Evaluating an Urban Installation - User Engagement, Shared Encounters and Spatial Influence

Bo Jacobsen
Aalborg University
Aalborg
bjacob14@student.aau.dk

Adam Kjær Søgaard
Aalborg University
Aalborg
asagaa14@student.aau.dk

Rune Lundegaard Uggerhøj
Aalborg University
Aalborg
rugger14@student.aau.dk

ABSTRACT

Introducing artefacts to the urban space can drastically change how that space is perceived and used by those who frequent the area. The fields of Smart- and Playable Cities respectively exploit this in an attempt to increase the efficiency of the city and the quality of life of its inhabitants, by introducing smart installations to the urban environment.

With that in mind, we set out with the intention of digitally augmenting the urban space to be more playful. A playful, interactive, urban installations with a with a focus on sonic elements was developed, based around the design of a pan flute, and evaluated through a series of deployments studies.

The results indicate that a recognisable design and adaptive modes of interaction are effective ways of creating an interactive urban installation that has a measureable impact on the space around it, as it was found to interfere with the natural flow of up to 17% of passers-by in a non-obstructive manner. Further iteration could likely increase its effectiveness, as modes of interaction become more congruent with the perceived affordances. .

1. INTRODUCTION

An artefact of any kind can change how humans act in an area drastically, take for example a slide on a public lawn. The lawn might only have been used to cross faster before introducing the artefact, but now both children and adults will occupy the area, with children frolicing and adults looking on. Similarly, a sign indicating the correct route to walk, will likely divert the foot traffic of that area. These concepts, if applied with modern technology, can be referred to respectively as Playable- and Smart Cities [10, 11, 23], and have in recent years risen in popularity and extent. With this rise has also come research, however, most of this has been within the Smart Cities field, meant to increase the effectivity with which the city functions. Digitally playable public installations is a much more scarcely researched subject, with no widely accepted guidelines for neither design nor evaluation, besides the underlying paradigms for usability and user experience.

With this in mind, there are notable considerations to be taken and variables to be addressed when designing and evaluating such an installation. Lures [16], affordances [24], and how well an installation fits onto a site are for example all important factors to be considered when designing it, while the evaluation requires fitting tools and methods to interpret the naturally chaotic data that is inherent with In the Wild deployments [29].

In this paper we attempt to design an arbitrary installation to be deployed publicly, while reflecting on, and evaluating the effects and elements identified in previous works, with the purpose of demonstrating an effective implementation and evaluation of an HCI based public installation. This is done by implementing a large pan flute, meant to attract users through both auditory and visual lures, and allow both group- and individual use. This is tested over three days of deployment, over a period of 19 hours, with a total of 263 use cases in groups of varying sizes. The data is evaluated from multiple perspectives, being spatial influence, user behaviour, and group use.

The results suggest that the installation is moderately intuitive in use, as users were quick to realise how to interact with and operate the device, and that groups are more likely to interact. The 263 use cases added up to a total of 13.56% of passers-by interacting in some level with the installation, which is argued to be a sizable influence on the surrounding space.

2. RELATED WORK

This section covers related work with regard to two subjects: work that details interactive urban installations, and work elucidating and applying methods for designing and evaluating public installations. Regarding the former, for this paper we decided to focus on installations that sonically enhance the urban space, as audio might have potential as an effective lure, since it is not as easily obstructed by the presence of other passers-by as visibility might be.

2.1 Interactive Urban Installations

As outlined in [13], sonic installations can be broadly separated into three categories:

Instrumental Installations - Users can directly manipulate the sonic expression of the installation through a tangible interface.

Non-instrumental Installations - The installation lacks a tangible interface but remains interactive, with the system detecting and reacting to the actions of users in the vicinity.

Non-interactive installations - The installation does not detect or react to the actions of users, but may remain context sensitive by e.g. changing the sonic expression in response to changes in temperature or light level.

This categorisation is somewhat widely applicable, thus for the purpose of this paper, it is extended to also apply to urban installations that do not primarily rely on sonic expressions.

The Sonic City project [14] falls into the non-instrumental category. Although it is not strictly speaking an urban installation but rather a wearable interface, it accomplishes several of the same goals of other playable urban artefacts using some of the same means. Specifically, it facilitates awareness of a user's surroundings by reacting to the city as the user moves through it, as well as reacting to gestures and motions of the user. It combines these two elements to generate electronic music for the user, giving them a source of personal enjoyment. Another non-instrumental installation can be seen in [6], although this one straddles the line between being instrumental and non-instrumental. It has an interface, but the interface is designed in such a way as to facilitate inadvertent interaction which in turn makes the user aware of the interface, allowing them to transition into more deliberate interaction. It uses proximity sensors to detect users, generating a myriad of sounds based on the user's position and gestures.

Instrumental installations seems the more common type of urban installation, as evident by their prevalence in the sources. Many examples exist of installations that are primarily sonic in nature [12, 13, 18, 21, 25, 28, 32, 33]. Prominent examples of these are the Hummingwall [21], a large scale installation which uses sound both as feedback and to lure in potential users, where users can collaborate to generate interesting soundscapes, and the Piano Staircase [32], which added sounds to individual steps of a subway staircase, adding a playful element to them which had a clear effect on the behaviour of pedestrians, causing more of them to use the stairs instead of the escalators.

As for instrumental installations making use of lights as primary feedback, Kick-flickable [11] was a collection of glowing artefacts which users could manipulate by kicking, picking up, and throwing. It was found that users were willing to incorporate the artefacts in creative play, coming up with their own ways of gaining enjoyment from the interaction. Another example of an installation primarily making use of light is Responsive Lighting [26]. This installation augmented a public area with coloured light which, depending on use-mode, either responded to pedestrian movement around lamp posts, or through direct input via a mobile application. Thus whether this installation counts as instrumental or non-instrumental depended on use-mode.

There are also successful installations making use of both lights and sounds as feedback, SwingScape [15] being one of those. SwingScape made use of swings, each having different notes and lights to create a playful zone, and using lights as safety measures, where depending on your placement your experience might differ.

2.2 Analysis Measures

The intended subjects for analysis are user engagement, shared encounters, and spatial influence, and as such, this section details concepts related to these.

Regarding users' interaction with, and use of, an installation, sources reveal several relevant concepts. First of all, one might observe users engaging in "championing", such as observed in [3]. Championing describes the act of users, entirely of their own volition, promoting an installation to other potential users and inviting them to engage with it, as well as teaching them how it works [2].
Another common method is to categorise the way users engage and interact with an installation, separating the different instances of use into use categories, such as in [5, 7, 17]. These categories can be as precise or broad as fits the purpose of the evaluation, and thus they can serve as a useful tool to outline and describe the overall use of an installation. This paper makes use of the use categories detailed in [17], which are the following: Active Participation (AP), Active Exploration (AE), Passive Exploration (PE) and Passive Observation (PO). Active participation is when a user interacts directly with the installation, active exploration is when users explore the installation up close, but do not make use of the functionality thereof, passive exploration is when users stop to observe, but do not approach, and passive observation is when users simply observe, but do not stop to observe or interact.

