Ubiquitous music in smart city: musification of air pollution and user context

Pedro Sarmento[1], **Ove Holmqvist**[2], **Mathieu Barthet**[1]

[1]Centre for Digital Music – Queen Mary University of London

[2]Holonic Systems Oy.

{p.p.sarmento, m.barthet}@qmul.ac.uk, ove@holonic.systems

Abstract. *As smart city technologies proliferate, new sources of data are becoming ubiquitous. The auditory modality appears promising to create experiences giving users new meanings about their environment. Building on the fields of sonification, ubiquitous music and mobile music, we set out in this paper a musical smart city ecology whereby environment and user-related data are used to drive the music creation process. We discuss several challenges related to this vision including how artificial intelligence may be used for data-driven music generation. Finally, we discuss air pollution and user context musification use cases illustrating applications to smart city contexts.*

1. Introduction

As smart city initiatives thrive, along with the development of the Internet of Things (IoT), data harnessing and interconnectivity between users, devices and decision-making agents arise. There is, however, a gap between sensed data and end users who often do not have direct access to data or the knowledge to make sense of it. In this work, we study the interpretation and contextualisation of smart city data through the lens of sonification, the rendering of data into non-speech audio [Kramer et al. 1999]. The audio modality as a means to convey (non-speech) information is still rather unexplored contrary to the visual modality [Ramos et al. 2018]. In previous work [Sarmento et al. 2020], we discussed interactions between musical creation and smart city based on the fields of sonification, ubiquitous music [Pimenta et al. 2009], and mobile music [Essl and Lee 2017]. We expand on this work by discussing artificial intelligence (AI) techniques for music generation [Briot et al. 2019] to assist in the creation of musical content driven by smart city data. Two prototypical use cases are also presented. In the first use case, sonifications of air pollution levels let listeners compare periods during COVID-19 lockdown with periods of normal daily traffic. The second use case presents Holonic Systems' mobile applications that generate music interactively based on user contextual data.

2. Background

2.1. Ubiquitous and Mobile Music

The field of ubiquitous music ("Ubimus") lies at the intersection of areas such as human-computer interaction, sound and music computing and creativity studies with a focus on mobility and social interaction [Keller et al. 2019]. Ubimus is influenced by ubiquitous computing concepts [Weiser 1991], leveraging the presence of technology in everyday objects [Mandanici 2019]. It aims to integrate components of a system as musical agents, each one adapting and responding to both the environment and other agents [Pimenta et al. 2009]. Mobile music focuses on designing mobile device interfaces for music generation and consumption [Essl and Lee 2017]. Examples include location-based music, where sound experiences are created according to user geolocation and/or movements within specific target zones [Joy and Sinclair 2008]. As stated by [Hazzard et al. 2015], locative, GPS-driven technologies support exploration and creation of novel music experiences, as demonstrated by their soundwalk experience in an outdoor location with adaptive mobile music. Further works explore the use of sensors embedded in mobile devices and their ability to produce novel means for musical expression, for example MoMU [Bryan et al. 2010], a mobile music toolkit for iPhone.

2.2. Sonification and Ubiquitous Music

As stated by [Worrall 2018], '*the idea that sound can convey information predates the modern era, and certainly the computational present*'. In our era characterised by data, this root function of sound is worth to be considered regarding information gathering and decision making. Relatedly, the field of auditory display investigates the use of sound to render information for human communication (e.g. speech, alarms). Sonification focuses on using non-speech sound not conflicting with verbal communication to render data [de Campo 2007], often with the purpose of exploration. [Barth et al. 2020] highlight parallels between sonification and ubiquitous music. The authors suggest that natural data (measurements of spatial or temporal patterns from human or nonhuman phenomena) can become a source of composition. They also note that the ability to record and access environmental signals is increasingly possible which provides an opportunity for music creation and artistic, educational, and scientific outcomes. Their sonification process that reveals the dynamics of Geyser data is linked with the ability of the auditory system to identify complex structures evolving over time which "might otherwise be overlooked or require significant processing ". Projects such as *The Decatur Civic Dashboard*, by [Winters et al. 2016], and *Datascaping*, by [Pigrem and Barthet 2017] propose sonification methods based on smart city data that capture relationships between users and their environment.

