Performative, extravagant, expressive, place-based experiences.

Betsworth, Liam George

Performative, Extravagant, Expressive, Place-based Experiences

Liam George Betsworth

Submitted to Swansea University in partial fulfilment of the requirements for the Degree of Doctor of Philosophy

Prifysgol Abertawe
Swansea University

Future Interaction Technology Lab
Department of Computer Science
Swansea University

2015

Declaration

This work has not been previously accepted in substance for any degree and is not being concurrently submitted in candidature for any degree.

Signed .... (candidate)

Date ....12/06/2015.................

Statement 1

This thesis is the result of my own investigations, except where otherwise stated. Other sources are acknowledged by footnotes giving explicit references. A bibliography is appended.

Signed .... (candidate)

Date ....12/06/2015.................

Statement 2

I hereby give my consent for my thesis, if accepted, to be available for photocopying and for inter-library loan, and for the title and summary to be made available to outside organisations.

Signed .... (candidate)

Date ....12/06/2015.................

Abstract

Many of the mainstream mobile location-based services that we have become used to in public spaces are good at delivering information privately, in context, but the opportunity to incorporate more engaging and exciting interactions is often overlooked – especially where the output is of shared public interested. Smartphones now offer us multimodal interaction, gestures and internet connectivity, all providing opportunities to interact in new, extravagant and expressive ways.

As members of *Cu@Swansea*—a multi-partner project leading the regeneration of the world significant Hafod-Morfa Copperworks—we have been tasked with designing a range of technologies that will attract people into the site, not to experience a finished, curated piece of heritage, but to bring the site to life, provoking discussion amongst the local community, stakeholders and other visitors. Instead of allowing people to pass each other, *"digitally divided"* [76], we focus on designing interactions that will start conversations, encouraging people to join together in a collaborative, public experience.

This thesis details the design, development and evaluation of a set of novel, extravagant, expressive mobile location-based experiences. We experiment with both audio and visual effects as a baseline. We then attempt to extend the framework, developing a remote mechanism that can be used to scale-up and direct audio-visual experiences. We consider the design and evaluation of our systems from a performative standpoint, attempting to optimise engagement between the perfomer—user of the system—and spectators – bystanders engaged in this performance. This was achieved through a range of amplified manipulations and effects [108]. During the design process, we organised community engagement events, meeting with stakeholders and holding a focus group with interested members of the community. These engagements concentrated more on deployment concerns, such as attitudes towards these technologies, and how they may be successfully integrated within the site.

Our main contribution in this work is a novel, performative mobile framework for more extravagant, expressive interactions in public spaces. Although our ultimate aim is to design and deploy these experiences for use within a heritage context, our findings suggest that these technologies could be utilised to promote a more social, active engagement in a range of public spaces.

Contents

Acknowledgements

Throughout this endeavour, I have met many individuals with whom I have exchanged differing concepts, opinions and perspectives. Connecting with and being directly affected by the local Hafod community, project stakeholders and other researchers has been an exciting and fulfilling process. In this respect, there are too many people to thank individually, though I would like to begin by thanking anyone who has come into contact with me and helped shape my research throughout the process. I would also like to thank Nitendra, Saurabh and Amit, my colleagues during an internship at IBM India Research Lab. Encounters, both big and small, have all had some effect on my work, and they make it what it is today. Thank you.

From a personal perspective, at times, the project has seemed overwhelming, with ecstatic highs and bitter lows in equal measure. Perhaps, the biggest thanks are in order for Matt Jones – my friend and supervisor. The support, advice and opportunities you continue to provide me with have been both invaluable and plenitudinous, consistently teasing the best qualities out of me. I could not have asked for a better supervisor. Additionally, Huw Bowen, you have been central to the existence of this entire project. Thank you for providing me with this opportunity and supporting me along the way. During my studies, I have been very lucky to have a group of good friends and colleagues that have taken the time to guide, support and socialise with me. Simon, Emma, Jen, Tom, Patrick, Ben, Elvis, Tim and Cameron – many of my best memories at university are with you lot, and I am sure we will all be friends for years to come.

Finally, I would like to thank my family. Without their backing and continued support, I would not be here. Sophia, Des, Jeff, Nick and Hazel – thank you for putting up with and supporting me. Thank you for letting me live my dream.

Chapter One

Introduction

This thesis is about a possible future; a future driven by a desire to surface the past. As members of a large, heritage-led, regeneration project, we were tasked with designing a number of mobile location-based experiences for an abandoned, unused heritage site. Our aim was to encourage the local community and visitors to make use of the area, starting conversations, sharing experiences and ultimately, building an appreciation for the space. Many existing mobile location-based experiences in such public spaces are designed to deliver media to visitors' devices, and also offer guidance within the local vicinity. However, these interactions are often very private and personal, not designed with collaboration or shared, social experiences in mind. Therefore, we argue that these kinds of mobile location-based experiences are perhaps less suitable for promoting engagement and conversation amongst people in public spaces. In locations such as the heritage site we are interested in regenerating, we believe visitors could benefit from more engaging, public location-based experiences. Here then, we suggest, lies an underexplored design space for mobile location-based services.

Much of the previous research in the area of mobile location-based services has been focused within a private, personal design space. In this thesis, we draw on novel technological and interaction design opportunities to examine the approaches and value of more performative, extravagant, expressive place-based services.

Extravagant: By *extravagance*, we aim to build bold, unconventional, conspicuous interventions that attract attention.

Expressive: This will be achieved through a range of *expressive* interactions, with amplified gestural manipulations and effects, and movement of users through space.

Performative: All the while, we consider the design of these experiences from a *performative* perspective, considering public spaces as a stage where performers, spectacles and spectators all play an important role.

Throughout the thesis, we report on the design, development and evaluation of a number of experimental, performative mobile location-based experiences for use in public spaces. Our goal is to delve into this underexplored design space, researching new, exciting ways for people to interact with others and their surrounding environment when using mobile location-based services in public spaces. There is a great irony in the increasing connectivity of our mobile devices actually disconnecting us from meaningful social interactions [143]. As smartphones continue to gain additional sensors and adoption increases, the kinds of technologies mentioned here could potentially help to reinstate these meaningful interactions.

1.1 Mobile location-based services

Over the last decade, location-based services have become increasingly prominent on mobile devices. This surge in use can be attributed to the introduction of the smartphone, along with cellular data, now seamlessly integrating location-based services into our everyday mobile experiences. Before widespread Internet availability, many location-based services relied on local databases containing geotagged information such as maps and points of interest [133]. With the growing number of WiFi hotspots and affordable cellular data, such applications are now less reliant on offline databases. Today, an equivalent mapping application such as Google Maps[1] downloads its data as and when needed via asynchronous web requests.

With widespread access to the Internet, mobile users are now able to access a whole range of different services, information and media types, meaning location-based services are no longer restricted to being offline utilities, merely providing

[1] Google Maps – `http://maps.google.com`

maps, points of interest and guidance. An example of a more modern and social mobile location-based service is Twitter's 'places' search functionality[2]. This allows a user to discover and read tweets near to their location. Using this same API, McGookin et al. [93] went a step further and converted nearby text tweets to speech, allowing users to wander round outdoors and listen to spoken tweets as they encountered their location of origin. Researchers have also experimented with sharing images in a similar way [75], linking discoverable pictures to points of interest in the real world. All of these examples, along with many other mobile location-based services, are concerned with delivering text, image and audio information to mobile devices, though in a private manner, solely to the user of the device. However, in the literature, there is less discussion and experimentation around public, collocated mobile location-based experiences.

1.2 Performative, extravagant, expressive experiences

Existing research that tries to facilitate enjoyment and engagement between visitors in public and semi-public places such as visitor centres has looked at themes such as contribution/sharing, gamification and augmented reality. Of the systems that allow for contribution and sharing of digital media, the main focus is often on how the media is shared or discovered – less so on how the media is finally experienced by the discoverer. The media is often experienced in a private, isolated manner, such as on a mobile device screen [75] or through headphones [123].

In museums and visitors centres, gamification is a popular method of encouraging visitors to move between points of interest, providing hints and clues for visitors to follow [55], or providing pieces to solve a larger riddle [32]. Our initial concern with this area of interaction was that it is aimed mainly at a younger demographic, and would perhaps not cater to the wider audience that we are trying to capture in our own experiences. Although ambitious, we are looking to develop experiences that span all age ranges, allowing young, old and middle aged people to experience our interactions together.

[2]Twitter Advanced Search – `https://twitter.com/search-advanced`

Augmented reality is certainly a growing area in interpretation, with smartphones and other mobile devices affording new possibilities. When attempting to bring a physical, public space to life–especially in the heritage context [95]–augmented reality has proven to be a useful and popular interaction technique. Very often, augmented or mixed reality experiences do indeed offer an immersive and engaging experience for the visitor [82]. However, although more immersive and engaging for a single user, there are no specific or obvious ways in which a designer is able to create engaging collaborative experiences for multiple people, both with friends and strangers. After initially considering a number of ways in which we could potentially design our interactions, we pondered the question: *How can we design our own interactive experiences in such a way that we are purposefully encouraging a more active, collaborative engagement between people?*

Before mobile internet-connected devices and location-based services, we had to look around and explore; now, as argued by Jones [76], perhaps the dangers of such systems are that we end up walking into *"public places we no longer connect with—where people pass each other, digitally divided."* With too many private, hidden and secretive mobile interactions, it is easy to see how we may end up cutting ourselves off from the world around us. Perhaps then, as he puts it, we should strive for more *"extravagant, expressive, place-based computing,"* where we attempt to employ interactions that are more mindful of our environment, paying closer attention to both our surroundings and those who populate it.

As our primary interest is in developing technologies that promote a more social, active engagement, we need to design our technologies in such a away that would draw people together, provoking curiosity and encouraging conversation. When attempting to develop a system to facilitate a more extravagant, expressive, experience, we believe that it makes sense to think of the design and evaluation from a performative perspective. According to Goffman, any interaction that we have with others in our everyday lives, even a one-to-one conversation, can be seen as a performance [57]. The fact that we use the majority of our mobile devices in front of others, in focused relation to our surroundings [42], makes the performance metaphor highly relevant to mobile location-based services.

In a performance—be it technologically mediated or not—a large factor in the impression that we give others can be derived from the observability of our actions and their subsequent reactions. Reeves et al. developed a taxonomy to classify a performance using the different levels of observability as a spectator [108]. In this taxonomy, a performance or interaction can be categorised into one of four classes. These are: secretive, suspenseful, magical and expressive. Each of these kinds of performance are classified by the observability of both its manipulations and effects, which range from hidden to amplified. In our research we are particularly interested in expressive experiences, where both manipulations and effects are amplified. We believe that such experiences are the most accessible to spectators, with the public, high visibility of manipulations and effects making both the performer and experience open and more approachable. We are also interested in magical experiences, where effects are again amplified, but manipulations are hidden. With these kinds of experiences, spectators can fully appreciate the feedback of a performance, though with a sense of curiosity as to how the performance is being controlled. With such interactions, performing perception also becomes less of an issue, as performers do not need to worry about their manipulations being judged [42].

So far, we have briefly touched upon the concept of more performative, extravagant, expressive place-based services. In the next section, we discuss our underlying motivations for focusing on this specific, underexplored research area.

1.3 Context & motivation

When Britain's industrial age went into decline in the early 20th century, many buildings and factories were demolished or, in a few cases, renovated to house new activities. Other industrial sites, however, were simply neglected and ignored. One notable site—which is the subject of this work—was completely abandoned and left exposed to vandals and the elements for over 30 years. This is the Hafod-Morfa Copperworks (see Fig. 1.1) in the Lower Swansea Valley, located on a $12\frac{1}{2}$-acre site just to the north of the city of Swansea, Wales. By 1890 the Hafod works was the largest copperworks in Europe, and it lay at the heart of a global network of supply. The site was, and still remains world significant.

When the doors of the works shut for the final time in 1980, the site was simply sealed off in a (failed) attempt to prevent any further access. While it is slightly

Figure 1.1: A crumbling building at the Hafod Copperworks (2012).

unsettling to see that there has been no attempt to maintain or develop the site for such a long period of time, it is necessary to look beyond the overgrown foliage and crumbling buildings. What has actually occurred at this site is the natural preservation of a set of historically significant buildings and structures, many of which are registered on the Statutory List of Buildings of Special Architectural or Historic Interest.

Due to the historical significance of this site, a number of organisations–led by Swansea University–have come together in a multi-partner project known as *Cu @ Swansea*[3], to raise awareness and explore the possibilities of regenerating the site. The primary focus of the project is to make the site a pleasant, usable area for the public, although in a heritage-led way that attempts to preserve the

[3] *CU @ Swansea* - `http://www.welshcopper.org.uk/en/Cu@Swansea.htm`

history that is there to be seen and interpreted. As members of this project, we have begun to explore a range of digital interactions that could be used to aid the regeneration of the site.

Visitor attractions such as heritage sites have come a long way with technology adoption in recent years. Many (e.g., Kew Gardens [82]) now offer a mobile application that acts as an augmented reality, interactive tour guide. This application, though, along with many others of the same kind, afford personal, hidden interactions.

At public, outdoor attractions, there is also a long standing tradition of visitors behaving as consumers. People visit a site expecting to play a passive role, learning the information that is present there. This kind of model has been thoroughly explored in standard location-based guides, with users discovering and receiving information as they move around a space. However, Bagnall [6] argues that consumption alone is not enough, arguing for the importance of playing a more active role in these visits. She suggests that at such sites, performance and performativity are key social practices, and that *"the relationship between visitors and the sites is based as much on emotion and imagination as it is on cognition."* It was claimed that these performative experiences helped stimulate reminiscence through emotionally engaging visitors. Though Bagnall [6] did not make use of technology in her performances, we think these same values would hold for technologically mediated performative experiences, playing a more active role, shaping not only their own experiences, but also others at the site.

Our overall goal at the site is to draw people together, not to view a finished piece of curated heritage, but rather, to start conversations about their memories and the significance of the site to them, and to discover what they would like to see at the site in the future. The technology we are producing is about engaging with the local community and stakeholders as groups to provoke discussion. This contrasts with previous uses of mobile guides which only attempt to be tourist aids. We believe that all of these discussed factors are particularly important to ensure the successful regeneration of the site – something that is desperately needed in this immensely important, but abandoned and under-appreciated area.

Engagements with the Hafod-Morfa Copperworks

During this project—for a large period of time—the Copperworks site was not accessible to the public. For this reason, many of our earlier experiments, studies and engagements happened outside of the works.

This research began in late 2011. Much of the first year was spent exploring different technologies and interaction possibilities. As we were not able to test these prototypes at the Hafod-Morfa Copperworks, we made an attempt to deploy the interactions in existing visitor centres and public spaces. Figure 1.2 presents a timeline, reporting the final progress in this project. The timeline includes project milestones such as experiments and studies, community engagement events and the changing physical status of the Hafod-Morfa Copperworks. Each milestone also contains the name of the location where that event occurred.

Figure 1.2: A timeline of the project milestones and accompanying locations.

The Hafod-Morfa Copperworks had an official opening event in June 2014. This was the first real opportunity to test our designed interactions in a naturalistic deployment, in-situ, with members of the public. A few months prior to this official opening, we were given access to carry out early evaluations onsite.

Although we were not able to test all of our interventions fully at the Hafod-Morfa Copperworks, we saw this as an opportunity, rather than a hindrance. Contrary to negatively impacting our research, we believe that this decision to test the technologies in different scenarios allowed us to gain additional insight as to what technologies could work in certain contexts.

1.4 Overview & contributions

We have identified an underexplored design space for mobile location-based services. Previous research in the area of mobile location-based services has been focused mainly on private, personal and secretive interactions with information. In this thesis, we draw on novel technological and interaction design opportunities, examining a number of different approaches to, and also the value of more extravagant, expressive place-based experiences that consider interactions outside of the device. The key contributing chapters in this thesis document a progressive journey through a number of prototypes, concluding with a chapter that discusses current attitudes towards these kinds of technologies, and how they may be successfully integrated to help facilitate regeneration. Our overall approach and methodology is discussed in *Chapter 3*.

All of the experiments and evaluations reported here contribute to forming a set of design recommendations for the development and deployment of similar mobile experiences in public spaces. Our main, overarching contribution here is a novel, mobile place-based framework to enable performances in public spaces.

Chapter 4 begins by considering audio as an effect. We conducted an initial experiment with spatial audio, exploring the engagement benefits of sounds originating from the surrounding environment. We then went on to focus on a number of more expressive, performative auditory experiences that are possible with interaction outside of the device. Situated loudspeakers were used, along with a range of interaction techniques to produce amplified feedback in public spaces. Amplified audio proved to be a very effective method for engaging bystanders. Control was shown to be an important aspect for performers, affecting comfort and enjoyment. The deployment location also had a major effect, highlighting the value of such systems being dependent on context.

Chapter 5 diverges from the auditory approaches in the previous chapter, exploring opportunities for more performative, visual mobile location-based services. A public deployment was carried out with a pico projection prototype, allowing users to project animated insects, animals and processes in a botanical garden. Similar to our previous experiments, the amplified, open nature of the interactions and feedback promoted a more active engagement among bystanders and users. Although not compared directly, the scale and brightness of pico projections appear to make them less effective than audio for this purpose.

Chapter 6 draws the previous contributing chapters together, detailing a performative framework for remotely directing large scale audio experiences in public spaces, harnessing visitors' mobile device speakers to create dynamic, ad-hoc, human soundscapes. There was again, a continuation of the kinds of bystander behaviour observed in previous chapters, with amplified audio prompting curiosity and conversation, though even more so due to users being forced into close proximity of bystanders. Due to the complex nature of real world settings, improvisation played a large part in the success of remote guidance. Important issues were also raised with the ethics and accountability of remotely guiding people.

When deploying the extravagant, expressive technologies demonstrated in this thesis, we learnt a great deal regarding their usability. However, less was known regarding the public acceptance and attitudes towards these kinds of interactions, and how they will permanently integrate within public spaces. The final key, contributing chapter—*Chapter 7*—explores and discusses these potential issues, and also provides suggestions on how to encourage users to cross the consumer-producer divide.

1.4.1 The author's contribution

The author of this work is very much an "artist-researcher" [15], linking practice, studies and theory – engaged in all parts equally. The vast majority of the work contained within this thesis was conducted solely by the author, some of which was completed as part of an internship at IBM Research India. The design and concept of each system was discussed with the author's supervisor and other colleagues, though the development of each system was carried out entirely by the author. When evaluating each of the systems in user studies, help was sometimes required for logistical reasons, such as setting up equipment in the field, recruiting users and data collection. All analysis of the study results was done by the author.

Much of the research contained within this thesis has been published in co-authored peer-reviews international conferences and journal papers. The details of these contributing publications can be found in Appendix A.

Chapter Two

Background

In this chapter we review previous related research and concepts, situating our work within the field. Our goal in this thesis is to investigate the underexplored design space of performative, extravagant, expressive place-based experiences.

Therefore, we begin this chapter by reviewing the origins of existing location- and context-aware interactions and services, progressing then towards future-looking research concepts that focus more on the *experience*, taking full advantage of modern lightweight, sensor-packed mobile devices. We then narrow our focus by considering the specific area of mobile guides, discussing different methods of information retrieval, navigation and control.

Moving beyond previous, conventional screen-based mobile interactions, we then begin to look at design approaches that allow opportunities for face-on, heads-up interactions. Considering the technologies and concepts introduced in previous sections, we then discuss the topic of socio-spatial interaction, looking for ways to further enhance social engagement through purposefully designing our technological interventions. We conclude with an overview of the related work in this area, summarising the key issues and areas that we build upon in this thesis.

2.1 Context-aware and location-based interactions

A naive definition of context-aware systems, would be: systems that contain some information about the current context, and are able to make decisions, providing the user with more useful or precise data, interactions or services based on that particular context. Abowd et al. [2] formally define context-awareness as *"the use of context to provide task-relevant information and/or services to a user."* Context, here, being,

> *"any information that can be used to characterize the situation of an entity, where an entity can be a person, place, or physical or computational object."*

Contextual information can include automatically sensed, as well as manually input data, allowing the system to *"react to changes"* [128] accordingly – though context is most often automatically sensed, increasing its usefulness.

ParcTab [151]—developed by Want et al.—was a wireless mobile communication device, which was arguably one of the first context-aware mobile devices. The device took into account a number of contextual factors, including the location of the user, the time, the presence of other users and also other nearby peripherals. All of this information combined, allowed the system to provide users with useful information and services in an office environment. Now, with the widespread use of mobile smartphones and cellular data, context-awareness has moved to the outside world [56]. By taking into account information about the device, user and local environment, such applications can make decisions for us in an attempt to make our everyday lives easier.

It would be easy to assume that the more context-awareness we have, the more we can, perhaps, enhance mobile user experiences. However, researchers have warned about over reliance on contextual information. In particular, Cheverst et al. [36] remark that trust plays a large part in these kinds of applications, with users becoming frustrated when systems are unable to predict intentions accurately. It is important, then, to ensure that designers understand such limitations, so that they may develop more useful context-aware interactions.

According to Schilit [128], context-aware applications can be broken down into a simple, three step process: discovery, selection and use. The experiences we are concerned with designing, developing and discussing in this thesis, will only be accessible in particular locations, or are somehow affected by the location of users in the experience. Therefore, we are interested in using proximate selection techniques [127], with users able to interact with things that are physically near to them.

Context-aware systems that are primarily interested in user's locations, are often referred to as location-based or location-aware services. The principal aim of a location-based service is to,

> *"integrate a mobile device's location or position with other information so as to provide added value to a user."* [129]

This is a broad, but useful definition, encompassing location-based services of all purposes, shapes and sizes – from mobile geocaching [101] to car satellite navigation.

2.1.1 Utilising location as context

Before digital location tracking was developed, obtaining a precise location was difficult. Progressively, over time—using a mixture of instruments—one could track and estimate their own position using sextants, chronometers, maps and compasses. Other people could also be tracked through token-based approaches, making a record of where a person was at a particular point in time. Only in recent history have we been able to precisely calculate our locations digitally, using a range of sensors.

One of the earliest attempts at tracking locations, involved using a network of infrared (IR) sensors in a building. The Active Badge Location System [150] required employees to wear a small IR emitting tag. The tag emitted short pulses every 15 seconds, which were then detected by the sensor network. This allowed the system to determine the location of every employee in the building. The main purpose of the system was for administration, allowing employees to be placed on a map, and for calls to be redirect to the location of the employee within in the building. This prototype later evolved into the ParcTab [151].

Beacon-based approaches are widely used for indoor navigation. Other researchers have looked at using radio-frequency (RF) to determine location [7]. Today, this can be achieved through harnessing existing WiFi hotspot infrastructure [52]. Bluetooth beacons have also been used in a similar way [51], using signal strength and triangulation techniques to calculate approximate locations. Further research has looked at an innovative combination of approaches [86], sharing known locations between devices using bluetooth.

Currently, one of the most popular and widely available approaches to location tracking is the Global Positioning System (GPS). Due to its nature, GPS only works outdoors, though this is well suited to most location-based services, as we tend to use them when we are out-and-about. GPS capabilities are built into most of todays smartphone devices – it has become something that people now expect, rather than an optional add-on. Although GPS can fairly accurately track outdoor positions, it is not without problems. It does not work particularly well in urban, built-up areas, nor does it work particularly well near the extremes of the northern- or souther-hemisphere. It can also be affected by cloud cover and other extreme weather. A user may expect the technology to work in these conditions, though go on to experience reduced sensing, leading to confusion or frustration [36]. Benford et al. [17] provide a number pointers to overcoming such expectations, including 'containing expected movements', 'communicating limits' and 'lowering precision.' It is also argued that such shortcomings can actually present opportunities, with designers perhaps wanting to reveal or exploit uncertainties instead of hiding them [12].

As well as GPS and beacon-based approaches, there are a number of other ways in which location can be determined. QR codes have become increasingly used as a means of accessing location-based content [99]. This approach works similarly to a 'you are here' sign, though the location or content is fed directly to the device by scanning the QR code. Although QR codes often require a smartphone application, they do not require a specific, bespoke application – any smartphone scanner will work. In our research, we consider this kind of approach useful for accessing or gaining control of content in specific locations.

2.1.2 Beyond services, towards place-based experiences

In this thesis, we put a particular emphasis on *place* rather than *location*, as we do not want to rely solely on a highly specific knowledge of a person's location. Rather, we focus on the fact that these experiences are only available to be accessed and used when visiting certain places. Whether acting as a user or a bystander, a person must be there to experience it. We also put an emphasis on *experiences* as opposed to *services.* Many of the examples that we have mentioned so far, have been concerned with providing a utility or service based on a user's context or location. Jones [76] raises an interesting point, commenting that, even though these technologies are 'undoubtedly useful', we must be weary of

> *"the dangers these innovations pose to the joys, surprises and even discomforts of exploring our cities, hills or beaches."*

Rather than purely providing information through a service, we are interested in providing users with engaging experiences – encouraging exploration and social interactions.

Computers were once extremely expensive. Much of their design and evaluation was driven by trying to be more efficient, though now, we are surrounded with devices that are not so expensive, and spend much of their time idle [49]. With mobile devices in particular, focus has gradually shifted away from utility and efficiency, towards facilitating valuable and enjoyable experiences. The following quote is a description of location-based experiences from Steve Benford:

> *"Moving beyond core information services, location-based experiences aim to provide the user with a richer experience that extends across a series of locations."* [11]

As later discussed by Benford, these kinds of experiences are well suited to guides, games and educational resources. A popular method for developing engaging location-based experiences such as these, involves designing them from mixed-reality performance perspective [10], introducing narrative and movement, whilst also augmenting the real world with audio-visual information [119].

Augmented reality has been used in many instances to promote learning in both the museum context [43, 64] and outdoors [38, 148] at sites of historical or archaeological interest. While some of these systems have been crafted intentionally to aid information access and retrieval, others have been built to engage younger audiences through a game like experience where the user(s) of the system has to achieve preprogrammed goals. This area of research that focuses its attention on teaching young children and students is one of growing interest, most likely due to the challenges of making history interesting and appealing to the younger audience. An example of a system developed for this very purpose is *Explore!* [5, 40]. The *Explore!* framework comprises of a mobile application that is run on the students' mobile phones and a master application that the game master uses to brief/debrief the students before and after the game. Each game is crafted to a particular site and will only work at that location.

Very often, mixed-reality location-based experiences encourage collaboration between users [13, 106], and are sometimes collocated [26], with digitally and physically situated users working together to achieve a goal. Two particularly interesting examples are *Can you see me now* [12] and *Uncle roy all around you* [18], both making use of location in differing ways. The former involved chasing digital characters in a physical environment. The latter was a slower, unfolding experience, where collocated players attempted to find a mysterious person in a city through communication and teamwork. These kinds of collocated location-based games provide interesting new experiences for users, where differing perspectives are exchanged.

Other interesting location-based experiences have looked at more distinctive opportunities, such as when cycling [33]. *Rider Spoke* gave cyclists the chance to explore a city freely, leaving audio comments at places of particular interest. Other riders using the system were then able to retrieve these messages by essentially riding through the location at which they were recorded in. The authors argue that cycling opens new opportunities for location-based experiences that are not available to pedestrians, with cyclists tending to travel at a much faster pace and cover larger distances than walkers.

Although researchers have begun to explore the possibilities of mobile location-based experiences, we believe that there are potentially many other, new experiences that could benefit from being designed in a more performative, extravagant, expressive manner – attempting to promote a more active social engagement among visitors in public spaces.

2.2 Mobile guides, navigation and control

In this research, we are particularly interested in public spaces such as the Hafod-Morfa Copperworks, where people are able to explore and learn about a space. There is a large and growing interest in the research area of location-based mobile guides, with attempts being made to develop fundamental principles [29] that these kinds of systems should adhere to.

Abowd et al. were one of the first to consider outdoor, location-based, context-aware mobile guides [1]. *Cyberguide* was a mobile guide, helping users understand and navigate unfamiliar public places. Many of the initial mobile guides followed a similar service-based design approach [133], providing maps and other on-screen information to aid people in understanding and navigation.

Over the coming years, some researchers opted for designing more engaging mixed-reality experiences, trialling new technologies of different shapes and sizes [130]. Some guides began offering collaborative experiences, mainly for children or young students [3, 5, 27]. However, when focusing on the majority of recent, mainstream, handheld mobile guides, we can see that many are still, perhaps, lacking in some of the more exciting interactions that are possible [81]. Kenteris et al. conducted a survey of mobile guides in 2011. Although the survey fails to take note of some of the later smartphone type experiences, it gives a good account of the period when mobile guides were becoming mainstream and widely available. The survey details that the majority of mainstream mobile guides feature a map, allowing a user to discover points of interest and then navigate towards them if they wish. Chittaro et al. [37] also claim, *"Supporting users' navigation is a fundamental feature of mobile guides."* It is clear then, that navigation is an important aspect of mobile guides that we must take into consideration.

2.2.1 Approaches to navigation

Before discussing mobile guides further, we will first begin by clarifying the separate terms navigation, guidance and control. *Navigation* implies that there is some known end location that a person is attempting to arrive at, and does so by ascertaining their own location and following a particular route to the destination. As previously discussed, there are many ways in which a mobile device can determine location, so very often, when using a mobile device, the user is simply left with the task of following a route. *Guidance* is the act of a technology aid or other third-party directing or positioning a user through providing a route. It is concerned mainly with the kinds of instructions that are received by the user, and how they are interpreted. Finally, in terms of remote guidance, *control* is the method by which the guidance is given. It is concerned with the commands that are sent, and how they are transformed into instructions.

Mobile navigation appears in many different forms; from usage in cars [142], to cycling [123] and walking in public spaces [117]. Although designed differently, they are all concerned with getting a user from A to B. This is the aim of most navigation aids, though we must remember that we are attempting to design place-based experiences. Although we intend on allowing users to navigate from A to B, we must also consider the experience of being guided. This includes using serendipitous discovery and selection techniques such as those described in [117, 123]. We must also consider what kinds of instructions the user will receive, and also how they will be controlled.

Traditionally, most mobile navigation aids provide automated guidance, calculated and controlled by a computer. This kind of interaction is desirable in many situations, as it often very fast, accurate and reliable. However, we must again consider here, that we are trying to promote a more active engagement amongst people in the same place. It is not particularly desirable, then, to have visitors passively staring at their screen, following automated routes that will take them from one point of interest (POI) to the next. Perhaps, with some thought, we can attempt make the navigation element more socially engaging.

We propose that in these situations, it may be beneficial to have a human that manually provides guidance. Humans have been used as controllers in remote guidance before [8], though mainly to aid visually impaired people. There are many reasons why human controllers could be beneficial to social engagement. Firstly, the human controller may have expert knowledge of a location, and be able to guide a user on different routes, to different POIs. There is evidence to suggest that the use of mobile guides encourages people to visit rare and obscure

POIs [141]. Secondly, humans are better at understanding social conventions, knowing how to provoke or excite people, and also when to avoid conflict or danger. We also argue that it is comforting when being guided by a person, based on the fact that quite often, we contact friends or family for directions when we are lost. Although there is an abundance of research and related work regarding navigation, very little of it focuses on one human user being controlled or guided by another.

Guiding techniques

As previously discussed in Section 2.2, many mobile guides are map-based. Typically, when guiding users, these maps display a suggested route, direction or path that the user should follow. Whilst maps provide a good representation of an area from a geocentric perspective, they do not always provide a good representation of what a user observes from an egocentric perspective [37].

There have been numerous attempts at multimodal navigation using feedback aimed at the user's perspective on the ground, e.g. [37, 94, 116]. Jones et al. [77] discuss a mobile application where a user is guided by a simple audio player, providing cues by panning music to the left and right ear. *Audio Bubbles* [94] was another attempt at audio navigation that uses simple non-speech sounds to signify proximity to a point of interest. The work builds on earlier audio navigation research [69, 136], though removes the need for stereo headphones, by instead using a one-dimensional Geiger Counter metaphor. Another method of eyes-free navigation from the users perspective is discussed by Robinson et al. [116], this time using the haptic modality. Vibrations were used when a user pointed in the correct direction to move.

However, for navigation over a small area, arrows have also been found to be a simple, unambiguous, egocentric instruction which are easy to interpret [37]. Here, the researchers compared three different approaches to displaying guiding instructions to user, using audio with a mixture of visual feedback such as maps, photographs and arrows. The results of this research concluded that participants performed significantly better when arrows and photographs were used, though *"the combination of map and photographs was highly preferred by users."* Arrows and photographs have similarly been used in museum navigation before [152], though projected from a handheld pico projector.

2.2.2 Discovery, selection and control mechanisms

In our location-based experiences we wish for people to be able to interact with certain objects, artefacts, monuments and people in the real world. Because there could potentially be many interactive entities in one space, we need a good way of determining which one we would actually like to interact with. Research exists on 'Interaction with Web Services through Associated Real World Objects' [23], although we focus less on the commercial aspect of things and discuss the best method of interaction for an outdoor guide. The idea of our exhibits being controllable and interconnected relates closely to the concept of the 'Internet of Things' [144], where the future Internet will consist of many different aspects. One of these aspects involves interaction with real world objects that are interconnected, some of which include objects that can read the world and think for themselves. In our quest to understand discovery, selection and control mechanisms, we discuss a number of example mobile applications that employ different mechanisms, providing different user experiences.

The *Ambient Wood* [119] project is a system that is entirely focused on facilitating learning for students through scientific investigation opportunities and augmented reality in a wooded area. The system comprises of two different types of augmented reality. The first is like many others we have previously mentioned; live information and graphics being displayed through a mobile device screen. The second type of augmentation is that of an audio nature, where sounds are played back from predefined loudspeakers situated throughout the wooded area. This augmentation could either be student initiated or automatically environmentally initiated. When conducting their user studies, Rogers et al. noticed that student initiated augmentation seemed to encourage more collaboration as opposed to environmentally triggered augmentation. This was due to the fact that students often missed the environmentally triggered sounds as their attention was elsewhere. The dangers of information overload were discussed as a possible problem with that particular implementation of the system. The researchers then tried another approach where students were alerted on their mobile devices and could choose whether they wanted to hear environmentally triggered sounds. This version showed a vast improvement in student recognition of environmentally initiated augmentation.