Then there is the concept of unexpected use [9]. This is described as users interacting with the installation in a way that was not conceived of - or intended - by the designers. It is argued that this unexpected use is not a fault of the design, and thus should not be seen as a negative. The reasoning behind this argument is that users would have to come up with the use, which they would not do if they had no interest in interacting with the installation in the first place.

Several sources, that have conducted in the wild studies with public installations, indicate that group use is more prevalent than single user use [18, 19, 21]. This is an important factor to keep in mind, especially if spontaneous group formation (SGF) is a goal. SGF describes a phenomenon where users, who otherwise had no plan to interact socially with one another - typically either because they are unaffiliated to each other or simply are in the space for different reason - form a social group triggered by the installation [19]. An installation should thus preferably have the option for multiple users to use it at once, especially as, if the installation is created with single user use in mind, users will still attempt to use it collaboratively [17].

With spatial influence, we refer to how an installation affects the surrounding area. This may be in the shape of changing the routes by which people traverse the area, or changing the general area in which people spend time. This can for example be measured in flow interference, which is a measure that describes the percentage of passers-by affected by the installation in question [11]. Example of measuring flow interference can be found in [2, 21, 22], and is shown to be a useful tool to supplement or be supplemented by use cases. It is, however, not an equally relevant measurement for all types of installations, as installations with a focus on engaging many users at once will likely have a higher flow interference than one made for fewer users.

3. DESIGN

Inspired by the concept of instrumental, sonic installations outlined in section 2, it was decided to create a playful instrumental design that sonically enhances the city. The final design can be seen in Figure 1. In this chapter we lay out the reasoning behind this design, and its technical implementation.

Figure 1: The final iteration of the pan flute.

3.1 Theory

The design of the installation has been considered with regard to several theoretical concepts. This section aims to explain what these concepts are, why they are relevant, and how the design was shaped around them.

3.1.1 Principles of Universal Design

Universal design is a method of design described by Preiser [27]. It outlines how you design a system while accommodating for a wide or nonexistent target group, through application of the following seven principles:

1: Equitable Use - Provide identical or equivalent means of use, without stigmatizing a particular user group.

2: Flexibility in Use - Provide a system where the user has an ability to chose how to interact.

3: Simple and Intuitive Use - Provide a design that is easy to understand, without unnecessary complexity while conforming to user expectations.

4: Perceptible Information - Effective communication of all necessary information, regardless of the user's abilities.

5: Tolerance for Error - Minimize the possibility for errors and hazards, and allow fail-safe conditions when such cannot be completely eliminated.

6: Low Physical Effort - Provide a design that is comfortable to use.

7: Size and Space for Approach and Use - Make users able to have the necessary space to interact, and make the design of a size where all users can reach the necessary elements.

These principles are relevant to the design of this installation, as a playful interactive urban installation conceptually is deployed in a fully public space, and thus is it difficult to control which kinds of people will interact with it. Moreover, the idea behind an installation in the field of the playable city is to improve the quality of life of a city's inhabitants, and thus it makes sense create a design that appeals to as many citizens as possible for maximum effect.

Touch input was chosen as the main method of interaction for the installation as this type of interaction caters well to principles 1 through 3. The amount of users excluded from interacting by touch is minimal, as all it requires are the use of your arms. As for flexibility, by spacing the touch surfaces equally the design is usable by left-handed as well as right-handed users. Moreover, not only is touch a common and well understood method of interaction, but designing the installation as a pan flute with clear, individual pipes helps communicate that each pipe does something different, ie. each pipe plays a particular note. As for principles 6 and 7, not only does touch input not require any particular physical strain, but the height of the installation was chosen such that the touch surfaces could be reached comfortably by adults as well as children; the mean height for children age 8 is roughly 130 cm [30, 31], thus the installation was designed to be 110 cm tall (final height was 108 cm).

The biggest drawback to touch as an input with regard to this particular design is that it is not directly congruent with how a pan flute usually works. It was considered to implement a way for users to interact by blowing air onto the pipes, however, this idea was discarded due to the intended placement of the installation. When positioned outside, a slight breeze might trigger the sensors, which - while it would arguably be an interesting art installation to observe - would result in the user not having full control over which pipes were playing audio. An alternative method of interaction which was considered, was drumming or tapping. This method has the advantage of being potentially more fun and engaging to observe, enhancing the honeypot effect - a phenomenon where passers-by observe users engaging with an artefact, and thus become interested in engaging with it themselves [8]. However, not only could frail people have a harder time interacting with the installation in this manner, but this method of interaction is even less congruent with how a pan flute work, than touch input is; a user needs to be able to maintain a continuous note as one would when applying an air stream to a pan flute, which would be difficult to do with a discrete input type such as tapping.

3.1.2 Entry Points

Of further relevance to the design of an interactive urban installation is the concept of entry points [16] - especially the idea of the progressive lure. This concept describes a string of lures, enticing a potential user to ultimately engage with the content (in this case interact with the installation) by drawing them from lure to lure.

With an urban installation of relatively small scale it is difficult to have a web of lures, one leading to the next, without putting up signs over a large area pointing in the right direction. That said, this design has its own small set of progressive lures. First of all there is the installation itself. Although this is the target towards which the lures should lead, it is in this case likely the first sign a passer-by will see of the installation. By designing the installation to look like an instrument with which the user is likely familiar, this first look should hopefully pique their interest. As the user then approaches the installation, they will enter an area - or point of prospect [16] - from which they can perceive the second link in the chain of progressive lures: light and audio. The intent behind these extra elements is to increase the odds that a potential user will approach the installation, even in the case that the user had already seen it but had no intention to explore. Moreover, to make sure that the light and audio not only enhance the potential honeypot effect, they were designed to animate even when the installation is not in use to draw in more passers-by.

This idle-animation of light and audio also helps explain what the installation is to potential users through more than one channel, accommodating the 4th principle of universal design. The idle-animation specifically demonstrates how the installation can be played, by drawing out the duration of the sounds, thus indicating that continuous touch creates a continuous sound. The current delay for how often the idle-animation triggers is 40 seconds, which is deemed to be frequent enough that users will likely be subjected to it during their pass through the area with the most effective points of prospect [20], ie. the area where the audio can be heard, and the installation can be clearly seen and identified. 'Likely' in this scenario refers to a likelihood of approximately 50%, which is calculated with the assumption of an average walking speed of five km/h, and the points of prospect extending in a radius of approximately 20 meters around the installation.