2.3. Musification

A research area within sonification that emphasises artistic creation is musification. Musification has been defined differently among researchers, either as a sonification technique that uses scales, chords, key and tempo changes [Barrass 2012], or as an organised sonification [Bonet 2019a], the latter stemming from Varèse's definition of music as 'organized sound' [Risset 2004]. There are many historical examples of composers who explore computational and algorithmic procedures driven by data [Vickers 2016]. Predating the advent of computing, Heitor Villa-Lobos' *New York Sky Line* (1957) depicts the skyline of the city of New York into notes for a piano piece, mapping skyscrapers and building heights, read from left to right, into a musical score [Bonet 2019a]. At first glance, this could be considered an application of parameter-mapping sonification, but the focus is on musical aesthetics rather than the accuracy of information transmission. Musification stems from the increasing use of sonification by both scientists and composers [Bonet 2019b]. As a practice, it is expected to find balance between the functional aspect of conveying information and the aesthetic values of a self-contained musical piece. While there is a risk of non rigorous data representation in the resulting sonic works, musification can have the upside to expose the artwork to broader audiences and foster engagement [Ballora 2014]. As stated by [Barrass and Vickers 2011], aesthetics is an important aspect to facilitate understanding, this notion becoming even more relevant when sonification (or musification) becomes embedded in everyday objects or activities (see Section 3).

2.4. Internet of Musical Things

The concepts of ubiquitous music and smart city also relates to the Internet of Musical Things (IoMusT), described in [Turchet et al. 2018] as the networks of musical things (e.g. smart musical instruments [Turchet and Barthet 2019] or wearables) supporting the production and/or reception of musical content. The IoMusT focuses on novel interactions between musicians and audiences made possible through context- and user-awareness computing with smart devices. Some of the IoMusT challenges are also relevant for the deployment of ubiquitous music in smart cities, namely those concerned with the interoperability between devices and reduced latency e.g. relying on 5G technologies.

3. Towards an Ubiquitous Musical Smart City Ecology

3.1. Context

Following the growth of smart cities linked with the evolution of urban settlements worldwide [Cugurullo 2018] we witness the exploitation of communication technologies to support the administration of the city and to improve citizens' experience [Zanella et al. 2014]. [Zanella et al. 2014] argue that the IoT is crucial for the deployment of such an intricate ecosystem. In the IoT paradigm, objects of everyday life have communication capabilities, connecting users and devices. This is supported by a network of ubiquitous and heterogeneous sensors, harnessing data from distinct sectors (e.g. environment, infrastructure, etc.), while integrating user information.

3.2. Vision

Pervasive computing [Satyanarayanan 2001] envisions digital spaces embedded within ubiquitous computing devices surrounding users. We previously proposed the concept of ubiquitous sonification as *type of sonification/musification that leverages ubiquitous computing environments* [Sarmento et al. 2020]. In the context of smart city environments, ubiquitous sonification can be used to deliver auditory experiences to listeners, through data-driven sonic exploration of the environment and/or active musical interaction based on their physical behaviors and the response from musical agents. In this sense, ubiquitous sonification is concerned both with musical aesthetics and information transmission.

Ubiquitous sonification can be applied to smart cities and their multiple heterogeneous data streams. Components for a musical smart city ecology are shown in Figure 1. Data providers gather smart city data, namely environmental data relating to user surroundings (e.g. data related to air pollution, energy consumption, crime rate, urban traffic), mostly gathered by *ad hoc* Wireless Sensor Networks (WSNs). User-related data can express locations, physiological data (e.g. heart rate, body temperature) and inertial data (e.g. hand gestures, head movement), and can be collected by smart devices (e.g. watches, bracelets, phones). Data can be streamed in *quasi* real-time pending latency issues, be stored and accessible in cloud computing environments, or be predicted (as is the