Yiannoutsou et al. [161] discuss two mobile museum games which catered for both children and adults. The simpler *Donation* game allowed young children to explore the museum, gaining clues from exhibits through RFID tags along the way. These clues would then point them to a final answer. This is similar to a game presented by Ceipidor et al. [32], though using QR codes instead of RFID tags. The adult game named *Museum Scrabble*, is of slightly higher difficulty, requiring users to find exhibits that are closely linked in the museum. The more closely linked two items are, the higher the points the person scores. This game once again employs selection through RFID tags. In both of these games, collaboration was encouraged, prompting users to share answers that they had found.

Although we have only provided a small number of examples here, there are in fact a whole range of ways to select, gain access to and control location-based content. In the following sections, we discuss some of these interaction techniques in more detail, weighing up their pros and cons.

RFID / NFC

NFC and RFID technologies have existed for a number of years, but are only recently being exploited on mobile devices. Due to the fact that these technologies are relatively new in the mainstream mobile market, there is often confusion over what the technology is, how to use it, or whether one's device actually has the capability of interacting. In an attempt to improve accessibility to these technologies, some researchers have run user studies to determine how best users may be guided and learn how to use such technologies [24].

Both RFID and NFC technologies have previously been used in the tourism context. One of the more interesting applications of RFID technology in the tourism domain was an interactive tour guide system, where each visitor was given pieces of paper which acted as the *Glue* [55]. In this experience, visitors traveled from place to place with paper clues containing RFID tags on them. The RFID tag could be scanned at different points to interact with exhibits, providing new clues along the way. Results showed that visitors were engaged with the task delegated to them and were able to make connections between points in their visit, showing a coherent experience. Another instance of using RFID tags to interact with exhibits is described by Mantyjarvi et al. [90]. The level of interaction in this research was then further developed by allowing users to tilt their mobile device to select, navigate or activate interface items that were available. The purpose of this research was to try and employ a natural feeling set of interaction techniques.

RFID and NFC technologies have many applications. The examples we have mentioned so far have involved handheld devices and objects containing RFID tags. Other research has explored the possibility of wearables, where the users themselves wear the RFID tags [28]. This would allow objects to seamlessly interact with users as they approach. The results of Cafaro et al's [28] research revealed that response time can become a large problem in wearable RFID interfaces, concluding that such issues needed to be overcome before the system is practical for use in the wild.

An example of an NFC system is *Touch & Interact* [65]. Here, NFC tags were used in an array formation under a projected display, where users were able to interact with different parts of a map using their NFC enabled mobile devices.

Each of the previous examples demonstrate that RFID and NFC are a viable way of determining a particular object that a user wants to interact with. However, unlike in other countries such as Korea [39], RFID & NFC technology is not hugely popular with mobile device users in the United Kingdom, and on the whole, support does not appear to be improving. For this reason, radio technologies such as RFID and NFC do not appear to be good solutions.

Pointing

Of the three interaction techniques we cover here, pointing appears to be the most common identification interaction technique in tourist/outdoor guide applications. Intuitively, when in the company of other individuals, we often point at entities we are trying to describe, identify or query. In terms of augmented reality or image recognition [5, 38, 44, 148], pointing is a very well suited interaction technique. The user looks at the device screen, at the same time the lens of the device facing the scene the user wishes to augment or capture. This is an obvious example of pointing, but there are also more obscure methods of pointing that fit our aim of trying to determine which objects a user wishes to interact with.

Researchers from Simon Fraser University and Dalhousie University of Canada, have explored using infrared technology as a way of determining device identity [138, 139]. The technology works by the means of an infrared emitter which is the pointing device, and an infrared receiver on the device the user wishes to identify. The intended aim of this project was to determine device identity easily, so that data may then be transferred to the identified device via wifi. This idea of using infrared for pointing is a good one, as infrared rays travel in straight lines (wherever they are directed to do so).

Delving into the some of the more experimental and not so well established pointing technologies, we find *VisionWand* [30]. *VisionWand* is described as an interaction technique for large displays, 'using a passive wand tracked in 3D.' The motivation for this project was to see if a new and improved kind of input technology could be developed for large displays. This solution, however, is rather unattractive to us as the wand contains no sensors of its own, requiring an array of cameras and additional technology to carry out the tracking. All of the movement detection is carried out through computer vision technology by two cameras situated behind and beside the wand user. Inherently, one of the issues with this kind of technique is the possibility of the wand user or other individuals breaking the line of sight from the camera. In terms of a solution to part of an outdoor tourist attraction such as the Hafod-Morfa Copperworks, this technology is probably not a good approach, as the area is likely to get crowded at times, rendering this technology completely unusable.

Another experimental technology that involves pointing—and also obstruction—is [35]. In this example, light is the pointing mechanism and is sourced from a handheld torch. The user turns the torch on and points at the projected screen, using the torch as a selection and manipulation device. The light position and behaviour on the projection plane is interpreted by a camera. This camera also detects obstruction and can draw shadow representations of the obstruction to the screen when needed. The technology was developed as a toy for child's play, although it's easy to see the attraction of this simplistic implementation in a tourist setting. However, this technology is again very bulky and resource heavy.

From studying previous point-based approaches on a mobile device, it would seem that the best solution—especially when outdoors—would be to use GPS to determine location and a digital compass to measure heading. This kind of interaction technique is commonly used, is lightweight and does not suffer from the obstruction problems we have mentioned here.

QR / Barcode

Barcodes and QR codes are very useful for storing and retrieving small amounts of basic data such as text, numbers and URLs. Whilst barcodes were designed to work with a laser scanner, QR codes are inherently more useful to mobile device users as they can be detected through a camera lens.

QR codes are essentially two-dimensional barcodes that allow for a much greater storage capacity. The capacity of each code is dependant upon the version (size of the grid), the character set used and the error correction level. Most commonly, QR codes are used to store URLs to websites. As previously discussed, QR codes are detected through image recognition. To detect a code, a user requires a mobile device with a camera and then an application that can interpret the photo. Although barcodes were originally designed to be detected by a laser, barcodes now too, can be detected using similar image recognition techniques.

In a research paper similar to the previously mentioned RFID paper clue game [55], Ceipidor et al. [32] developed a QR code clue game where young students had to answer historical questions by finding the related artefact and scanning the QR code below it. When the correct code was scanned, the answer to the question was revealed to the student. With each clue, the student would also gain a secret letter. These letters would then form a word at the end of the game. The student's final goal was to find this word.

With an increasing number of mobiles featuring cameras—providing the capability of reading QR codes—it would seem a useful interaction technique.

2.3 Face-on interaction

In this research, we are interested in developing technologies that promote a more socially active engagement between visitors. For this reason, we would like to divert attention away from the technology and mobile device screen as much as possible. Rather, while engaging with our interventions, we would prefer if visitors were able to focus their attention on their surroundings and the people who populate it. We would like visitors to interact, face-on. We believe that there are a number of ways in which we can design our interventions to encourage this kind of behaviour. Here, we present a mixture of existing eyes-free and heads-up interaction research, concluding with considerations on how these findings may help promote future face-on interaction.

2.3.1 Exploring eyes-free & heads-up interactions

There are many reasons why designers may want to design for eyes-free, heads-up interactions. Designers reasons range from safety critical [89]—focusing visual attention in the right places—to experience-based [113] – allowing users to look around at the world and *"get on with their lives whilst using the technology"* [21]. Robinson's experience-based angle sets out a similar vision to ours, believing that,

> *"our focus should not be constantly and solely on the device we are using, but fused with an experience of the places themselves, and the people who inhabit them."* [113]

By developing and evaluating a range of new and novel 'eyes-off' interaction techniques, Robinson was able to demonstrate how these kinds of interactions can support exploration and improve the user experience.

Eyes-free modalities

Considering first, eyes-free interactions. Oakley and Park remark that, *"The fundamental motivation for eyes-free interaction is that as it leaves visual attention unoccupied, users are free to perform additional tasks."* [100]

Both the audio and haptic modalities are well suited to eyes-free interaction, as they require little or no visual attention from the mobile device. This is true of both modalities, as inputs and outputs. While visual feedback is often considered a necessary, corner stone to mobile computing—with the audio and haptic modalities seen as complimentary [160]—there are in fact many examples of mobile systems that consist purely of eyes-free, audio and haptic interactions [78, 156].

Audio has been used for a range of basic, eyes-free feedback mechanisms in user interfaces. These include things such as task progression [41] and notification events [121]. More recently, commercial mobile devices have become capable of rendering spatial audio [145]. Vazquez-Alvarez and Brewster ran a user study with one of these early devices, testing the user perceived error of directional sound. Each participant was asked to listen to a number of sounds and then point to the direction the sound originated from immediately after. The study showed that participants could quite accurately detect 45° intervals of sound in the front 180°, proving that mobile devices were indeed beginning to simulate three-dimensional audio to a fairly acceptable standard. The advent of spatial audio led to a number of new experiences, including soundscape exploration [135], new eyes-free navigation techniques [136] and improved conference calling [47].

In this thesis, we are particularly interested in explorative, experience-based approaches. A good example of such an approach is the Roaring Navigator [135]. Here, users were able to explore a real zoo, hearing the physical locations of animals through headphones. Users could then decide whether they wanted to walk towards the origin of the sound, discovering the animal. Other researchers have recreated similar environments, presenting a virtual zoo that users can walk around [66]. This, and many other previous systems have used a digital compass on the user's headphones to measure their orientation, though Heller et al. [68] argue that it is just as useful—and perhaps, simpler—taking device orientation measurements, as users tend to direct their devices in front of themselves.

All of the systems we have described so far have made use of headphones. There are, however, also opportunities to use loudspeakers in eyes-free interactions. Ardito et al. developed *Explore!* [4], placing a speaker in children's backpacks, which were then activated by changing locations when walking around. The children had to try and map the area based on these sounds and other clues. Another educational mobile game experience—*Ambient Wood* [119]—embedded speakers in the environment, which were then triggered during activities.

Similar to previously discussed interface-based audio research [41, 121], other researchers have also looked at haptic feedback methods for eyes-free interactions with touch screen interfaces, rendering progress bars and scroll bars through piezoelectric actuators [85]. Again, similarly to audio-based research motivated by eyes-free interactions, haptic feedback has also been used for target finding and navigation. In [114] Robinson et al. investigated the difference in finding a target using both visual and haptic navigational cues. They discovered from their study that haptic feedback had very similar target finding results to visual feedback in 2/3 cases, proving it as a good modality to aid navigation. The researchers also concluded that the eyes-free navigational method of haptic cues allowed users to carry on communicating and avoiding without ever needing to look at the mobile device screen. This research was later further developed [117] by allowing users multiple choices of what route to take. Instead of a user following a vibration at a single point, if multiple routes were detected, the device would vibrate over a larger area indicating to the user that they may head anywhere in that direction. As well as mobile haptics, researchers have also looked at eyes-free interaction techniques for wearables, demonstrating an *"effective alternative to visual-centric interface designs on mobile devices."* [22]

Considering, again, experience-based approaches that are more similar to our mobile guide scenario, McGookin et al. [95] developed *Virtual Excavator*, a multimodal location-based experience. In this system, previous archeological finds were geolocated at an unstewarded heritage site. Users of the system were encouraged to physically explore the site and uncover these finds. When a user stumbled upon the location of an artefact, the device would vibrate. Further haptic input was then encouraged, with users required to shake the mobile, simulating digging and uncovering the find.

Heads-up displays

When referring to the term 'heads-up' interaction, researchers are most often referring to visual interfaces. In many of these cases, researchers are not attempting to remove the visual modality entirely, rather, visual feedback is moved away from a device screen and is placed within the user's line of sight. Simply put, the user is able to continue carrying out a task without having to glance away at a screen. While many of the previous audio and haptic approaches we have already discussed could indeed be considered as heads-up, our specific focus here is on heads-up visual interactions.

The earliest mention of heads-up displays (HUD) appears in aviation [154], where pilots requested that vital information be displayed in front of them. This is a safety critical design, which has also been applied in other areas of transportation, including cars [89]. In both of these situations, the user's attention is best placed on the direction they are travelling in, so it is desirable for important information to be displayed in the user's direct line of sight.

It has been argued by many researchers that mobile devices introduce 'fragmentation of attention' [102], with users slowing down to write text messages and withdrawing from less important tasks. Designers have looked at addressing these concerns in a number of ways, including showing a live camera image behind a text message, allowing users to *Walk and Text*[1]. More recently, technologies such as *Microsoft HoloLens* [96] and *Google Glass* [60] have appeared. Unlike many wearable displays that attempt to immerse a user in virtual reality [62], *Google Glass* and *Microsoft HoloLense* are lightweight wearables that augment the real world, allowing a user to carry on their daily life, uninhibited.

Perhaps, a more established and common form of heads-up display is digital projection. Just like wearables, projectors can be portable, and can project onto a whole range of different surfaces. Bert Bongers [19] was one of the first to realise the potential of the projector, taking it outside of its intended context, projecting 'gateways to alternate realities' in the streets. Now, many years on, public outdoor performances often include the use of large-scale projections in a

[1] `https://play.google.com/store/apps/details?id=com.incorporateapps.walktext` - Walk and Text (Type n Walk)

similar way (see `projectionartworks.com`).

In our scenario, there are multiple benefits of using a projector as opposed to a wearable heads-up display. The main advantage is that instead of a single user witnessing the visualisations, many more people can witness a projection at once, allowing for shared, collaborative experiences. Of the range of commercially available projectors, pico-projectors are both lightweight and affordable, making them particularly useful for such purposes. Research has already begun to realise the potential of handheld projectors in a collaborative context [115]. The focus of *Picotales* was to allow users to collaboratively author stories through drawing and animating projections by moving the projections around accordingly. There is also a possibility of allowing multiple users to stream image or video through a peephole system similar to that mentioned in [88]. With many different users, each with a pico projector, one could quickly build collaborative, large-scale displays. This has already been done with conventional mobile displays [131], and also more experimental displays [34]. Integrated projection may be widely available to mobile device users in the near future [159], making such scenarios a compelling use case for heads-up displays.

2.3.2 Encouraging face-on interactions

To foster a face-on experience amongst visitors at the Hafod-Morfa Copperworks, we need to move the focus away from the technology by ensuring that interactions do not happen around a small device screen – we need interaction to occur outside of the device. Cabrera et al. [27] remark that guided tours and similar technologies are often built around the mobile device, and not around the physical space they are intended to be used in. If there is too much focus on the technology, it may detract all attention from the actual exhibit, people and location [43]. We want our experiences to be deeply embedded in the place, with as little focus as possible on the technology. It is clear, then, that we need to employ a range of eyes-free and heads-up interaction techniques to encourage a face-on experience.

We have already seen from a number of previous examples [4, 95, 115], that moving attention away from a device screen can be a good way of facilitating collaboration. Large-scale visual displays such as projection, and spatial audio interfaces both appear to be popular ways of trying to divert attention outside of the screen, encouraging users to look around them. Haptic feedback such as vibrations also allow users to accomplish tasks without a mobile device screen. Gestures have then been used a method for eyes-free, heads-up input, allowing exploration of physical spaces [113] and interaction with interfaces [95]. We will make use of the range of modalities that we have, considering also possibilities of interacting with internet connected peripherals, outside of the device.

2.4 Socio-spatial interaction

In this thesis, we are interested in the use of technology to promote engagement in public spaces. In our particular case, we focus on the digital regeneration of the Hafod-Morfa Copperworks, encouraging people to visit and use the site, and to start conversations when they arrive there. As argued by Bederson [9], one of the main reasons people go to museum is for a social experience. To enhance our chances of designing successful interventions, we must first gain some understanding of socio-spatial interaction, and how certain theories and related work may be able to help us to develop such interventions.

Ishii and Ulmer presented an early vision of interacting with the world when we are out-and-about, rather than interacting directly with a computer interface [71]. Weiser also had a similar vision of ubiquitous computing [155], though with computers invisible to the eye – hidden from view. The term 'embodiment' has been used by Dourish [49], to draw together the two similar themes of tangible and social computing. Embodiment plays a particularly important role in phenomenology – explaining how we perceive, experience and act in the world around us.

Hornecker and Buur [70] draw on previous research, presenting four themes to help evaluate tangible interactions in physical spaces. The framework discusses ways of evaluating tangible manipulations, spatial interaction, embodied facilitation and expressive representation. Spatial interaction and embodied facilitation are useful for providing insights relevant to the broader area of embodied interaction, useful for understanding and supporting social interaction. The remaining themes address social interaction in a more indirect way. Their definition of the tangible design space is purposely left broad, encompassing such characteristics as 'tangibility and materiality', 'embodied interaction' and 'embeddedness in real space.'

Technology designers have been concerned now for some time with social interactions mediated by technology in public spaces. An early example of such an experience is Bederson's collaborative audio exhibit in a museum [9]. The technology allowed groups of visitors to walk up to an exhibit together, which then triggered an audio description. After visitors had walked away from the exhibit, the audio description stopped. A similar approach has more recently been used in a Denmark art museum, also attempting to facilitate a more social experience. *The Sound of Art* [83] encouraged visitors into a small physical space—marked out by a circle—to hear audio descriptions. The sound was directed from above in such a manner that only visitors standing in that particular circle could hear the audio. Both of these examples are concerned with being part of a shared experience in a small, confined space.

Researchers have also commented that in publicly accessible experiences, there is a potential for strangers to come together and collaborate [92]. Marshall et al. suggest that in these scenarios, users are not dependent on any particular person in a collaborative system, and that the goal is not too personal that users will shy away. In contrast to this view, however, Benford et al. [16] recommend *"deliberately and systematically creating uncomfortable interactions as part of powerful cultural experiences."* They suggest that designers may want to include discomfort through visceral, cultural, intimate and controlled experiences, *"not to cause long term suffering or pain, but rather to underpin positive design values related to entertainment, enlightenment and sociality."*

In the following sections, we discuss existing methods of community engagement and regeneration using technology, and also explain our reasoning for designing and evaluating our interventions from a performative perspective – creating spectacles and trajectories, and encouraging visitors to take a more active role during visits.

2.4.1 Community engagement and regeneration

During its heyday, the Hafod-Morfa Copperworks and surrounding industry was once the heart of the Lower-Swansea Valley community. The works would have been visible from miles away, with plenty of noise and activity occurring at all times of day. Now that the copperworks have been closed for some time, the area occupies an eerie feel. An area that was once so densely populated and full of life is now desolate.

At the Hafod-Morfa Copperworks, as well as building experiences that can serve tourists the information and interactions they request, we are also interested in building and facilitating a framework for communication amongst the local community. There are a number of such research projects where attempts have been made to rebuild communities through the use of novel technologies [107, 110, 160]. The reasons for each of these projects range from engaging younger individuals in community developments, to trying to build or rebuild distant communities. Our motivation stems from trying to attract people into and make use of an important area of a community that has been long forgotten.

The first of the aforementioned projects [107] focused on the design and development of a system to support local community communications. The reason for this was due to the community's organisation and communication usually being kept private and out of view from the public. Due to the nature of this private communication, information is often lost and no other record of communication is kept. For community issues, the researchers believed that there should be a central system where all local community related communication could be carried out. This would ensure that the local community were kept up to date with goings on and public announcements. Through an iterative design cycle, situated public touch screen displays were deployed, along with a mobile interface design for phones. Both of these methods allowed the local community to read and leave quick messages. A more complex system was then provided for internet users on their home computers, allowing the community to moderate the system and gain a deeper experience of the ongoing communication. At first it seemed that the situated display was fairly popular, although after the first email digest was sent to members of the community, the web usage statistics soared far beyond the display's usage ones, and remained steady. The situated display's initial burst of usage was most likely due to the novelty of the technology, which quickly wore off. An analysis of the messages left by users on the display showed that they were often either private or relevant to a small group of people. The messages left on the display were noted to have a different nature to those left online and were often not very serious. From this research we can conclude that situated displays can potentially offer fun interaction, but are perhaps less useful for serious community building and communication.

Another similar project is presented by Yamada et al. – the *YeTi system* [160]. Just like [107], this project focused largely on the concept of community building, with its main intent on discovering who is using the system and how. Unfortunately, unlike the previously mentioned user study that was conducted in-the-wild, the user tests conducted in [160] were lab studies containing mainly staff from technology departments. Yamada et al. did provide statistical usage of the large display, although we consider these statistics less useful, as the user-base in this experiment is not a good reflection of our target audience. It is also worth noting that the remote communities that the study attempts to bring together are ones that are located in Japan and the USA. Techniques that we will consider will focus more on bridging the community in a more collocated manner.

Other research has discussed the potential use of mobile blogging to articulate community messages [110]. The aim of the *Moblogging and Belonging* project was to encourage young people to think of their media practices as part of their community rather than just personal experiences. Richardson et al. concluded from their investigation that mobile blogging did indeed have the potential to be a good medium for social inclusion. Since this paper was written, many developments have occurred, such as the introduction of smartphones, increased connectivity and cross-compatible social network applications. These developments have all helped boost the profile of mobile blogging, and with the widespread availability of localisation in smartphones, the idea of local and community driven posts is becoming more popular. Social networks have a huge influence on the way we currently communicate online. If we were to pursue the idea of bridging the community remotely, it would be so through social networking. A physical link to these web services at the Copperworks could boost community communication.

Considering these previous examples, it is clear that we must choose the way in which people use our technologies to communicate very carefully. We can think of the means of communication that we choose as a 'Social Proxy' [46]. As DePaula discusses, we have to make design decisions very carefully in a 'community-centered design' approach, where we not only focus on the human-computer interaction but also the human-human interaction that is mediated by the technology. For the design to be successful, we must follow an iterative design approach that involves field studies of potential system users.

2.4.2 Performative engagement

To begin, the terms 'performative' and 'performance' have been interpreted in many different ways by different groups of people. This problem has been discussed in some depth by Spence et al. [134], who have attempted to classify the field of 'performative experience design' by reviewing the human-computer interaction literature that focuses primarily on performance. The researchers remark,

> *"Until we sort out what we mean when we invoke 'performance', we will fail to realise more than a tiny fraction of performance's enormous potential to contribute to and engage with the work in this area, or to identify new fields in the spaces between HCI and performance."* [134]

Following Spence et al.'s performative experience design taxonomy, we consider that our work here is primarily concerned with *staging* and *engagement*. Similar to other works in this area, we identify those involved in the experiences as 'performers', 'spectators' and 'bystanders' [14, 48, 108]. When in front of another person, we become part of a performance [57], where we can choose to acknowledge each others actions and engage in focused interaction [58], or continue in unfocused interaction by blending in. Users who interact directly with the system are referred to as performers. When others are aware of the situation and are in focused interaction [58], we refer to them as spectators. Otherwise, we refer to them as bystanders, who may become engaged at some point in time. This complex relationship, and the transition between each state has previously been explored by Sheridan et al. [132], differentiating between witting and unwitting spectators (audience and bystanders), and performers and participants (in control and taking part). The relationship between the user, system and spectator in a performance has also been discussed in depth by Dalsgaard and Hansen [42]. Here, the term *performing perception* is used to describe how the user is *"simultaneously engaged in three actions"* - the act of interacting, the act of performing and the act of perceiving. These three acts relate to being *"the operator of the system, the performer for people present, and the spectator of the action in her immediate surroundings"* respectively.

Much of the literature regarding the design and evaluation of artistic installations focuses on the theme of performance. The *humanaquarium* [140] is one such project, where the designers themselves became performers, improvising music from an acrylic box. Bystanders were able to engage with the performance, controlling the output by making contact with the touch sensitive front. Over the course of a year, the designers were able to *"gradually re-imagine content and interaction strategies in response to a deepening understanding of the design space gained through performing the work in public"* [140]. The design of other ubiquitous, performative experiences is a topic that has been discussed in depth by Jacucci et al. [73].

Extravagant, expressive performances

Technologically mediated performances appear in many forms – from discreet to extravagant. Thanks the Reeves et al., we are able to classify a performance based on its observability to spectators [108]. This is particularly useful to us, as we would like to design experiences that will draw in and engage bystanders. In this taxonomy, a performance or interaction can be categorised into one of four classes: secretive, suspenseful, magical and expressive (see Fig. 2.1).

Each class is determined by the observability of both the manipulation and effect with regard to the spectator. A performance with hidden manipulations and effects can be considered as *secretive*, where as a performance with amplified manipulations and effects can be considered *expressive*. We are particularly interested in expressive performances, where we would like to draw bystanders

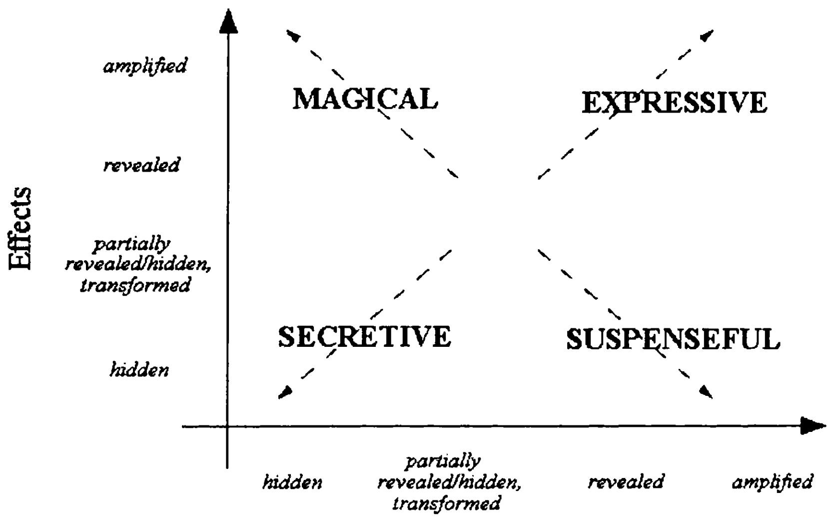

Figure 2.1: Reeves et al. spectator observability taxonomy. [108]

into focused interaction with the performers and performance, effectively making them spectators. To make such experiences accessible to both performers and spectators, Hornecker and Buur stress the importance of a *"clear relation between actions and their effects"* [70]. In most cases, short-lived performative interactions can be thought of as a single class, but Benford et al. [13] demonstrated that more complex, longer interactions can successfully cover all four classes in the taxonomy.

The acceptability of gesturing in public is a topic that has been covered in depth by Rico et al. [111]. Further work has also touched on the social acceptability of multimodal performance in public spaces [157].

From consumer to producer

At public, outdoor attractions such as heritage sites, visitors are often encouraged to behave as consumers. When visitors arrive at a site, they expect to play a passive role, exploring the space and absorbing the information that is present there. While this kind of model is very popular among standard location-based guides, it is perhaps lacking from a participation perspective. Bagnall [6] argues that when visiting these sites, consumption alone is not enough. She argues for the importance of playing a more active role in these visits – with performance and performativity appearing as *"key social practices."* Bagnall studied a number of performative experiences at heritage sites, with actors re-enacting past events for visitors. It was claimed that such performative experiences helped stimulate reminiscence through emotionally engaging visitors. She later remarks that *"the relationship between visitors and the sites is based as much on emotion and imagination as it is on cognition."* When designing a performative experience, Bagnall suggests,

> *"It is the physicality of the experience, the capacity of the sites to engage and stimulate a whole range of physical and sensory experiences, and the way the sites engage visitors on an emotional level that is important."* [6]

Though the performances Bagnall studied were of a more traditional, theatrical nature, we believe that these same values hold for technologically mediated performative experiences. By playing a more active role at the Hafod-Morfa Copperworks—performing with imagery and sounds of the past—we hope to encourage reminiscence in a similar way. Additionally, by making these experiences both extravagant and expressive, we hope that not only will visitors shape their own experiences, but also the experiences of others around them.

As identified by Sheridan et al. [132], getting people to wittingly participate and perform in a performance can be a difficult task. In an attempt to ease this transition, we will set the entry level low, with visitors able to access and control the performance easily.

2.5 Conclusions

In this chapter we have reviewed the previous related work in this area, situating our research within the field. Our goal in this thesis is to investigate the underexplored design space of performative, extravagant, expressive mobile place-based experiences, which encompasses a number of different themes.

With the advent of smartphones, GPS and widespread internet-connectivity, a range of useful location-based services and mobile guides have began to appear.

While these services provide a useful utility, most are lacking in some of the more exciting interactions that are now possible with mobile devices. Researchers have been quick to realises this, with many now designing location-based experiences – moving beyond core information services. Rather than providing hard information, location-based experiences focus on providing a fun and enjoyable experience for users, interacting in new and different ways. In this thesis, we attempt to extend this design space of location-based experiences, introducing new and novel experiences that make use of a number of existing mobile and internet-connected technologies, combining, applying and presenting them in new ways.

The main issue that this thesis aims to address, is that the majority of mobile location-based experiences still place a large amount of attention on the screen or headphones, failing to cater for more social experiences. We argue that one of the main reasons people visit places such as heritage sites and museums is to have a social experience [9]. These people want to look up at the world and converse with each other – they do not want to be divided by technology. To combat these concerns, we have looked at a number of heads-up, eyes-free interaction techniques. The reason for using these techniques is to remove the focus from the technology, and to place it back on the environment, and the people contained within it. People should be able to enjoy face-on, shared, collocated experiences, uninhibited by a mobile device screen or headphones.

Approaching from our Hafod-Morfa Copperworks regeneration perspective, we have looked at a number of past and existing community interventions that attempt to bridge communities and foster social engagement. We then looked to a theory-based approach. To encourage a more active social engagement while face-on, we have considered the potential design and evaluation of our mobile interventions from a performative perspective. We have discussed our reasoning for extravagant, expressive performances, and have also argued why performance and performativity are particularly relevant in the heritage context.

After reviewing the related work—to our knowledge—there appears to be little mention of mobile controlled performative experiences in public spaces. That is, using a personal mobile device as a control, manipulating public, extravagant, expressive effects. This thesis is a first step to understanding these kinds of interactions, and how they may encourage a more social active engagement.

Chapter Three

Approach & Methodology

As active members of the *Cu@Swansea* project (see Section 1.3), we have been heavily involved in much of the process, attending meetings, workshops, open-days and a range of other events related to the project. Throughout our research, the entire project has been a lens through which to develop ideas, engaging with and shaping the final outcome of this work. The Hafod-Morfa Copperworks is a kind of living laboratory, where things are continually changing and evolving. At the beginning of this project there was no interpretation on-site. Even now at the time of publishing, the site is by no means a finished, curated piece of heritage. Designing our interactions within an evolving space has allowed our ideas to coevolve with and shape the project, and although none of our interventions up to this point have been permanent installations, they will have had some impact on the way visitors experienced or now perceive the area.

The methodology and approaches outlined in this chapter are heavily influenced by the fact that the work in this thesis forms a multidisciplinary PhD, crossing boundaries between modern history, computer science and social science. In this work, we have an important duty to be respectful and preserve the significant history at the Hafod-Morfa Copperworks, at the same time, conceptualising novel and engaging technologies that use this rich heritage to its advantage – encouraging conversations and social regeneration.

The following chapter describes our design approach, conceptual framework and outlines the selected research methods and approaches applied during the experimentation and evaluation of the prototypes within this thesis.

3.1 Design, experimentation and evaluation

In this thesis, we followed an *exploratory design* approach, making use of a mixture of well understood design traditions. The initial interaction concepts and designs were born from our *conceptual framework* (see Section 3.2), which in turn, was heavily influenced by a previous well established framework [108]. Each design followed these strict guidelines, attempting to build a performative, extravagant, expressive experience.

The design of each system was equally *novelty-led*, attempting to identify new and unexplored interaction techniques. Before any development commenced, we first carried out extensive literature reviews, discovering key gaps in the literature. From this, we were able to identify a set of opportunities for new and novel interactions. These opportunities helped fuel a number of designs – many of which, we implemented as fully working prototypes. For the development of each interaction technique, we chose to employ an *iterative design* [98] approach. Early testing in controlled settings and pilot studies allowed us to learn of any potential weaknesses or defects with our proposed interaction techniques. The redesign and implementation of each system then continued until they were ready to be publicly deployed for further, in-depth evaluation. Most of these deployments could be considered as *in-the-wild* [120], testing the interactions and experiences in their intended context. In many of these deployments, participants were pre-recruited, though in one particular deployment, we were able to test the experience more naturally with unsuspecting visitors and users.

Along the way, we also carried out a number of community engagement activities. These ranged from focus groups within the local community, to group discussions with project stakeholders. These focus groups and discussions proved to be invaluable engagement activities, allowing us to gain additional perspectives on our design concepts. These meetings also became a place for debating the potential benefits and drawbacks of introducing more extravagant, expressive place-based services to the Hafod-Morfa Copperworks (see Section 3.4 for more details). Although the local community and stakeholders were not involved throughout the entire design process, these *participatory* and *user centred design* approaches played some part in the final concepts and prototypes.

3.2 Conceptual framework

When trying to develop and understand more performative, extravagant, expressive mobile location-based experiences, it is first important to map out and understand the design space. Building upon Reeves et al.'s spectator observability taxonomy [108], we highlight the distinct areas of a performance that we were interested in exploring: manipulations; effects and visibility:-

Manipulations: Gestures (point, tilt); Movement through space

Effects: Visual; Audio; Movement through space

Visibility: Amplified

The aforementioned manipulations and effects make use of a number of modalities that are widely available and compatible with existing mobile smartphones. With these modalities, we were interested in trying to develop expressive experiences, with amplified manipulations and effects. We hypothesised that these would provide the most accessible and open experiences for visitors in public spaces.

We achieved amplified **visual** effects through using pico-projectors. These were small, smartphone sized projectors attached to the mobile device, allowing the user to project content into the environment instead of viewing it on the screen. By bringing this content to the environment, it was hoped that it would be seen and experienced by bystanders who were nearby. Amplified **audio** effects were achieved through using collections of loudspeakers. These ranged in size, from onboard mobile device speakers to 10-inch portable loudspeakers. **Movement** through space was achieved by guiding users with live directions, sent remotely by another user from their own device. With this remote guidance mechanism, users are also able to disseminate sounds or projections as they walk around, acting as portable media hotspots. We could have instead chosen to experiment with other more complex effects, such as controlling and producing smoke through a mix of olfactory and visual techniques, though we decided against this, as we felt that audio and visual technologies alone provided a good baseline for understanding the value of mobile performative technologies. We discuss more complex effects as an avenue for future work in Chapter 8.