3.1.3 Collaborative Use

Looking at papers with comparable goals or installations, e.g. [21], it was observed that groups have a higher likelihood of interacting with this type of installation. As such, considerations were made to accommodate this. Primarily, the pan flute design of the installation was scaled to a degree where multiple users can interact with the installation simultaneously without getting in each others' way, and interaction is possible from both sides of the installation, which further strengthen the installation with regard to the 7th principle of universal design.

3.1.4 Sites

The intention with this installation, was to deploy it in areas where users are relaxed and open to such types of experiences, and thus the nature of the site is important. For this purpose, we group different sites into three categories: Semi-public Spaces, Transit Spaces, and Leisure Spaces. A semi-public space is a space where the demographics of those present are controlled to some extent, e.g. inside a university or a large company. A transit space is a space which is primarily used for traversal, for example sidewalks and bus terminals. Finally, leisure spaces describe areas where people can relax and allow themselves to be distracted, e.g. in a public park. By these definitions, a leisure space is likely the best option for the deployment of this installation.

3.2 Implementation

This section documents the implementation of the pan pipes, and which materials were used.

3.2.1 Hardware

As mentioned in section 3.1, instead of using a blowing or drumming method of interaction, touch was used, due to various factors. To implement the touch input, the installation makes use of Grove touch sensors [1], with the sensing surface enlarged through the use of copper tape attached to the underside of an acrylic glass surface (see Figure 2). This allowed for sensing without the user having actual physical contact with the components, and was a positive both aesthetically, as well as with regard to durability. Two Arduinos were used to split up the sensors in two groups, one with six sensors and one with five (Figure 3 shows one of these Arduinos). This was done in order to shorten the wires between the Arduino and the sensors, as it was found that long wires (40 cm and above) overlapping each other led to a significant increase in noise.

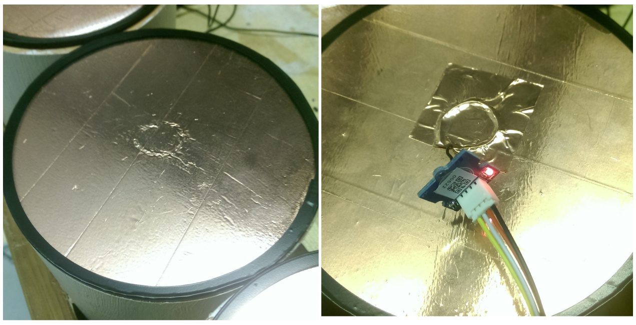

Figure 2: The Grove sensor attached to the underside of an acrylic glass surface, on copper tape.

The Ardunios were connected to a Raspberry PI 3B through USB connections. Two speakers were connected to the Raspberry through its jack port. The speakers were placed in each outermost pipe so that, through implementation of stereo sound in the software, the installation can emulate the appropriate origins of the sounds. Specifically, the audio of each pipe emits - depending on how far left or right the touched pipe is - either only from the left or right speaker, or proportionately between them.

LEDs were installed on the underside of the frame (see Figure 4) with eight for each pipe such that touching a pipe will result in light emitting from that pipe. However, it

[1] https://www.trossenrobotics.com/p/grove-touch-sensor.aspx - Last Accessed: 21/05/2019

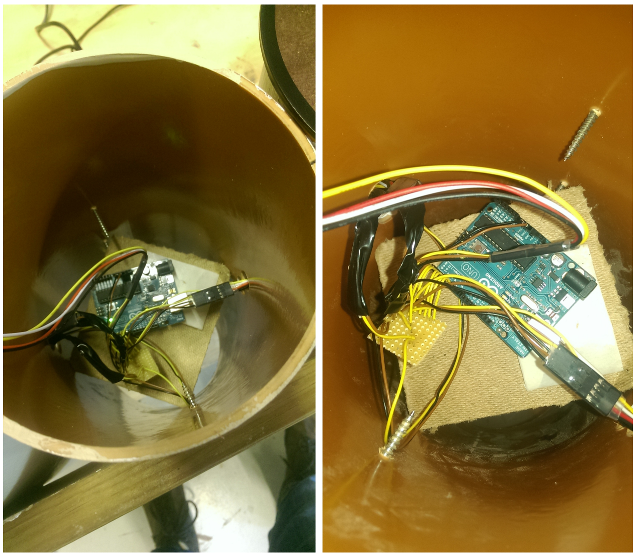

Figure 3: One of the Arduinos, in its position in a pipe.

was found that lighting all LEDs simultaneously required a greater current that the Raspberry is capable of outputting, resulting in brief loss of power on the Arduinos, crashing the software of the Raspberry. To overcome this, the number of LEDs for each pipe was reduced to six, 66 in total, allowing for all LEDs to be lit at once without the system shutting down.

Figure 4: The LED lights mounted to the pan pipe installation.

3.2.2 Software

The software of the installation was comprised of a python script on the Raspberry, communicating with the two Arduinos through serial.

The Arduinos continually read the signal from the sensors and keep track of their current state in a string of 'y's and 'n's. When a sensor fires (or stops firing) the corresponding character in the string will update, and the Arduino will then check if the new string is different from the previous string, and only if it is will it write it to the serial connection. After this step, the Arduino then iterates through the characters of the string it just wrote to the serial, turning on all LEDs that corresponds to a touched pipe.

Once the Arduino writes a string to the serial, this string is read by the Raspberry in the python script. Here, the strings from each Arduino are concatenated, and the script then iterates through the characters in this new, longer string. For each 'y' it plays the sound corresponding to that characters' index, as long as that sound is not already playing, and for each 'n' it stops the playback of that sound. The audio itself was sampled, and stored in the form of .wav files.

Additionally, the python script keeps track of for how long a user has been interacting with the installation. When interaction seizes, the duration is logged to an external file. Finally, the script also keeps track of how long has passed since the last interaction with the installation, and if more than 40 seconds has passed, the script will initiate the idle-animation, playing a sequence of notes and lighting the corresponding LEDs. This functions as one of the installations progressive lures.

Technical Difficulties

Despite the Grove sensors being intrinsically reliable, a combination of the length of wires running from the sensors to the Arduinos and the density of said wires created interference between the sensors, which caused some of them to fire sporadically, resulting in noise in the input. This noise resulted in many false positives, and the majority of these occurred while a user was interacting with the installation, not while it was idle. After attempts to eliminate the noise by further insulating the wires were unsuccessful, a solution was found by implementing a degree of artificial latency in the system. The Arduinos were given a variable determining for how long a sensor would have to fire continuously (in milliseconds) before the Arduino would accept it as input, rather than regard it as noise and ignore it. It was found that all of the noise was eliminated with an artificial latency of 250 ms. This was not an optimal solution, as not only can too much latency be disruptive and unpleasant to the user, but this artificial latency could result directly in false negatives if the user touches a sensor for less than the specified duration. Due to this, it was decided to carry out a lab test to explore how much noise vs. latency users were willing to accept.