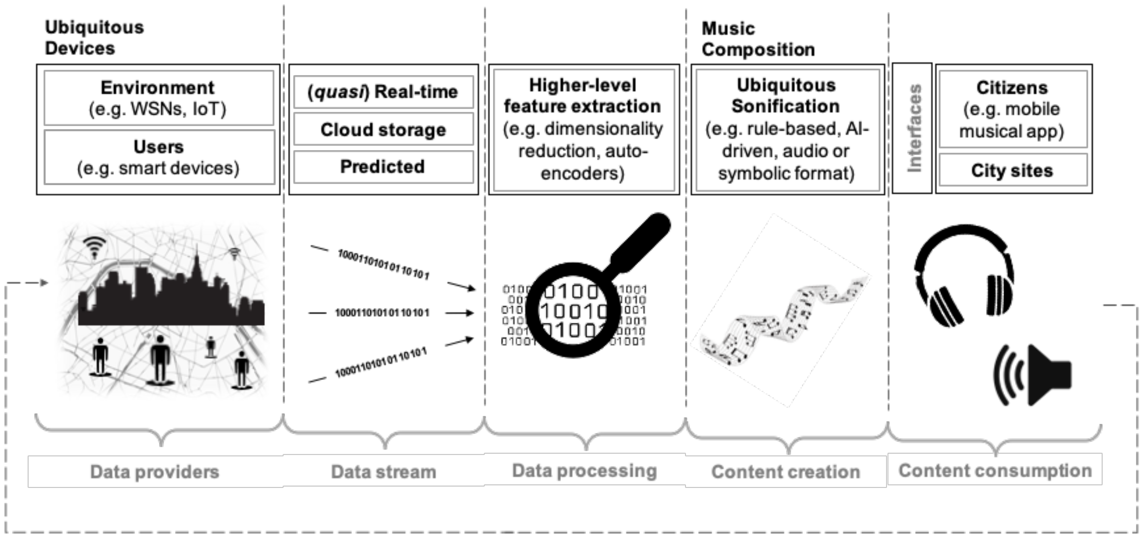

Figure 1. Conceptual diagram of a musical smart city ecology.

case with most weather-related applications). Considering large datasets with multiple features, higher-level feature extraction could support ubiquitous sonification. This could involve filtering operations, where the most relevant features are selected, or, in datasets of higher ambiguity, dimensionality reduction techniques could be pursued (e.g. Principal Component Analysis, deep auto-encoders). Musical compositional models are core to the process (see Section 3.3). Interfaces enable users to experience the generated sonic works. Mobile applications are examples of interfaces (See Section 5).

3.3. Compositional Models and the Use of AI

To support a musical smart city ecology, our objective is to develop data-driven methods for automatic music generation focusing on experience and meaning-making [Bødker 2015]. Musical works generated from the user context bares several values, for example, didactic (to become aware and learn about intangible aspects of one's environment), and yields an aesthetic experience, characterised in [Goldman 2006] as "the simultaneous challenge and engagement of all our mental capacities—perceptual, cognitive, affective, imaginative, even volitional — in appreciation of the relations among aspects and elements of artworks". We classify sonification models for composition into two broad categories: **knowledge-based (or rule-based) sonification**, the most traditional type of approach, guided by rules and manual parameterization of relations between data points and sonic output; **AI-driven sonification**, in which the sonic generation leverages e.g. deep generative approaches. Rule-based sonification is well documented. With this approach, a high degree of control is possible, by mapping data characteristics to aspects of the musical composition. Common practices link data to music generation parameters (e.g. pitch, loudness and timbral features). Recent years have seen a rise in research that explores the use deep learning (DL) models for music generation [Briot et al. 2019], either creating content directly in the audio domain [Zukowski and Carr 2017] or in a symbolic format (e.g. MIDI [Ferreira and Whitehead 2019], ABC notation [Sturm et al. 2015], GuitarPro format [Sarmento et al. 2021]). We propose to use DL models to generate music conditioned by (inferences from) IoT data. To our knowledge, this has not yet been addressed in the sonification literature. In their project *Hearing Artificial Intelligence*, [Winters et al. 2019] sonified the penultimate layer of a neural network used to detect skin lesions, following parameter-mapping sonification principles. Often deep generative approaches for music generation are seen as *black boxes*, offering a small degree of control over the output besides the choice of the dataset used by the model to make inferences. However, recent research evidences that conditioning can allow for a fine-tuning of the musical result. [Kaliakatsos-Papakostas et al. 2018] present an RNN-based system that composes polyphonic MIDI music inspired by a previous learned style, while integrating user-defined features, namely rhythmic density and pitch register, with possible real-time control. Furthermore, [Meade et al. 2019] explore conditioning techniques with PerformanceRNN, a language model used in a musical setting capable of generating expressive piano compositions [Oore et al. 2018]. More specifically, the generation is performed at the symbolic level (MIDI) and uses features such as velocity, tempo, musical key, also including information about specific composer styles. Generative music systems that support a wide range of musical expression and styles while being conditioned by contextual data would enable to produce ubiquitous sonification works of varied aesthetics matching listener musical tastes. For musical smart city applications, the ability to adapt the generated music to user preference is deemed important to foster engagement with the system. This could be dealt with by training deep learning models on different datasets (either by musical style, composer, emotion) and tailor the output to user intent.

3.4. Challenges and Insights

We discuss in this section five major challenges.