From a control perspective—when using a mobile device—one of the most obvious, amplified manipulations a person can engage in is physically moving and waving their device. For this reason, we chose to use **gestures** as a means of amplified control. The gestures we chose to explore included pointing and tilting. Again, it was believed that the enhanced visibility of gestures made it obvious that that particular person was controlling the performance, making them both visible and approachable.

3.3 Deployments

When each of our prototypes was at a robust enough stage, public deployments were carried out. As stated by Oulasvirta [103], *"field experiments are required when phenomena do not fit in the laboratory or cannot be simply staged there in a convincing manner."* We chose to carry out public deployments instead of more controlled lab studies, as we believed that to achieve reliable usability results, we had to introduce the systems to more public contexts, similar to where they would finally be used. Conducting a controlled study in a lab setting may certainly have taught us fundamentally whether systems worked or not, but it would not have given us a real sense of how the system and its users perform in a real world, public context, where things are ever-changing and less predictable. As well as being able to measure usability, public deployments also allowed us to observe how bystanders behaved and reacted towards our interventions, something that would not have been possible in a controlled setting.

Similar to other evaluations of location based experiences [95], we compared a number of our systems with comparable, existing technologies. In other, more naturalistic deployments, and where a comparison was not suitable, we followed a similar approach to Benford et al. [15], studying and making sense of the interactive experience in-situ. By carrying out these field experiments, we were able to learn a large amount regarding our interaction techniques from a usability and usefulness perspective. We were also able to begin to understand the social impact they have, and whether these systems are indeed good for promoting more social engagements in public spaces.

In the early days of the Hafod-Morfa Copperworks reopening to the public, there were very few visitors, making it difficult to carry out recruitment and measure any kind of social impact. In one instance, participants were pre-recruited off site and taken there for a study. For other deployments that were not conducted onsite, we chose to try and conduct them in similar visitor centre contexts, including the National Botanic Garden of Wales and Oystermouth Castle. Further experiments that were not conducted in any of these locations were done so in other public contexts with bystanders present, such as Swansea University and a shopping mall.

3.3.1 Study technique

Apart from a small number of controlled usability experiments, all studies were deployments of prototype systems in public spaces. Each of the systems were designed for the particular context they were to be deployed in. The majority of experiments involved providing participants with a mobile device to take part. This is common practice in the evaluation of mobile location-based experiences [43, 119]. In one deployment, however, no devices were provided, with visitors able to freely discover and use the system using their own mobile device. We believe this kind of deployment is closer to reality, as visitors are often expected to bring and use their own mobile devices to interact with public interpretation.

Throughout the deployments, we used a mixture of pre-recruited participants and participants recruited in-situ. Participants were sometimes pre-recruited due to the difficulty of being able to reliably recruit people in public spaces. Similar to other location-based experiences for visitor attractions [4, 54, 95], we focused on recruiting groups and pairs, though for certain experiments we also recruited individuals. We also ensured that different participants were used in each study – none of which had any experience with the systems before taking part.

Prior to conducting any study, demographic information was captured for each participant, allowing us to provide an accurate representation of our pool of participants. In all instances, upon beginning the study, participants were briefed on how the system worked and what their specific task was. When participants understood the system and were ready, they began to carry out the task. In studies with multiple tasks, participants would return to the investigator after completing each task for additional briefing. During each study, depending on the task, we captured a range of data, including participant and bystander observations, device logs and questionnaire responses. All studies finished with an interview, debriefing participants and digging deeper into their experiences. Depending on the suitability, most studies involved the capture of observations via video recordings, and interview responses via audio recordings. In all studies, participants were given a monetary incentive to take part. This ranged from £5 to £15 depending on the length of the study. The only exception to this was our studies conducted in a shopping mall, where there was no monetary incentive for participants.

3.3.2 Data collection and analysis

During our deployments, we followed approaches in other studies that have measured the effect of mobile interventions in public scenarios [147, 156, 157], recording observations, conducting questionnaires and interviews, as well as capturing device logs.

Observations were useful for analysing participant and bystander reactions, behaviour and movements. We took a similar approach to that mention by Ardito et al. [4], taking a mixture of observational notes and video recordings, though ensuring not to get too close as to influence participant behaviour [59]. Observations were especially useful when video recorded, allowing us to play back and observe situations from a distance, making note of things that may have been missed when conducting the experiment. All observations were transcribed, which then allowed us to carry out thematic analysis [20]. We also used video observations to measure task timings in a controlled experiment. Video recordings proved to be a very useful form of deconstructing and analysing study sessions.

Log data proved equally helpful in this regard. Prior studies have already shown the value of log data [147, 156], enabling the capture of data that cannot be easily observed or accurately measured. By recording such things as the locations the user visited and how they interacted, we were able to build a detailed view of each session. Capturing individual button presses, gestures and other instructions also allowed us to determine how individual features were being used by participants.

In some experiments, we used questionnaires with Likert-type scales, capturing self-reported low to high ratings [87]. Statistical analysis was conducted on these scores to learn of any statistical significance, the result of which allowed us to compare users opinions on different systems, or individual elements within each system. Though borrowed from psychology, Likert-type scales are commonly used in human-computer interaction research, with 45.6% of the CHI 2009 proceedings making use of them [79].

All of the interviews we conducted were semi-structured. This allowed us to blend conversation and discussion in with questioning, adapting the interview based on prior answers and providing additional insight into the investigation themes in a friendly manner [97]. Some interview questions were posed in an explanatory fashion [112, p. 31], prompted by observations or answers given by participants in questionnaires. Others were exploratory based, probing participant's own observations, opinions and feelings – unknown to the researcher prior to asking. By audio recording these interviews, we were able to go back and re-listen to each interview in more detail, capturing the dialogue, subtle tones, laughter and other important emotions that may have been missed. All interviews were transcribed. Thematic analysis [20] was again used when analysing these interview transcriptions.

By taking advantage of methodological triangulation [45] and combining the analyses of these individual data points, we were able to build up a clear, more reliable picture of each study session.

3.4 Community engagements

As well as deployments, we carried out a number of community engagement events, where the technologies and underlying concepts were demonstrated and discussed with project stakeholders and the local community. While engaging with individuals and smaller groups yielded very personal and individual opinions, focus groups were chosen as our primary engagement as they provided homogeneous perspectives from both the community, and also stakeholders [63].

Both the stakeholder and community focus group engagements were prearranged, though no incentives were provided to members of these meetings. Other meetings with individuals and smaller groups were moulded around deployment studies. This gave us an opportunity to demonstrate additional interaction techniques and concepts, gaining further useful insights from participants. At each community engagement, at least one system or interaction technique was demonstrated – the feedback of which, helped inform some later changes and new design concepts. As much as participants critiqued the individual designs and encapsulating notion of mobile, performative technologies, they also commented on their personal concerns regarding the Hafod-Morfa Copperworks – sharing their own thoughts and vision on how the site should be.

The community engagements followed a variety of formats. In more formal events such as our stakeholder and community focus groups, a presentation was given, along with a brief demonstration of the technologies being discussed. Members of these meetings were then able to trial these technologies and discuss their thoughts amongst themselves and with the researcher. As the community focus group was a particularly large group, participants were asked to leave any additional parting comments on paper slips which were collected at the end. On other occasions such as select deployments, smaller groups and individuals who had already been exposed to one technology—by participating in a deployment—were given a brief introduction to an additional interaction concept, and then given the opportunity to discuss the technology, along with the notion of performative, extravagant, expressive place-based experiences. These participants were given monetary incentives for taking part – mentioned previously in the deployment section. All community engagements were audio-recorded, allowing us to later transcribe these conversations and deduce themes through thematic analysis [20].

3.5 Summary

We have defined the performative, extravagant, expressive mobile conceptual framework that we were interested in exploring, and based the design and development of our interaction techniques on key gaps we discovered in the literature. We carried out the evaluation of our interaction techniques and concepts through a mix of public deployments and community engagements. The deployments gave us an opportunity to place the experiences in their intended context, allowing us to gain valuable usability feedback and begin to understand the kind of impact that these interaction techniques have on bystanders. Our community engagement focus groups presented a chance to introduce these interactions to directly affected parties, discussing them and their impact in greater depth, also debating the underlying pros and cons of performative, extravagant, expressive technologies.

As we followed an exploratory based approach, our research here is not presented in chronological order. Rather, it is presented as groups of ideas and concepts that follow on from one another. For a true roadmap of project developments, please refer to Appendix B.

CHAPTER FOUR

Auditory Performance

As discussed in our methodology (see Section 3.2), in this research, we aim to explore audio and visual effects as a baseline. In this chapter, we begin by exploring the use of audio to create more performative, extravagant, expressive place-based experiences. We started by investigating the value of audio as en engaging technology, appearing to emanate from the environment. This was simulated using headphones and spatial audio techniques. We then went on to explore a range of performative interaction techniques in a public space, placing interactive loudspeakers in the surrounding environment. We go on to discuss the implications of using each of our designed interaction techniques in a public scenario. Considering these findings, further tests were conducted through a public, naturalistic deployment, where visitors were able to turn a number of loudspeakers on and off at will using their own mobile device. From all of these experiments, we are able to draw a number of important design implications for performative auditory experiences in public spaces.

4.1 Introduction

When listening to digital audio in public spaces—whether enjoying music or an audio book whilst sitting on the train or going for a run—we do so predominantly, privately, using headphones. These kinds of experiences are perhaps personal to us, and things that we would not like to share or be made public. However, there are other contexts where this may not be the case – where digital audio is of collective interest to groups of people in the same space.

In public contexts such as visitor centres and outdoor attractions, headphones are widely used [31, 124], offering a collection of curated audio tours and soundscapes. We argue that in these kinds of contexts, visitors could perhaps benefit from more shared, collaborative experiences. Shared audio experiences in such places give visitors time to reflect and start discussions, ultimately enhancing the experience through promoting a more active engagement – making the visit social, as most are intended to be [9]. One attempt at collaborative audio experiences led Heller et al. to designing *Corona* [67], an interactive museum soundscape where audio feedback was synchronised between devices. Similar approaches have also been used outdoors, with large soundscapes that users can openly explore and experience together [109, 135]. Fosh et al. have tried to foster a more social audio experience in a different way, through designing specific trajectories that encourage discussion and reflection [54], and gifting experiences to one another through personalised interpretations [53].

Turning now to the underlying motivation for many audio tours and soundscapes – audio augmented reality is *"an attempt to combine our real world interactions with the richness of computational information without isolating people from each other"* [9], though still, in the previous research we have seen, researchers remain fully aware of the isolating effect of headphones [54, 67]. Perhaps, then, headphones are not the best way of facilitating a more social experience. A different approach that focuses not on the synchronisation of audio or design of trajectories, is the use of loudspeakers for public, amplified feedback. With loudspeakers, the sound is further reaching, and visitors are not forced to distance themselves from each other by wearing headphones. An additional benefit of this approach is that people who do not have a mobile device or headphones can still be a part of the experience, just through standing nearby. There are a small number of examples of mobile location-based experiences that have used loudspeakers before [4, 119]. Rogers et al. report on the design and evaluation of *Ambient Wood* [119], experimenting with new forms of digital augmentation for learning outdoors. One particular piece of technology in this experience involved situated, wireless loudspeakers playing environmental sounds – something participants of the experience were not able to directly control. A similar learning experience—*Explore!* [4]—allowed groups of users to walk around a heritage site, experiencing sounds through a loudspeaker placed in one of the user's backpacks. Different audio snippets were triggered depending on the GPS location of the tracking device.

4.1.1 Extravagant, expressive auditory place-based experiences

Although loudspeakers have been used in mobile location-based experiences before, in this research, we are specifically interested in experiences where users are given the opportunity to directly trigger and manipulate the public, amplified audio. Traditionally, in audio mobile location-based services, a user is only ever given control over the playback on their own device. Here, using our performative, extravagant, expressive interaction techniques, we encourage users to publicly manipulate the audio, directly affecting bystander's experiences. The motivation for directly affecting other people's experiences is to create a more social experience – to create a sense of curiosity, evoking emotions and sparking conversations.

To ensure that the experience remains expressive, the audio must be amplified and public. We do not suggest a particular size of loudspeaker for creating amplified audio, though the output must be audible to those in the local vicinity. Any kind of speaker, from an inbuilt mobile device speaker to a high powered situated loudspeaker, can be considered more amplified than headphone feedback. To complete the expressive experience, the manipulations or control mechanisms must also be amplified. As discussed previously, one of the most amplified kinds of manipulations that can occur on a mobile device is a gesture.

Therefore, we imagine a performative, extravagant, expressive auditory place-based experience as one which makes use of large, high powered embedded loudspeakers, and can be controlled by using gestures on a mobile device.

4.2 Exploring the value of spatial audio

Before embarking on a substantial development and evaluation project—investing large amounts of time and money in creating and testing a performative audio experience with loudspeakers—we first decided to explore the notion of spatial audio, and its value in large, unfamiliar public contexts.

Spatial audio is a fairly well established area in the field of HCI. Much of the original work to do with spatial audio with mobile devices focused on GPS navigation by panning sounds in the left and right ear, also using a range of auditory icons to differentiate certain instructions [69, 78, 136]. A number of different approaches have also emerged, including a Geiger counter metaphor where beep frequency increase as the user walks closer to a point of interest [94]. In recent years, vast improvements in mobile computational power have meant that these devices can now also render much more complex three-dimensional scenes. By combining an embedded compass with head related transfer functions (HRTF), developers have been able to create soundscapes that rotate around the user, not only for navigation purposes, but also for exploratory, entertainment purposes [66, 135]. Vazquez-Alvarez et al. [147] provide a useful comparison of these existing rendering techniques—among others—discussing the benefits and limitations of each.

Although our end goal is to harness embedded loudspeakers in the environment, producing a range of amplified sounds, our initial experiments here are an attempt to explore the fundamental concept and value of audio emanating from points of interest in the environment. Similar to some of the aforementioned examples, we attempt to simulate spatial audio using a mobile prototype that produces a binaural representation of spatial audio using the HRTF abilities packaged with the OpenAL framework. By using the embedded compass in the mobile device to estimate which direction the user faces [68], we are able to produce a live and fairly accurate representation of the soundscape to the user. However, unlike conventional spatial audio location-based services that trigger multiple sounds as the user walks around, this prototype focuses on passing sequential suggestions to the user, avoiding increased cognitive load [146]. We also use this as a technique to simplify selection in an audio display, and to promote serendipitous discovery – an important aspect of way finding in unfamiliar places [25].

This section reports on the design, development and evaluation of *Audvert*, a prototype that attempts to provide users with a sense of a place using only audio. *Audvert* is an eyes-free, heads-up audio display, allowing the user to look around, reinforcing the sounds they hear against what they see in the environment. We first conduct a controlled study to prove that our prototype is an accurate simulation of spatial audio originating from the environment, and that users are able to infer direction fairly accurately. This controlled study also provided an opportunity to measure the effectiveness of different kinds of audio delivery and refine our prototype. We then gave the system to users in an exploratory usability test, allowing participants to use the system freely. We report on our findings, discussing the value of spatial audio in an unfamiliar public space.

4.2.1 Prototype design

Audvert is a prototype system designed to give users a better sense of the place that they are in using only spatial audio. To use the system, a pair of headphones and the *Audvert* application is required. When using *Audvert*, sounds appear to originate from their actual physical locations. A user should be able to infer proximity through amplitude—closer is louder—and direction through the panning of sounds. Spatial audio feedback in the system is constantly changing due to lightweight interactions such as the user walking around (changing their location in the system) or the user changing direction (changing their listening orientation). *Audvert* uses a digital embedded compass, so as the user rotates, sounds continue to appear as though originating from the same physical location in space. With regards to location tracking, we decided upon a Wizard of Oz study methodology for *Audvert*. Users are not told that their location is entered manually elsewhere, so it is fair to say that users believe that their location is being automatically tracked. We use a separate purpose built system where a researcher may tap a map to update the location of a user. This Wizard of Oz approach was chosen as it offers a high accuracy with little complication. GPS was not a viable option as the system was used indoors. In reality, this kind of system would most likely work indoors with WiFi localisation.

When considering the user interaction with the prototype system, we decided that it should facilitate three separate functions, all relative to tourist way finding [25]: to aid serendipitous discovery of points of interest (wandering), to deliver useful information about them, and to offer users the ability to navigate towards them if they wished to do so (homing). To facilitate serendipitous discovery, *Audvert* randomly chooses points of interest in the local environment, playing them back one at a time to the user. Useful information about the points of interest is spoken in the audio recordings (sales, offers and general information about the shop). Once a user has selected a point of interest, *Audvert* facilitates navigation through continually playing spatial audio for that point of interest.

The following figure—Figure 4.1—and scenario depicts a user using *Audvert*.

Scenario. *Fred is at a large shopping mall but is unsure about which shops to visit. Putting on his headphones and opening the Audvert application, he holds his phone out in front of him (1). The Audvert app randomly picks a store from the mall and begins playing an audio clip containing the name and description of the goods and services available there. Fred can hear this information on his right side indicating that this particular store is in that direction (2). Deciding he is not interested in this shop, Fred waits a few seconds for the next random selection to be played. After listening to the name and short description of the second store, which is now being played in his left ear, Fred decides he wants to know more and signals this via a shake gesture with his phone (3). Fred now hears additional information and uses the direction the sound is coming from to guide him to the store (4). As he approaches the store, the amplitude increases.*

Figure 4.1: User interaction in *Audvert* (explained in scenario). The user hears a collection of random shops spatially rendered. The user selects a certain shop and heads towards the sound source.

4.2.2 Controlled pilot study

The first study was conducted to test the spatial audio rendering techniques employed by the system, identifying if users of the system were able to comprehend the spatial audio feedback given by the system. As a comparison, we used two different types of audio clips to determine which type allowed users to more easily understand the spatial audio feedback. One system played short intermittent audio clips and the other played long continuous audio clips. We define short intermittent as audio clips that do not have continuous background sound and are short in length (i.e., purely spoken word, 5-6 seconds long). Long continuous clips contain continuous background sound and are longer in length (i.e., spoken word with a musical backing track, 25 seconds long). Our goal in this preliminary evaluation was to determine which of these two types of audio provided users with the most accurate path to the audio source location. We chose these two test cases, as long continuos audio closely resembles a television or radio advertisement, where as short intermittent audio resembles receiving a prompt, useful snippet of information, similar to when asking a stranger for directions.

This study was conducted in a shopping mall. Participants of the study were members of the public that resided within the shopping mall. In total 24 participants were recruited (Male: 18, Female: 6), with an average age of 24 (Min: 16, Max: 37). A between groups method was employed for the study, where half of the participants used the system with short intermittent audio, and the half used the system with long continuous audio. There was a similar age and gender distribution between each system. No incentive was offered for participants who took part.

Task

So that every participant had a clear idea of how spatial audio worked, a training exercise was developed. This training exercise was a collection of four animal noises that the participant could listen to. Participants were told the locations of animal sounds on the map so that they could understand the link between the source placement and audio feedback experienced (see Fig. 4.2). The placements

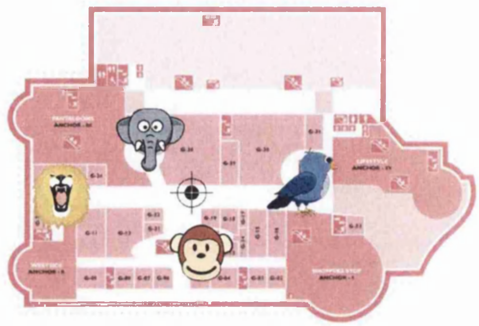

Figure 4.2: Animal locations laid on top of the mall map for the training system. The target is the location of the user.

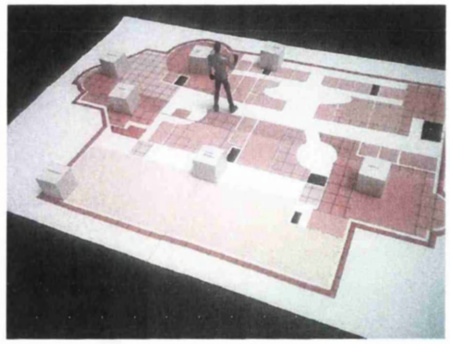

Figure 4.3: Example cube placements on the mall map. Pink regions signify the existence of a shop at that location.

were in front, behind, to the left of and the to right of the participant. All animal sound sources were varying distances from the user.

For the actual study session, five different sound sources were used. Each sound source signified the existence of a shop at a certain point in physical/virtual space. Each sound source was played one after another with a small pause in-between. For the short intermittent audio system, the name of the shop was played and a very short description of the goods or services offered at that shop (~5 seconds). For the long continuous audio system, the name of the shop was similarly played, but this was then followed by a long description of the shop's goods and services with a backing music track (~25 seconds). As well as having a longer duration, the audio contained music and was continuous – unbroken.

During the study, participants were asked to place cubes on a shopping mall map to signify where they thought the sound of each shop originated from (see Fig. 4.3). To ensure that this study was as fair as possible, two additional conditions were added to the study. The first condition was that all shops included in the placement task should not be in the participants direct line of sight. This would prevent participants from glancing at the names of nearby shops. The second condition was that participants should not have prior knowledge of the location of any of the shops used in the placement task. This too, would give users an unfair advantage in placement tasks. To combat this second issue, all

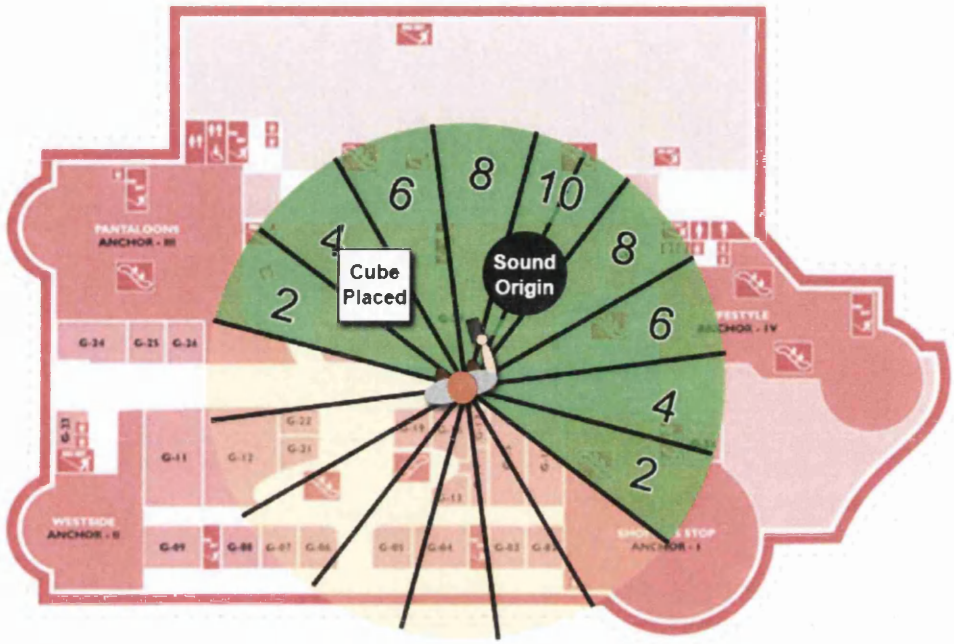

Figure 4.4: Segmented scoring chart – to take a measurement, the wheel is centred on the user location, then rotated until the dotted line goes through the actual sound origin. The placement is scored on which segment the cube resides in. In this particular example, the participant would score 4/10 for the placement.

participants were asked if they knew the location of the five shops used in the study before taking part. If a prospective participant did know the locations of the five shops, they were not able to take part in the task.

Data Collection

We measured the physical placement of each cube on the shopping mall map. The map was A3 in size and each cube was 2cm^3. Direction scores were calculated depending on the cube placement angle of error. For the direction measure of each cube, a moveable circular segment chart was placed over the map and scores from 10 to 0 were given for $22.5°$ segments (see Fig. 4.4). Proximity scores where calculated depending on the cube placement error in centimetres (distance from user). Observations and informal feedback were also recorded.

Findings

Overall, the system proved successful in allowing participants to determine the general direction and proximity of sounds. However, participants sometimes got confused between whether audio was in front of or behind them. Unfortunately, this is a side effect of using off-the-shelf head related transfer functions that have not been trained to specific users. This limitation would not be present when using loudspeakers situated in the environment.

	Average	
	Distance Error (cm)	**Direction Score (/50)**
Short Intermittent	5.7 (σ 4.4)	21.6 (σ 8.2)
Long Continuos	3.6 (σ 2.7)	33.0 (σ 6.5)

Table 4.1: Average distance errors and directional scores for short intermittent and long continuos audio.

The results of our evaluation (see Table 4.1) showed that on average, participants using long continuous audio performed better in terms of perceiving both proximity ($p < 0.01$; Unpaired t-test; 2.1cm average difference) and direction ($p < 0.01$; Mann-Whitney; 11.4/50 average score difference). We therefore decided that short intermittent audio should be used as a brief alert of a shops existence (when a shop is unselected), and long continuous audio should be used to give a person more information about a shop and allow a user to navigate towards it if they wish (when a shop is selected).

4.2.3 Experiment: exploring the value of spatial audio in unfamiliar public spaces

After demonstrating that our prototype was a suitable representation of spatial audio, our second study focused on gaining a better understanding of the usability and user experience offered by interactive sounds emerging from the environment. We wanted to learn how participants felt about these sounds appearing to come from physical locations in a public space, and also discover if participants could use these sounds to navigate towards shops they had not visited before.

We explored these points by deploying *Audvert* in the same shopping mall, allowing mall visitors to trial our system while shopping there.

Participants

Participants of this study were again, visitors to the shopping mall. We recruited 19 participants (10M, 9F, 17-47 years), none of which had taken part in the controlled pilot study. Approximately 3 were regular visitors to the mall. All participants could speak and understand English. No participant reported hearing problems. All participants had used a mobile phone before, but not all had used a smartphone.

Procedure

Participants were recruited individually as they walked through the shopping mall and were given a short introduction to the *Audvert* application. For the task, the only prompt participants were given was to explore the shopping mall using the system. As participants moved around the shopping mall, a researcher stood around 20 metres away and updated the participant location every few seconds. This was achieved using the prior mentioned 'Wizard of Oz' technique, with a researcher manually entering the participant location on a second device. While exploring, participants were free to interact with the system as they wished. At the end of each study session, the participant was asked a series of questions in a semi-structured interview regarding their experience with the system. Again, no incentive was given for completing the study.

Data Collection

A log of each participant's movements and interactions was recorded every second by the mobile device. Qualitative data on the user experience with the system was gathered through a post-study interview. Questions asked in this interview were:-

Q1 Has the system given you a better or worse understanding of the shop locations?

Q2 How would you describe your sense of place after using the system? (Level of comfort and familiarity with the place)

Q3 Did you hear any shops that you have not visited before?

Q4 *(If Q3 yes)* Would you consider visiting one of them now?

Q5 How would you sum up your experience of using this system?

4.2.4 Findings

The following section details the responses given to the questions asked in the post-study interview. One participant did not complete the entire study, leaving midway through the interview. Some participants were also undecided, and did not provide answers to all questions. Questions that participants did answer have been included.

	Response		
	Better	**Unaffected**	**Worse**
Q1	79% (15)	5% (1)	16% (3)
Q2	80% (12)	7% (1)	13% (2)
	Yes	**Maybe**	**No**
Q4	63% (10)	31% (5)	6% (1)

Table 4.2: Quantitative results from questions asked in the study.

Informal Observations & Feedback

It became apparent in the study that not every participant was selecting a shop by shaking the device. By avoiding this function, the participant was not able to gain extra information about a shop (a long description) – the system however, is still usable without this function. As *Audvert* provides no visual feedback of where the user is in the system, one participant was unsure whether a shop was selected or unselected at any point in time. This participant carried on for a further two minutes and after this period appeared to be much more comfortable with the system. One participant that successfully used the shake gesture to activate a shop was seen attempting to navigate towards the selected shop. This participant attempted to follow the direction of sounds, and did so successfully by finding the shop. The participant mentioned that they knew they were going in the correct direction because they could hear the sounds getting louder as they got closer to the shop. This person said that they had never visited the shop before, and that they had no prior knowledge of its location. There was also another situation where a participant used the varying volume of sounds to navigate towards a shop. This participant had selected a shop, walked in one direction and then realised that the volume was decreasing – inferring that they were heading in the wrong direction. The participant then turned around to face the correct direction and successfully found the shop that they had selected.

4.2.5 Discussion

Most participants enjoyed the experience, claiming that the system appeared to be 'very accurate' and that the direction sounds came from was the same direction that they could see the shops. A large portion—79% (15)—of participants said that they had a better understanding of the location of shops after using the system. One participant who claimed a worse knowledge of shop locations said that it was *"difficult to differentiate between the distances of each shop."* Another participant agreed, finding it *"difficult to tell the distance to a shop"*, but claiming the *"direction was good."* Although 79% of participants said that their understanding of shop locations had improved, a third of these (5 participants) claimed that they could either infer direction or distance, but not both. In most of these 5 cases, participants could only infer direction (4 participants). This is an interesting result, as without knowing both, one would imagine it being difficult to pinpoint an exact location. These participants obviously believed that knowing only direction or proximity was sufficient for knowing the location of a shop. A similarly high percentage—80% (12)—claimed that they had a better sense of place after using *Audvert*. There were 5 participants who did not answer this question, unable to differentiate between knowing the locations of shops and a sense of place. One participant said that the system gives a *"good sense of the area. I could tell what kinds of shops and products were available."* Another participant argued that the system gave a bad sense of place as there were *"lots of shops and I only heard a few."* Although a similar percentage said they had a better sense of place and an improved understanding of shop locations, not every participant who had an improved under-standing of shop locations had a better sense of place, showing that the two were not regarded as the same question by participants.

80% of participants heard a shop they had not visited before, of which, 63% (10) said they would definitely consider visiting that shop and 31% (5) said they might visit now. Only one participant decided that they would not visit a new shop they had discovered, saying that they are the *"kind of person that comes to the mall for one thing."* There does appear then, to be a compelling change in behaviour, where the system seems to encourage participants to visit new, serendipitous discoveries that are made. Generally, most participants explained that they would like to visit these shops because of sales and discounts mentioned in the audio recordings. Participants who said they might visit a new shop often said that they would do so if more detailed information about a shop was available. Judging by this feedback, it seems that the content is an integral factor of this change in behaviour. The majority of participants liked the shake gesture. Two participants explicitly mentioned that they would *"prefer the shake over a button."* In contrast, two participants said that they were not comfortable with the shake gesture, preferring a button. One participants concern came from not wanting to shake their expensive mobile phone. To satisfy all users, a multimodal approach could be introduced, where the user chooses whether the interaction occurs onscreen or through a gesture.

We have introduced *Audvert* to illustrate the potential of a lightweight system that facilitates navigation and serendipitous discovery of large indoor spaces using spatial audio. While the idea of an audio display for navigation and discovery is not a new one, we argue that the novelty of our approach lies within its indoor application and the feedback and selection techniques employed. We use amplitude and directional feedback to manipulate useful, spoken audio information about a point of interest in real time. *Audvert* has also introduced a simplified selection technique for spatial displays through using a single active element and a simple shake gesture to select.

Our user evaluations have shown that *Audvert* was an accurate representation of spatial audio feedback, and that it is capable of helping users navigate towards and discover new points of interest. We saw multiple examples of users successfully arriving at a shop and for the majority of participants, we witnessed a change in behaviour with participants claiming that making new, serendipitous discoveries made them want to go and visit there.

4.3 Exploring performative audio interactions

After conducting our initial experiment and the results presenting a compelling argument for spatial audio in unfamiliar contexts, we went ahead with our exploration of more performative place-based experiences using audio. To understand the true, comparative effect of extravagant, expressive audio interactions in public spaces, we developed a number of different systems. These ranged from highly expressive, to secretive and suspenseful interactions similar to *Audvert*.

In the following section, we discuss the design rationale behind each of our proposed interaction techniques, and introduce a deployment we conducted within the grounds of Swansea University. A within-subjects approach was followed, where each participant was asked to use all system variations. From this deployment, we discuss our findings, outlining the key benefits and drawbacks that each interaction technique presents.

4.3.1 Prototype designs

Surround You is a system that attempts to look at spatial audio in a new way. There are three main components to the system, two of which are clients and one of which is the main web server that handles all requests and carries out all of the computational work. The first client is the mobile application that the user sees and uses. This is the controller that is used to actuate each sound point. The second kind of client is the individual sound points. For our experiments, we used netbooks as clients. To each of these sound points, we connected a portable

Figure 4.5: A sound point – a netbook connected to a battery powered loudspeaker. The netbook communicates with the central server via the Internet, returning the correct amplitude to play the sound through the loudspeaker.

battery powered loudspeaker so that the sound would be audible from a distance (see Fig. 4.5). A web server was then used to handle communication across client devices.

The main difference between *Surround You* and other embedded speaker systems is that the user has complete control of the output using a mobile device. Also, it is developed in such a way that n speakers may be connected to the system, working as independent audio streams and m users may also connect to the system with the ability to control any of these audio streams simultaneously.

Interaction

When designing *Surround You*, we decided that the user manipulations should be amplified and that the effects should also be amplified. This way, in terms of the performance from a spectator's perspective [108], the interaction would be an expressive one. By making an expressive interaction, we hoped that this would be the most engaging experience for bystanders and would encourage them to become involved in the experience.