4. CONTROLLED EVALUATIONS

Two lab tests were conducted. One with the purpose of identifying the relationship between latency and sensor noise (leading to false positives) most preferred by users, and a second one with the purpose of measuring general usability.

4.1 Latency Test

The latency test was conducted prior to the usability test. The test was conducted on a moderately warm day, in a courtyard, as the wind would otherwise interfere negatively. The reason the test took place outside, was to properly simulate the circumstances of an actual test, as the reverberation of the sound and the capacity of the sensors are affected by the environment.

Method

The test followed a within-subject, repeated measures design. The independent variable was the length of the artificial delay, of which three different settings were used: 70 ms, 160 ms & 250 ms. The dependent variable was the order in which the user preferred the settings.

Participants
The participants for the test were recruited from amongst students at a university building. A total of six people partook in the test, 5 of whom were male. The age-range was approximately 20 to 30 years. These students were all familiar with this type of installation, on account of their study direction.

Procedure
The participant was introduced to the prototype, being told it was a digital instrument and instructed in its basic functions. They were told that they were to interact with the prototype in three sessions of 25 seconds, in between which the facilitators would change some setting. The participant was not told which settings were changed in between sessions. The order of the three latency settings were balanced between test participants using a latin square method. After the final session the participant was asked to rank the three sessions by personal preference, by arranging cards labeled 1-3, after which they were asked why they ranked the sessions in this particular order. If latency or noise was not brought up by the participant they were asked directly if they noticed the difference in latency. Finally, they were asked for any closing comments.

Results
The results of the rankings of the three latency settings are shown in Table 1, and have been scored by weight of rank. The results do not show a clear preference for any of the three settings, although they do suggest that the participants prefer to trade some noise for latency based on the 70 ms setting being ranked at half the score of the other settings. At the same time no preference between medium or high latency is evident. Based on this lack of preference, it was decided to aim for a middle ground of medium latency and medium noise.

Table 1: Latency Test Results

	1st	2nd	3rd	Score
70 ms	0	3	2	3
160 ms	2	2	1	6
250 ms	3	0	2	6
Weight:	2	1	0	

The frequency of rankings for the three latency settings, and the weighted score of each setting.

4.2 Usability Test
After making a choice regarding the latency, a usability test was conducted. The test took place outside on a warm day, once again to be as comparable as possible to an actual deployment of the installation (see Figure 5).

Method
The test followed a between subject, repeated measures design, where no independent variables were being manipulated. Success was measured qualitatively and quantitatively, via semi-structured interviews and participants' answers on a complete System Usability Scale (SUS) [4].

Figure 5: The installation as positioned for the usability test.

Participants
The participants were university students, recruited in a university building. A total of 15 people partook in the test, 13 of whom were male. The participants were in the age range of 20-29 ($\mu = 23.5, \sigma = 2.4$)

Procedure
The participant was introduced to the purpose of the test, and was informed that it was a usability test. They were then told to interact with the installation until they felt they were comfortable with the system. They were not introduced to any of the functions of the installation. Once the participant expressed that they were comfortable with how the system worked, they were asked to fill out a questionnaire containing the SUS, as well as questions about their age and gender. Once the questionnaire had been submitted, the participant was asked if they had any general comments on the installation. Finally, if had not been brought up by the participant themselves, the facilitator would ask if they had recognised the installation as a pan flute.

Results
The results of the SUS showed a general acceptable usability of the system, with a mean score of 84.67 ($\sigma = 8.23, median = 85$). This is above the threshold of 81, which is argued to mean that the system being tested is usable, without major flaws or shortcomings. This suggests that the installation is usable in its current configuration, which is likely a result of the simplicity of the design. However, observations from the test as well as comments from the participants highlights an issue: the affordances of the installation do not reflect the intended input method. During the test, 10 of the 15 participants (66.66%) initially thought that drumming or tapping was the correct way to interact with the installation, and 4 of those 10 never fully realised that touch was the intended method of interaction. Two of these four became confident that hovering their hand above the pipes was the intended method (it was possible to trigger the installation in this way due to the high sensitivity of the sensors), with one of them commenting that they could not figure out at what distance from the pipe their hand had to be to trigger the sound. Finally, two participants initially looked for a way to trigger the installation through blowing air on it, assuming the method of interaction was the same as that of an actual pan flute. All of this combined suggests

that the affordances should optimally be made clearer, to properly communicate the correct method of interaction.

One thing that was clear, however, was the intended look of the installation. All participants recognised it as a type of musical instrument, and only two thought it was meant to be something different from a pan flute (xylophone and piano-like, respectively).

4.3 Iteration

Following the two lab tests, several iterations and considerations were made regarding the installation, some of which were directly informed by the results of the lab tests.

4.3.1 Latency

A modification was made to the circuits of the installation, which allowed for a significant reduction in the artificial latency. A capacitor was soldered onto each sensor, connecting the copper tape directly to the ground on the sensor (see Figure 6). This meant that, whenever a sensor was touched, part of the capacitive signal would go directly to the ground, rather than into the input of the sensor, making the relative change in capacity smaller, thus resulting in lower sensitivity. Moreover, an idea was conceived to place a wooden board on top of the installation whenever power was first connected. Doing this meant that the sensors performed their initial calibration with a higher capacitive threshold, again lowering the sensitivity of the sensors. While this method can be seen as a suboptimal solution, it was deemed acceptable for the purposes of a prototype, as the sensitivity of the sensors cannot be manually manipulated directly in the chip.

Combined, these two changes meant that the threshold for an artificial latency that would eliminate all noise was lowered from 250 ms to 120 ms. This new lower threshold was selected as the level of latency for the installation going into the deployment tests.

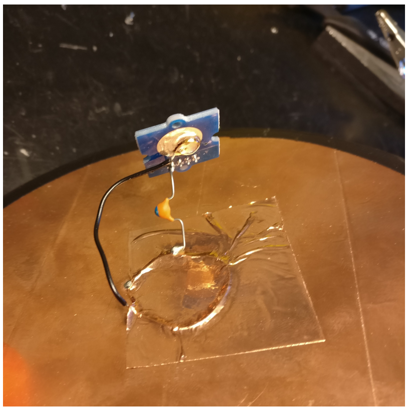

Figure 6: A Grove sensor with added capacitor connecting the capacitive surface to GND.