[1] How to harness and make sense of data within the context of a smart city: Relevant data streams and domains must be identified for musical smart city applications. This should take into consideration large, fuzzy and often incomplete streams of data, coming from a myriad of distinct sensors. To accommodate this, dimensionality reduction techniques may be explored in order to extract higher level features of said data (e.g. PCA, Deep Auto-Encoders).

[2] How to apply AI for musification of data: An important issue is the choice of musical datasets to train models, including their format, either in the audio or symbolic domain. The symbolic domain has the following two advantages: the models are in principle easier to train and handle, and the output should be easier to be conditioned on data, to reach satisfactory musical results. The choice of datasets can be guided by multiple parameters such as musical style, number of instruments, mood/emotion. However, with the symbolic format, timbre quality variations due to performer interpretations [Barthet et al. 2010] are not directly taken into account, and the synthesis/rendering of MIDI to audio is separate. The type of mappings for the data conditioning can be informed by previous sonification research and user feedback. With the symbolic approach, pitch-related features are a reasonable choice, as well as note density, loudness and temporal characteristics.

[3] How to balance aesthetics and informational objectives: Music and related aesthetic experience is often one of the concerns in auditory display as illustrated by [Neuhoff 2019] stating that the word 'music' appears in 74% of the ICAD publications. In his paper on sonification viewed from translation studies perspectives, [Lepri 2020] proposes a three-level categorization for the relations between input data and sonic rendering: literal translation, a sonification that focuses on the functional aspect; semantic interpretation, where there are imaginative and emotional allusions in relation to the input data, sacrificing a more systematic approach; critical interpretation, with an emphasis on the artistic side of the work. A musical smart city application would fall into the semantic interpretation, aiming for a musical output with aesthetic value, where the transfer of information into sound is not a direct rendering of data, instead guiding the listener's awareness to new references.

[4] How to evaluate interfaces and outcomes: The hurdle of evaluating a musical smart city system relates to the two-level nature of the assessment: its efficiency to convey information and the aesthetic merits of the musical content. Based on a survey on evaluation criteria of sonification and musification, [Bonet 2019a] proposes a framework using five dimensions: *Gain*, concerned with how much the aesthetics of the work serves the purpose of information transmission; *Intuitivity, clarity and learning effort*, how efficient is the process of information transmission; *Sonification method*, assessing if mappings between data and musical parameters do serve a musical purpose; *Aesthetics*, focused on the aesthetics valences of the work and also on how immersed in the musification experience a listener is; *Audience feedback*, which includes all the previous criteria from the perspective of the audience. Self-report questions for each of these dimensions can be used to quantify the performance of a musification work using scales ranging from 1 (lowest) to 5 (highest). Results are mapped into an evaluation chart for musification that comprises both artistic and functional values. A similar approach could be used to assess musifications produced from smart city data. Furthermore, for an AI-driven compositional process, an iterative design of the system can be envisioned by integrating feedback received by listeners.

[5] How to consider environmental challenges and impact: The use of AI techniques requiring computing resources and IoT technology yields an environmental challenge. Questions such as how to devise efficient DL architectures minimising the impact on the environment, and how to use sustainable IoT technologies that consume less energy [Vieira et al. 2020] should be addressed.

4. Use case 1: Musification of Air Pollution Data

4.1. Background

Musical smart city systems aim to raise awareness about our environment. In this use case, we developed techniques to characterise air pollution through musification. Reports suggest that a reduction in air pollution levels occurred during the lockdown related with COVID-19[1]. One of the motivations for this use case was to inspect if such differences

[1] Available at: `www.eea.europa.eu/themes/air/air-quality-and-covid19`

could be conveyed in an auditory way. According to the World Health Organization[2], Nitrogen Dioxide (NO_2) is the main source of nitrate aerosols and, in the presence of ultraviolet light, of Ozone. The major sources of emissions of NO_2 are combustion processes from heating, power generation, and engines in vehicles and ships. Here NO_2 data from a measuring site located on Mile End Road, London, was used[3]. The data are comprised of hourly mean values from a week in 2019 (18^{th} to 25^{th} of April) and a matching week in 2020 (16^{th} to 23^{rd} of April). Data files were not subjected to any processing stage. As a proof-of-concept, data fetching was not done in real time.