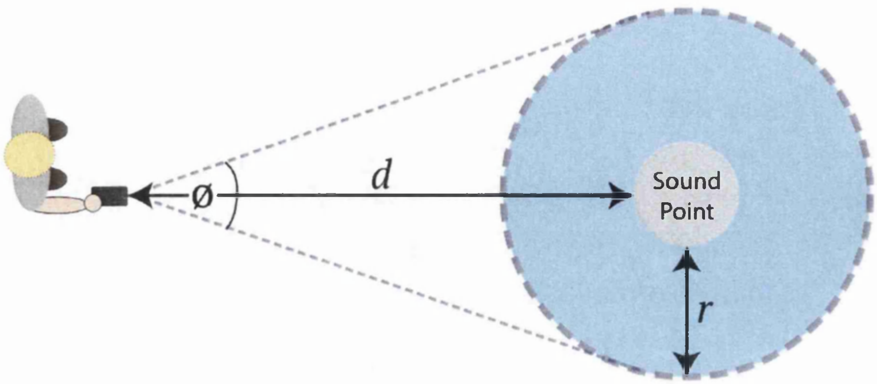

Figure 4.6: A user can activate a sound point by pointing directly at it. The valid selection range (θ) is based on the distance (d) between the user and sound point, and also the exaggerated activation radius (r) around the sound point.

To activate a sound point, the user must point the mobile device in the direction of a sound point (see Fig. 4.6). To compensate for GPS inaccuracies and the relatively small size of the electronic sound point equipment, each sound point is programmed to think that it takes up much more physical space than it actually does. For example, if one of these sound points were to be placed inside an old building at a heritage site, we would assign the diameter of the building to the sound point. If the user then pointed at any point of that building, the sound point would activate and begin playing. It would not be feasible to have sound points that have a valid selection range of a few degrees, as it would be almost impossible for users to select them.

In *Surround You*, each speaker is pre-programmed with its location, making it location-aware. The smartphones we used were GPS enabled, allowing continual calculation of the bearing from the user to the sound point, which would be considered as the valid selection bearing to 1°. Previous research shows that a valid selection range should be at least 20° to achieve 90% accuracy [91]. The web server actively checks the location and orientation of every user in the system, calculating valid playing ranges and using this information to determine whether any of the sound points are currently being pointed at (see Fig. 4.6).

When a user interacts with a sound point by pointing the device, we use a tuning-in and tuning-out metaphor. When the user points directly at a sound point, the sound point is at maximum amplitude. The further away the user points from the centre of the sound point, the more the amplitude decreases. It does this until the user is no longer pointing at the sound point, at which point the amplitude level returns to zero. We used this metaphor so that the audio would gradually fade in and not startle spectators or those standing close to the loudspeakers. In the current implementation of *Surround You*, if more than one user points at the same sound point, the loudspeaker takes the mean amplitude and plays at that volume. The framework is built in such a way that we can quickly and easily change the effect that multiple pointing users have on the system. Possible effects that multiple users have could include increased amplitude or access to different audio streams that are not available to single users.

Control systems

In comparison to *Surround You*, a standard style audio guide system would use headphones instead of loudspeakers and would require the user to enter a unique number at each of the sound points using a number pad. As a control system, we recreated this standard style audio guide to see how our system fared against existing, conventional audio guide style systems to gain information from surroundings. In terms of manipulations and effects, both are more hidden in this interaction, creating a secretive experience [108] from the point of view of the spectator.

To select a sound point using the number entry manipulation, a user must enter the unique number assigned to that sound point. With the addition of headphones, the loudspeaker feedback is substituted for headphone feedback that only the user can hear.

Two further systems were created with variations of the different manipulations and effects mentioned. The first of these systems uses the pointing manipulation and gives feedback to the user through the headphones, creating a suspenseful experience [108]. Our last variation uses the number entry manipulation and the effect is heard through the loudspeakers positioned in the environment,

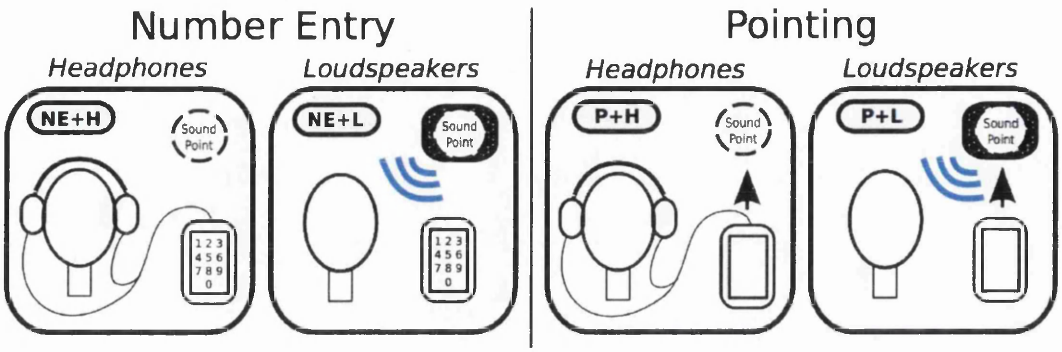

Figure 4.7: The four different interaction techniques used for evaluation. From left to right: Secretive (NE+H : *Number Entry & Headphones*), Magical (NE+L : *Number Entry & Loudspeakers*), Suspenseful (P+H : *Pointing & Headphones*) & Expressive (P+L : *Pointing & Loudspeakers*)

providing a magical experience [108] for the spectator. The four systems can be seen in Figure 4.7.

4.3.2 Exploratory pilot study

Before embarking on a full study of *Surround You*, we conducted an exploratory pilot study that involved the deployment of the system in a public setting - Swansea University campus. The aim of this pilot study was to test a number of varying prototypes and study methods. We recruited 8 participants (5M, 3F, 21-40 years) from the University, all of whom where staff members of different departments. Of the 8 study participants, all had used a smartphone before, but only 7 actually owned one.

We set up 5 different sound points in close proximity on the University campus and assigned contextual sounds to them that one would not expect to hear in that context. For each of the sound points, we used sounds such as the rainforest, cackling geese and monkeys. We used these distinctive sounds so that users would know they were part of the experiment.

Each participant carried out the experiment individually. Participants were given very few instructions, only being told that there were 5 sound points located in the local environment and that they had to discover them. After discovering all 5 sound points, the participant was asked to return to the researcher for questioning. For each of the participants, we tried slightly different methods of evaluation, including questionnaires, interviews and asking participants to attempt sound localisation, marking a map with the location the participant thought each sound was originating from. In some cases, participants were also given different kinds of output to test, such as headphones or the mobile device's speaker as opposed to the loudspeakers used in *Surround You*. Pointing was the only selection technique used in this study.

The study helped shape the full experiment (see Section 4.3.3) by providing a better grounding for the future procedure.

Findings

Using loudspeakers in the environment produced some interesting behaviour from bystanders. Most would turn round to look what was happening, some would stop and watch and others would go as far as approaching the participant and asking them *"what are you doing?"* or *"how are you doing that?"*

Some participants used the system to directly affect spectators, with one participant saying, *"One woman jumped when she heard the monkeys. It was funny."* This was a recurring theme throughout the study, where participants mentioned that spectators often appeared shocked or surprised.

These implications seem to suggest that expressive audio interactions do indeed gain the attention of bystanders, and that the participants were happy to assume this role, playing an active part in directly affecting bystanders and spectators. In the full study, we will attempt to record the effect that the performance has on bystanders.

In this pilot, it was also decided that the study location was overcrowded with five individual sound sources. For this reason, sound points in the full study will be fewer, and will be more uniformly spaced.

4.3.3 Experiment: exploring performative audio interactions

We conducted a study over eight days on the grounds of Swansea University, Wales. During exploratory testing, we discovered that the behaviour of both performer and spectator were interesting, so our study focused on interpreting the behaviour of both participants and bystanders. As we did not question spectators in this study, particular emphasis was placed on finding out the performers enjoyment, comfort and the perceived effect of their performance on bystanders between systems.

Participants

For this study, all participants who were recruited were affiliated with Swansea University. A total of 16 participants (11M, 5F, 18-50 years) took part in the study. This included 2 members of staff and 14 students (6 postgraduates and 8 undergraduates). All participants except one had used a smartphone before, with 13 of the 16 participants owning one. All 16 participants reported having no hearing difficulties prior to taking part in the study.

Before each study session, participants were also asked a set of questions regarding any group activities they regularly partake in (i.e. Team sports, singing, dancing, lecturing etc.) and whether they saw themselves as an introvert, extrovert or ambivert. A short dictionary definition was given to users to help make their choice of personality trait. These questions were asked so that we could gain a brief understanding of how used to performing our participants were. Upon self assessment, 4 participants claimed to be introverts, 1 extrovert and 10 ambiverts. There was one participant who did not wish to answer this question. In terms of group activities, 6 participants reported not partaking in any regular group activities, with the remaining 10 taking part in one or more regular group activities. By observing this data, we concluded that the majority of participants were fairly comfortable performing in front of others, with the remaining quarter of participants having little experience of performing around others.

Procedure

The study was spread over 8 days, with each session lasting for around 30-45 minutes. For most of these days the weather was relatively cold, cloudy and windy. There was only one study session where it began to rain lightly, but the participant decided to carry on with the study. Although the Swansea University campus is not a visitor attraction as such, it was again chosen for the main study as it is a busy public location that provided a good network infrastructure upon which to operate our systems.

When each participant arrived outside at the study location, they were asked to read and agree to the ethically approved research consent form and were given a demographic questionnaire to complete. After completing the short questionnaire, the participant was given a short briefing regarding the capabilities of system, including the different selection (number entry & pointing) and feedback (headphones & loudspeaker) techniques. The participant was also shown the physical location of the four sound points used in the study (see Fig. 4.8).

A within-subjects design was used, where each participant was required to use all four of the different systems (see Fig. 4.7). The order systems were used in was counterbalanced to reduce the effects of presentation order on results. For each condition, the participant was first given instructions as how to use the prototype. For the pointing selection, participants were told that all they were required to do was point at the individual sound points to activate them. For the number entry selection, participants were told that they were required to visit the number displayed at each sound point and enter it into the mobile device. To avoid participants learning the numbers for each sound point and not having the visit them again, two random double digit numbers were displayed on each of the sound points. Each sound point in our experiment consisted of a numbered label and a netbook with a portable speaker (see Fig. 4.9).

Once participants had been introduced to the first interaction technique, they were asked to explore the area for 2-4 minutes, returning to the researcher once they had activated and listened to all four sound points. During this time, the researcher took notes of participant and bystander behaviour. When participants returned to the researcher, they were then given a system with the same manipu-

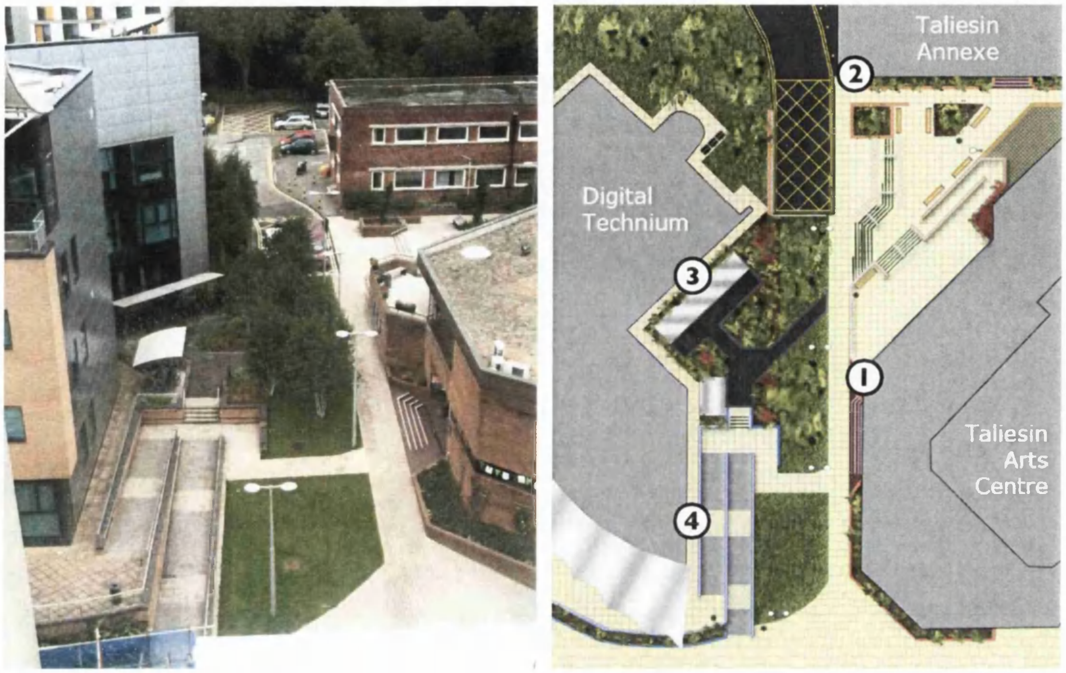

Figure 4.8: *Left:* The location the study was conducted in. *Right:* The locations of each of the four sound points used in the study.

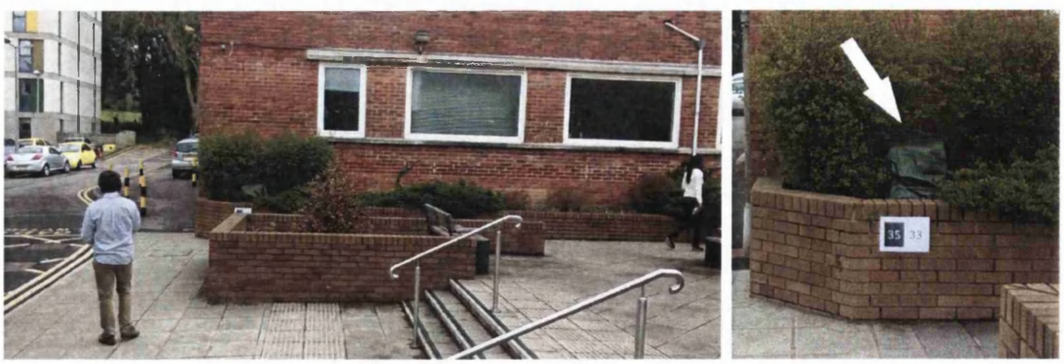

Figure 4.9: *Left:* A participant pointing at a sound point. *Right:* A close up of the sound point - a portable speaker and netbook protected by a waterproof cover. Also, a numbered label for the number entry systems.

Session 1 & 9	Session 2 &10	Session 3 & 11	Session 4 & 12
NE+H	NE+L	NE+L	NE+H
NE+L	NE+H	NE+H	NE+L
P+H	P+L	P+H	P+L
P+L	P+H	P+L	P+H
Session 5 & 13	**Session 6 & 14**	**Session 7 & 15**	**Session 8 & 16**
P+H	P+H	P+L	P+L
P+L	P+L	P+H	P+H
NE+H	NE+L	NE+L	NE+H
NE+L	NE+H	NE+H	NE+L

Table 4.3: Order of using systems in each study session. NE+H - *Number Entry & Headphones*, NE+L - *Number Entry & Loudspeakers*, P+H - *Pointing & Headphones*, P+L - *Pointing & Loudspeakers*

lation technique (pointing or number entry), but now with either headphones or loudspeaker (depending on which one had not been used). As before, participants were asked to explore the four sound points for a further 2-4 minutes. After using this second system, participants returned to the researcher and were asked to answer some comparative questions regarding the two systems in a questionnaire.

Participants were then asked to repeat the experiment, but using the new input technique (pointing or number entry) with the two different outputs again (headphones & loudspeakers). Once participants had used all four systems and answered the same comparative questionnaire again, participants were then interviewed with questions regarding their entire experience. Participants were thanked for taking part and given a voucher once they had completed the post-study interview.

Table 4.3 depicts the order in which the systems were used in each study session.

Data collection

Data was collected in the form of a pre-study demographic questionnaire, two mid-study questionnaires given to the participant after using both variants of each input system (number entry & pointing), informal observations and a post-study interview.

The pre-study demographic questionnaire focused on gaining background information from a participant such as their experience of mobile phone use, whether they have a hearing impairment, whether they see themselves as an introvert, extrovert or ambivert and if they regularly partake in group activities. These last questions were asked to gain a sense of the personalities of our participants, and whether they were used to performing. The two questionnaires given to participants during the study, required participants to rate their enjoyment and comfort level on a 7-point Likert scale after using each set of systems. One questionnaire was given to the user after using the first two systems with the same selection method (e.g., pointing). The same questionnaire was then given when the user had used the final two systems with the other selection technique (e.g., number entry). At the end of both of these mid-study questionnaires, the participant was asked by the researcher, "What did you think of these two systems?" This question was deliberately left open and allowed participants to compare the systems or give an account of their experience however they pleased. During the study, informal observations were made of both participant and bystander behaviour. In terms of participants, observations of general behaviour when using each system were recorded in note form, along with any incidents that the researcher deemed interesting. For bystander observations, the number of people in the local vicinity was recorded, along with the number of those that stopped to look or reacted to the system in any way. The investigator kept count of bystanders and spectators in a tally chart, making the process quick and simple. In some cases, detailed accounts of bystander behaviour was also recorded. The post-study interview gave participants an opportunity to reflect on their entire experience.

4.3.4 Findings

The following section includes the results from the study, including the two questionnaires given to the participant after each interaction technique, the post-study interview and observations made of participants and spectators during each study session.

	Number Entry		Pointing	
	H	**L**	**H**	**L**
Comfort (1-7)	6.44 (σ 0.70)	5.31 (σ 1.31)	5.69 (σ 1.26)	4.44 (σ 1.32)
Enjoyment (1-7)	5.81 (σ 1.01)	5.63 (σ 0.85)	5.38 (σ 1.05)	4.56 (σ 1.27)

Table 4.4: Average score for *comfort* and *enjoyment* level for each system (1-7 Likert-type scale; 7 being high).

Questionnaire between interaction techniques

Participants were asked about their comfort and enjoyment levels when using each of the interaction techniques (see Table 4.4).

Pointing and loudspeakers was rated lower than all other interaction types for enjoyment, showing statistical significance against each of the interaction techniques ($p < 0.05$; Wilcoxon signed-rank test). It was also rated as the most uncomfortable, which also showed to be statistically significant against all other interaction techniques ($p < 0.05$; Wilcoxon signed-rank test). Although number entry and headphones was ranked both the most comfortable and the most enjoyable interaction technique, it did not register significantly higher than all of the others. A few participants mentioned that they were more comfortable with using loudspeakers the second time around. As well as the increased control and accuracy of number entry over pointing, some participants also mentioned that they liked the idea of being able to activate a sound point and walk away.

Post-study interview

Most participants preferred number entry as a selection technique as opposed to pointing. Participants that chose number entry said that they did so because it worked well and gave them greater control over what was playing. The general consensus among all participants was that pointing was convenient and fun, but again number entry gave more control. Regardless of selection techniques, most participants preferred using headphones as they provided a more personal and less invasive experience, along with the ability to manually control the mobile device volume. Participants that preferred the loudspeakers did so because it was easier to differentiate between audio streams, and they liked the idea of spectators being able to link their pointing gesture to the loudspeaker output.

The vast majority of participants decided that the number entry and headphones interaction offered the most solitary experience. When asked to comment on their views on a solitary experience in this context, participants' replies were universally positive, with one participant saying, *"I was more comfortable, wasn't worried about others and could concentrate on the info."* One or more of these individual aspects were mentioned by most participants. When asked about number entry and headphones as individual factors, the view of most was encapsulated by a single participant who said, *"people didn't seem to take much notice, it's just normal behaviour."* At the other end of the spectrum, most participants thought that the pointing and loudspeaker interaction offered the most sociable experience, with one participant explaining, *"they can see the interaction and listen."* Participants commented that *"some [spectators] were looking for where sound was coming from"* and *"I think someone stopped their conversation to look and see what was going on."* A few participants however, raised concern with the loudspeaker feedback, saying, *"the speakers got people's attention, but it depends on the context if it's ok"* and *"if they want to listen then fine. Otherwise, not so much."* The majority of participants did not notice a change in spectator behaviour when pointing, although some participants said that *"a couple of people noticed"* and *"some gave strange looks."*

When asked, five people said that they had purposely tried to affect the actions of others or gain reactions from spectators. Of these five, some tried harder than others, claiming that they purposely tried to scare people that were walking by. An example of this was when one participant said, *"towards the end, I waited until a guy was in front of it [the speaker]. It didn't scare him, but he jumped a bit."* These 5 participants were all seen to be playing around with spectators, trying to make them look for where the sounds were coming from. The other 11 participants claimed that they tried to be as inconspicuous as possible. One participant said, *"I tried to stop the sounds as soon as possible. I Tried to avoid interrupting anyone."*

Participants voted the two loudspeaker systems as the most performative ones. Opinion was divided between whether a system being performative in a

	Number Entry		Pointing	
	H	**L**	**H**	**L**
Most Solitary	15 (94%)	0	1 (6%)	0
Most Sociable	0	4 (25%)	0	12 (75%)
Favourite Interaction	10 (63%)	3 (19%)	1 (6%)	2 (12%)
Least Favourite Interaction	0	1 (6%)	4 (25%)	11 (69%)
Most Performative	0	5 (31%)	1 (6%)	10 (63%)

Table 4.5: Quantitative results for questions asked in the post-study interview.

public place was actually a good or a bad thing. One participant explained, *"I think it could be fun. I don't get embarrassed in public but some would."* Another participant remarked, *"I don't want to be a performer. It wouldn't come naturally, but the system works well for a performance. Others always look and listen."*

The quantitative results gathered in the post-study interview are provided in full in Table 4.5.

Observations

Participants did not give away many reactions during the study. The main reaction that was observed from participants was smiling and laughter when they were looking at the reaction of spectators (when using loudspeaker versions). There were, however, some participants that appeared a little embarrassed when using some of the systems and attempted to use each interaction technique as quickly as possible. Participants who stood further away from loudspeakers when selecting them generally appeared more comfortable. Overall, participants appeared to stand closer to sound points when using the loudspeaker versions, even though they could be heard from far away. This was also true for both pointing systems, where participants also appeared to approach sound points more closely than when using the number entry systems.

Throughout the entirety of the study, the researcher recorded 528 bystanders who were standing nearby. These were not participants of the study. Of this number, 216 were deemed to have played an active part as a spectators, being in focused relation to the performance [58]. Spectator acknowledgements ranged

	Number Entry		Pointing	
	H	**L**	**H**	**L**
Potential Spectators	104	130	84	210
Spectators in Focused Relation	0 (0%)	73 (56%)	4 (5%)	139 (66%)

Table 4.6: The number of spectators that were in focused relation with each interaction technique.

from small reactions to stopping to look and listen to what was going on. Table 4.6 shows the number of spectators that acknowledged the system and/or performer for each interaction technique:

Through all of the study sessions, not a single potential spectator acknowledged the performer or system when using the number entry and headphones interaction technique. In some cases, potential spectators were standing very close to the performer but still noticed nothing. The headphone and pointing interaction technique was acknowledged by 5% of spectators. These spectators appeared intrigued by the pointing interaction. The pointing and loudspeaker interaction technique gained the most attention, with 66% of potential spectators acknowledging the system or performer in some way. In most of these cases, spectators would begin by looking for the sound source. Once they had found the sound source, some would then attempt to look for the cause of the sound. In some cases, the performer purposely hid or stood further away so that they could not be seen. Acknowledging spectators showed a range of reactions; those standing close to the speakers sometimes gave a shocked or annoyed expression and began to walk away, some were seen smiling and laughing and there were also some spectators that appeared confused.

4.3.5 Discussion

Generally, different participants liked individual aspects of each of the interaction techniques, but there was a clear preference among participants for the more secretive, number entry and headphones interaction technique. In contrast, there was a clear dislike among participants for the more expressive, pointing and loudspeaker interaction technique. The accuracy and level of control for selection—especially when feedback was public—appeared to be a large factor in participants' bias towards number entry. The privateness of the headphone feedback and the ability to control the volume then drove many participants towards a preference for headphones. Many participants said that they were uncomfortable affecting others with loudspeakers, but it is possible that participants only had this opinion due to where the study was conducted. To know whether this is true, we would have to try the system out in a number of different contexts. It could be, that our findings do not hold true in the heritage site context and people may find *Surround You* more enjoyable and more comfortable. Despite this, according to Benford et al. [16], there are some potential benefits to developing uncomfortable interactions, but this was not our aim when designing these interaction techniques.

Interestingly, when participants were asked to give an example of the best usage scenario for the number entry selection and the headphones feedback, suggestions for both were indoor related. In contrast, suggested usage scenarios for pointing and loudspeakers were both primarily outdoor based. Although participants believed that pointing and loudspeakers—rated as the most sociable and performative interaction—should be used outdoors, it was still rated as the least favourite, least enjoyable and least comfortable when used outdoors. This seems to strengthen the context argument, that participants may find the loudspeakers more acceptable elsewhere. If our findings held true in for instance, a visitor centre context, we would argue that there may be some reticence to our new approaches. Even if our findings held true in this context, our loudspeaker interaction techniques might still provide the most sociable and engaging experiences.

In our study, the main reason the pointing selection did not always work as users had hoped was due to a combination of GPS inaccuracy, network lag between communication devices and more generally, users lack of familiarity with the technique. Throughout the study, users were observed standing very close to sound points, expecting to be able to point at a very small, specific point that we had marked out by the physical locations of each loudspeaker. From this study, we have learnt that one should not pinpoint a very small, exact location in space where something can be activated when pointed at. In exploratory testing prior to the study, participants were not told the location of the loudspeakers. These participants were happy to explore and did not seem to mind if the valid pointing location changed. There was no reference point as to where the pointing should work, so it did not bother participants when it moved around. In future, we suggested that pointing may be more suitable when used with bigger targets in a larger area, where the targets are more spaced out.

Although the only interaction techniques that gained any real spectator attention used loudspeakers, they did tend to grab the attention of at least half of the spectators that passed by. When using loudspeakers with both selection techniques, interesting behaviour was observed in a number of participants. Around a third of participants purposely attempted to gain reactions from spectators. Most of the remaining participants claimed that they tried to act as inconspicuous as possible. Those who tried to gain reactions tended to stand far away from sound points when activating them. All of these participants found it rather enjoyable confusing or startling people who were nearby. Although some spectators appeared confused or shocked at first, most reacted with a smile or a laugh when they saw the performer and/or the location of the loudspeaker. A few spectators however were witnessed looking annoyed and walking away. In a future experiment, it would be interesting to interview spectators and ask them about their feelings regarding this.

4.4 Public deployment in visitor center

After conducting our comparative study of performative interaction techniques, it became apparent that although successful at gaining attention and generating curiosity, extravagant, expressive interactions were not as popular with users as we had first imagined. Many performing participants reported feeling uncomfortable when triggering these sounds, not wanting to draw too much attention to themselves or to disturb others. Additionally, there was no statistically significant difference between the attention garnered by hidden or amplified manipulations, though there was also a clear preference among participants for the hidden, screen-based manipulations – providing a greater level of control.

Based on the feedback from this study, we developed a new variation, continuing to use amplified audio produced by loudspeakers, though exchanging the gestural interaction for a more definite, screen-based on-off switch. Interested if our findings from the previous study were also context based, we then decided to carry out a further deployment in a more public, naturalistic setting – the Hafod-Morfa Copperworks. During the deployment, all visitors were able to use the system, and we were interested in speaking with those who did, and also those who did not use the system. It was hoped that this new deployment context, improved interaction variation and discussions with bystanders would provide additional, useful insights into performative mobile location-based experiences.

4.4.1 Prototype design

As participants in the performative audio comparative study enjoyed the definite control and anonymity of the number entry interaction, the prototype for deployment in a visitor centre focused on the similar magical design concept [108], where manipulations are hidden and effects are amplified. Instead of using numbers to identify each sound point like in the previous magical system though, quick response (QR) codes were used. QR codes were chosen as an access point, as they are already commonly used as a portal to information for visitors in public spaces[1], and would hopefully be seen as an acceptable point of entry to the system for visitors with compatible mobile devices. Although previous research suggests that QR codes are perhaps less favourable than visual recognition and number entry [153], QR codes are the only technology here that work without requiring a site-specific application. For visitors who were not able to scan QR codes—but could still access the internet—a text web address was provided (see Fig. 4.10).

When visitors scanned a QR code or manually entered the URL, the corresponding web page would load for that point of interest (POI). On this webpage was the title of the POI, a short description and an on/off switch (see Fig. 4.11). The need for a short description was prompted by our experiment in Chapter 5, where our findings suggested using a mixture of performative and informative

[1] History Points QR Codes – `http://www.historypoints.org`

Figure 4.10: Large QR codes were placed at each point of interest that had an interactive sound point. A web address was also added for visitors who did not have a reader installed. These signs were branded with site-specific graphics to align with other interpretation available on-site.

design. The switch in the web page controlled the amplified audio at that particular POI. The switch state was synchronised between all devices – if 5 users had a particular POI open in their web browser and one user turned the switch on, the switch would change to 'on' for all devices.

The audio snippets that could be triggered at each POI were carefully tailored to the specific location where they were to be deployed. In the case of the following experiment, the audio snippets that were used were; a canal, a river, an engine house, a locomotive shed and a works-entrance. Each audio snippet was a personally curated, rich collection of sounds, attempting to recreate a fairly realistic soundscape.

4.4.2 Study: public audio in a public context

We conducted a study on an open day at the Hafod-Morfa Copperworks in Swansea, Wales. QR codes were placed outside of points of interest so that visitors could scan them and gain access to the control page for that sound point (see Fig. 4.10 & Fig. 4.11).

As well as our sound exhibit that had been placed on-site for the day, other interpretation that was available included information placards with text and

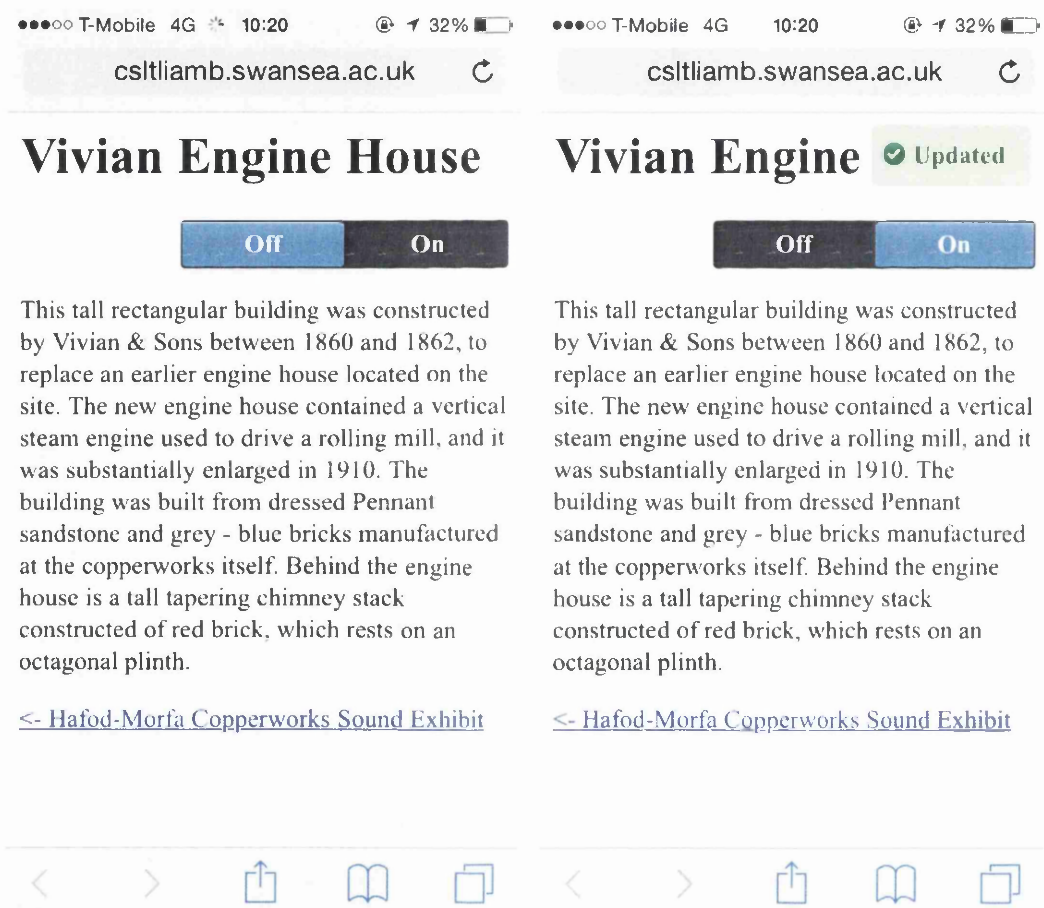

Figure 4.11: The interface for turning a POI on and off. *Left:* POI in the **off** state. *Right:* POI in the **on** state. An 'updated' pop-up also appears to acknowledge a successful response from the server.

Figure 4.12: Existing interpretation available on-site. *Left:* Information Placards *Center:* Worker Statues *Right:* Wind-up Character Sounds

images describing points of interest, statues of workers commisioned by an artist and also wind-up sound boxes that played back the voices of individuals who were historically involved with the site (see Fig. 4.12). This existing interpretation was relatively new to the site, though not related to our research.

Participants

In the hope of achieving a true reflection of visitor opinions, participants were recruited on-site at the Hafod-Morfa Copperworks. These were visitors, all with the intention exploring the site on its official opening day. Estimates from the organisers put the footfall at around 5,000 for the day.

In total, we recruited 13 participants (2M, 11F), forming 5 separate groups (see Table 4.7). Participants were not asked their explicit ages, though considering their age brackets, 4 were children, 2 were young adults, 5 were adults and 2 were older adults. Around half of the participants had visited or seen the site previously, with the remaining participants visiting the site for the first time. None of the participants had been exposed to the technology on display prior to the day of the study. All participants in the study had walked around and explored the site prior to taking part in the study.

One of the participants in Group 5 was an artist involved with other interpretation at the same site, though had no involvement in this research. Group 4 was not actually a group but a single participant. The participant is still referred

Group	Age Group	Gender
1	Adult	M
	Adult	F
	Child	F
	Child	F
	Child	F
2	Old Adult	M
	Old Adult	F
3	Adult	F
	Child	F
4	Young Adult	F
5	Adult	F
	Adult	F
	Young Adult	F

Table 4.7: The constituents of the 5 groups recruited at the open day.

to as a Group 4 in the evaluation as not to cause confusion.