4.3.2 Affordances

As discovered in the usability test, there was a twofold issue with the perceived affordances of the installation. Many users believed tapping to be the appropriate method of interaction, however, it remains unclear why they initially thought this to be the case as they did not offer any clarification during the interviews, and the facilitators did not think to ask about it directly until after all usability tests had been concluded. However, it is likely a result of the users perceiving the top of the pipes to bear a resemblance to percussive instruments, even through they also recognised it as a pan flute. Several methods for overcoming this incongruence between perceived affordances and intended interaction were discussed, such as putting up a sign explaining the use of the installation (a method used by the authors in previous projects [18]), or by attaching icons communicating the use, e.g. a variation of the common icon for touch (see Figure 7). However, no such change was implemented, partly because a touch icon could potentially be misinterpreted to mean tap, thus not solving the issue, and partly because both adding icons and signage was not compliant with the authors' aesthetic idea of the design, and would work against the 3rd principle of universal design by adding additional complexity, as it would result in more information a user has to digest before interacting. Additionally, it can be argued whether this is actually an issue, a line of reasoning supported in [9]. If a majority of users interact by tapping rather than touching, intended or not, this suggests that the installation simply has more than one mode of use - a mode of use that, if prevalent, will likely only become more prevalent as it is observed by other passers-by, which may be enhanced by the honeypot effect.

Figure 7: A variation of a common icon for touch interaction on smart surfaces. [1]

5. DEPLOYMENT STUDY

The following section covers the evaluation of the installation, which was an in the wild deployment study. The purpose of the test was to see how the installation would affect a given area in regards to spacial influence, user engagement, and shared encounters. The design of this study follows the framework outlined in [17].

The deployment study took place over three days in mid May. The installation was introduced to two distinct sites, get a broader understanding of the effect on the city, as well as to allow for comparison between the two sites.

Environment

The first and second days of testing were conducted at Østre Anlæg, which is a park relatively close to the city center. The park is largely used by people out for walks, runs, or walks with dogs. Additionally, as the surrounding area is highly populated, the park is used for barbecuing and other social activities, weather permitting. On these two days of testing the weather was mild and relatively warm for the season, and the tests were conducted from 11:00 to 16:00 on the first day, and 11:00 to 17:00 on the second. The installation was placed alongside a trafficked pathway. Figure 8 - *Upper*, shows the primary points of prospect around the installation at its deployment in the park, ie. from where the audio could clearly be heard. As can be seen, the path north of the installation did not fall within this area, and the flow in that area was therefore not calculated as part the spatial influence (nor was the area marked in darker blue, as this area was blocked by bushes). The installation could be moved slightly further north, but doing so would have meant that less of the southern path would be covered, and as that path was estimated to be more trafficked than the path further north, this area was chosen as preferable.

On the third and final day the installation was deployed to the harbourfront (see Figure 8 - *Lower*), which traffic-wise was similar in type to the previous deployment, but with a larger amount of people. The weather on this day was also mild and warm, and the deployment was from 12:00 till 20:00. The installation was placed outside a university building, but was as much on the path as it would be in any other place along the harbourfront.

Both of these sites were categorised as leisure spaces, ie. a site where users have time to relax and can allow themselves to be distracted, which was optimal considering the nature of the installation.

Procedure
All three days of testing followed the same overall procedure, albeit within different time frames. The installation was set up, and the facilitators positioned themselves in a spot from where they could easily observe the installation, while remaining inconspicuous. A camera was used to film the entirety of the tests, and it was positioned such that it could not be seen from the installation. The footage was used to identify demographics, categorise use instances, extract use duration and to count both the total number of users as well as total number of passers-by. Figure 9 show the cameras field of view for the two different locations.

5.1 Results
This section covers the results from the three-day deployment tests.

Use Categories
The two days of testing in the park saw a total of 378 use cases, while the final day of testing on the harbourfront saw 594 use cases. Figure 10 shows the proportions of use instances separated into the four categories outlined in section 2.2. It is notable that in both sites there were more instances of AP than AE, and more instances of AE than PE. Moreover, a Fisher Exact test showed a significant difference in the proportion of users who engaged in PE between the two sites ($p = 0.007$). There are no significant differences in the proportions of AP, AE and PO between the two sites.

Use Duration
The use duration was calculated for every use instance, and the means within each use category extracted. As can be seen in Figure 11, users averaged a much higher use duration when engaging in AP, than they did in AE and PE. Moreover, Mann-Whitney U tests showed a significant difference in the use duration for AE and PE respectively, between the two sites ($p < 0.001$ and $p < 0.001$), meaning users who engaged in either AE or PE on the harbourfront generally did so for a shorter period of time than in the park.

User Behaviour
As was discovered in section 4.2, users had a tendency to tap

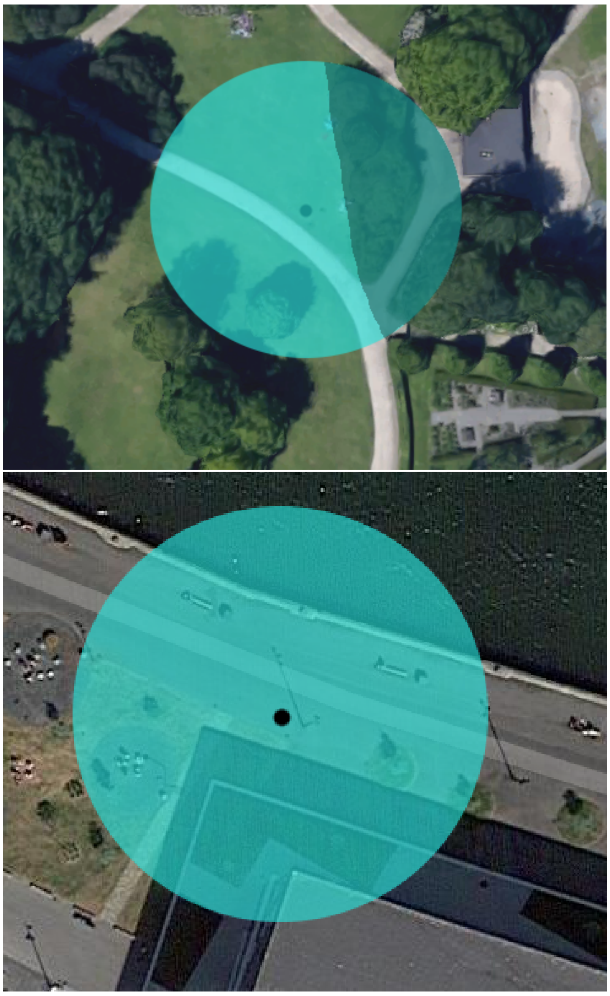

Figure 8: *Upper:* The location during the first two days of testing. *Lower:* The location during the third day of testing. In both images the black dot indicates the location of the installation, while the blue area indicates the approximately 20 meter radius within which the installation could easily be heard.

the pipes rather than touch. This was quickly discovered to also be the case during the testing. This seemed to be the case for most users, regardless of demographic.

Group Size
Users who engaged in AP in groups of 2+ had an average use duration of 36.44 s, whereas users who engaged in AP alone averaged 21.49 s. A Mann-Whitney U test showed this difference to be significant ($p = 0.001$) suggesting that groups overall engage in AP for longer periods of time than individuals.