4.2. Design and Implementation

This use case explored musical borrowing in the sonification process, where a piece of music (or an excerpt of it) is used as source material to drive a new compositional process. Rarely explored in musification, this technique can be defined as '*a deliberate evocation within a composition of a different musical work*' [Bicknell 2001]. Previous examples of this technique map deviations of the value of the pound Sterling after Brexit to shifts in pitch applied to Britain's national anthem, *God Save the Queen*; other renditions use predicted data of floods in US cities in the year 2045, mapping values to changes in tempo of *Ice Ice Baby* (1990), a song by Vanilla Ice [Bonet 2019b]. As stated by [Bonet 2019a] regarding her piece *Wasgiischwashäsch*, it is important to note that the musical borrowing practice does not constitute plagiarism because the compositional goal is to invoke the audience's *musical literacy* and knowledge about the originality of the piece. Musical borrowing leverages the knowledge that a listener has about a composer's work. This presumably makes any modulation caused by the data more explicit since the original work acts as a reference.

The piece 'Air on a G String', a Wilhelmj's arrangement of a Bach composition (2nd movement of the Orchestral Suite No. 3 in D major, BWV 1068) was *musically borrowed* for the process of sonification[4] . An interactive system was created using Bela, an embedded computing platform desgined for ultra-low latency and high-quality audio [McPherson and Zappi 2015], and audio signal processing code was written in C++. Bela was chosen for its audio capabilities and the possibility of embedding the platform in smart city devices. Also, its sensoring capabilities enable a degree of interaction that was considered interesting. The musification process was done by controlling the parameters of an *ad hoc* spectral delay effect. Here, 10 bandpass filters were used, the output of each was stored into 10 delay buffer lines. This rendition is considered a knowledge-based musification, as mappings are directly created and manually tuned.

A flow diagram of the whole process is depicted in Figure 2. Holding a button down switches from data from 2020 to data of 2019. Data points were evenly spread for the duration of the piece, and linear interpolation of data values was performed to account for changes at sample level. Subsequently, data values controlled spectral delay parameters such as cutoff frequencies of every filter, feedback and gain amount of every delay line. The output of every filter-delay line was summed. Regarding mappings, lower cutoff frequencies inside every band are linked with higher data values; higher values of

[2]Available at: `www.who.int/news-room/fact-sheets/detail/ambient-(outdoor)-air-quality-and-health`

[3]Retrieved from: `https://www.londonair.org.uk`

[4]Retrieved from: `http://www.orangefreesounds.com/air-on-a-g-sring/`

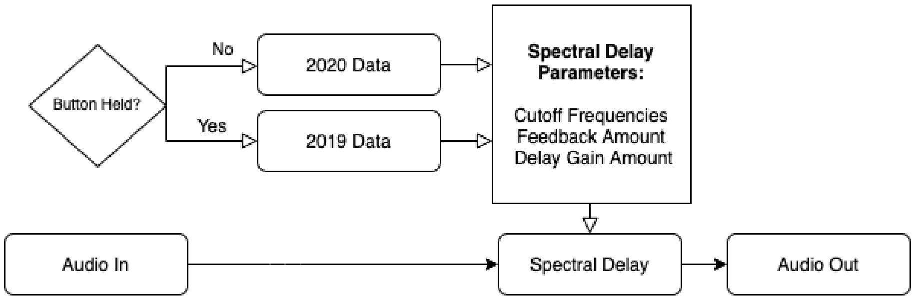

Figure 2. Flow diagram of the sonification process.

data correspond to higher values of feedback of the delay effect; higher data values are mapped to higher delay gain. This is so that low pollution/NO_2 levels, approximate the output to a 'clean' rendition of the piece, whilst high pollution/NO_2 levels 'pollute' it with increasingly high amounts of the delayed signal. A demonstration video, as well as all the code and resources, are made available online[5].

4.3. Discussion

Some of the design decisions for this use case were driven by practical constraints (e.g. limiting the number of filter-delay lines to stabilize CPU usage on the Bela board). Others were the result of experimentation and subjective listening evaluation by the designer in an iterative fashion (e.g. choice of quality factor, range of cutoff frequencies, delay parameters). Aesthetic decisions were taken considering both aesthetics aspects and the ability to convey information. The choice of musical borrowing as a basis technique stirs the project precisely in that direction, privileging engagement with the musical work instead of information transmission accuracy. However, intrinsic musical variations of the original piece may mask data-driven modulations, so a careful choice of borrowed material is required. This effect is reduced when there is a solid familiarity with the piece, allowing the listener to distinguish between what is changing due to data from changes that are part of the original piece. In order to assess the didactic value and aesthetic experience of interactive musical works generated with the proposed method, evaluations should be conducted with listeners. Expansion of this work include its integration with real-time pollution data streams in a smart city environment.

