancy evaluation should not be considered a method to be used in special cases. In a sense, Sensor City automatically provides a post-occupancy evaluation, with identical measurements taken before and after a hypothetical new installation. In fact, post-occupancy evaluation is generally not performed for urban installations because it is often not clear what to look for. Sensor City provides the flexibility to decide metrics on the fly or well after measurement.

Consider the case of an urban park, presented in Figure 3. It has been established that urban parks provide restorative benefits to users and that the perceived access to a park from a person's residence improves the daily indoor and outdoor soundscape (Gidloff-Gunnarsson & Ohrstrom, 2007). In the first drawing, the park lacks good access because it does not have good crosswalks. If researchers can identify an opportunity like this, sensor data such as unique conversations and park visits can be counted and a questionnaire could be circulated to park neighbors and users to understand the current park usage and soundscape evaluation. An identical study performed after the installation of new crosswalks will tell researchers the efficacy of that installation. If later, researchers realize that the new crosswalks may have slowed cars, sound level and magnetometer data can be gathered from before and after the installation to confirm whether this is true.

LONG-TERM IMPLICATIONS FOR ASSEN AND URBAN DESIGNERS/SPECIALISTS

The authors engaged in a literature search for urban design and architectural resources on the pedagogy of both adopting scientific approaches in architecture, such as evidence-based design, and soundscape design. They also performed a manual search by speaking to architecture and design colleagues and the overwhelming result is that no such education exists. While there is a recent trend for evidence-based methods in healthcare infrastructure, these lessons have not been transferred to more general cases of urban infrastructure. Currently, urban designers build their designs but have no responsibility to evaluate the performance of their designs – the final product to the designer is an image of the space before it opens to the public. The unfortunate meaning of this culture gap is that the "scientific" and evaluative approach that Sensor City is taking is potentially unmeaningful as a feedback mechanism for designers. Therefore, it will also be the responsibility of Sensor City researchers to formulate meaningful research outcomes for designers that can result in real change.

Politically, soundscape evaluation does not receive a lot of weight either, despite the evidence that it should. Some studies have demonstrated that while soundscape concerns are not considered with priority in political or residential considerations, decision makers still have a responsibility to provide good soundscapes (Gidlof-Gunnarsson and Ohstrom, 2007; Raimbault and Dubois, 2005) because poor soundscapes can adversely affect individual public health, happiness, and productivity. With proper knowledge management, Sensor City is capable of providing a real, tangible and meaningful product for designers with automatic ecological validity, negating designers' concerns about the shortcomings and validity of laboratory testing. In addition, the municipality of Assen gets real-time feedback on the effects of policy, installations, construction/demolition, special events, wildlife changes, and more both spatially and temporally.

Finally, once appropriate lessons have been gleaned from the Sensor City Project, simply reporting them is not the end of the story. The language used to convey findings must be audience-specific. Such a consideration is critical because the expertise of those who make decision about the functioning of a city comes from many sources – architects/designers, scientists, engineers, contractors, and politicians (who are further influenced by the lay public.) Raimbault and Dubois (2005), for instance, revealed through a semantic analysis that even experts fail to agree on basic technical vocabulary to describe usual sound events. At the same time, the already limited technical vocabulary of experts is totally unused by the "city users" who opt for "vocabulary of comparison" and "human noise descriptions" much more than urban planning experts. Also, experts' descriptions of phenomena showed richer and deeper categorical structure. Thus, we must distinguish the sounds of 'bicycle wheels' from 'people on bicycles' if we are to communicate meaningfully to our audience with Sensor City. It is clear that a one-size-fits-all tool will not be the best solution for conveying our findings.