Observations of group interaction revealed a difference between the behaviour of groups who engaged in AP or AE. AP group use typically began with one individual approaching the installation and interacting with it, while the rest of the group would observe for a while before then approaching to interact with it themselves. AE use generally began

Figure 9: *Upper:* The field of view of the camera during the first two days of testing. *Lower:* The field of view of the camera during the third day of testing.

in the same manner with a single user exploring the system, however, the rest of the group would only rarely approach to explore it themselves, instead hanging back and observing.

Spontaneous Group Formation
The first two days of testing saw two cases of spontaneous group formation (SGF). The first one being lasted roughly 2 minutes and 31 seconds, where two adults, a woman and a man engaging in AP use were joined by a second woman, after which they left the installation together. The second case saw an adult man engaging in AP use, when an adult woman walked by and engaged a conversation with the man, while also engaging in AE. This SGF lasted for 12.5 seconds.

The last day of testing at the Habourfront, saw three cases of SGF. The first being between two groups of elderly people, where a group of two men and two women were engaging in AP use, and a second group of elderly people, one man and one woman, engaged in AE use and exchanged words with each other. This SGF lasted for five seconds. The second case of SGF lasted for 17 seconds, where a woman was joined by a second one, both engaging in AP use, after which they split up again. The last SGF lasted for 14 seconds, where two men were interacting with the installation and were then joined by two women, after which the women left and the men continued to interact with the installation for some time.

Flow Interference
Flow interference was calculated as the percentage of all passers-by who engaged in AP, AE or PE. The first day showed 17.57% flow interference, with 78 out of 444 passers-

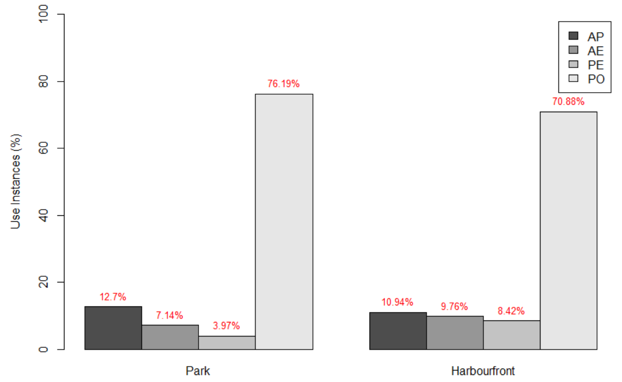

Figure 10: The number of use instances in each of the categories: Active Participation (AP), Active Exploration (AE), Passive Exploration (PE), Passive Observation (PO).

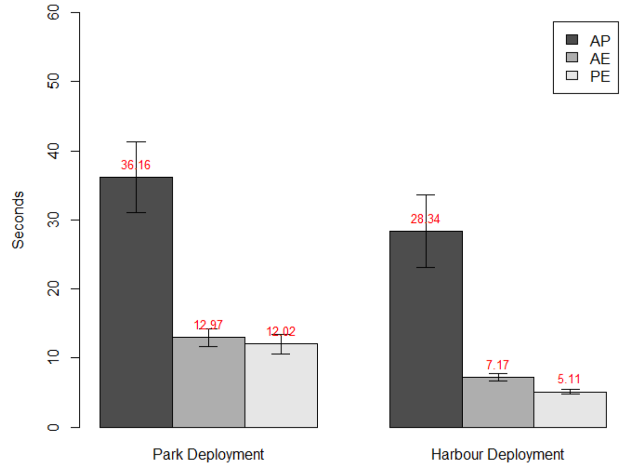

Figure 11: The mean use duration in seconds for Active Participation (AP), Active Exploration (AE) and Passive Exploration (PE) for both sites.

by interacting, while the second day showed 15.06%, with 80 out of 531 passers-by interacting, for a total flow interference of 16.21% in the park. The third test on the harbourfront, saw a total flow interference of 12.51%, with a total of 2429 passers-by. These numbers average out to a total of 13.56% flow interference over the 19 hours of testing, not including passive observers. The two areas, however, show to have a significant difference in flow interference ($p = 0.003$), calculated with a Fisher's Exact Test, suggesting that the installation has greater spatial influence in the park, than on the harbourfront.

The numbers suggest that bicyclists have a low interaction rate, however, as it is not non-existant, discarding these numbers from the equation would be misleading, even though the cyclists make up more than 12.72% of the flow in the park, and 34.83% on the harbourfront, while making up only 0.31% and 1.15% of the interactions, respectively.

6. DISCUSSION

This section covers a discourse on the results, puts validity and reliability of this study under scrutiny, and outlines potential future work.

User Engagement

Upon investigation of the frequency of the different non-PO use cases, it is notable that AP is the use category with the most instances, particularly in the park. Earlier work with different installations by the same authors shows that this is not unheard of, but nor is it the norm; in [18, 21] the majority of use was participatory in nature, but in [7, 17, 19] the majority of use was exploratory. Moreover, three of these installation were supported by a sign explaining how they worked, which was not the case for the installation presented in this article. This suggest that users have a higher likelihood of engaging in AP with this installation than could be expected, extrapolating from previous work. This likely stems from the installation being modelled after a pan flute, recognisably so (as shown in section 4.2), and that its use and purpose is thus intuitive to the user, even if the perceived affordances do not conform with the intended input method.

The measurements of use duration showed that users who engage in AP do so for longer than AE- or PE-users. This conforms to the results of [17, 18, 19], and thus suggest that the use duration heavily depends on the category of use, and whether or not users are part of a group.

Spatial Influence

The results showed a relatively high flow interference compared to other installations. Papers referenced in related works and tested in similar settings, such as [17, 18], saw lower flow interference rates of 10.65%, and 7.61%, respectively, compared to the 13.56% of this installation. This is again despite the fact that these were supported by an additional lure in the form of a sign, whereas this installation was not. This suggests that the installation in itself is more effective at luring in passers-by than those from earlier works, and thus that its progressive lures are overall more effective. Additionally, in [11], it is suggested that 5% flow interference is an acceptable level, and this level is significantly exceeded with the 13.56% observed in this test, confirmed with a binomial test ($p < 0.001$).

As shown in section 5, the flow interference observed in this paper was significantly higher in the park than it was on the harbourfront. This may be caused by several different factors, but the most likely one is that the harbourfront is more akin to a transit space than the park would be considered. Additionally Hummingwall [21], which was deployed to an area no more than 100 meters away, found that users were more likely to interact on weekdays, rather than weekends, and since the final deployment was the day prior to a national holiday, it could be considered to effectively be a Friday. Going by these results, deploying the installation to the harbourfront the day prior may have resulted in a higher flow interference, and conversely for the deployments in the park.

As shown in Figure 11, there was a significant difference in the use duration of AE and PE use between the two sites. This is likely once again a result of the harbourfront being more alike a transit space than a leisure space - a fact that is supported by bicycle traffic accounting for a higher percentage of the overall flow - which suggests that passers-by will be less likely to stop and explore the installation for a prolonged period of time.