5. Use case 2: Holonic Systems - Become Sound, Musification of User Context

5.1. Background

Holonic Systems[6] is a company which started in 2016. The apps from Holonic Systems produce music resulting from user activities and interactions with the environment. Certain musical elements are structurally coupled to phenomena that can be traced within a smart city ecology, such as proportions of vehicles, weather, or time of day. The approach aims to ensure that the musical output of the participants (users or audiences) is related to the activity at hand.

[5]Due to anonymity, available at: `http://drive.google.com/drive/folders/1hb6-uW3m3a27yN8-mD5RaS5H6ZaljpR6?usp=sharing`

[6]Available at: `https://www.holonic.systems/`

5.2. Design Principles

A key feature in the design of Holonic Systems interactive music apps is the reliance on passive, even non-volitional, operation that is based on the mapping of accompanying user gestures in everyday activities. This choice was taken to counter the physical fatigue that active control would present if used during prolonged times. It was also chosen to limit potential theatrical gesticulation for use in everyday environments and encourage a more natural experience. It relies on a so-called Zero UI paradigm, in which user input is minimal. The assumption behind the decision was that limiting screen time increases immersion towards the generated music.

5.3. User Context Musification with Holonic Systems Apps

We describe here two mobile applications, a generative music app, Holon, and an authoring tool called Holon.ist. Both apps are native, universal iOS applications written in Swift and can run on iOS mobile devices.

5.3.1. Holon.ist App

The Holon.ist app lets users create their own mappings from input sources to outputs. Its interface is depicted in Figure 3. Inputs include onboard sensors and external data, such as the weather service Open Weather Maps. Inputs can also come from hardware sensors, such as the Apple Watch (Series 3 or newer), Apple AirPods Pro, Estimote iBeacons and the smart sensor platform Movesense. Outputs are either in OSC or MIDI formats.

Figure 3. Interface of the Holon.ist app.

Input sources include data relating to postures. These can be learned in a way similar than using a MIDI Learn function in a digital audio workstation. Custom locations can also become data sources by either learning the current GPS location, or by entering longitudinal or latitudinal values. The distance to or azimuth of a given location can be used as a modulation source, thereby indicating the proximity and direction of a physical place. Entering a location can be used as a switch to trigger a new state in the music composition. Similarly, relative sun and moon positions are calculated from the current user location and are made available as data sources. Sunrise and sunset times are also calculated according to user position and the amount of sunlight is expressed at dusk and dawn as an indicator of twilight. As the sun has a measurable physical effect (e.g. body temperature, circadian phase), this can be used for musical effect in a number of ways.

5.3.2. Holon app

The consumer app Holon, shown in Figure 4, is a packaged Pure Data patch running on the embedded libPD engine on the Holon.ist platform. No samples are used, instead, sound is synthesised entirely from first principles. Holon includes 85 subpatches that contain features like adaptive activity-based routing, harmonic rules and various voices for each context; stationary, walking/running, vehicular and unknown. The voices use different synthesis methods ranging from subtractive synthesis to FM synthesis and Karplus-Strong.

Figure 4. Interface of the Holon app.

The engine adapts to user input devices automatically. For example, one can use the existing onboard sensors of the iPhone or iPad: accelerometer, gyroscope, magnetometer, barometric air pressure or altitude, as well as GPS. However, when a compatible sensor, such as an Apple Watch, is connected, the inputs are replaced with the ones coming from the wearable. Using external wearables adds additional dimensions of movement and expressivity. For example, user heart rate replaces an average default resting heart rate, which is then mapped to the master tempo. The sonified heartbeats can be heard when a user is sitting down. The sitting posture is determined from device orientation, although it requires users to place the phone in their pocket as instructed. When the user is stationary and sitting down, it is assumed that attention is inward, towards the self. When mobile, attention is external and tempo is derived from step rate, or in the case of vehicular sonification, speed derived from the GPS. The vehicle mode features a subtle synthetic engine sound, which is pitched towards the current generated root note determined by speed. This was added to provide drivers with a reference as a safety feature to determine speed and acceleration. Holon also supports ascent/descent sonification. A pitched sound can be heard when travelling vertically in a rapid manner, for example in an elevator. This is possible by using the barometer, which also reacts to doors closing, or even the column of air in front of an approaching subway train. Time of day is also a music composition factor, as is orientation towards the sun. Facing the sun activates a synthetic string section, which is panned according to solar position relative to the user.