The individual groups of this study later went on to take part in more open discussions regarding the sharing of personal media using mobile performative technologies – reported in Section 7.4.

Procedure

Prior to the open day, the investigator placed 5 loudspeakers across the Hafod-Morfa Copperworks site, hidden around different outdoor structures and points of interest (see Fig. 4.13). Using the interaction described in Section 4.4.1, all visitors to the site were given the ability to turn these loudspeakers on and off when visiting each point of interest. During this period all usage data of the sound exhibit was logged, and intermittent, informal observations were made of visitor behaviour.

After the event had been running for an hour, participant recruitment began, discussing visitors' experience of visiting the site. Before taking part in the study, participants were first asked if they had explored the site on that day. This was integral to the study, ensuring that visitors first had the chance to potentially

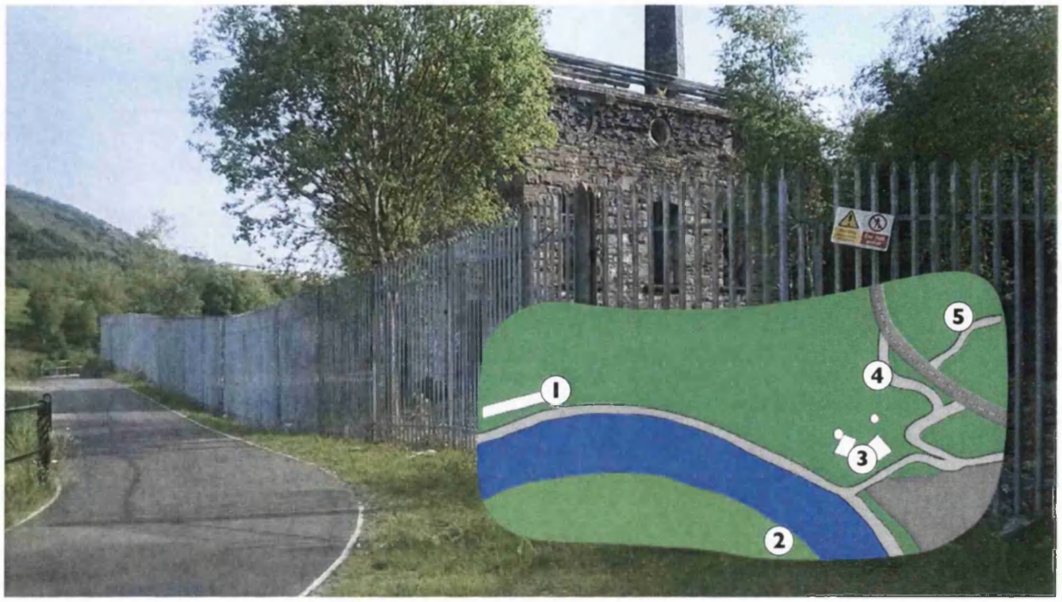

Figure 4.13: *Top:* the Vivian Engine House. *Bottom:* a map of the POI locations 1–5 at the Copperworks; *1:* Locomotive shed *2:* River quay *3:* Engine house *4:* Works-entrance *5:* Canal basin.

see, hear and interact with the sound exhibit. Visitors who had explored the site were asked if they would like to take part in a short, semi-structured group interview regarding their experience of exploring the site. Visitors that were happy to take part were then asked to read and agree to an ethically approved consent form detailing the experiment terms. Where individuals such as children were involved, parents or guardians gave consent on their behalf.

After consenting, participant groups took part in a 15 minute semi-structured interview regarding their experiences of exploring the site. Upon completion, each group was given a £5 Amazon voucher as a token of appreciation for taking part in the study.

Data collection

Data was collected in the form of informal visitor observations and a semi-structured interview that gave visitors the chance to express their opinions of the sound exhibit.

The informal observations were captured by an investigator in-between conducting semi-structured interviews. During this time, the investigator walked around the site and blended in as a visitor, observing others' behaviour. General notes were taken on visitor behaviour that related directly to the sound exhibit on site. This included both visitors who interacted directly with the exhibit, and also those that passively experienced it by standing nearby.

The semi-structured interview data from each group was audio-recorded to ensure that important information and quotes were not missed by the investigator. The interview questions were all unbiased, with most being open-ended in an attempt to provoke thoughtful responses. While most of the questions provided qualitative data, there were also some that required straight yes or no answers. These questions were an attempt to understand whether visitors had understood, seen, heard or were even aware of certain aspects of the exhibit.

Open usage logs

During the open day, all 5,000 visitors were free to interact with the audio exhibit. Each time a visitor interacted with a sound point, log data was captured. This data included how many unique visitors had loaded the control pages, when sound points had been turned on and off and also by who. The log data was captured to give an idea of how many visitors were actually using the sound points, and how they were using them.

4.4.3 Findings

The following section includes the results from the study, including the open usage log data, informal observations of visitor behaviour and the semi-structured interviews that some visitors chose to take part in.

Open usage log data

When distinguishing individual users by distinct IP addresses, there were a total of 22 unique users. Throughout the day, the QR code signs received a collective 91 scans. Users carried out a total of 133 actions between all sound points, with 67 (50.4%) 'on' actions and 66 (49.6%) 'off'. This meant that over all users, there

	Scan (Count)	Action (Count)	
		On	Off
Hafod Quay	20	14	14
Locomotive Shed	14	10	10
Engine House	30	20	21
Works Entrance	17	16	15
Canal Basin	10	7	6

Table 4.8: Scan and action count for each sound point.

were an average of 1.46 actions per page load. The total scan and action count is contained in Table 4.8.

However, upon observing the log data more closely, it is clear that of the 91 page loads, 25 (27.5%) resulted in no actions. Considering the remaining 66 (72.5%) page loads, the average number of actions per page load for users who actually committed an action was 2.02. For users that intermittently switched between the on and off states repeatedly, the log data suggests that on average, they left around a 10-20 second gap before the next action. Of the 22 unique users that made scans, 7 (31.8%) users did not carry out any actions at all.

Although the number of unique users was fairly low, we did see some conflict in control. At 13:41:26, a user turned the works entrance on. Then, 28 seconds later at 13:41:54, a second user turned the works entrance off. The first user then came back 4 seconds later, turning the works entrance on again. At 13:44:45—when the first user had left the scene and began triggering sounds at the canal basin—the second user then came back and turned the works entrance off again. The log data shows that there were 6 separate occasions where this kind of pattern was exhibited throughout the day, with two to three individual users switching between control in a period of less than 5 minutes. This behaviour happened across all sound points except for the canal basin, which was the quietest and most secluded sound point.

Informal observations & feedback

The majority of the visitors at the site were of an older demographic, though this was not unexpected at a heritage site. While walking around the site, the investigator got into a conversation with a visitor who was scanning the codes at sound points. The visitor did not realise that the on and off switch on the page would control the ambient sound at that point of interest. After discovering this, they were very excited and explored the site a second time, activating all of the sounds as they stopped at each sound point.

When sounds were triggered, they transformed fairly quiet areas into industrial cathedrals, appearing the bring the place back to life. This shocked many bystanders, especially those who did not understand what had triggered the sound. Upon speaking to a number of visitors to the site, they were not aware that anyone was triggering the sounds, or that the option to trigger sounds was even available on-site. These visitors thought that the sounds were automatically triggered on a timer, and had no human intervention.

When sounds had been playing constantly for some time, visitors would often pay a large amount of attention to the points of interest, looking for where the sounds were coming from. These visitors, however, were unable to locate the sounds sources, as they were purposely well hidden in an attempt to remove the obvious digital element.

Semi-structured interviews

All five of the groups that took part in the semi-structured interview reported hearing the 'digitally recreated sounds' as they moved around the site. Three groups reported hearing 4-5 separate sounds, with the remaining two groups hearing 2 sounds each.

When asked what their initial reaction was upon hearing these sounds, reactions between the groups were mixed. Groups 2 and 3 reported a sense of awe. In Group 2 in particular, this was exhibited by a participant remarking, *"I thought you had fired it up. I thought, how on earth have you done that?!"* Groups 1 and 4 were more reflective, discussing that the sounds made them ponder what the site would have been like in its heyday. The two adults from Group 1 conversed, *"[It] sort of brought it to life didn't it. [...] Tried to imagine what it was like, back in time. [...] What the workers had to put up with. [...] The noise."* One participant from Group 5 was also reflective, but not before being shocked, exclaiming, *"My initial reaction was 'Good ...! What is that?!' Because it just seemed to come out of nowhere. And someone was pointing in the trees, and it was like, is there going to be some big bulldozer or something come over the top?"* Another participant in this group had the complete opposite reaction, saying *"I thought it was normal [when] walking past the engine house. [...] For that couple of seconds, it sounded right."*

Groups 1, 3 and 4 spent their time actively looking for the sounds. In Group 1, the parents reported that their children enjoyed searching for the sounds. Group 2 and 5 did not actively seek out sounds. When describing the introduction of sounds into the site, two groups commented that the sounds enhanced the experience, with the other three individually referring to the sounds being meaningful, a source of immediate information and generally being effective in conveying a sense of place. A participant in Group 5 gave an interesting analysis of the sounds, saying, *"I didn't actually enjoy the noise. I don't like machine noises, I like bird noises. [...] I suppose I enjoyed the canal a bit more because it was more decipherable – you know, you could actually listen to it and go, 'Oh yeah, I get that. I understand what those noises are.' Where as down by the river, I didn't feel like I could deconstruct the noise in any way."*

Groups 3 and 4 were the only ones that reported trying to activate or control the sounds. Both of these groups liked the fact that the sound points were controllable, commenting that it felt 'good'. The single participant in group 4 summarised the positive points of control, saying, *"If you don't like a sound, you can turn it off. You have a sense of power. [...] It's also good, say if, you wait for people to come around you and you switch the sound off, and they're not using the app. [...] And then, you just randomly turn it on and they're like, 'Oh my ...!'"* Neither of these groups reported a conflict of control at any point. Both groups commented that the activation of sounds evoked shock and awe among bystanders, prompting them to then move closer to the points of interest. A participant from Group 3 recalled, *"[They] came running... [laughing] Towards, towards definitely."* Both groups also suggested that bystanders did not know it was them who was controlling the sounds. The same participant from Group 3 went on to say, *"We found it quite amusing down by the Quay, to be turning it on and off. Because people on the other side didn't know we were controlling it, so that was quite fun [group laugh]."* The single group 4 participant similarly remarked, *"It was pretty cool as well that nobody knew"*. They then continued to say that for this reason, they did not feel "embarrassed."

Turning now to the groups that did not attempt to activate or control sounds. Groups 1 and 2 were not aware that the sounds were controllable until the interview. Group 1 reported scanning some codes, though one participant admitted, *"I seen a switch, but I thought that meant turn the QR code reader off."* No participants in Group 2 had a device capable of accessing the internet. The remaining group—Group 5—said that they were aware that the speakers were controllable, though didn't attempt to control them as they were already activated. At the end of the interview, the participant in Group 4 suggested using signs or a poster, *"explaining that you can turn the sounds on and off."*

4.4.4 Discussion

During the deployment, of the estimated 5,000 attendees at the site, we recorded only 22 unique users – 7 of which did not trigger any of the sound points. Although this appears to be a very low number, the majority of visitors that we spoke to on-site had heard the sounds produced by the system at some point during their visit. We believe that there are a number of reasons for the low numbers of visitors who interacted directly with the system. Many of the visitors to the site were of an older generation, and did not have a device capable of interacting with the QR codes or accessing the Internet. Something else that became apparent during observations and the semi-structured interviews was that many visitors were actually unaware that these sounds could be manually triggered. This was even the case for some participants who scanned the QR codes, thinking the on/off switch controlled a setting on their own device. Perhaps, as one participant suggested, if the ability to control these sounds was better advertised, visitors would have been more willing to try it out. We believe the fact that many visitors—even those who scanned the codes—remained completely unaware of the functionality, shows that this is not the expected kind of interaction that people are used to in public spaces. In mainstream location-based services, users are used to being consumers, passively receiving information – not taking part as a performer.

The two groups of participants in the semi-structured interview that had turned sound points on and off admitted to trying to provoke reactions among bystanders by unsuspectingly triggering sounds when they were nearby. In both cases, the participants said they thought that bystanders did not know it was them that was controlling the sounds. The participant in group 4 went on to say that this was the main reason for them not feeling embarrassed. Although bystanders did not interact directly with any of our participants that triggered sounds, participants did report that the sounds affected bystander behaviour, with visitors *"running"* towards sounds and talking about them. Also, although none of our participants experienced a conflict in control, we believe that the conflict in control recorded in our log data may well have created some social conflict or discussion amongst the users involved.

Overall, participants in this study portrayed a very different attitude to those of our previous experiment in the University. All were upbeat and very excited by the technology potential, including both those who controlled, and those who experienced the sounds passively. Those who controlled were not worried about disturbing people or becoming a spectacle. This could perhaps, be attributed to the ability to trigger sounds confidently, with control from afar. We believe that the different context may have also played an important role, where the sounds were of collective interest to the visitors. Participants in the interviews all thought that the sounds used at each POI were highly representative of what the site would have sounded like in its heyday, though one participant had concerns with the composition of the sounds. This participant gave an interesting analysis of the audio used, saying that it was more enjoyable when it could be broken down. Although we made an attempt to simulate reality closely, this participant's point is something to perhaps consider in the future development of public soundscapes, if sounds are intended to be enjoyable and a pleasure to listen to.

4.5 Conclusions

In this chapter we began by exploring the value of spatial audio for discovery and navigation in large, unfamiliar places. We then carried out a deployment study on a number of performative audio experiences, comparing the effects of each. Due to the numerous unanswered questions born from our findings in this second study, we conducted further research on a refined prototype in a more naturalistic, public deployment. Throughout this chapter, we have demonstrated a number of novel interaction methods—magical to expressive—that allow users to trigger amplified audio in public spaces using their mobile device as a control.

In our initial experiments, we discovered the engaging qualities of spatial audio in unfamiliar surroundings, allowing listeners to gain a better understanding and sense of place within a very short period of time. Our early prototype helped visitors discover new points of interest they had not visited before, and also provided a clear enough representation of spatial audio that users were able to successfully navigate towards these sounds without visual cues.

Following Reeves et al.'s [108] spectator observability taxonomy as a guide, we then developed four different interaction techniques; secretive, suspenseful, magical and expressive. Each of these interaction techniques was tested in the same public setting, with participants able to compare each approach. There was a clear dislike among participants for the most expressive interaction technique—involving pointing at loudspeakers—with many preferring the conventional secretive approach of screen-based number entry and headphone feedback. This cause for dislike was attributed to a lack of definite control with the pointing gesture manipulation, as well as the University context, with performing perception [42] most likely playing a part. Participants creating a spectacle of themselves or causing a disturbance in this particular context may be seen as antisocial, or perhaps less acceptable.

With these findings, we altered our interaction technique, providing a more definite control through an on/off switch, which was then made publicly accessible through a web page. We then decided to carry out an additional experiment in a more naturalistic deployment, where the context was different—Hafod-Morfa Copperworks—and the feedback was of shared public interest. With this new system and context, users appeared to be much more comfortable using the system, enjoying the triggering of sounds.

When using *Surround You*, our expressive audio interaction, we were surprised to see that bystanders did not approach the performer. Similarly, when using the magical interaction at the Hafod-Morfa Copperworks, no bystanders were reported to have communicated with users who controlled the points of interest. Although bystanders did not always approach the performer directly as we had first imagined they would, curiosity and mindfulness were prevalent among bystanders in both of these experiments, with the majority of spectators reacting positively to amplified sounds, looking for the audio sources and discussing them with friends.

4.5.1 Designing for performative audio place-based experiences

From the design and evaluations of prototypes in this chapter, we suggest a number of factors that are critical to the success of an engaging performative audio place-based experience. Many of these factors are applicable to the entire performative mobile place-based experiences design space.

It is clear from the experiments conducted in this chapter that, amplified audio proved both engaging and immersive. Both expressive and magical interaction techniques gained a large amount of attention from bystanders, turning them into spectators. The visibility of manipulations had less of an effect. Amplified loudspeaker audio made users feel fairly uncomfortable in a University setting, though this was not the case when amplified audio was used at the Hafod-Morfa Copperworks. For this reason, we suggest that the acceptability of these kinds of amplified, public interactions is heavily dependent on the context. When using our expressive prototype, pointing gestures gave users a feeling of less control. A more simple and definite on/off was preferred, also allowing users to stand further away from the effect, offering users anonymity – openly welcomed by some. In our later study at the Hafod-Morfa Copperworks, many visitors claimed that they did not have a compatible device or were not aware of the interaction possibilities. We recommend making interaction points clear, physically and digitally – many people were not aware the technology was controllable, some even after scanning. Although a very low number directly interacted with the system, a much greater number were still able to be a part of the experience, making it very inclusive.

Affordance Ensure that the technology, access points and functionality are all clear. Upon arrival at the interactive space, we suggest some kind of prompt—physical or digital—that there are sound points in the area that can be actuated. Physical markers should then be placed in the environment at each sound point, either as an access point or a prompt to trigger a nearby sound point. In the user interface, next to the trigger or switch, we recommend making the functionality explicit and clear for users.

Visitors to public spaces are used to being consumers, passively receiving information. It is important these prompts and affordances exist so that users will know what the technology is, as well as how and when to interact.

Control For the experience to be usable and enjoyable, there needs to be a high level of definite control when triggering and stopping sounds. Our amplified gestural pointing mechanism worked, though was perhaps too expressive, with too many affecting variables. This sometimes led to triggering a different speaker to the one the performer had intended on interacting with, or even triggering multiple speakers within a direct line of sight. There seems to be an overwhelming preference for being able to accurately choose a sound point and then explicitly decide when it should be on or off.

Visibility Amplified effects had much more of an impact on bystanders than amplified manipulations – at least on the mobile scale. Compared to effects, differing levels of manipulation visibility garnered barely any attention at all. Although this was the case, performers who used the magical interaction techniques reported that they enjoyed their manipulations being hidden, where it was not always obvious that they were the ones who were controlling.

Engagement In the case of magical and expressive interaction techniques, both sparked curiosity and conversation amongst bystanders, though no spectators engaged directly with the performer. We conclude that perhaps people are drawn towards effects, and not manipulations. In our experiments, the amplified effect was placed far away from performers, and performers were not approached. If designers desire engagement between spectators and the performer, we suggest placing the effect with or around the performer.

Context The social acceptability of these interventions is heavily dependent on the context. Amplified audio was enjoyed impetuously at the Hafod-Morfa Copperworks, though was deemed less acceptable on a University campus. We conclude that these kinds of interactions are perhaps more acceptable where the sounds are of shared public interest, and are also carefully used to compliment the surrounding environment, rather than append to it.

Chapter Five

Visual Performance

In this chapter we focus on visual performance, exploring the use of visuals to create more performative, extravagant, expressive place-based experiences. We began by thinking of the more expressive kinds of mobile visual interactions – the result of which culminated in the design and creation a pico-projection prototype, able to project a set of small-scale images and videos. As with our performative audio prototypes, it was again our goal to develop an experience with amplified manipulations and effects, making the experience as accessible as possible to bystanders. We carried out a deployment study of our pico-projection prototype, comparing its expressive and performative nature to a more conventional, screen-based interaction. We report on the results of this experiment, discussing the use of pico-projections for performative interactions. The chapter concludes with a set of design implications for performative visual experiences in public spaces. Our visual experiments were the first that wes carried out, and therefore, do not build upon the findings of our other research.

5.1 Introduction

Many of the mainstream location-based information services to date have focused on delivering images, multimedia or text in situ, direct from the internet to a mobile device screen. While this sort of interface does indeed provide valuable location-based information, we feel that it is lacking in some of the more exciting or participatory interactions that could be possible. Small mobile screens are good at displaying private and personal information, but this also makes them very poor at displaying public information to groups of people.

The *Augurscope* [130] was an early outdoor mixed reality prototype for groups of people at an attraction. The system used a wheeled tripod base, and a laptop to display historical augmented content. While far from a current mobile portable device, the system allowed groups of people to share their experiences in situ.

In an attempt to draw people away from the screen, we have looked at developing publicly visible content using projectors. Bert Bongers was one of the first to realise the potential of taking a projector outside of its intended context [19]. Before the introduction of pico projectors, Bongers took to the streets at night with a full sized projector to project gateways to alternate realities. Over time, situated projections have become more interactive, with certain examples allowing users to control the content that is projected by external means [84, 126]. In L.A.S.E.R. Tag [84], a user is able to use a laser pointer as a paintbrush. The projections are aimed at large outdoor objects, giving the user a canvas to paint over the real world. Full-size projectors have been used on many occasions to augment buildings as part of a multimedia performance, using 3D video mapping to make the buildings appear to come to life (e.g. see `http://nuformer.com/`). Sauter developed an intrusive but playful artwork named Light Attack [125], where a white silhouette of a person was projected onto the walls and sidewalks of streets from a moving car. A further example of situated, but moving projection is demonstrated in [72], where projections are dynamically generated and displayed on the banks of a river from a riverboat.

Since the introduction of pico projectors, although there has been interest in their uses in public space, there has been little research regarding their usage as mobile guides – in particular, where a person is able to explore and learn about the place that they are in. Pathlight [152] is a mobile guide designed for a museum. Its main purpose is to give information about exhibits, but it also doubles up as a navigation guide within the museum, directing the user to different exhibits using projected arrows. Researchers, have been quick to take advantage of pico-projectors in mobile information scenarios, using projectors for collaborative learning [74] or gaming [158]. Our prototype is inspired by these previous designs, but we focus on the collaborative, shared viewing experiences possible when using public mobile projection, such as those discussed by Greaves et al. [61] and Robinson et al. [118].

In this chapter, we report on a prototype system to display pico projected images and animations. By using a handheld projector to view digital content, instead of the typical touch screen on a modern mobile device, we aim to allow for a more expressive visual mobile experience, with performance and collaboration in mind. The novelty of this work lies with the use of pico projection to augment exhibits at a visitor attraction. While some of the findings have been seen in prior works, the performative perspective of this research sheds new light on the possibilities.

5.1.1 Extravagant, expressive visual place-based experiences

Although projections have been used in public spaces before, in this research, we are specifically interested in developing a mobile visual experience that promotes a more active social engagement between visitors. Where other kinds of mobile augmented reality (AR) lens applications overlay information over a camera image, here, using pico-projection, we take this type of AR lens application one step further – overlaying the real world itself. The amplified visual effect here, then, is the projection – taking the content out of the screen, enlarging it and placing it in the environment to be viewed publicly.

In terms of amplified manipulation, our first choice here is again gestures and movement. Both are well suited to pico projectors, as they are lightweight, portable, and do not require complex tracking. In particular, there was a large emphasis on the user being able to control the projections in a free-form way, positioning projections using a red dot, preventing the need for a sophisticated tracking system. It was hoped that the free-positioning of projections—without the bounds of tracking—would allow for more amplified manipulations.

Therefore, we image a performative, extravagant, expressive visual place-based experience as one which makes use of publicly visible projections, and can be expressively controlled by moving the mobile device, gesturing and walking through space freely, projecting without sophisticated tracking.

5.2 Pico projections for performative interactions

Following an in-depth review of the literature and a discussion of a suitable mobile concept, we proceeded with the design and development of a performative visual prototype. In the following section, we introduce our pico projection prototype, designed to encourage a more playful and active engagement amongst visitors. As a baseline, we also introduce a comparison system, which follows a more traditional, mainstream screen-based information approach. Both techniques were deployed in-situ in a visitor centre, and compared in a between subjects study. The results of this study are presented, and we discuss their implications.

5.2.1 Prototype design

We developed a mobile prototype to demonstrate performative, place-based projection. Since we wanted to test the system with real visitors at a local attraction, we worked with the curators of the National Botanic Garden of Wales, carefully crafting and tailoring the types of performative projection specifically to the context of their visitor attractions. One of the key reasons for working with this attraction was their previous interest in being early adopters of emerging technologies. For example, the gardens (both outdoor sections and displays in the large indoor glasshouse) are viewable on Google Street View[1], allowing visitors an online 3D tour of the attraction before visiting. Our system reuses QR codes and some curated content from a previous project at the National Botanic Garden of Wales (similar to [122]). We use QR codes for indoor location awareness and as a visual indicator to show users that they can interact with that exhibit.

For this research, our main aim in using a pico projector is not as extra screen space, but specifically to allow a visitor to augment the real plants and objects with digital content, acting out or performing actions with the elements they project. Our approach contrasts with text projection onto objects [80], as we do not project text information into the environment, but visual animations

[1]GoogleMaps Street View : National botanic garden of wales (`http://goo.gl/maps/exaFC`)

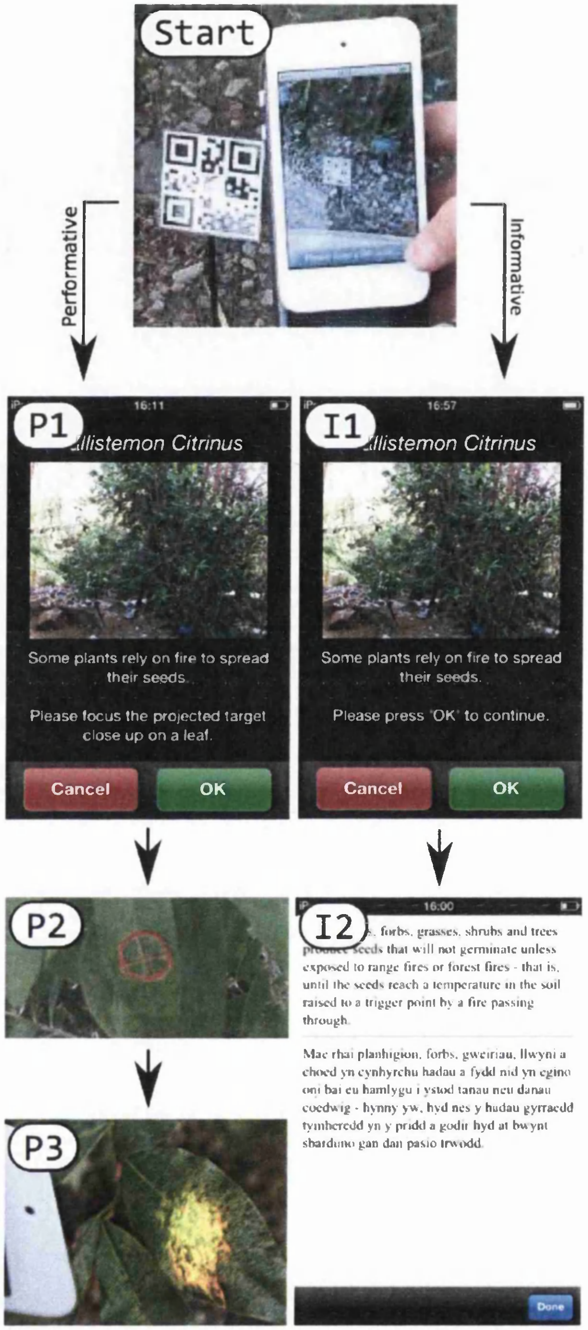

Figure 5.1: Using both prototype systems: scan QR code with iPod touch (Start) *Performative:* prepare (P1 & P2; confirm and point target) and project (P3; a fire) or, at right, *Informative:* confirm scan (I1) and display text (I2).

Figure 5.2: Projecting images and animations in *performative* mode. Clockwise from top left, projections show: a leaf on fire; a sunbird flying to a plant; raindrops falling on a leaf; and, a witchetty grub on a tree trunk.

that attempt to overlay and bring the environment to life. Where spectators are present, the moving and positioning of the projector and the effect can create an illusion of AR.

Rather than previous lens-based approaches, however, the augmentation is projected directly onto the physical elements to which it refers. Our system uses an iPod touch attached to a pico projector (see Fig. 5.1). The iPod is used to scan QR codes situated next to eight exhibits around the gardens' visitor centre. After scanning, an image and sentence of context about the exhibit are shown onscreen, along with a prompt to focus a projected target on the object. The user presses a button when ready, and imagery or animation is then projected. As the prototype was built for use at this specific attraction, the imagery used is of insects, animals or environmental factors that are related to the plants and other displays at the botanic gardens. Figure 5.2 shows several such examples, where the projection appears next to or on top of the related artefact. Apart from the initial QR scan, the system does not implement any additional tracking. This allows users of the system to project freely onto objects in an attempt to promote performative and playful behaviour. For this prototype, content and QR codes are stored and recognised locally, as the device is designed specifically to augment this particular attraction. For use with a wider set of attractions, we imagine an online repository of content paired with displays in other visitor centres.

Informative, screen-based system

We also built a second, alternative mode into the system, allowing us to compare traditional screen-based location information with the projected content approach. After recognising the QR code and showing the same initial content on-screen as in the performative system (i.e., a sentence of context and an image), upon pressing a button it then displays a page of textual information about the object, instead of a projection (see Fig. 5.1).

5.2.2 Study: a comparison of performative and informative interactions

We conducted a study over six days at the National Botanic Garden of Wales. The aim was to test both systems with real visitors in situ. We had two research questions:

RQ1: How do perceived learning and enjoyment through performance with projections compare to perceived learning and enjoyment with text-based information?

RQ2: How does the performative aspect of the projector system affect involvement or interest from non-participant visitors when compared to the informative system?

Participants

Twenty groups of participants were recruited as they entered the building. Ten groups used the informative system and ten used the performative system. A total of 58 participants took part, with 34 people using the informative and 24 using the performative system. Participants' ages ranged from 3–80, with 29M (Male), 29F (Female) overall, and similar gender distribution between systems. The average group size was around 3 participants, with 3.4 (σ 1.26) and 2.4 (σ 0.84) on average for the informative and performative systems, respectively. None of the participants had used a pico projector before, and very few had used QR codes. Most participants were repeat visitors to the garden. Groups consisted mainly of couples or families looking for something interesting to do on a day out.

Measures

To gather users' opinions of the system, a short survey was built into the prototype. After scanning a QR code and either projecting or reading the related content, the prototype prompted the group to give feedback. Groups were instructed to give feedback collectively. The survey questions asked:

Q1 *How many other visitors stopped to look at what you were doing?*

Q2 *Rate your enjoyment of this particular QR code experience.*

Q3 *Rate the value of this particular QR code experience as a learning resource.*

Q4 *How has your understanding been affected by this QR experience?*

Questions 2 and 3 allowed participants to select a rating from 1 (low) to 5 (high). Q4 allowed a selection from 'decreased,' 'unaffected,' and 'increased.' In addition to the survey, participants answered a short semi-structured interview at the start and end of each session. For one of the six days, two additional researchers observed groups' behaviours from a distance while they used the prototype, being careful to avoid intruding on the experience. In total four groups (13 visitors) were observed, with three using the performative and one using the informative system. In addition, during that day, many visitors who were not participating in the study were also observed.

A pre-study interview gathered information regarding participants' previous experience with QR codes, projectors and smartphones. Participants were also asked about any previous visits to the botanic gardens, and their ages, genders, and the number of participants in each group were recorded.

1. *Have you been to the Glass House before? (if so, when?)*
2. *Have you used QR codes before? (if so, where?)*
3. *Have you used the QR codes here in the Gardens? (if so, when?)*
4. *Have you used a Smartphone before?*
5. *Have you used a Pico Projector before? (if so, where and when?)*
6. *Group Size, Gender(s) & Age(s)*

A post-study interview focused on the effect the prototype had on participants' and other visitors' experiences during their visit. Participants were also asked to report any surprising, exciting or annoying experiences with the system, and how other people around them reacted when they saw it in use. Finally, participants were asked to suggest any improvements they might like to be made to the prototype or the experience.

1. *How has your visit to the glasshouse today been affected by the QR system you've just used?*
2. *Was there anything that surprised or excited you during the experience?*
3. *Was there anything that angered or annoyed you during the experience?*
4. *How did people react around you when you were using the system?*
5. *On a scale of 1 to 7, how would you rank the usability of the system?*
6. *How did you cope with finding the QR codes?*
7. *Can you suggest any improvements to the (projector/information) experience?*

Procedure

The study used a between-groups method where each group used either the performative projection system or the informative textual system. Participants did not use both systems to avoid learning artefacts. At the start of each session the study procedure was briefly described to the group, and they were asked to confirm that they wished to participate by reading and signing an ethically reviewed research consent form. In an attempt to minimise the possibility of separate groups of participants encountering each other, the researcher conducting the study left a 20 minute interval between recruiting new groups.

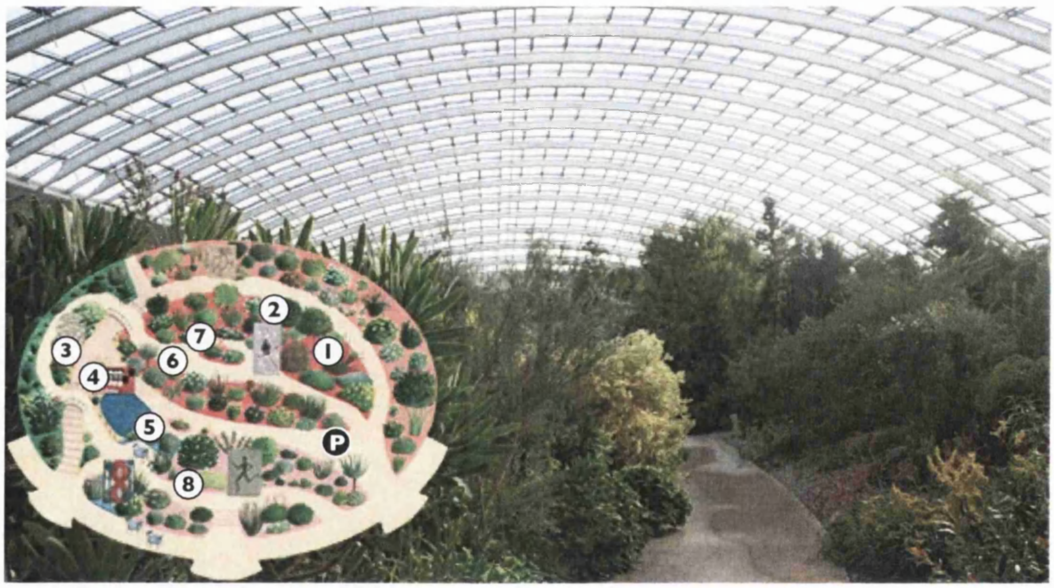

Figure 5.3: *Top:* the large glasshouse in which the study took place. *Bottom:* the garden map with exhibit locations 1–8 indicated.