Shared Encounters

It was observed that one of the common ways for groups to interact with the installation was where one or two people interact directly, while the rest of the group would observe for a while before engaging themselves. This likely helps explain why groups averaged a higher use duration than individuals. When an individual had had enough of the interaction they would leave, whereas in a group, a single user might have had enough, but their co-users might not.

As mentioned in section 5.1, 5 of the 263 (1.9%) use instances were observed resulting in SGF. This suggests that the installation is not effective at triggering this kind of social encounter, when compared to 9 out of 161 (5.59%) of use instances resulting in SGF in [19], an installation designed specifically for that purpose.

Spaces Within Spaces

As mentioned in section 5, the installation was deployed to leisure spaces. However, a pattern was noticed during the first two days of testing in the park. Even though the park was itself a leisure space, the immediate area around the installation in effect also functioned as a transit space; a notable portion of passers-by were jogging or biking past the installation. Moreover, aside from two exceptions there were no park-goers who spent any prolonged time relaxing in the area, unrelated to the installation - although it remains possible that this is in fact a result of the installation being present and thus causing park-goers to prefer other locations. This suggests that spaces can have areas within them in which citizens behave more like they do in other types of spaces, resulting in spaces within spaces. This potentially makes the choice of location not as straightforward as simply picking a location that overall conforms to your expectations of space, but that different areas within the space should be considered in detail. Although it would likely be a time-consuming process, this could be done e.g. through measuring a baseline of the flow at different areas inside the space, and identifying key activities of citizens.
In the recordings, no notable amount of honeypot effect was observed, which may add to the hypothesis that the areas chosen were leaning more towards transit spaces, rather than leisure spaces. This, or the flow was simply not consistently high enough for a group of people to gather around. It may also be a combination of these factors, or simply, that the installation does not have the magnetism required.

No Interviews

Interviews were neglected primarily to avoid giving the users the impression that they are observed, as this may reduce interaction, as shown in [34]. Should interviews have been done, they could have provided some information regarding why people chose to transition from one use category to the next, but would have reduced the validity of the data regarding the flow. This is a trade-off, and the choice was made to not include the interviews, in order to get a more even spread of data, as no specifc data that interviews could gather was sought after for this specific study.

Entry Points & Lure
There is no statistical significance between the flow interference observed in [18] and this paper, however, it is worth noting that that study had a sign with instructions as an additional lure in order to draw in participants, which this study did not. This may suggest that the inherent lures of the pan flute are effective to a high enough degree on their own. Additionally, earlier in the paper, the points of prospect were described as a sphere around the installation, which, although accurate, does not describe the entirety of the situation. Specifically, the progressive lure, which extends far further, depending on the curve of the road, and the angle at which the installation is placed. In the park, the angle was such that a user would be able to see the installation clearly when they entered line of sight, however this is not as much the case on the harbourfront, as the angle was closer to parallel to the road, meaning it is more difficult to recognize as a pan flute. This may also have had an effect on the abovementioned interference levels.

Thermal Imaging and Flow
For this paper, thermal imaging was not available for practical reasons. It could, however make measuring flow significantly easier, as was shown in [21], and especially if time permits for baseline readings. A full day of reading baseline with the thermal camera, followed by a full day of testing with both thermal and regular camera. In this case, a very visual demonstration of the change in the flow of people would be available. Although even in cases where thermal imaging is unavailable, measuring a baseline is advantageous, however in the case of this paper, the days with appropriate weather were few in the period, and could therefore not be spared.

6.1 Future Work

This sections details potential future work, and why this work is relevant in the context of this study.

Mode of Interaction
Given the amount of users that elected to drum on the pipes as the main method of interaction, this can be considered the current primary mode of interaction, especially as the prototype is easily usable this way, because of its inherent flexibility of use (see section 3.1). Rather than trying to discourage this through change in affordances or signalling the information otherwise, it might simply be worth leaning into. This could be done in multiple ways, such as giving a fading or echoing pipe sound that extends beyond the initial strike, when struck, or implementing a drum sound when struck quickly, but maintaining the pipe sound when holding for a touching period of time. There are many options, and there is no particular reason the pipes should have a congruent sound upon interaction, given the options available with a digital implementation. However, should any of this be implemented, the issue of false negatives should be addressed. Specifically, when the installation is used in the above-mentioned manner, multiple consecutive strikes to the same pipe will lead to falls negatives, as the capacitive sensor becomes satiated. This interferes fairly notably with the use of the installation in this manner, and makes anything that requires quick movements very difficult. A part of this issue also stems from the implemented delay, which also means quick touches becomes difficult to manage.

Confirmation of Lures
Currently, the installation has been tested exclusively in daylight. This means that the additional lure and effect, the LED lighting, has likely had no effect as it is very difficult to see in sunlight. Testing during nighttime could garner information regarding the effectiveness of this part of the lure. In addition to this, it could reveal how much the sound affects users, in the sense that it will be harder to see the pan flute in the dark, except for the lights, which will drastically change the nature of the progressive lures.

Permanent Installation
Should this concept be developed into a permanent installation for a site, this could have a number of advantages. A finished installation would have the advantage of looking finished, which the current prototype does not in all cases, with things such untreated wood being obviously visible, as well as screws being visible, and a few paint spots here and there. These are issues that may affect users to be less likely to interact, however to truly find out whether this is the case, more information is needed.

A further developed installation could also have implemented proximity sensing, such as was implemented in [6]. This type of lure could, especially if the points of prospect are identified, allow for a much more consistent idle animation, and possibly for a non-zero amount of non-instrumental interaction.

7. CONCLUSION

This paper set out with the intention of digitally augmenting the urban space to be more playful. For this purpose, a playful, interactive, urban installations with a with a focus on sonic elements was developed and evaluated through a series of deployments studies.

Results indicated a high proportion of users actively engaging with the installation, compared to related works, suggesting an intuitiveness in its use, and the installation was found to cause a high degree of flow interference - interfering with as much as 17% of passers-by. Moreover, results showed that the installation was suitable for group use, as it was not uncommon to have multiple users of the same group engaging and observing simultaneously. .

Despite shortcomings in the design with regards to poor affordances and visual lures that have not been effectively tested, analysis of the results suggest that this installation is a successful example of a playful, interactive, urban installation, capable of notably affecting the environment that surrounds it. This could be an indication that recognizable design and adaptive modes of interaction are effective ways of creating an impactful urban installation through an iterative design process. Iterating further could likely lead to even better rates of interaction, as the modes of interaction become increasingly congruent with the affordances.

References

[1] Touch icon. Made by https://www.freepik.com/ - Last Accessed: 21/05/2019.

[2] Akpan, I., Marshall, P., Bird, J., and Harrison, D. Exploring the effects of space and place on engagement with an interactive installation. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems* (2013), ACM, pp. 2213–2222.