6. Conclusion and Perspectives

In this article, we presented a vision towards ubiquitous music in smart cities, where AI is involved in the creation of musical content conditioned by IoT data. We discussed several approaches for the design of musical smart city tools, emphasizing on possible musification methods. Finally, we discussed two use cases that focus on different aspects of the smart city data spectrum, namely an interactive musification of air pollution and consumer mobile music apps driven by user contextual data. In order to respond to some of the challenges of a musical smart city described in this paper, future work can address the creation of AI-based models suitable for musification. A promising direction concerns the development and assessment of music generation models driven by user contextual data using deep learning.

7. Acknowledgements

This work is supported by the EPSRC UKRI Centre for Doctoral Training in Artificial Intelligence and Music (Grant no. EP/S022694/1) and partly supported by Holonic Systems.

References

Ballora, M. (2014). Sonification, Science and Popular Music: In Search of the 'Wow'. *Organised Sound*, 19(1):30–40.

Barrass, S. (2012). The Aesthetic Turn in Sonification Towards a Social and Cultural Medium. *AI and Soc.*, 27(2):177–181.

Barrass, S. and Vickers, P. (2011). The Sonification Handbook, Chapter 7: Sonification Design and Aesthetics. In *The Sonification Handbook*, pages 145–171. Logos Publishing House, Berlin.

Barth, A., Karlstrom, L., Holtzman, B. K., Paté, A., and Nayak, A. (2020). Sonification and animation of multivariate data to illuminate dynamics of geyser eruptions. *Computer Music Journal*, 44(1):35–50.

Barthet, M., Depalle, P., Kronland-Martinet, R., and Ystad, S. (2010). Acoustical correlates of timbre and expressiveness in clarinet performance. *Music Perception*, 28(2):135–154.

Bicknell, J. (2001). The Problem of Reference in Musical Quotation: A Phenomenological Approach. Technical report.

Bødker, S. (2015). Third-wave hci, 10 years later—participation and sharing. *interactions*, 22(5):24–31.

Bonet, N. (2019a). *Data Sonification in Creative Practice*. Phd thesis, University of Plymouth, School of Humanities and Performing Arts.

Bonet, N. (2019b). Musical Borrowing in Sonification. *Organised Sound*, 24(2).

Briot, J.-P., Hadjeres, G., and Pachet, F.-D. (2019). *Deep Learning Techniques for Music Generation – A Survey*.

Bryan, N. J., Herrera, J., Oh, J., and Wang, G. (2010). MoMu: A Mobile Music Toolkit. In *Proc. of the 2010 Conf. on New Interfaces for Musical Expression*, pages 174–177.

Cugurullo, F. (2018). The Origin of the Smart City Imaginary: From the Dawn of Modernity to the Eclipse of Reason. In *The Routledge Companion to Urban Imaginaries*, pages 113–124.

de Campo, A. (2007). Toward a Data Sonification Design Space Map. *Proc. of the 13th Int. Conf. on Auditory Display*, pages 342–347.

Essl, G. and Lee, S. W. (2017). Mobile Devices as Musical Instruments - State of the Art and Future Prospects. In *Proc. of the 13th Int. Symposium on CMMR*, pages 364–375, Portugal.

Ferreira, L. N. and Whitehead, J. (2019). Learning to Generate Music with Sentiment. In *Proc. of the 20th Int. Soc. for Music Information Retrieval Conf.*

Goldman, A. H. (2006). The experiential account of aesthetic value. *The Journal of Aesthetics and Art Criticism*, 64(3):333–342.

Hazzard, A., Benford, S., and Burnett, G. (2015). Sculpting a Mobile Musical Soundtrack. In *Proc. of the 2015 Conf. on Human Factors in Computing Systems*, pages 387–396.

Joy, J. and Sinclair, P. (2008). Networked sonic spaces. In *ICMC*.

Kaliakatsos-Papakostas, M., Gkiokas, A., and Katsouros, V. (2018). Interactive Control of Explicit Musical Features in Generative LSTM-Based Systems. In *ACM Int. Conf. Proceeding Series*.

Keller, D., Schiavoni, F., and Lazzarini, V. (2019). Ubiquitous Music: Perspectives and challenges. *Journal of New Music Research*, 8215:1–7.

Kramer, G., Walker, B., Bonebright, T., Cook, P., Flowers, J., Miner, N., and Neuhoff, J. G. (1999). Sonification Report: Status of the Field and Research Agenda. Technical report.

Lepri, G. (2020). Sonification as Negotiation - Learning from Translation Studies. In *Proc. of the 17th Sound and Music Computing Conf.*, pages 212–219, Torino.

Mandanici, M. (2019). Sounding Spaces for Ubiquitous Music Creation. In *Procedings of the 14th Int. Symposium on Computer Music Multidisciplinary Research*, pages 675–684, Marseille, France.