After groups agreed to participate, a short training session was conducted to demonstrate usage of the system to the group. The group was then given the prototype (in either performative or informative mode), and an information sheet in case they needed further guidance. This sheet also incorporated a map showing the approximate location of eight QR codes to scan (see Fig. 5.3). Because the study was conducted in such a brightly lit area, the location of each of these codes was carefully selected so that projections would have the best visibility possible. These areas were shaded and darker than the rest of the glasshouse.

After receiving these instructions, the group then left the researcher, finding and scanning each separate code and completing the five survey questions after viewing the content associated with each display. Discreet observations were made of participants during some study sessions. At the end of each session, the group were debriefed in a short post-study interview, thanked for participating and given a gift voucher.

5.2.3 Findings

Firstly, we will consider the data gathered by the mobile application after each exhibit (see Fig. 5.4). For Q1, the average numbers of non-participant visitors that were reported were 1.39 for the informative system (σ 1.66) and 1.88 for the performative system (σ 2.33). There is an overall significant difference in participants' rating of whether their understanding of an exhibit was affected (Q4), with the informative system seen to be more beneficial in that respect ($p < 0.002$; Mann-Whitney). Turning to the ratings of enjoyment (Q2) and perceived learning (Q3), there was no significant difference between systems.

In the post-study interview, all participants indicated that they had noticed interest from other non-participant visitors around them. A common sentiment was captured by one participant, who said: *"if people were around they looked."* In some instances, other visitors were curious enough to ask participants what they were doing. Three groups using the performative system reported that they demonstrated the system and engaged with non-participant visitors. One of these said that their performance involved 13 visitors who became interested in what was happening.

Feedback from the post-study interviews for both systems suggests that in nearly all cases groups felt that using either of the systems had enhanced their experience. Participants often commented that the system they used added interest to their visit, with one participant claiming that the performative system gave *"an extra dimension."* Some of the groups with children (using either of the systems) noted the enjoyment in seeking out and scanning the QR codes

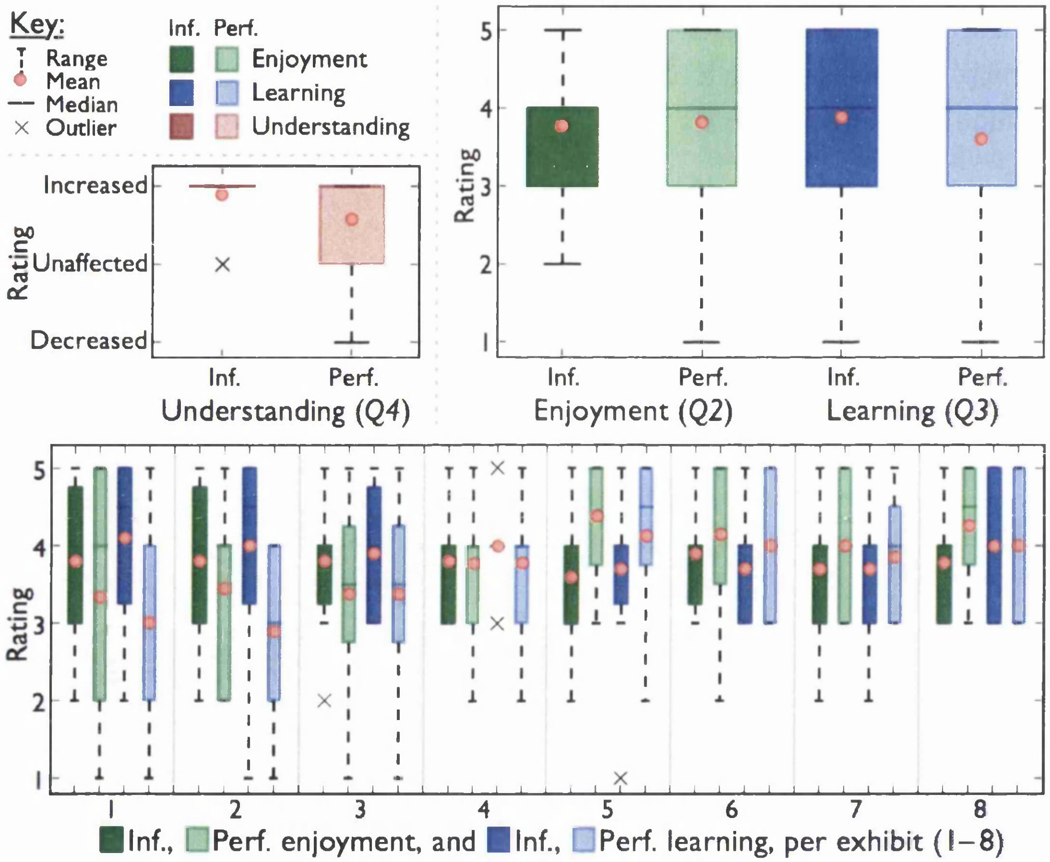

Figure 5.4: Box plots of survey results (Q2, Q3, Q4) for both systems. Top: overall results between systems: understanding is significantly increased when using the informative system, but there is no evidence of an effect on enjoyment or learning. Bottom: individual results for enjoyment and learning separated into exhibits 1–8.

themselves. One participant using the informative system explained this, but noted: *"the children love to find the codes and scan them but they're not interested in reading any of them."* One group mentioned in the post-study interview that people would sometimes get in the way of their projections. This would seem to suggest that participants were pointing projections at a specific object, strengthening the argument that participants were using the performative system as intended – to project onto surfaces that the device instructs. Several participants commented that the brightness of the projector was sometimes an issue. One participant went as far as returning to find a researcher during the study to ask *"how am I supposed to rate this [projection] if I can't see it?"*

Observations

Considering first the group observed using the informative system (four adults; one child) – in general this group gathered closely around the system after scanning each QR code. No single individual in this group took control of the prototype; instead, participants took it in turns to scan each QR code. In some cases one participant read aloud to the rest of the group; for other exhibits individuals read to themselves instead, huddled tightly around the device. While other visitors were aware that the group were doing something unusual, they were not seen to experience the information the group was reading.

With the performative system, where three groups were observed, there was evidence that projection encouraged participation beyond the device itself. Participants were not gathered around the device, but were seen to be focused on the projections rather than the prototype. In one group (two adults; two children), an adult held the device and let the children direct his hand, pointing the projections at plants while visitors stood by and watched.

In terms of the non-participant visitors that were observed, most wandered slowly and casually, occasionally glancing at their surroundings. We only observed one group of visitors (one adult; two children) who seemed more actively engaged in the exhibitions. This group moved directly between information points, paying attention mainly to the signage and other content, such as drawings showing the plants' native habitat.

All groups were asked in the pre-study interview if any participant had used a handheld projector previously. Out of 58 participants, none of them had ever used a handheld projector before. Initially we were worried that this lack of experience would affect the participant's experience of the performative system. Upon reviewing the usability scores (1-7) that group's presented in the post-study interview, it was evident that such a problem did not exist. On average, the performative system received a score of 6.1/7 (min:3; max:7; σ 1.28) and the informative system received a score of 6.7/7 (min:6; max:7; σ 0.48). One possible explanation for the slightly lower performative system score may be that 90% of informative groups had prior experience with a smartphone like device, whereas only 70% of performative groups had prior experience with a similar device.

5.2.4 Discussion

The higher rating given to the informative system in terms of 'understanding' is not surprising given that the system provided detailed textual content for each exhibit, in contrast to the performative system's images and animations. We might have expected a higher rating for perceived 'learning' in the informative version for similar reasons; and, conversely, a higher rating for 'enjoyment' for the performative. However, no significant effect was apparent. For this reason, then, we may speculate that both types of system provide benefits in these respects – allowing for both informative and performative modes in future designs would seem a sensible approach.

Clearly image quality, particularly brightness, impacts on the efficacy of projection. Visibility seems to have played a part in participants' opinion of the performative prototype as a learning resource. For example, the significantly lower rating for exhibit 2 in the performative system may have been because a darkly coloured projection was used in a very light part of the glasshouse. No significant difference was found in the numbers of people reported as stopping to watch by participants using both systems. However, post-study interviews and group observations suggest that bystanders had a more active engagement with the performative prototype.

5.3 Conclusions

In this chapter we began to explore the design space of extravagant, expressive visual interactions for mobile devices. We made the decision to use pico projections for two reasons; projection is an amplified effect, and should therefore be more publicly accessible than a private mobile screen, and also, the projector is easy to move around in a freeform way – allowing for amplified, gestural control without complex tracking. Considering our design choices, we went on to build a projection prototype which we believed was a good representation of what a performative, mobile visual experience may be. We then carried out an in-situ deployment of our performative projection prototype, comparing it to a more conventional information guide that one may expect to see in a visitor centre.

The results of our deployment seem to suggest that a pico projection system—such as the one employed here—may encourage people to engage with their surroundings rather than focus on signage or, if using a conventional mobile device, the device itself. Furthermore, there is some evidence that projection might allow groups to enrich their shared experiences and to draw in bystanders. The choice of locations and attractions for performative projections can clearly impact on the effectiveness of the approach. Future work could explore how to camouflage QR codes to enhance the fun we observed some children having during the hunt for markers. Forcing visitors to stand in *"disruptive"* locations to project content—for example, changing the flow of others along a pathway—may encourage spectators. Careful stage-craft is needed, though, to avoid annoying bystanders or embarrassing performers.

Pico-projection brightness will remain an issue for some time. To accommodate this, and to further use digital output to prompt physical engagement, we might consider providing more stage direction to users. For example, in the garden context, instead of simply asking people to target the beam on an exhibit, the group could be asked to stand round the plant (providing shade), with one of them cupping their hand around a leaf (further darkening the object) before animations begin. Or perhaps larger scale projections.

The novelty of this work lies with the use of pico projection to augment objects at a visitor attraction. In particular, there was a large emphasis on the user being able to control the placement of the projections in a free-form way. While some of the findings have been seen in prior works, the performative perspective of this research sheds new light on the possibilities. Taking into consideration our previous auditory performance findings, along with the current limitations of pico projectors, we have decided that in the case of performative mobile place-based experiences, audio is perhaps a more accessible effect for bystanders. For this reason, we have chosen not to pursue the use of pico projections any further in this research.

5.3.1 Designing for performative visual place-based experiences

From the design and evaluation of our performative projection prototype in this chapter, we suggest a number of key design implications that will encourage designers to develop engaging performative visual place-based experiences. Some of these factors are also applicable to the wider performative mobile place-based experiences design space.

In this chapter, we demonstrated the possibility of a performative projection prototype without the need for complex tracking or control. Users were literally able to point the projections where they wanted them to appear, providing playful and intuitive control. The low-visibility of the dimly lit projections appeared to play a part in some of the more critical feedback received from participants. This could perhaps be overcome with more powerful projections, or careful direction and instruction to users, providing tips on where or how best to project. Additional stage-craft could also play a part in the experience, encouraging further social interaction through the clever placement of exhibits, directly impacting the movement and flow of visitors. Concluding, the performative and informative approaches ranked very similarly among participants, suggesting that both approaches have their merits. Some combination of both approaches, would perhaps, be the most desirable solution.

Control Users can be given control of pico projections without the need for complex tracking. With a small amount of priming/direction, users will point the projection where it should be. In our study, this free-form gestural control allowed users to manipulate the projections expressively. The sheer simplicity of this interaction makes the connection between the manipulation and effect clear, perhaps, giving users a high sense of control – which was proven desirable for performative, expressive audio (see Section 4.5.1).

Stage-craft There is only so much that a user can be expected to do with a performative system, without further direction. We believe that stage-craft, such as strategically placing exhibits to intercept the flow of visitors and prompting users to carry out additional tasks such as physically interacting with exhibits, could all enhance the performance further. This would be achieved through careful planning and curation, signage and further instructions to the user.

Visibility No significant difference was found between the number of bystanders who stopped to look at each system. Clearly, here—although we had encouraged visitors to project in dark places—the brightness of the pico projector did not have an amplified enough effect. Designers should be very careful in ensuring that the brightness of projections are suitable for the context.

Variety In our study, the informative and performative approaches both scored fairly equally. When designing a performative place-based system, there may, perhaps, be too much emphasis on the performance, with the system not providing enough useful information. We suggest that designers are mindful of this, and should cater for a combination of approaches, providing more in-depth information, allowing users to make sense of the performance. This finding prompted a description to be added to our prototype in our audio experiments at the Hafod-Morfa Copperworks (see Fig. 4.11).

CHAPTER SIX

Directing the Performance

In this chapter we build upon the findings of our previous investigations into performative auditory and visual interactions, extending the framework to enable remote, large-scale, dynamic performances. We refer to this concept as portable media hotspots (PMH) – where users and their devices can be commandeered, receiving directions to follow, and amplified media direct to their mobile device. Our motivation stemmed from connecting experiences, and utilising multiple mobile devices to create large displays, rather than using embedded peripherals – as we had used in our previous audio research. We began by designing a prototype remote guidance mechanism, allowing a controlling user to remotely direct a hotspot user through the use of tilt gestures. A controlled study was carried out, measuring the system's guiding accuracy and affect on preferred walking speed. Encouraged by the findings of this first study, we developed a second prototype – retaining and improving all of the original features. We deployed the second prototype in a public setting amongst castle ruins, offering small groups the opportunity to remotely direct each other around and trigger contextual sounds at will. From our experiments, we outline a number of key design implications for remote performative experiences such as portable media hotspots.

6.1 Introduction

While walking through a busy city centre, one can expect to find a whole range of ambient audio and visual displays – screens and speakers providing the latest news and offers, or simply setting the mood. Predominantly, the kinds of displays we seem to encounter in public spaces are static and rooted to the ground. This is usually the case as these kinds of displays have a high power consumption and require a constant, wired supply of energy. Also, large public displays are often fairly expensive and need to be secured accordingly.

A large percentage of the global population own a mobile phone [137], many of which are smartphones. Mobile devices such as smartphones can be used in a similar manner to retrieve and display location-based and context-aware information, though as we have discussed before, the results are often kept private. Although mobile devices are not as visually or aurally extravagant as large, situated public displays, they are far more abundant. Perhaps, then, these mobile devices can be used together in unison, to create large, dynamic ambient displays. We have already learnt from our previous studies that placing the effect directly on or around the performer encourages spectators to engage directly with the performer, promoting a more active engagement. As well as promoting a more active engagement, mobile devices are low-cost, low-energy and can be moved freely around a space, or even be directed towards specific positions in space. Bearing all of these qualities in mind, we have a vision – to use Internet connected mobile devices as scalable, controllable outdoor displays.

The idea of using mobile devices as a collective screen-based visual display is not an entirely new one [131], though less research has been conducted on the use mobiles as a collective audio display. *Samsung Group Play*[1] is one particular example of a mobile application that attempts to use multiple mobile devices as a collective audio display, though its functionality is fairly limited, allowing a user to share and transmit only one audio track at a time to other devices. Each device must also be connected to the same wireless network for the application to function. *MoPho* [149] is an example of a more creative scalable mobile audio display, where each mobile phone is used as an instrument in an orchestra.

Something that is not usually considered in collective mobile displays is how each device should be directed into position. There appears to be no common mechanism of sending and receiving instructions. A remote guidance mechanism is especially useful if the instructions are to be sent to many devices, and

[1]Samsung Group Play App - `goo.gl/O94hXV`

over a large geographical area. Some researchers have experimented with remote guidance involving humans controllers before [8], though mainly to aid visually impaired people. In this case, the controlling user sat at a computer, viewing the scene through a camera and sending verbal instructions to the impaired person via a GSM voice call. As we are interested in sending instructions to many devices at once, this is perhaps not the best mechanism to use here. For navigation over a small area, arrows have been found to be a simple, unambiguous, egocentric instruction which are easy to interpret [37]. Arrows have been used in museum navigation before [152], though projected from a handheld pico projector. A similar approach was used to that mentioned by Chittaro et al. [37], where a photograph of a landmark was overlaid with an arrow.

6.1.1 Directing extravagant, expressive place-based experiences

We were interested in developing a remote guidance mechanism for the positioning and control of portable media hotspots (PMH). In contrast to previous examples that use the mobile device for collaborative input [50], we focus on the use of these devices as an extravagant, scalable, fully controllable, distributed display for use in public spaces. Throughout the chapter, we refer to the mobile devices that form this display as *portable audio hotspots* (PAH). Portable audio hotspots are a subclass of portable media hotspots, focusing on the audio as the amplified effect. The location of, and sound produced by each PAH should be able to be controlled individually, or as a group. This task is overseen by a *controlling user* who has control over all connected hotspots. While PAH devices receive and play sounds accordingly, location and directional instructions are to be interpreted by the human user. For such a framework to remain usable, the direction mechanism must be simple, effective and robust for both the controlling user and the hotspots. Considering guidance techniques that allow the specification of a destination and/or explicit directions, we present a novel framework for the positioning and control of portable media hotspots in outdoor public spaces.

Our novel direction mechanism makes use of a simple set of tilt gestures. This method was designed the give the controlling user a high level of granularity, providing a specific direction to move a hotspot in as opposed to an end location. The acceptability of gestures in public is a topic that has been discussed in depth by Rico et al. [111]. We have have chosen to use the tilt gesture specifically as it is a subtle but noticeable gesture, and we believe its egocentric nature provides a strong correlation between the direction the device tilts and the direction the user receiving instructions should move in. It also provides an eyes-free method of direction, where the controlling user can focus their full attention on the hotspot.

Unlike many mobile location guides [69, 78, 94], we have chosen not to use audio as a means for navigation, but purely as a means of creating a public soundscape. For directional instructions, we have chosen to use arrows, though unlike previous work [37, 152], we do not overlay the arrow over an image – we display an arrow that can rotate 360° onscreen, always pointing towards the direction to move in. Robinson et al. discuss haptic feedback for use in navigation [116], though we chose to use haptic feedback to alert the user of a new instruction – long vibration, and when one ends – short vibration bursts. This means that hotspot users only need to look at the screen once per instruction.

To illustrate our imagined use case for this technology, we present the following scenario – a user taking part in a remote controlled soundscape. A controlling user orchestrates the experience, with visitors acting as performers in a soundscape.

Scenario. *John decides to visit a medieval castle for the day. Upon arrival, he registers his interest in being part of an ongoing interactive soundscape exhibit at the site. Later on, John's device awakens with a buzz, presenting an arrow with a direction to walk in. He follows the instruction. While walking through a crowd, his mobile device begins to emit the sound of a raging battle. John continues to follow the arrow presented on the device. After following the directions for some time, John arrives at a courtyard. The directions stop, and the sound on John's device has now changed to a medieval market place. John stands and listens as other visitors walk by, their mobile devices producing a whole range of other sounds. At this point, John notices that some other visitors are being directed to stand near him. Suddenly, a trumpet fanfare erupts from these visitor's devices. Along with others, John then receives the same next instruction. They proceed to walk through the castle grounds together, emitting a fanfare as they pass-by.*

6.2 Developing a direction mechanism

This first section in the chapter discusses the design and evaluation of our initial remote guidance prototype, *Director*. We begin with an explanation of the prototype, from both the controller and hotspot's point of view. We then report on an experiment we conducted to test the accuracy of the guiding approach, and also its effect on preferred walking speed.

6.2.1 Prototype design: phase I

Director is a mobile application that runs on the Android platform. To guide a PAH, the controlling user must carry out tilting gestures (see Fig. 6.1). We chose simple gestures as a control as they feel like a natural mapping, and can be done without looking at the screen, allowing the controlling user to focus their full attention on hotspots they are currently controlling. Hotspots receive directions as an arrow on screen (see Fig. 6.3a). An arrow also appears on the controlling user's device, though only in the direction of the instruction they are currently sending.

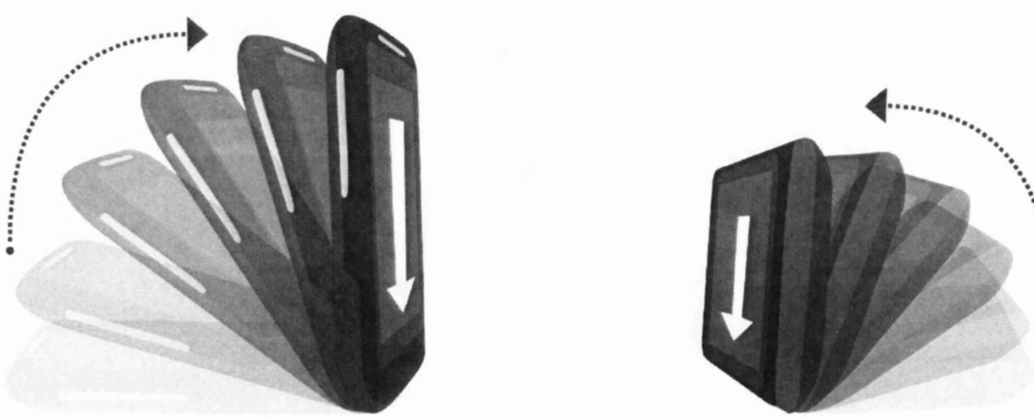

Figure 6.1: *Left:* Tilting towards - sending an instruction to walk towards the controlling user. *Right:* Tilting left - sending and instruction to walk left of the controlling user. The controlling user can tilt their device backwards or forwards to make a hotspot walk towards or away, and left or right to turn left or right.

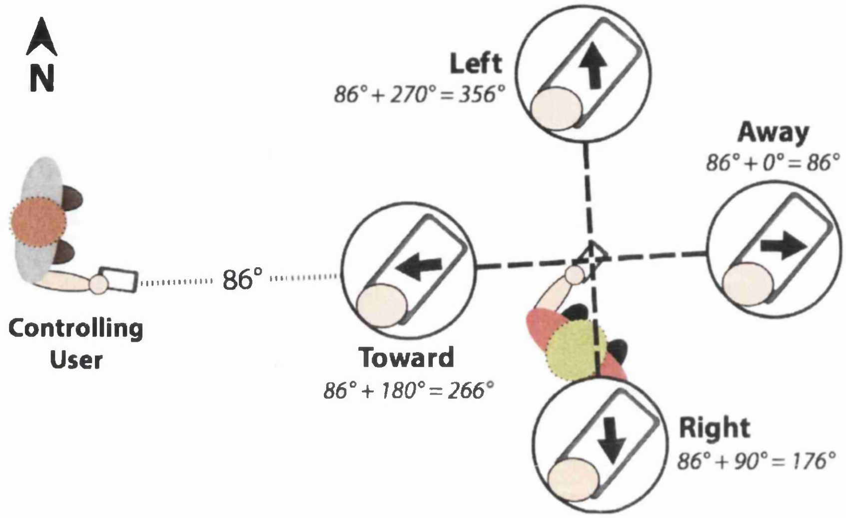

Figure 6.2: Tilt gesture directions the hotspot receives are based on the bearing between the controlling and hotspot devices. All instructions are egocentric to the controlling user.

The tilt gesture instructions work by first calculating the bearing from the controlling device to the hotspot device, for example 86°, and then adding 0°, 90°, 180° or 270° depending on whether the instruction is away, right, toward or left (see Fig. 6.2). Then, after taking the hotspot device orientation into consideration using the digital embedded compass, it displays an arrow pointing towards the resulting direction.

To send sounds to a PAH user, the directing user must click on the controller icon seen in Figure 6.3a, top left. This reveals a list of sounds that the controlling user can send (see Fig. 6.3b).

To test our first *Director* prototype, we conducted a controlled experiment. The focus was on testing the new tilt gesture guidance mechanism, measuring the guiding accuracy and effect on preferred walking speed. We also briefly discuss a deployment pilot study, testing the full system in a public context before the development of a second prototype.

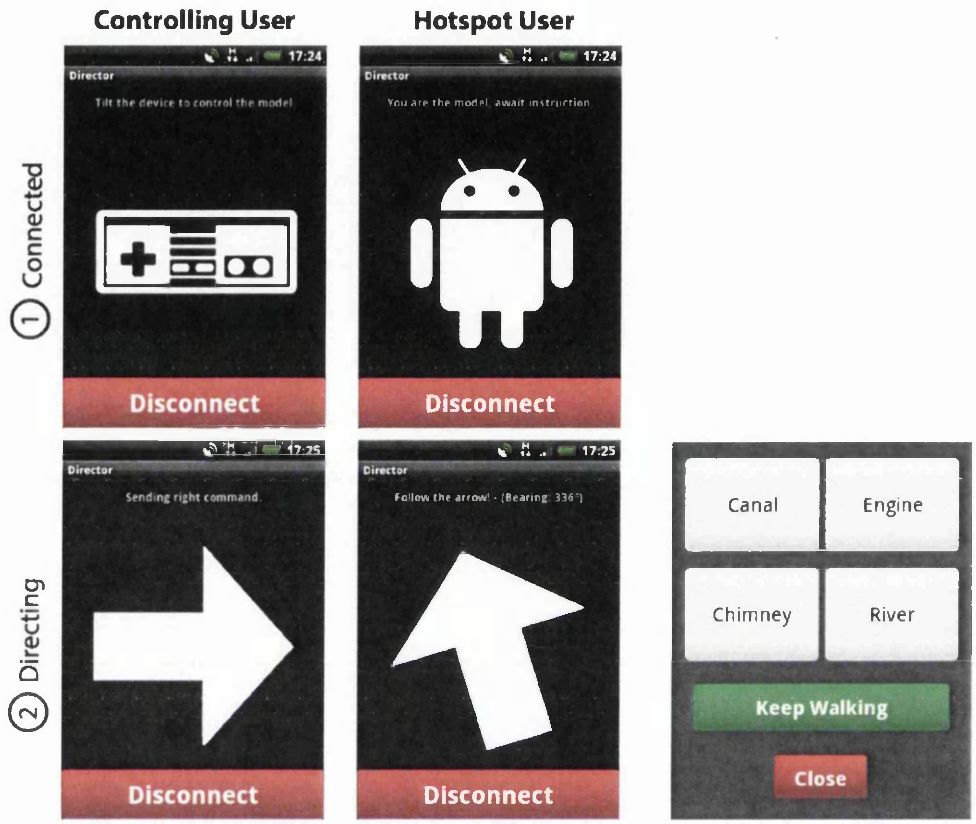

(a) Remote guidance. (b) Triggering audio.

Figure 6.3: The *Director* mobile user interface. *Left:* The flow of interaction while sending/receiving remote guidance in the initial prototype. *Right:* The popup interface containing a list of sounds that the controlling user can send to the hotspot user.

6.2.2 Experiment: system effect on speed and accuracy

As our tilt gesture guidance mechanism with arrows was a novel, untested concept, our first study focused on determining how well this kind of navigational approach could work.

There have been numerous cases where researchers have measured the capabilities of mobile guiding technologies when walking [59, 104, 105]. Goodman et al. [59] give a number of important pointers on using field experiments to evaluate mobile guides. They suggest following participants from a distance as not to influence them – something we consider in our experimental design. Another important suggestion is "the extent to which the use of a device disrupts normal walking." This can be measured in percentage preferred walking speed (PPWS), taking into account first the participant's preferred walking speed. PPWS has been used in a number of different experiments [104, 105], one of which evaluated the effect of a gestural audio interface on walking speed by timing laps of a circuit mapped out by cones. Bearing all of this in mind, we developed a collection of tasks that would allow us to determine the PPWS of hotspot users and the accuracy to which they could be controlled.

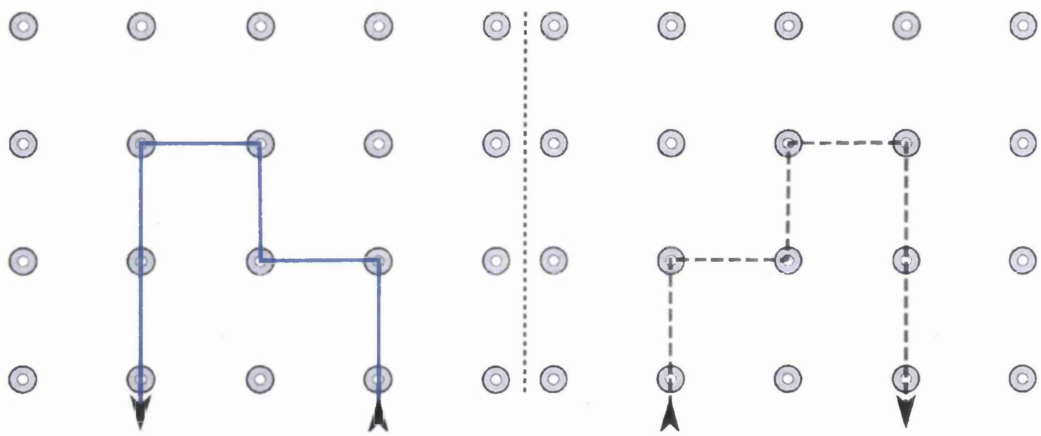

(a) Blue ribbon path (b) Mirrored *Director* path

Figure 6.4: Paths used to determine PPWS when being guided by gestures.

To determine user's PPWS, we first had to capture users normal Preferred Walking Speed (PWS), as well as their walking speed when using the system. To determine a user's PWS, we arranged a set of sports cones in a 4x5 grid formation, 4 metres apart from each other and defined a path through them with a blue ribbon. The study was conducted on a flat field and the distances between each of the cones was accurately measured between each study session. As the path traversed 6 sides inside the grid, the ribbon path was 24m long. Each participant was asked to walk along the blue ribbon from start to finish at their normal walking speed (see Fig. 6.4a).

The second part of this task involved introducing the gesture guidance mechanism to participants. Participants were briefed on how the system worked and were asked to assume a specific starting position on the grid. Participants were asked to begin the task as soon as they received the first instruction. The second path was a complete mirror image of the first, ensuring that the path complexity remained the same, but without any learning effects (see Fig. 6.4b). Participants were not told that this path was a mirror of the first.

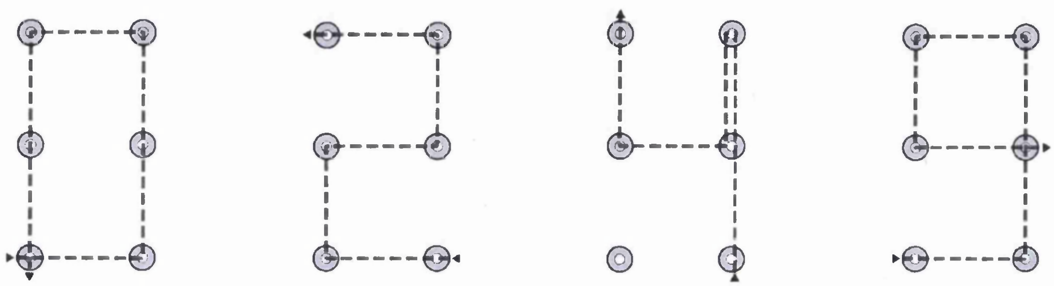

(a) Zero-path (b) Two-path (c) Four-path (d) Nine-path

Figure 6.5: The four 7-segment numbers used as paths. An extra layer of cones around the edges allows errors to occur.

The third and final task was an attempt to understand how accurately participants could be directed around complex paths in a small area. The paths we used were the sides of 4 different numbers from a 7-segment display (see Fig. 6.5). In this task, we measured deviations in two different ways. When participants deviated less than a cone away from the correct path, but then righted themselves by rejoining the correct path, we recorded a 'self-correction'. When participants incorrectly followed an instruction by walking to the wrong cone, we recorded an 'error'. In the instance that participants made an error, they were redirected back to the last correct cone on the path, and continued from there.

In an attempt to keep the level of guidance skill fair throughout all participants, the same expert user was used as the controlling user in all tasks. This user was an investigator. For each of the sessions, the controlling expert user stood around 30m away from the grid (see Fig. 6.6), increasing the distance between both devices in an attempt to compensate for GPS inaccuracies.

When participants finished the tasks they received a questionnaire, followed by an interview to discuss the overall experience. All tasks in this study were video recorded for later in-depth analysis of timings, errors and general user behaviour.

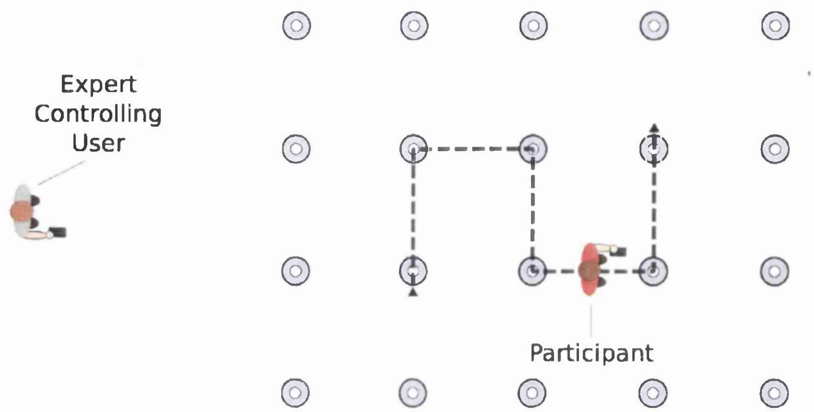

Figure 6.6: Expert Controlling User guiding a participant around the two-path. The numbers 0, 4 and 9 were also used as paths.

In total, we recruited 20 participants (10M; 10F) with an average age of 21 (Min: 18; Max: 26). All participants were University students from a large range of subject areas. Each study session lasted around 30 minutes, and participants were given a £5 Amazon voucher for taking part.

6.2.3 Findings

The results have been divided into four subsections: the effect on walking speed, guiding accuracy and user feedback.