[3] Balestrini, M., et al. Jokebox: Coordinating Shared Encounters in Public Spaces. *Proceedings of the 19th ACM Conference on Computer-Supported Cooperative Work & Social Computing - CSCW '16* (2016), 38–49.

[4] Bangor, A., Kortum, P., and Miller, J. Determining What Individual SUS Scores Mean: Adding an Adjective Rating Scale. *J. Usability Studies 4*, 3 (5 2009), 114–123.

[5] Behrens, M., Valkanova, N., gen. Schieck, A. F., and Brumby, D. P. Smart Citizen Sentiment Dashboard: A Case Study Into Media Architectural Interfaces. *Proceedings of The International Symposium on Pervasive Displays - PerDis '14*, September 2013 (2014), 19–24.

[6] Birchfield, D., et al. Interactive public sound art: a case study. In *Proceedings of the 2006 conference on New interfaces for musical expression* (2006), IRCAM—Centre Pompidou, pp. 43–48.

[7] Claes, S., Wouters, N., Slegers, K., and Vande Moere, A. Controlling In-the-Wild Evaluation Studies of Public Displays. *Proceedings of the 33rd Annual ACM Conference on Human Factors in Computing Systems - CHI '15* (2015), 81–84.

[8] Cox, T., et al. Uncovering the honeypot effect: How audiences engage with public interactive systems. In *Proceedings of the 2016 ACM Conference on Designing Interactive Systems* (New York, NY, USA, 2016), DIS '16, ACM, pp. 5–16.

[9] Dalsgaard, P., and Halskov, K. Designing urban media façades: Cases and challenges. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems* (2010), ACM, pp. 2277–2286.

[10] Dameri, R. P. *Smart City Implementation.* Springer International Publishing, 2017.

[11] Fischer, P. T., and Hornecker, E. *Creating Shared Encounters Through Fixed and Movable Interfaces.* Springer Singapore, 2017, pp. 163–185.

[12] Franinovic, K., and Visell, Y. Recycled soundscapes. In *Symposium on Designing Interactive Systems: Proceedings of the 5 th conference on Designing interactive systems: processes, practices, methods, and techniques* (2004), vol. 1, pp. 317–317.

[13] Franinovic, K., and Visell, Y. New musical interfaces in context: sonic interaction design in the urban setting. In *Proceedings of the 7th international conference on New interfaces for musical expression* (2007), ACM, pp. 191–196.

[14] Gaye, L., Mazé, R., and Holmquist, L. E. Sonic city: the urban environment as a musical interface. In *Proceedings of the 2003 conference on New interfaces for musical expression* (2003), National University of Singapore, pp. 109–115.

[15] Grønbæk, K., et al. Designing playful interactive installations for urban environments – the swingscape experience. In *Advances in Computer Entertainment* (Berlin, Heidelberg, 2012), Springer Berlin Heidelberg, pp. 230–245.

[16] Hornecker, E., Marshall, P., and Rogers, Y. From entry to access: How shareability comes about. In *Proceedings of the 2007 Conference on Designing Pleasurable Products and Interfaces* (New York, NY, USA, 2007), DPPI '07, ACM, pp. 328–342.

[17] Jacobsen, B., Svendsen, M. U. K., Søgaard, A. K., and Uggerhøj, R. L. Designing an evaluation framework for public interactive installations. `https://projekter.aau.dk/projekter/files/294398532/MTA18931_P9_Report.pdf`, 1-12.

[18] Jacobsen, B., Svendsen, M. U. K., Søgaard, A. K., and Uggerhøj, R. L. Interactive furniture for urban intervention: Sonically augmenting the urban environment. `http://projekter.aau.dk/projekter/files/259994395/Paper___MTA17639.pdf`, 2017.

[19] Jacobsen, B., Uggerhøj, R. L., Svendsen, M. U. K., and Søgaard, A. K. Urban Intervention - Spontaneous Group Formation, 2018.

[20] Lidwell, W., Holden, K., and Butler, J. *Universal principles of design: 125 ways to enhance usability, influence perception, increase appeal, make better design decisions, and teach through design*, second edition, revised and updated ed. Rockport Publishers, 2010.

[21] Morrison, A., et al. The humming wall: Vibrotactile and vibroacoustic interactions in an urban environment. In *Proceedings of the 2016 ACM Conference on Designing Interactive Systems.* ACM, 2016, pp. 818–822.

[22] Müller, J., et al. Looking Glass: A Field Study on Noticing Interactivity of a Shop Window. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems* (New York, NY, USA, 2012), CHI '12, ACM, pp. 297–306.

[23] Nijholt, A. *Playable Cities*, 1st ed. Springer Singapore, 2017.

[24] Norman, and Norman, D. A. *Design of Everyday Things.* 2002.

[25] Peeters, M., Megens, C., van den Hoven, E., Hummels, C., and Brombacher, A. Social stairs: taking the piano staircase towards long-term behavioral change. In *International Conference on Persuasive Technology* (2013), Springer, pp. 174–179.

[26] Poulsen, E. S., Morrison, A., Andersen, H. J., and Jensen, O. B. Responsive lighting: The city becomes alive. In *Proceedings of the 15th International Conference on Human-computer Interaction with Mobile Devices and Services* (New York, NY, USA, 2013), MobileHCI '13, ACM, pp. 217–226.

[27] Preiser, W. F. E., and Smith, K. H. *Universal Design Handbook, 2nd Edition*, 2 ed. McGraw-Hill Professional, 2011.

[28] Rasamimanana, N., et al. The urban musical game: using sport balls as musical interfaces. In *CHI EA'12 Proceedings of the 2012 ACM annual conference on Human Factors in Computing Systems* (2012).

[29] Siek, K. A., Hayes, G. R., Newman, M. W., and Tang, J. C. Field Deployments: Knowing from Using in Context. In *Ways of Knowing in HCI*, J. S. Olson and W. A. Kellogg, Eds. Springer New York, New York, NY, 2014, pp. 119–142.

[30] Sundhed.dk. Growthcurve, boys 0-20. https://www.sundhed.dk/borger/patienthaandbogen/boern/illustrationer/tegning/vaekstkurve-drenge-0-20/. Last Accessed: 21-05-2019.

[31] Sundhed.dk. Growthcurve, girls 0-20. https://www.sundhed.dk/borger/patienthaandbogen/boern/illustrationer/tegning/vaekstkurver-piger-0-20/. Last Accessed: 21-05-2019.

[32] Tan, L., and Chow, K. K. Piano staircase: Exploring movement-based meaning making in interacting with ambient media. In *IFIP Conference on Human-Computer Interaction* (2017), Springer, pp. 282–291.

[33] Tanaka, A., and Bongers, B. Global string: A musical instrument for hybrid space. In *ICMC* (2002).

[34] Waern, A. The Ethics of Unaware Participation in Public Interventions. *Proceedings of the 2016 CHI Conference on Human Factors in Computing Systems - CHI '16* (2016), 803–814.