CONCLUSIONS

This paper presents a cognitive sensor network capable of powerful urban measurements. The unique setup allows researchers to expand their ideas and pursue research questions they might not have asked. Our lab is pursuing the idea that soundscape can be rated according to its appropriateness for a certain situation. The appropriateness measure takes into account both the urban morphology, by suggesting that a certain type of place might call for a certain type of soundscape that affords certain types of activities to take place there, and the human, by suggesting that determining appropriateness is necessarily a cognitive task. We will investigate this idea by conducting laboratory and *in situ* questionnaires and observances about how people use certain spaces for certain kinds of activity and whether the soundscape of that space appropriately matches their expectations. Soundscapes that violate listener expectations by being inappropriate would be considered unsuccessful and, we speculate, would discourage use. For example, a noise-exposed park would fail to attract visitors wishing to exercise, while a quiet marketplace might discourage people from browsing and shopping. Soundscape appropriateness nicely balances the expectations of automatic recognition with the imposition of human-like cognition and could ultimately lead to fungible success measures agreed on by both the scientific and design communities. Such corroboration would ultimately lead to a better cycle of design and evaluation across fields of practice and potentially lead to practices in the field of the built environment that consume fewer resources to do a better job.

REFERENCES

Botteldooren, D., Dekoninck, L., & Gillis, D. (2011). The influence of traffic noise on appreciation of the living quality of a neighborhood. International Journal of Environmental Research and Public Health, 8(3), 777-798.

Brown, A. L., Kang, J., & Gjestland, T. (2011). Towards standardization in soundscape preference assessment. Applied Acoustics, 72(6), 387-392.

Dubois, D., Guastavino, C., & Raimbault, M. (2006). A cognitive approach to urban soundscapes : Using verbal data to access everyday life auditory categories. Acta acustica united with acustica, 92(6), 1060-1070.

Gidlöf-Gunnarsson, A., & Öhrström, E. (2007). Noise and well-being in urban residential environments: The potential role of perceived availability to nearby green areas. Landscape and Urban Planning, 83(2), 115-126.

Guastavino, C. (2007). Categorization of environmental sounds. Canadian journal of experimental psychology, 61(1), 54-63.

Guastavino, C. (2006). The ideal urban soundscape: Investigating the sound quality of French cities. Acta Acustica United with Acustica, 92(6), 945-951.

Guastavino, C., Katz, B. F., Polack, J. D., Levitin, D. J., & Dubois, D. (2005). Ecological validity of soundscape reproduction. Acta Acustica united with Acustica, 91(2), 333-341.

Korte, C., & Grant, R. (1980). Traffic noise, environmental awareness, and pedestrian behavior. Environment and behavior, 12(3), 408-420.

Krijnders, J. D., Niessen, M. E., & Andringa, T. C. (2010). Sound event recognition through expectancy-based evaluation of signal-driven hypotheses. Pattern Recognition Letters, 31(12), 1552-1559.

Maris, E., Stallen, P. J., Vermunt, R., & Steensma, H. (2007). Evaluating noise in social context: The effect of procedural unfairness on noise annoyance judgments. The Journal of the Acoustical Society of America, 122, 3483.

Niessen, M., Cance, C., & Dubois, D. (2010). Categories for soundscape: toward a hybrid classification. In INTER-NOISE and NOISE-CON Congress and Conference Proceedings (Vol. 2010, No. 5, pp. 5816-5829). Institute of Noise Control Engineering.

Raimbault, M., & Dubois, D. (2005). Urban soundscapes: Experiences and knowledge. Cities, 22(5), 339-350.

Shenai, K., & Mukhopadhyay, S. (2008). Cognitive sensor networks. 26th International Conference on Microelectronics, MIEL 2008, (pp. 315-320). IEEE.

Steele, D. L., & Chon, S. H. (2007). A perceptual study of sound annoyance. Audio Mostly 2007, 9-24.

Steele, D., Luka, N. & Guastavino, C. (2012). Closing the gaps between soundscape and urban designers. Proceedings of Acoustics 2012, joint meeting of the 2012 IOA annual meeting and the 11th Congrès Français d'Acoustique, April 23-27, Nantes, France.

Truax B. 1999. Handbook of Acoustic Ecology. 2nd ed. (CD-ROM). Cambridge Street.

Yong Jeon, J., Jik Lee, P., Young Hong, J., & Cabrera, D. (2011). Non-auditory factors affecting urban soundscape evaluation. The Journal of the Acoustical Society of America, 130(6), 3761-3770.