Percentage preferred walking speed (PPWS)

When timing participants' walk along the 24m blue ribbon path, the average time taken was 18.86s (Min: 14.2s; Max: 24.5s). Using a simple $speed = \frac{distance}{time}$ calculation, the average PWS for participants was 1.3m/s (Min: 0.98m/s; Max: 1.69m/s).

When timing participants' walk along the 24m path delivered as instructions through *Director*, the average time taken was 24.32s (Min: 18.2s; Max: 31.3s). Participants' average walking speed when using *Director* was calculated as 1m/s (Min: 0.77m/s; Max: 1.31m/s). When using the *Director* system, the average participant's PPWS was calculated as 78.35% (Min: 59.1%; Max: 92.68%). Comparing each participant's results, their walking speed when using *Director* was significantly lower than their PWS ($p < 0.0001$; paired t-test).

Of the six instructions given, on average, each participant paused around two to three times. Here, participants were either waiting for the next instruction to arrive, or confirming it before beginning to walk again. If we take these pauses into account, focusing only on the time that participants were walking, the average time it took for participants to complete the path with *Director* was 21.26s (Min: 17.3; Max: 29s). With these new times, participants' average walking speed was 1.15m/s (Min: 0.83m/s; Max: 1.39m/s). This means that if we only consider the time that participants spent moving, the average PPWS for participants was 89.16% (Min: 71.36%; Max: 111.22%). These individual results were also significantly lower than participant's PWS ($p < 0.0001$; paired t-test). The results are summarised in Table 6.1.

	Average		
	Time (s)	**Speed (m/s)**	**PPWS (%)**
Ribbon	18.86	1.3	100
Director	24.32	1.0	78.35 (σ 9.58)
-without pauses	21.26	1.15	89.16 (σ 9.15)

Table 6.1: Results of the PPWS task.

Accuracy of approach

Of the four 7-segment number paths used, every participant managed to complete the number zero-path and two-path with no errors. For the number zero-path, we did observe one deviation, though it was self-corrected. This meant that after 120 instructions, 1 (0.83%) had been mis-interpreted at first, but the participant corrected their direction after walking a step or two. We observed a similar self-correction rate with all paths, with 2% of instructions on the two-path being self-corrected, 3.33% on the nine-path and 4% on the four-path.

The only deviations that were not self-corrected–which caused errors–occurred on the four-path and nine-path. When following these paths, we witnessed 2 errors out of 100 instructions (2%) and 3 errors out of 120 instructions (2.5%) respectively. In each of these instances, participants left the correct path. From our results, considering all instructions, there were 16 deviations in total (3.64%). Of these deviations, 68.75% were self-corrected and 31.25% became errors. These results are summarised in Table 6.2.

		Deviations	
	Instructions	**Self-corrections**	**Errors**
Zero-path	120	1 (0.83%)	0
Two-path	100	2 (2%)	0
Four-path	100	4 (4%)	2 (2%)
Nine-path	120	4 (3.33%)	3 (2.5%)

Table 6.2: Total number of instructions and deviations over all participants.

Feedback

In the questionnaire, enjoyment, usability, comfort and effectiveness as a navigation technique all ranked highly among participants (see Table 6.3).

	Avg. Score (/7)
Enjoyment	5.8 (σ 0.70)
Usability	6.1 (σ 0.79)
Effectiveness	5.7 (σ 0.86)
Comfort	6.25 (σ 0.79)

Table 6.3: Results of the 7 point likert-like scale post-study questionnaire; 1 being very low and 7 being very high.

In the interview phase, participants commented that it was fun and simple following the arrow, not knowing what direction would be presented next. One participant remarked, *"it made me feel like a kid again."* Regarding the usability of the application, one participant went as far as to say, *"it's just an arrow on the phone, you can't get much more simple than that."* Another participant claimed that the system would be easy for someone who is terrible with directions, saying, *"This is a like a live thing, telling you to stop and start. I'm terrible with directions and managed that, so it must be pretty good."* Participants were very keen on the idea of remote guidance, with one claiming that it was comforting that another person on the other end was directing you. One participant did mention however, that having to look at the device screen a lot impacted adversely on their enjoyment. The same participant went on to voice their concerns regarding the possible dangers of affixing ones attention on the device completely.

Many participants remarked that the directions given by the application did not always line up directly with the cones, so they were forced to make sometimes difficult decision on which cone to walk to. An issue was also raised regarding the timing of instructions, with some arriving just before cones and others arriving just after, leaving participants debating whether to stop walking or continue to the next cone. A number of participants suggested priming the user with the next instruction to expect, just a few seconds before.

All participants believed that the system was effective in achieving simple navigation, though some questioned it's ability in the real world, where instructions are *"not just left or right."* Another issue raised was that of the direct line of sight being broken between the controlling and the hotspot user. In these cases, the controlling user would be blindly directing hotspots. Some participants suggested ways in which this limitation could be overcome, including a map of user locations to make reference to, or even a similar solution to that seen in [8], where a live video feed from the hotspot user's mobile camera is streamed through to the controlling user.

6.2.4 Discussion

Generally, all participants enjoyed the experience of being directed, finding the system comfortable, usable and effective as a guiding technique. The main concerns from participants were regarding a delay in receiving instructions, and directions sometimes not lining up with the cones. Although this posed a slight problem in our task, it would not present such a problem in the real world as there are no such visual markers. One participant raised concern with the 90° directions, commenting that there would not be enough control in real life situations. In the next iteration of *Director*, we will introduce increased granularity, featuring 45° increments.

When looking at the effect of the system on PWS, we can conclude that there is a large effect (1.61; Cohen's d). There were times during the task where participants remained stationary at a cone. This was attributed to either waiting for the next instruction to arrive—sometimes due to connectivity lag—or pausing to confirm before they committed to the next instruction. Though we cannot omit lag entirely, we can speculate user's PPWS if instructions were received and followed simultaneously. If we negate the time where participants were stationary during the task, on average, PPWS was nearly 90%. Considering participants followed spontaneous sequential directions, we believe this small disruption to users PWS, even with the pauses (78.35%), is an acceptable figure.

The results of the accuracy task show a very low number of deviations, most of which were self-corrected by the participants themselves. After studying the errors in detail, it appears that every error occurred by participants continuing to walk in the same direction as the last instruction. We believe this was also due to lag, with participants still following the previous direction. Although the error rate was low, we believe that again, if the intermittent lag did not exist, we would be able to achieve an error rate much closer to zero.

6.3 Public deployment in a visitor centre

In our controlled user study, we learnt that our guidance technique was a fairly accurate remote mechanism. We also discovered a number of potential improvements that could be made to the technique.

In this section, we discuss our second iterative prototype, adding new functionality, as well as retaining and improving all of the features in the original. We deployed this second prototype in a public visitor centre, exploring the experience of directing a performative experience in a public space. Our experiment explored two different aspects; the usability of the system and the effect on people's experiences during the performance, including the controlling user, hotspot users and bystanders. We conclude with a discussion of our findings, and a set of design recommendations for remotely controlling portable media hotspots.

6.3.1 Pilot study at the Hafod-Morfa Copperworks

Before refining our system further for public deployment, we first conducted a pilot study at the Hafod-Morfa Copperworks. By conducting this pilot study, we hoped to gain a better understanding of how the system and its instructions would be interpreted in a real world scenario. This pilot study involved 8 participants (2M; 6F), testing both the audio and guidance mechanism together. The task involved an expert user directing participants on a guided audio tour of the heritage site (see Fig. 6.7).

Figure 6.7: Participants being directed in the pilot study.

In total, we conducted four study sessions, each of which involved an expert user directing pairs of participants on a path around the site (see Fig. 6.8). At each point of interest, the expert user would trigger the matching audio snippet.

These tests suggested that the sound produced by the Android device loudspeaker was not audible enough in public spaces. This was demonstrated by participants holding the devices up to their ears to listen. It was also apparent that more flexible control was required, with the ability to send hotspots off to specific locations without defining a full path through gestures.

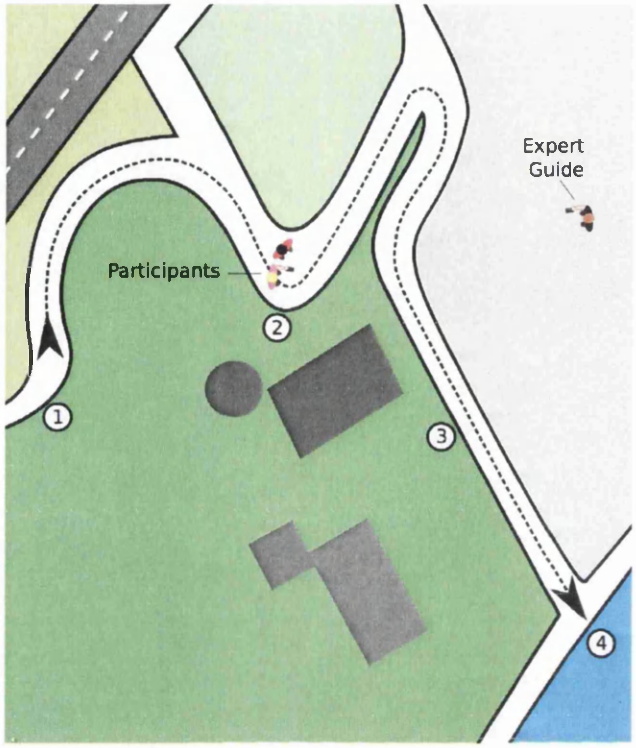

Figure 6.8: The four points of interest that participants were guided through: *(1)* Canal bridge *(2)* Large chimney *(3)* Engine house *(4)* River quay

6.3.2 Prototype design: phase II

Our second prototype was an iterative version of our initial prototype (see Section 6.2.1), retaining and improving all of the features. Considering participants' opinions and concerns from the first study, a number of changes were made. To compensate for the background noise apparent in our pilot test, we chose to attach small loudspeakers to the PAH devices (see Fig. 6.9).

In terms of gestural control, 45° intervals were introduced as opposed to 90°. These shorter intervals were added to give the controlling user a greater range of directional control. Participants in the first study were also concerned about controlling hotspots out of their direct line of sight. For this reason, the controller interface was changed to an interactive map, marking the PAHs (see Fig. 6.10).

In the controller interface, each available PAH appears as a grayscale android figure on the map. To carry out any kind of interaction with a hotspot, the

Figure 6.9: The loudspeaker attached to PAH devices.

controlling user must select the hotspot on the map. Once selected, the android icon turns green. Any number of PAHs can be selected at the same time. The interface details the number of connected and selected PAHs at the top of the screen. This information is especially useful if hotspots exist outside of the local map view. The audio snippet selection menu was also updated to include more options (see Fig. 6.10b). The menu is accessed by tapping the blue dot – the current location of the controlling user. Once open, the user selects the speaker, and then chooses an audio snippet to send. These snippets are presented as small icons, each representative of the sound they trigger. For example, the sword icon triggers a battle audio snippet. All snippets continue to play on loop until they are instructed to stop. This can be done by selecting the mute button on the radial menu. Each PAH has the ability to play the same, or different sounds simultaneously, depending on which ones are selected.

Another feature implemented was the ability to drag and drop PAHs on a map (see Fig. 6.10a). This was introduced to allow the controlling user to send hotspots on longer, unsupervised paths to specific locations. This is achieved by first selecting a PAH on the map, and then dragging it to the end destination. Only one hotspot can be dragged at a time, though all can be dragged sequentially. Hotspots that have received these kinds of instructions appear with a green trajectory towards their intended destination. Trajectories can be cancelled by selecting a PAH and then the cancel location icon in the radial menu. The tilt gesture mechanism works as in the first prototype (see Fig. 6.2), though only when a PAH is selected. With the tilt gesture, all selected hotspots can be directed simultaneously.

For hotspot users, a green coloured background was added to denote when they had been selected. This was an attempt to alert them that they were likely

(a) Drag and drop trajectory - the green line denotes the path to the destination.

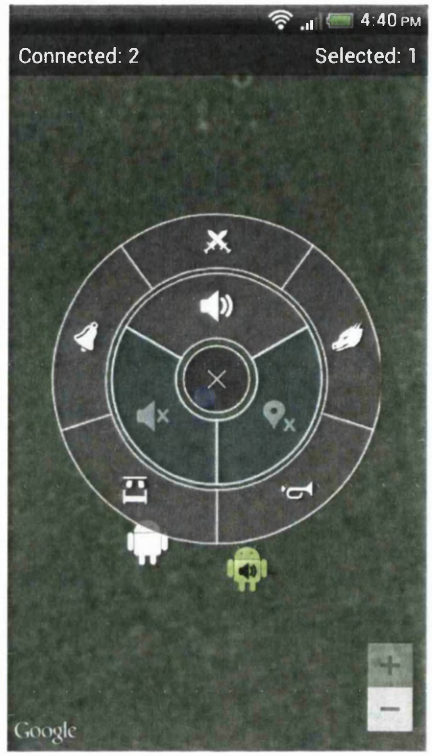

(b) Audio snippet menu - a battle, a dragon, a fanfare, a market and a church.

Figure 6.10: The controlling user's interface. Selected PAHs appear green, and their current state is overlaid as an icon – on the left, a place marker is shown to indicate that an instruction is being sent and on the right, a speaker is shown to indicate that an audio snippet is being sent.

to receive a new instruction soon. This was something that participants requested in the first study. Apart from this change, the PAH application remained very similar to the first prototype (see Fig. 6.3a).

6.3.3 Study: testing the full system in a public visitor centre context

Our second study was an exploratory study, focusing on how the *Director* system would be used in a public context. Specifically on the audio front, we were interested what kind of effect the triggering of sounds in public spaces would have on controlling users, hotspots and bystanders. We were also interested in how hotspot users react to and deal with social situations that may arise from navigating crowds and emitting different sounds. From a technical perspective, we were interested to see how our guidance techniques would stand up in a real world scenario, where there are crowds, obstacles to negotiate and other uncertainties.

The study was designed to simulate a group of friends visiting an attraction, all taking part in an interactive soundscape, similar to that in our scenario (see Section 6.1.1). The site chosen was the grounds of a 12th-century Norman castle, on a free-entry day to ensure that the location was busy with other visitors. Each session involved one group of 3-4 participants.

Upon arrival at the site, participants were given a device each and introduced to *Director*. Participants were told that there was no specific task, other than to use the system to create a soundscape. Unlike the previous study, participants were given a chance to become the controlling user. Once the group of participants had chosen which one of them would be the controlling user, the other participants dispersed around the area, participating as PAHs. When ready, the controlling user proceeded to guide the hotspots around, triggering and stopping the audio snippets at their leisure. Every group repeated the session once again, but with a different participant as the controlling user. This gave more than one participant in each group the chance to be the controlling user.

In total, we recruited 11 participants (8M; 3F) with an average age of 29 (Min: 20; Max: 53). These participants were organised into three groups, meaning there were six different sessions. Two of these groups were pre-recruited University students and staff, and the other was a group of visitors recruited in-situ. Of the participants, 7 reported listening to sounds in public places using their mobile device, all but one using headphones. This participant used the device loudspeaker to watch videos on the web. Each participant was given a £10 voucher for taking part. None of the participants had taken part in the first study.

Similar to other public exploratory studies [157], we observed participant and bystander behaviour with the help of video, concluding each session with an audio-recorded semi-structured interview. Like other location-based services studies, we also logged device usage data [116, 147], allowing analysis of usage.

6.3.4 Findings

The results of this second study have been divided into three subsections: the observations made during each study session, the device log data captured and the semi-structured group interviews that ensued after each session.

Observations

During each study session, it became apparent that controlling users were placing a large amount of attention on the device screen. Throughout each session, there were, however, two recurring situations when controlling users would look up from the device. The first was when a PAH was near to its intended destination, and the controlling user switched to the tilt gesture guidance for fine control. The second was when the controlling user was about to trigger an audio snippet, and would look towards the PAH and other bystanders for a reaction. Much of the time, however, controlling users sent hotspots out of their direct line of sight, using the map as their reference. Around half of the hotspots followed directions exactly, jumping over walls and down grass verges. Others took a more improvised approach, using paths and alternative routes to avoid obstacles. Bystanders also sometimes caused an obstruction. In one session, a hotspot was directed towards a group of people sitting, eating a picnic. The participant looked at the controlling user and shook their head, refusing to walk closer.

Some participants fully embraced the experience, engaging in performative exchanges with one another. For hotspots, this was mainly apparent when they were unable to follow an instruction any further due to obstacles. In these circumstances, hotspots would often raise their hands in the air in a theatrical manner, as if to show the controlling user that they could not follow the instruction, and that they had given up. One particular hotspot participant handled this issue in a different way, purposely walking into the obstacles, perhaps for comedic value – ultimately, attempting to convey the same 'hopeless' message to the controlling user. In this same session, the controlling user also exhibited performative behaviour, exaggerating their tilt gestures so that the hotspots were able to see.

Much of the time, the controlling user sent PAHs on different paths to different locations, with different sounds. We did however, witness one occurrence where the controlling user trigged the same audio snippet for all hotspots and sent them to the same location. All of the hotspots converged on the destination playing the church audio snippet. To bystanders, this would have sounded like a group of monks coming together and praying. There were numerous occasions in each session where participants crossed paths briefly. When close enough, these participants acknowledged each other, often engaging in conversation regarding the experience. This same behaviour was also exhibited when hotspots crossed paths with the controlling user. Most of the time this occurred among hotspots, though hotspots also crossed paths with the controlling user and exhibited the same behaviour. This happened whether participants knew each other or not. Although the controlling user was ultimately responsible for what directions and sounds were sent to hotspots, in one particular session—when crossing paths with the controlling user—a hotspot was seen taking over the performance, requesting certain sounds and locations to travel to. In this case, the controlling user obliged.

Bystanders did not tend to pay much attention to hotspots being directed around, though they did pay attention to the audio snippets. During a session, one participant commented, *"People are probably wondering, what are these nutters walking around with sounds coming out?!"* In another session, a hotspot was playing the church audio snippet, at the same time being directed through bystanders. One group of bystanders was overheard saying, *"Oh, that's the church*

Figure 6.11: Participant discussing the system with a bystander.

over there is it?" At which point, the hotspot walked past. The group of bystanders then realised that the sound was coming from the device and started laughing. In most instances, bystanders would acknowledge the audio snippets by turning to look, sometimes stopping if walking. There was one instance where a child began to follow a hotspot for around 10 seconds. The same child–along with a group–later stopped the hotspot to ask what they were doing. Similar behaviour was observed in most sessions, with one or two bystanders approaching hotspots and asking questions (see Fig. 6.11).

Device logs

Each session lasted an average of 8 minutes and 56 seconds. Of this time, hotspot participants spent an average of 7:20 (82%) receiving either an audio snippet or a direction, and 1:36 (18%) receiving neither. Individually, on average, hotspots received audio snippets for 4:50 (54%) and directions for 2:48 (31%). The fanfare audio snippet was chosen the most often by the controlling user, and this sound was also played for the longest period (see Table 6.4).

	Frequency	**Duration (s)**
Battle	7	490
Dragon	9	666
Fanfare	15	1654
Market	7	840
Church	11	988

Table 6.4: Total number of audio snippets triggered and play duration over all sessions.

In total, the drag and drop guidance mechanism was used to send 86 instructions, 33 of which were rerouted paths. Of the 53 completable paths, 43 (81.1%) were completed successfully and the remaining 10 (18.9%) were cancelled by the controlling user. The gesture guidance mechanism was used to send 405 instructions in total (see Fig. 6.12). In terms of how many hotspots the controlling user had selected at any one point in time, 1 hotspot was selected 117 times, 2 were selected 63 times and 3 were selected 12 times. The paths of hotspots from one session can be seen in Figure 6.13.

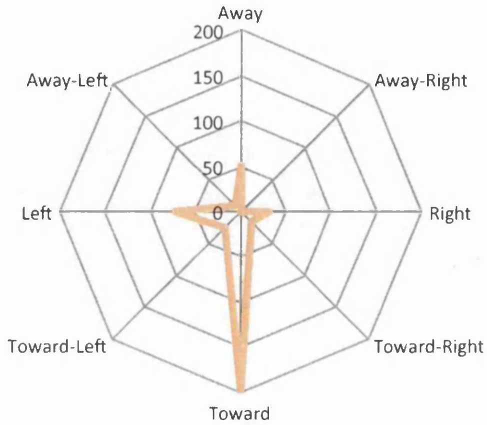

Figure 6.12: Gesture frequency distribution.

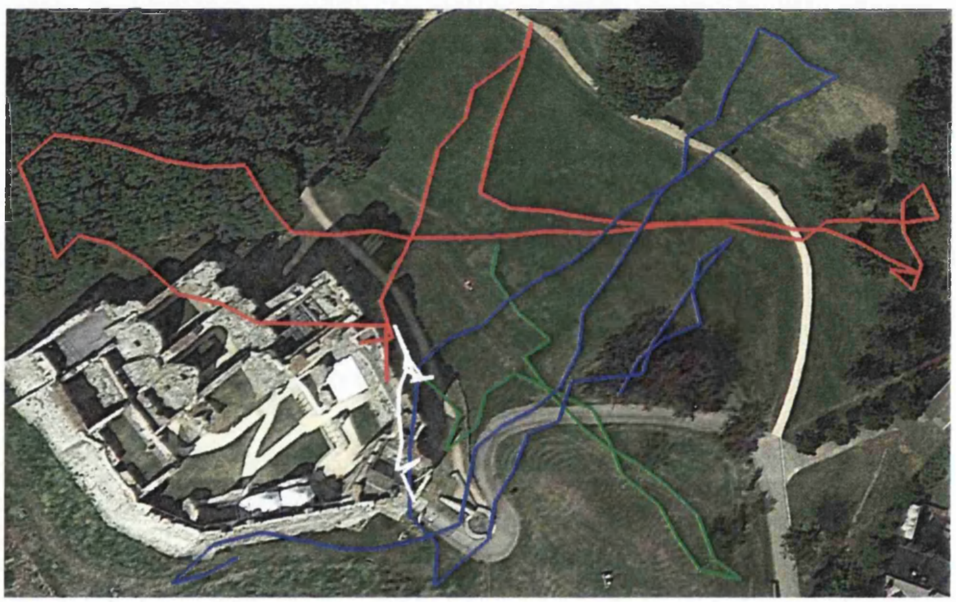

Figure 6.13: PAH paths for one session - red, green and blue. Controlling user path in white. The controlling user generally stood very still, and hotspots were directed extensively, often out of sight.

Semi-structured post-study interviews

All of the hotspot participants said that they were comfortable and confident in the directions given, commenting that they were simple and easy to follow. A number of participants did however mention that there were some instructions they were simply unable to follow due to obstacles. In these cases, participants said they improvised where possible, taking alternative routes. In one group, a participant suggested a mechanism for communicating with the controller when these situations arise. This was contested by another member of the group, arguing, *"I actually prefer it the way is. [..] I mean, you know where you have to go, and you just have to find your way to that point. You will end up there one way or another."*

All participants who acted as controlling users claimed that the hotspots were 'easy' to control. Two of these participants said the most enjoyable thing about the experience was seeing bystanders reactions to the sounds produced by the PAHs. One of these participants remarked, *"I try to send them [hotspots] to crowded places and see how they react... see how other people react around them."* This was something that a number of hotspots enjoyed, one of which mentioned that their favourite part of the experience was, *"Not knowing what sound was going to come out of the phone next."* Participants noted that most bystanders were aware that something going on, though not exactly what. Many participants saw the sounds as a *"trigger for curiosity."* One reported, *"I think they looked up to the sounds. I don't think they minded us walking around, but then like, the bells and the monks. [everyone laughs] Yeah, I got a few funny looks."* Another remembered, *"At one point, these people were having like a stereo experience, with music coming from both sides of them. [..] They looked bemused and then wandered off."* This fascination and enjoyment of gaining bystanders' attention was encapsulated by one participant's comment, remarking, *"You attract attention obviously and then people start talking to you. Which I kind of like, because usually you're very isolated. It makes it a bit more open."*

Role preferences were spread equally among participants, with some preferring the simplicity of being directed and others preferring the 'power' of control. Multiple controlling users confessed to using the system mischievously, with one directing a PAH in circles and another directing a PAH up and down a hill several times. In general, this only occurred between friends and more extroverted individuals, and the hotspots did not mind this. Two hotspots went as far as to say, [P1] *"I quite liked following the arrows..."* [P2] *"...yeah you can just sort of switch off can't you..."* [P1] *"...almost be a child isn't it..."* [P2] *"...yeah, the lack* *of responsibility."* All participants said the experience enhanced their visit to the site. At least one participant from each group pointed out that the directions took them places they would not usually go, with one participant referring to their favourite part as, *"Being sent around the back of the castle [...] It's not that bad, it's quite nice!"*

6.3.5 Discussion

Controlling users spent a large amount of time looking at the onscreen map. We believe this is not necessarily a bad thing, as it gave them a quick overview of the positions and states of all hotspots. Without the map, we presume hotspots would have been much more difficult to locate and interact with, especially when out of sight. It also appears from the results that controlling users most often selected a single hotspot when sending instructions. This and the fact hotspot participants spent 18% of the time standing around with no instructions or sounds, seems to suggest controlling users may have been overloaded, even with three hotspots.

The simplicity of the directional arrow was liked by all, and most of the time, hotspots were pointing in the correct general direction, attempting to follow instructions closely. Improvisation, however, played a large role in public spaces, with obstacles and impassable objects. As many hotspots discovered though, if they walked around obstacles and followed the general direction, they arrived at the same location. In most instances, hotspots did indeed manage to arrive at the intended destination using this technique. This ability to improvise was made possible by the fact that participants saw the arrow as a general guide, which always rerouted based on their current GPS location and device heading.

Director proved a powerful tool for orchestrating an audio performance. The most popular sounds were the fanfare, church and market place – all of which were fairly neutral in comparison to the violent nature of the battle and dragon. This, perhaps, made them more enjoyable sounds to listen to. We witnessed some situations where hotspots were reluctant to fully commit, avoiding social conflict, though on the whole, many participants were happy to play a full, active role in the performance. Two participants discussed the fact that they would not usually walk around randomly, producing these sounds, but the experience gave them the entitlement, comparing it to being a child again. We saw a number of participants exhibiting signs of more expressive, performative behaviour in their actions, with exaggerated tilt gestures, purposely walking into obstacles and expressing helplessness and discontent.

Additional examples of participants using the system to create a performance include triggering the church audio snippet for all hotspots, making them converge together at the same point, and also when hotspot participants reported surrounding a bystander, giving them a 'stereo audio' experience. Many of these encounters did not happen by chance. As controlling users admitted in the post-study interview, they had tried to get reactions from bystanders, directing hotspots near to them. One hotspot said they enjoyed the open nature of the experience, enabling them to start discussions with strangers. These events were caused by the guidance mechanism encouraging participants to cross paths with bystanders, and the amplified audio drawing attention.

During all of our experiments, participants joked about using the system for mischievous purposes. We witnessed one hotspot being directed up and down a hill repeatedly, and another around in circles. We also witnessed controlling users unintentionally direct hotspots off ledges, down grass banks and over walls. The area we conducted our study in was fairly safe, though nothing is to say that hotspots will not be directed towards roads or other hazardous situations in other public spaces. The dangerous potential of this system raises some interesting ethical issues regarding accountability. Before using the system, users should all be made aware of these potential dangers.

6.4 Conclusions

We presented the design and evaluation of a novel, low-cost, low-energy, lightweight and scalable portable audio display using existing mobile devices. Our work over both studies now provides a baseline for handheld, performative, remote guiding technologies such as Portable Audio Hotspots.

Both studies identified the key potential benefits and drawbacks of using *Director*, in a controlled setting and an open public space. In our first study, we discovered a number of important issues with our design, which we were able to address before taking our system into a real world setting. Our gesture guidance mechanism was proven to be a very accurate remote guidance technique, though connectivity sometimes played a part in delayed instructions. We also discovered that mobile device speakers alone were not audible enough in public spaces, so a miniature loudspeaker was added to each of our devices for amplified effect.

Our public deployment study showed that *Director* could be used to create a performative experience, with many participants happy to assume this role. Generally, participants enjoyed the experience and got used to playing the different roles very quickly. In terms of the controlling and hotspot roles, preferences were split equally among participants, with some enjoying the power of control, and others enjoying the social experience and lack of responsibility when being controlled. The dynamic and amplified nature of the system directly affected the experience of many spectators. It also promoted a very activate engagement among participants and bystanders, with many approaching to ask questions.

Although we have only touched upon the use of mobile devices as portable audio displays, this kind of framework could be extended to include many other mobile modalities, such as the pico projection used in Chapter 5, or haptic feedback. We define this broader design space as Portable Media Hotspots (PMH), and believe that there are many more interesting displays left to experiment with.

6.4.1 Designing for remotely controlled performative place-based experiences

From the iterative design and evaluation of the prototypes presented in this chapter, we suggest a number of key design implications when designing remotely controlled performative place-based experiences. Although we experimented with audio as an effect, we believe that these design implications are applicable to the wider portable media hotspot design space.

In the experiments reported in this chapter, we have once again seen the engaging and immersive qualities of amplified effects in public spaces. However, the remote direction and movement of multiple hotspots provides an extra dimension to mobile performative experiences, further enhancing enjoyment and engagement. Controlling users enjoyed directing hotspots around extensively, provoking reactions from both hotspots and bystanders. Hotspots enjoyed the lack of responsibility when receiving and following directions, visiting new areas and engaging directly with bystanders. However, the added direction mechanism introduced a new potential issue, with the possibility of hotspots blindly following directions into hazardous or dangerous situations.

We have categorised the results of our studies into four key findings. These findings form key design implications for scalable, performative mobile experiences that employ remote guidance.

Control Simple tilt gestures proved to be a useful eyes-free method of control, especially when directing hotspots near to their destination – allowing a finer level of control. However, gestures alone were not enough. The map and drag-and-drop functionality were vital to the controlling user being able to keep track of the overall situation, and also directing hotspots more extensively on unsupervised paths.

To further reduce cognitive load for controlling users, we suggest increasing the number of controlling users as the number of hotspots increases. We propose that each controlling user would then have their own mobile device, or perhaps share a tablet between multiple users.

Improvisation Use live directions that afford improvisation. In our controlled study, participants were asked to follow the directions to the next cone – which was not a problem as the study was conducted on an open field. In a busy public space however—with impassible obstacles and hazardous situations—hotspots were not able to follow all instructions. In this respect, an egocentric, live and responsive arrow was a simple direction, allowing users to see the general direction they should head in. If it was not possible to head in the exact direction the arrow was pointing in, participants were able to improvise by rerouting.

Engagement Similar to findings in our performative projection experiment, we have confirmed that placing the amplified effect on or around the performer promotes a more active engagement between bystanders and the performer.

However, the introduction of a direction/guidance mechanism saw an additional factor for promoting a more active engagement. The dynamic movement element drove hotspots into bystander's personal space. This, combined with the amplified audio element, made the experience even more accessible to bystanders. We encourage designers to be mindful of these kinds of interactions, in the hope of facilitating a more social and engaging experience for all.

Ethics There is an inherent danger to this kind of remote guidance technique, with controlling users potentially directing hotspots into hazardous situations – both intentionally and unintentionally. If a hotspot leaves a controlling user's field of view, or if a hotspot follows a controlling user's directions verbatim, hotspots could quite easily find themselves in hazardous situations. A mutual understanding is needed, requiring confidence from the controlling user, trust from the hotspots and a certain level of cautiousness from both. Although we do not have a definitive solution for this issue, designers of these kind of guidance techniques must be aware of and carefully consider the ethics – which may differ depending on the location and context.

Chapter Seven

Attitudes, Integration & Regeneration

In this chapter we discuss communal attitudes towards our technologies and performative approach. Along with our previous findings, we take on board participant opinions, considering how these technologies may be integrated into a public space such as the Hafod-Morfa Copperworks to facilitate regeneration.

We began our investigations with a focus group comprising of members of the community interested in the *Cu@Swansea* project. In this focus group, we demonstrated our performative audio and visual prototype interactions. Participants were encouraged to discuss the interactions amongst themselves and with the researcher, making comments and asking questions where appropriate. We later held another focus group, this time with the *Cu@Swansea* project stakeholders. Again, a number of performative technologies and concepts were demonstrated within this meeting, giving the project stakeholders an opportunity to see and hear of the kinds of experiences we had been designing, also allowing them to voice any opinions and potential concerns.

Further public engagements were carried out during deployments, with visitors taking part in open-ended discussions on the use of performative technologies, touching on the concept of sharing personal media. Considering our combined engagements, we conclude with a set of guidelines to help promote the social regeneration of the Hafod-Morfa Copperworks, and also similar public spaces that could benefit from an enhanced, more active engagement between visitors.

7.1 Introduction

The classic way in which heritage sites—similar to the Hafod-Morfa Copperworks—usually operate, is to either preserve the ruins as they are, or to attempt to restore them and turn them into facilities. In most cases, visitors to the site can then acquire additional information through textual, visual and audio interpretation guides that may be available. In our case, our goal was not to develop a set of interpretation applications, but to design a set of performative interactions and experiences that would help encourage a more active engagement among visitors to the site – to make the site a pleasant, enjoyable and usable space for the public. By developing these technologies to be more performative, we hope to give visitors a more expressive medium to evoke memories, enabling them to connect with the site and with others.

At the Hafod-Morfa Copperworks, we started with a blank canvas. Throughout this research, this canvas has continued to evolve, with new interpretation, developments and events. We have been engaged with all parts of the *Cu@Swansea* project throughout, providing us with opportunities to involve stakeholders, the local community and the wider public in the design and evaluation process. This has given individuals and groups the ability to express their own opinions of what they would like the site to become. While the studies we have reported on so far have given us an idea of the usability of the interactions and how they directly affect people, we know less about what people generally think of the idea of extravagant, expressive computing. We also know less about how these technologies will integrate within the site, and if or how people will use them – after all, at such sites visitors are used to being consumers, not producers [6].