McPherson, A. P. and Zappi, V. (2015). An Environment for Submillisecond-Latency Audio and Sensor Processing on Beaglebone Black. In *138th Audio Engineering Soc. Convention 2015*, volume 2, pages 965–971.

Meade, N., Barreyre, N., Lowe, S. C., and Oore, S. (2019). Exploring Conditioning for Generative Music Systems with Human-Interpretable Controls.

Neuhoff, J. G. (2019). Is Sonification Doomed to Fail? In *The 25th Int. Conf. on Auditory Display*, number June, pages 25–28, Northumbria.

Oore, S., Simon, I., Dieleman, S., Eck, D., and Simonyan, K. (2018). This Time with Feeling: Learning Expressive Musical Performance. *Neural Computing and Applications*, 32(4):955–967.

Pigrem, J. and Barthet, M. (2017). Datascaping: Data Sonification as a Narrative Device in Soundscape Composition. In *Proc. of the 2017 Audio Mostly Conf.*, pages 1–8.

Pimenta, M. S., Flores, L. V., Capasso, A., Tinajero, P., and Keller, D. (2009). Ubiquitous Music: Concept and Metaphors. *Proc. of the Brazilian Symposium on Computer Music*.

Ramos, F., Trilles, S., Muñoz, A., and Huerta, J. (2018). Promoting Pollution-Free Routes in Smart Cities Using Air Quality Sensor Networks. *Sensors*, 18(8).

Risset, J.-C. (2004). The Liberation of Sound, Art-Science and the Digital Domain: Contacts with Edgard Varése. *Contemporary Music Review*, 23(2):27–54.

Sarmento, P., Holmqvist, O., and Barthet, M. (2020). Musical Smart City: Perspectives on Ubiquitous Sonification. In *Proc. of the 2020 Ubiquitous Music Workshop*.

Sarmento, P., Kumar, A., Carr, C., Zukowski, Z., Barthet, M., and Yang, Y.-H. (2021). DadaGP: a Dataset of Tokenized GuitarPro Songs for Sequence Models. In *Proc. of the 22nd Int. Soc. for Music Information Retrieval Conf.*

Satyanarayanan, M. (2001). Pervasive Computing: Vision and Challenges. *IEEE Personal Communications*.

Sturm, B. L., Santos, J. F., and Korshunova, I. (2015). Folk Music Style Modelling by Recurrent Neural Networks with Long Short Term Memory Units. *Proc. of the 16th Int. Soc. for Music Information Retrieval Conf.*

Turchet, L. and Barthet, M. (2019). Smart musical instruments: Key concepts and do-it-yourself tutorial. *Foundations in Sound Design for Embedded Media: A Multidisciplinary Approach*, page 136.

Turchet, L., Fischione, C., Essl, G., Keller, D., and Barthet, M. (2018). Internet of Musical Things: Vision and Challenges. *IEEE Access*, 6:61994–62017.

Vickers, P. (2016). Sonification and Music, Music and Sonification. In *The Routledge Companion to Sounding Art*, pages 135–144. Routledge, London.

Vieira, R., Barthet, M., and Schiavoni, F. L. (2020). Everyday use of the internet of musical things: Intersections with ubiquitous music. In *Anais do X Workshop de Música Ubíqua (UbiMus 2020)*, page 60.

Weiser, M. (1991). The Computer for the 21st Century. *Scientific American*, pages 94–104.

Winters, R. M., Kalra, A., and Walker, B. N. (2019). Hearing Artificial Intelligence: Sonification Guidelines & Results From a Case-Study in Melanoma Diagnosis. In *The 25th Int. Conf. on Auditory Display*, pages 23–27.

Winters, R. M., Tsuchiya, T., Lerner, L. W., and Freeman, J. (2016). Multi-Modal Web-Based Dashboards for Geo-Located Real-Time Monitoring. *Proc. of the 2nd Web Audio Conf.*

Worrall, D. (2018). Sonification: A Prehistory. In *Proc. of the 24th Int. Conf. on Auditory Display*, number ICAD, pages 177–182.

Zanella, A., Bui, N., Castellani, A., Vangelista, L., and Zorzi, M. (2014). Internet of Things for Smart Cities. *IEEE Internet of Things Journal*, 1(1):22–32.

Zukowski, Z. and Carr, C. (2017). Generating Black Metal and Math Rock: Beyond Bach, Beethoven, and Beatles. In *31st Conf. on Neural Information Processing Systems*, number NIPS.