The experiences we have demonstrated and deployed so far have been expressive and magical [108], allowing people to take control of amplified feedback. We are also interested in taking the performative concept further, allowing people to share their own media instead of using curated content. This would enable visitors to contribute expressive audio, video, thoughts and memories that they think are relevant to the space, as well as then being able to experience them in an expressive and engaging way. We hypothesise that this emphasis on contribution will promote an enhanced feeling of ownership, and may possibly encourage people to revisit an area to discover new content. The acceptability of this approach is explored at the end of this chapter in our public engagements.

The main focus of this research has not been on evaluating the long-lasting regenerative effect of these interventions, though here in this chapter, we ask the question; how might these technologies be successfully integrated within the site, so that visitors will want to use them and begin to enjoy the effect they have?

7.2 Community focus group

On the 7th March 2013, a quarterly gathering called 'Friends of the Hafod' was held. The purpose of these meetings is for the main bodies involved in the regeneration of the Hafod-Morfa Copperworks to provide an update on the progress of their individual projects. It is also a good chance to network with people who are interested in taking an active role in the regeneration of these sites. On this day, a focus group was held to demonstrate two performative experiences that may potentially be deployed at the Copperworks to facilitate visitor interaction. It was hoped that this session would provide valuable feedback on the acceptability and usability of the prototype technologies.

7.2.1 Focus group

Technologies demonstrated

Two prototype systems were brought to the session – a version of the performative projection system from Section 5.2.1 that we named the *Hafod Torch*, and *Surround You* from Section 4.3.1. The underlying concept of the *Hafod Torch* is similar to our previously mentioned performative projection prototype, but for this session, we removed the QR code functionality from the system and allowed users to choose their animation by pressing and holding a button. In this demonstration application, four animations were included. These animations were chosen carefully and were relevant to the things a visitor would have been able to see at the Hafod-Morfa Copperworks when it was fully functional. The animations that the *Hafod Torch* projected were smoke, fire, sparks (see Fig. 7.1) and water.

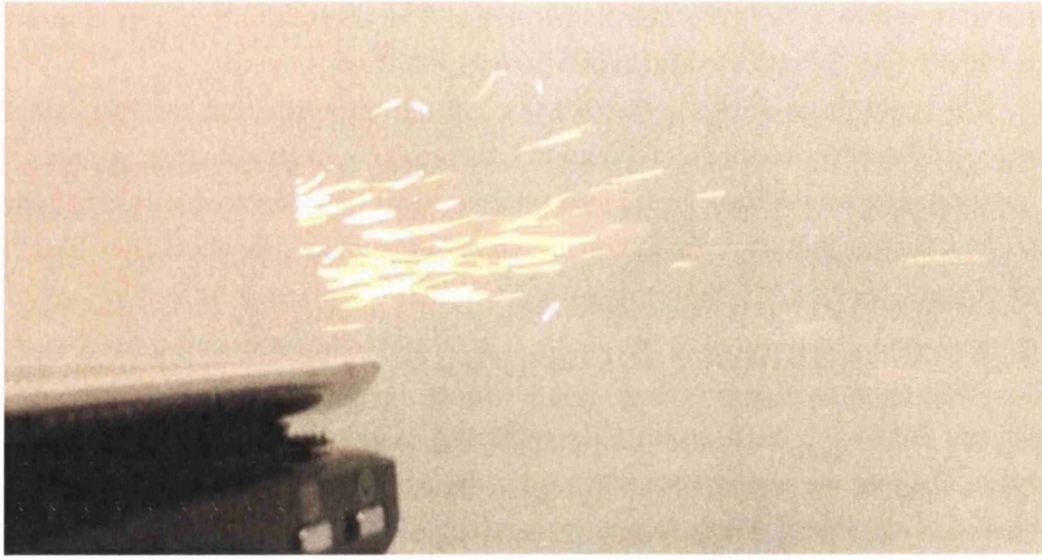

Figure 7.1: The *Hafod Torch* – projected sparks flying from left to right.

Sounds that were used with *Surround You* were closely matched to the kinds of sounds that would have been heard at the Hafod-Morfa Copperworks during its heyday. These sounds included: a busy locomotive shed, an engine house, barges on a canal and boats navigating a river. If this system were to be deployed for public use, these individual speakers would be placed in their respective locations on site.

Participants

During a networking break, members of the meeting were invited to attend a focus group where they could learn about and experience the potential kind of prototype systems that could be integrated into the Hafod Copperworks. Of the 80 people that attended the meeting on the day, around 20 stayed for the focus group session. These participants were generally middle to old aged and came from a range of backgrounds. There were participants who were academics, technologists, engineers and archeologists, but the majority of participants were general members of the public.

Feedback from these participants was gained through informal verbal comments and questions, and also post-it notes that participants were asked to write feedback on at the end of the session.

Procedure

For the duration of the session, a presentation was given. There were two break off points during this presentation where each of the systems were given separate demonstration time. Participants in the session were asked to come and stand at the front of the room so that they could play an active role in the demonstrations.

The first system that participants were shown was the *Hafod Torch*. Each of the animations was shown to the group and participants were told why these specific animations were chosen. During the demonstration, participants were asked for suggestions on what other kinds of things could be projected. A conversation then ensued regarding possible changes and additions to the system to improve it.

The second system that participants were shown was *Surround You*. For this system, a demonstration video of possible usage at the Hafod Copperworks was shown. During the video, participants were given an explanation of how the system was working and why particular sounds were coming from certain buildings. After the video had finished, participants were given a real life demonstration of the system working indoors. Again, continuing with the theme, industrial noises were used. During the course of this second demonstration, participants asked questions and actively contributed suggestions.

Participants who were interested in taking a closer look at the systems were given a chance at the end of the session. At this point, participants were given the opportunity to test the systems for themselves. In total, the session lasted for around 40 minutes.

Findings

It became apparent during the demo session that many participants were interested in being able to create their own media for these systems. Instead of only being able to use animations and sounds that had been supplied with the systems, participants claimed that it may be beneficial to allow users to create and share their own animations and sounds through these technologies. One suggestion of the kind of media that could be captured and shared by users at the site was sightings of plants and animals. This participant also suggested being able to record animal noises.

After seeing the *Hafod Torch*, without seeing *Surround You*, the first suggestion that many participants made was to incorporate sound into the same system. In terms of *Surround You*, after watching the demo video, most participants thought that the sounds came from the phone. Participants were quite shocked and excited when they saw the system being used in real life and realised that the sound came from situated speakers. When the real life demonstration was given of *Surround You*, a number of individuals in the room that were not a part of the session came to see what was going on. In terms of audio choice, participants said that they thought it would be a good idea to hear the voices of workers at the site and for people to be able to share their own stories. One participant explained that they liked both experiences because they were in control of the output. Generally, all participants responded very positively to both systems.

Use for disabled users was mentioned by one participant saying, *"you have to think of these things near the beginning. I've seen these new technologies at sites before and people don't think about how disabled people will use them."* Another participant commented that both systems were *"innovative ideas, easily adapted for those who are visually/auditory impaired."*

Concerns that participants had were with regards to the cost of the equipment and whether the equipment was resilient enough to work at the Copperworks. Another point that a participant raised was the question of how a user who had an older device would interact with these technologies, and if a visitor without a mobile device could be a part of these interactions at all.

7.2.2 Discussion

Although the majority of participants were of an older generation (50+ years), both the *Hafod Torch* and *Surround You* prototypes were embraced by all that were present. All participants that used the systems found the technologies easy to use with little instruction. The ability for users to create and share their own media seemed to be a popular view held amongst most participants. Continually fresh user contributed content may indeed be a good way of engaging visitors and getting them to come back to the site. The only possible difficulty with this contribution mechanism is how the content is regulated. As the Hafod Copperworks site will be publicly accessible to those of all ages, it would be important for inappropriate content to be dealt with somehow.

The fact that many more people began to join the session when the live demonstration of *Surround You* began is another indication that the system is good at drawing bystanders in. The audio in the system is fairly loud and when a person cannot see the audio source, they are bound to be drawn in by curiosity. When designing a performative system for public use, the needs of those with disabilities should be taken into account. This could possibly extend as far as creating remote experiences and interactions for those who are not able to visit certain parts of the site.

Regarding the resilience of the technology at the site, careful considerations will need to be made when designing a final product so that these technologies will continue to work under the tough conditions at the Hafod Copperworks. The site has no internet access or power which may cause potential problems.

During the session, one participant spoke of the possibility of a visitor not having a device capable of interacting with these technologies. As long as a visitor is in the presence of a performer who is using the system, we argue that they are still a part of the experience, but as a spectator. If a visitor does not have a device capable of interacting and there is no one else around that is performing, we will have to think very inventively of how they may remain a part of the experience.

7.3 Stakeholder focus group

To engage the *Cu@Swansea* project stakeholders, we organised a focus group meeting where we could demonstrate the technologies and concepts that we had been researching. It was hoped that this would give us a different perspective compared to that of the community focus group, potentially raising different concerns on a higher, project level. The focus group was held on the 12th May 2014, inviting stakeholders to come along and discuss interpretation at the Hafod-Morfa Copperworks, based on our research to date. In this section we outline the procedure and technologies demonstrated at the focus group meeting, concluding with a discussion of our findings.

7.3.1 Focus group

Technologies demonstrated

In this focus group, we brought our three main prototypes; the *Hafod Torch* (see Section 7.2.1), *Surround You* (see Section 4.3.1) and an early *Director* (see Section 6.2.1) prototype. The *Surround You* prototype was demonstrated in the same proof of concept video used in the community focus group.

Projections used in the *Hafod Torch* were the same as previously mentioned in the community focus group, and the *Director* prototype shown to participants was a display of the ability to connect to and send directional instructions and sounds between devices (see Fig. 7.2). Sounds used in the *Surround You* prototype were all industry based, including a canal, an engine house and others.

Participants

Participants were recruited through an email sent to all *Cu@Swansea* project stakeholders, inviting them to a focus group to discuss the experiences we had designed and developed. The majority of stakeholders agreed to attending the meeting. In total, 9 stakeholders (8M; 1F) attended, averaging 47 years old (min: 23; max: 54). Participants included a project co-leader, the project manager, design and contract admins, interpretation and events planners and other contributing team members. All participants had visited the Hafod-Morfa Copperworks before, and did so on a regular basis.

Procedure

All participants were asked to sign an ethically approved consent form, detailing the aims of the focus group and that the entire process would be audio-recorded.

Figure 7.2: Directional instructions were sent between devices to demonstrate the guiding functionality.

The focus group followed a very similar procedure to our community focus group, with a presentation describing the technologies and break off points to demonstrate the technologies to participants. The first prototype that was shown to participants was the *Hafod Torch*. Then, participants were shown the *Surround You* proof of concept video. The remainder of the meeting focused on *Director*. Participants were encouraged to interrupt the presentation or demonstrations at any time to discuss or ask questions, though much of the feedback came after all of the concepts had been demonstrated. The focus group lasted for around 40 minutes in total.

Findings

From analysing the audio-recorded discussions, we were able to identify three main themes that the project stakeholders were interested in discussing:-

Performativity: Crossing the consumer-producer divide, including teaching and encouraging people to use these performative technologies.

Relation: A comparison of performative and conventional interpretation technologies, with arguments for and against each.

Sustainability: How the prototype technologies may be integrated within the site in a sustainable way so their use may continue.

Beginning with performativity – one stakeholder had initial concerns with the concept of performative technologies. Speaking about *Director* specifically, this participant asked, *"to what extent do you have to brief the user before hand?"* They went on to say, *"you have to have two people, you need the director and the directed, you need to explain what people are doing, you need someone to explain what happens next."* Another participant suggested that this particular remote guidance approach would work well with children being directed by a teacher.

The stakeholder with the initial concerns later went on to say *"it's purely about engaging with the site in a different way ... this isn't about understanding or interpretation is it ... I'm just thinking about people turning up cold to use it, and how you would encourage them to get going with this."* Another participant again interrupted, claiming that this was not such a big problem. Their own past experience suggested that people who visit these kinds of sites do research prior to arriving, allowing them to prepare before hand. This participant continued *"for somebody that wanted to engage with something technical like this, they would probably be aware in the first place and that might be the actual thing that draws them to the site."* The concerned stakeholder later responded, *"it's quite a simple thing isn't it, but it requires a massive leap of the imagination from people using it to actually engage with it properly. It's not about receiving information, it's about doing stuff to the site yourself."* All stakeholders agreed that this technology *"opens the site up to a group of people who may not have visited."*

Participants in the focus group were very keen on comparing the technologies with existing interpretation and experiences they had witnessed elsewhere. Turning to participants discussion on *Director* again, two participants were in favour of removing the remote guidance element, replacing it with a more conventional, automated GPS approach that triggered sounds and events automatically – removing the need for a human controller. One participant retorted, arguing, *"it's all about the interaction though isn't it ... it becomes very passive."* This was a turning point in the conversation, after which, participants recalled a number of more performative experiences that they had encountered, including an unfolding experience through text messages in the Titanic Quarter in Belfast, projections of a swimmer in the Roman Baths in Caerleon, interactive projections in a local leisure centre and other interactive and immersive experiences in the Rhondda Heritage Park. Participants agreed that these technologies allowed visitors to experience these places in a *"playful way."* A stakeholder concluded this part of the discussion saying, *"I think it's great to have the experimental stuff, otherwise it's quite predictable."*

The final theme that stakeholders were interested in discussing was the integration and sustainability of these technologies within the site. Referring back to the participant who mentioned the Rhondda Heritage Park – this participant commented, *"they've got some fantastic stuff there. It worked for the first couple of years and they've never had the money to fix it. It doesn't work anymore and people have lost that as part of the experience."* There was much talk of the robustness of the technologies, being able to continuously work outdoors in different weather conditions. Another large concern participants had was with security, with multiple participants discussing the possibility of peripherals and devices being stolen from the site. Finally, there was talk of refreshing the experience, giving visitors a reason to keep coming back to the site. One stakeholder suggested engaging with schools from a history and technology perspective, capturing sounds and images that could then be experienced at the site through these performative technologies.

7.3.2 Discussion

Crossing the consumer-producer divide presented a large problem for one stakeholder. As discussed throughout this research, many visitors are not used to performing or actively contributing, often expecting such visits to be passive. In an unstewarded site such as the Hafod-Morfa Copperworks, this is perhaps more troublesome, as there will be no one on-site to inform people of the technologies available, or how to use them. To ensure the successful adoption of these technologies, the interactions will have to be well advertised, and also provide clear instructions on how to use them – perhaps through signage or other digital prompts. However, another stakeholder repeatedly made the point that many visitors to such sites plan their journeys, and may already know what technologies to expect at the site, providing an opportunity to prepare themselves. This same participant went on to say that these technologies may even encourage people to visit the site, especially those of a younger generation who would not usually consider visiting.

Many participants were able to recall existing performative experiences they had engaged with in other public spaces. These participants recognised the benefits of a non-passive experience, with one in particular arguing against a more passive, automated guidance interaction. While participants saw the potential advantages of extravagant, expressive experiences, one participant raised the topic of sustainability. With any technology intervention, it may work for a number of years, but will most definitely require additional funding and maintenance to keep running. This is something that is perhaps often overlooked, and should be carefully considered in grant applications and budgets.

Robustness was regarded a highly important factor in sustainability. Our current implementations are only prototypes, so are perhaps less robust than they could be. Further development will be able to fix this issue, addressing any vulnerabilities. Security was also important to many stakeholders, with the possibility of mobile devices, peripherals and other equipment being stolen. We can attempt to secure on-site peripherals and equipment in secluded locations and in secure boxes. If we were to deploy these technologies on-site, we assume that visitors would bring their own devices, meaning we do not need to be concerned with visitors returning mobile devices. Fundamentally, both robustness and security are at the peril of funding. Something else stakeholders mentioned—that would cost far less money—is encouraging the local community and schools to contribute images and sounds to the site, keeping the media available fresh. This would perhaps encourage visitors to revisit the site, and could easily be incorporated into our framework.

7.4 Further public engagement

Our engagements so far have focused on asking both the community and project stakeholders on their views of how performative technologies may fit within the Hafod-Morfa Copperworks. A common, recurring theme within these engagements has been the ability to share media at the site. The community wished to share media, perhaps allowing them to leave their mark on the site, creating a feeling of ownership. Stakeholders suggested the same functionality, though to create fresh content and to keep drawing visitors into the site. Freshly contributed content may indeed have the potential to provide the community with a sense of ownership, while at the same time getting people to revisit the site and also draw in new visitors.

In an attempt to better understand visitors opinions on sharing performative media, on 14th June 2014, we conducted a short discussion with a number of groups in-situ at the Hafod-Morfa copperworks (see Fig. 7.3), asking visitors what they thought about the concept of being able to share their own images, videos and sounds for others to experience at the site.

Figure 7.3: A public engagement event encouraging visitors to discuss their thoughts and experiences.

7.4.1 Engagement

Participants

In the hope of achieving a true reflection of visitor opinions, participants were recruited on-site at the Hafod-Morfa Copperworks. These were visitors, all with the intention exploring the site on its official opening day. Estimates from the organisers put the footfall at around 5,000 for the day.

In total, we recruited 13 participants (2M, 11F), forming 5 separate groups (see Table 7.1). Participants were not asked their explicit ages, though considering their age brackets, 4 were children, 2 were young adults, 5 were adults and 2 were older adults. Around half of the participants had visited or seen the site previously, with the remaining participants visiting the site for the first time.

Prior to these discussions, these groups also took part in a semi-structured interview regarding a performative auditory deployment (see Section 4.4.2).

Group	Age Group	Gender
1	Adult	M
	Adult	F
	Child	F
	Child	F
	Child	F
2	Old Adult	M
	Old Adult	F
3	Adult	F
	Child	F
4	Young Adult	F
5	Adult	F
	Adult	F
	Young Adult	F

Table 7.1: The constituents of the 5 groups recruited at the open day.

Procedure

Participants were first required to have walked around the site and heard the performative audio exhibit that was deployed on that day (see Section 4.4.2). After conducting a short semi-structured interview regarding their experience of this deployment, participants were then asked to engage in an audio-recorded open discussion on the sharing of personal media at the site.

Firstly, participants in each group were introduced to the *Hafod Torch* prototype. This was to show participants a potential way of experiencing imagery in a performative manner, as well as the performative audio they had already experienced on-site on that day. Once participants had seen the projection prototype, the researcher began to discuss the possibility of being able to share personal media at the site to be experienced by others through these performative mediums. The main points the researcher used to direct the discussion were:-

- The acceptability and attitudes towards sharing media at the site.
- The specific kinds of media people would share at the site.

Findings

Of the five groups, two were very excited by the idea of being able to share their own personal media at the site, making clear that they would like to contribute images and sounds. Another group mentioned that they would not personally contribute, though could imagine other visitors willing to do so. One group were not explicit on whether they would contribute media to these experiences. The remaining group clearly stated that they would not share media, with one participant commenting, *"[it] never entered my frame of mind."*

Two of the five groups were quick to mention that they would expect to see strict filtering on the media that is uploaded, ensuring that unsuitable items are removed. Most of the groups had a strong feeling that many people would be interested in contributing, with one particular group commenting, *"people in the Hafod seem very favourable to that. They're all willing to share."* A number of groups saw the opportunity to capture historical accounts and archive data from the local community and past workers at the site. As well as the contribution of stories, testimonies, historical and archive material, many participants were also in favour of contributing more 'contemporary' media and interpretations of people's experiences in the area. The same participant from the previous group again commented, *"you look at twitter and Facebook, people are open to sharing their personal experiences and activities."* To deal with different kinds of contributions, a number of groups touched on the topic of categorisation, attempting to sort the contributions by theme. One particular group favoured the idea of separating contributions into the *"past, present and future."* The past would contain historical sounds and images, the present would capture people's experiences at the site in its current state, and the future would be a forum for people to share their visions of what they would like the site to become.

In terms of specific sounds that participants themselves could imagine sharing at the site, suggestions were presented mainly in the form of personal anecdotes and stories. Some participants suggested contributing their own industrial sounds, though one participant said it was important that *"human sounds, as well as machine sounds"* were included at the site. This participant believed that human sounds were more emotionally engaging.

Imagery that participants could see themselves sharing were mainly personal images such as their children playing at the site, their family and old photos of the local area. One participant mentioned that they would like to have contributed images and videos of when they personally worked at the site, showing processes such as *"rolling copper."* This participant, however, commented that they were not allowed photography in the works, so this kind of media would be hard to come by. Approaching the issue from a more contemporary perspective, one participant was interested in sharing the work of local artists, filmmakers and community projects, utilising the site as a kind of exhibition space, demonstrating the work that is going on in the community.

7.4.2 Discussion

There appears to be a clear divide between visitors who would personally share media and those who would not, though most participants did believe that others would want to contribute media to the site in some kind of way. Potential contributions that participants suggested ranged from historical archive material, to present experiences and future aspirations for the area. A number of participants were keen on screening contributions, ensuring that the media remains clean and suitable for the range of audiences that visit the Hafod-Morfa Copperworks. In such instances where visitors are given powerful authoring tools, designers may indeed need to consider measures to counteract malicious or offensive content. A simple method to regulate content would be to have a verification step that is carried out by a trustworthy person, though this may become arduous as the number of contributions increases. Another solution would be to use a more democratic approach, where visitors themselves can vote content up or down depending on its likability. Although regulators should aim to remove undesirable content, a set of simple guidelines should be produced to ensure that there is a coherent definition for what constitutes 'undesirable content.' It may be the case that the heritage site welcomes more general contributions from the community, some unrelated to the site.

We also welcome the idea of categorising and making contributions searchable. The simplest method of accessing contributions would be to let visitors have access to the entire archive. Although this would be the easiest solution to implement, it is not necessarily the best. It is quite likely that over time, the archive would become too large to access in its entirety. Simple filtering techniques could be used on the meta-data to keep content relevant to visitors' specific context, such as time, place and theme. Visitors should also perhaps, have the ability to tag content, linking it to other relevant media contributions.

7.5 Conclusions

We conducted two focus groups, introducing the concept of performative mobile experiences through showcasing a number of prototype interactions. One of these focus group included members of the local community, and the other comprised of stakeholders of the *Cu@Swansea* project. Our findings from these focus groups prompted us to carry out a further, in-situ investigation into visitors opinions of performative interactions. Collectively, the results of these investigations have raised some important points regarding attitudes towards and the acceptability of performative mobile interactions.

Conducting two separate focus groups was very useful, as our community and stakeholder groups came from two different perspectives, motivated by separate aims, objectives and outcomes. Members of the community focus group were mainly interested in the live interaction, and how it would directly affect the experiences of people at the Hafod-Morfa Copperworks. Stakeholders, on the other hand, were interested in how to encourage people to use these technologies, and ensuring they remain sustainable. Both focus groups were very positive regarding the prototypes. One subject that both focus groups showed a combined interest in was content sharing – allowing visitors to contribute sounds and images when visiting the site.

In our further public engagement investigating performative sharing, we learnt that there was a clear divide in visitors who did, and did not want to share media at the site. However—similar to findings in both of our focus groups—participants of this public engagement continued to believe that there are sufficient people who would contribute media at the Hafod-Morfa Copperworks. The vetting of contributions was a major concern for some, with participants recommending the regulation of content. We discussed a number of ways in which this could be done. We also discussed the topic of filtering and categorisation of contributions, keeping content relevant to those who access it through search and context.

We do not know whether performative sharing would indeed provide a sense of ownership or continue to draw people into the site, but it is a feature that has been continually requested, and is perhaps worth exploring further in future work. We believe that this kind of contribution mechanism would at least compliment the mobile performative framework, allowing visitors to not only experience, but also share sounds and images in a performative, extravagant and expressive manner.

Chapter Eight

Discussion & Conclusions

We set out in this thesis to investigate the underexplored design space of performative, extravagant, expressive place-based experiences. Many of the mobile location-based services and experiences to date have focused on private, secretive interactions that leave people disconnected from one another – digitally divided. In the Hafod-Morfa Copperworks—the subject of this work—we would like visitors to have the opposite experience. We would like to encourage people to visit the area, to start conversations and have a social, more meaningful experience when they arrive there. This digital regeneration project—*Cu@Swansea*—is the underlying motivation for our work here.

Approaching the issue from a broad perspective, in this thesis, we have looked at developing interaction techniques that may help facilitate more active social engagements amongst visitors in public spaces. Armed with a vision, we began to explore potential interactions from a performative perspective, designing interactions to be as expressive and as extravagant as possible. By creating a spectacle with visible performers, it was hoped that such interactions would encourage visitors to approach and engage with one another. Previous work has also shown that performativity is particularly relevant in the heritage context, with this more active role promoting reminiscing and participation amongst visitors.

At the beginning of this thesis, we outlined that we would be investigating the use of audio and visual effects as a baseline for performative mobile interactions. We achieved a performative audio experience by placing interactive loudspeakers in the environment, and we achieved a performative visual experience through encouraging users to overlay the environment with pico projections. After experimenting with these performative audio-visual interactions, we then considered a technique for scaling-up and orchestrating these performances, making them remotely collaborative. Each performative experience was deployed in a public space and studied with potential users. Although our studies focused mainly on the usability of our designed interactions, we were able to gain some early insights into their potential social benefits. Throughout this work, we made a conscious effort to include both the Hafod local community and project stakeholders in the design process. This was realised through organising a number of focus groups and community engagement events.

In this thesis, we have provided a number of examples of novel, performative designs to promote a more active social engagement amongst visitors in public spaces. Each contributing chapter has studied and evaluated these concepts, drawing key design recommendations from the findings. Our focus groups have then given a different, broader perspective, concerning the opinions of different parties involved in a large regeneration project. We believe that our findings will help designers, developers and other researchers who are interested in pursuing performative, public, mobile controlled interactions.

In this chapter, we bring together our key findings, summarising and presenting them in a way that will be helpful for the design and evaluation of future prototypes. We hope that this framework will inspire and assist people in building and researching the next generation of prototypes.

8.1 Contribution summary

At the beginning of this work, we discovered an underexplored design space – peformative, extravagant, expressive place-based experiences. Our main contribution here is a novel, mobile place-based framework to enable performances in public spaces.

In *Chapter 1*, we began by looking at existing location-based services, arguing for the need for more performative, extravagant and expressive experiences that draw us together in public spaces. We also spoke of our underlying context and motivation—the *Cu@Swansea* project—which entails the digital regeneration of the Hafod-Morfa Copperworks. *Chapter 2* reviewed previous literature, situating our work within the field. We began by introducing the broad theme of context- and location-aware computing, then narrowing our scope towards mobile guides, navigation and control. We continued by discussing face-on interactions, and how we may facilitate these through a range of eyes-free and heads-up interaction techniques. Continuing the theme of social experiences, we then probed deeper into socio-spatial interaction. In our literature review, we considered past attempts at community engagement and regeneration projects, and how performative computing may help facilitate a more active social experience for collocated visitors.

Chapter 3 included a discussion of our methodology and research approach, documenting what we planned to research and how. We decided to conduct the majority of our research by designing and deploying public interventions, allowing us to get first hand experience of how mobile performative interactions would work in-the-wild. Additionally, we stated that we would determine the needs and concerns of both the local community and project stakeholders through a number of focus groups and community engagement events.

Chapter 4, the first of our contributing chapters, explored audio as an effect in a performative mobile experience. We conducted an initial experiment with spatial audio, exploring the engagement benefits of sounds originating from the environment. We then explored more extravagant, expressive experiences that are possible with sounds and interaction outside of the device. Situated loudspeakers were used, that could be triggered by movement, gestures or visiting certain points of interest. Amplified audio was very effective in turning bystanders into spectators. Control was shown to be an important aspect for performers, directly related to comfort and enjoyment. The deployment location also had a major effect on comfort and enjoyment, highlighting that the suitability of such interactions depends on the context.

Chapter 5 then explored the visual modality as an effect in a performative mobile experience. A pico projection prototype was developed, where projections could be controlled in a free-form way, simply by moving the device. The system was built to encourage users to project into and focus on the environment as opposed to looking at their screen. We carried out a public deployment in a visitor centre, where users could project insects, animals and processes. Similar to our audio experiments, the expressive nature of the interactions promoted a more active engagement amongst bystanders and performers. However, although not compared directly, the scale and brightness of projections appeared to make them less effective for this same purpose.

Chapter 6 drew the previous contributing chapters together, detailing an interaction for directing large scale, performative mobile experiences. By harnessing users' mobile device speakers, we were able to create a mechanism for orchestrating a life-size human soundscape. A controlling user could trigger sounds on individual users' devices, and also position them in space with a novel guidance mechanism. Again, we saw a continuation of the kinds of behaviour observed in previous chapters, with amplified audio prompting curiosity and conversation, though even more so due to users being forced into close proximity of each other. When remotely directing, the complex and continually changing nature of real world settings meant improvisation played a large part – which our simple feedback technique was able to cater for. Issues were also raised with the ethics and accountability of remotely guiding and controlling users in public spaces.

By building and deploying a number of extravagant and expressive mobile performances, we were able to learn a great deal regarding their usability. To learn more about the public acceptability and attitudes towards these kinds of experiences, we conducted a number of focus groups and community engagement events with interested parties. The final key, contributing chapter—*Chapter 7*—reported on these investigations, exploring and discussing potential issues and ideas raised by participants. The chapter concluded with a discussion of potentially encouraging visitors to contribute content and play a more active role in experiences in public spaces, further crossing the consumer-producer divide.

8.2 Place-based experiences framework

Throughout this thesis, we have learnt a great deal regarding performative place-based experiences. By studying our designed interactions in a range of controlled experiments and public deployments, we have discovered a number of key factors related to both their success and failure. Here, we draw these key findings together, providing a simple framework to aid in the successful design and evalu-

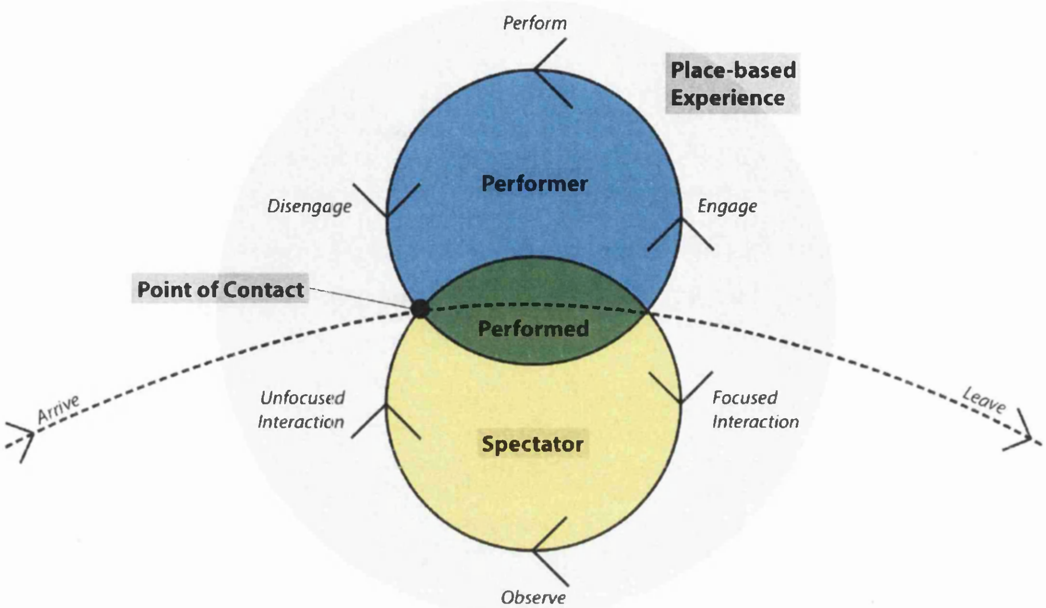

Figure 8.1: Place-based Experiences Framework trajectory with the five concepts: *Place-based Experience*, *Point of Contact*, *Performer*, *Performed* and *Spectator*.

ation of future performative, extravagant, expressive place-based experiences (see Fig. 8.1). The following five, essential areas—along with their suggestions—of the *Place-based Experiences Framework* (PEF) can assist with the evaluation of designs, encouraging more socially active place-based experiences:-

Place-based Experience: The experience itself, including the physical location and space in which the place-based experience occurs.

Point of Contact: The physical point of contact with the experience. This concerns *if* and *how* visitors come into contact with the experience.

Performer: A visitor who has chosen to actively participate by controlling and manipulating the experience, sustaining a performance.

Performed: A visitor who has chosen to actively participate, though acts through instructions and effects caused by another performer's manipulations.

Spectator: A bystander who has acknowledged an ongoing performance, and chooses to engage in focused interaction, observing the experience.

The framework is represented as a kind of trajectory, where a visitor must first arrive at the experience, and then later leaves. All visitors first arrive in a neutral state as a bystander, behaving in unfocused interaction. When a bystander reaches a point where they are able to experience or take part in the performance (*Point of Contact*), they have three options; to remain as a **bystander** and leave, become a **performer** or **performed** by engaging and playing an active role in the performance, or finally, become a **spectator** by coming into focused interaction and observing the performance. The trajectory is designed in such a way that it allows visitors to enter any state, and allows those in contact with the experience—spectators and performers—to alternate between states. As in similar work [42], we believe that a performer is always a spectator to their own actions, however, in our framework, we differentiate between performer and spectator as visitors who are either actively participating and controlling the performance, or just simply observing. The distinct terms bystander, spectator, performer and performed are useful in describing visitors' behaviour, and how they come into contact with each other.

Similar to other performance related work [54], we have chosen to use a trajectory to represent the framework, as we believe it provides a simple, easy-to-understand representation of the five main concepts that are critical to performative place-based experiences. Designers wishing to apply this framework should consider the five concepts in the model, and the proposed suggestions that come with each. Central to the framework, is determining a suitable **setting**, **manipulations** and **effects**.

8.2.1 Place-based experience

> *The experience itself, including the physical location and space in which the place-based experience occurs.*

Context

One of the first things designers should consider is the context into which they are placing the experience. The social acceptability of extravagant, expressive interventions is highly dependent on context. A single experience may be greeted with very different reactions if situated in two different physical locations or scenarios. The importance of context has been discussed in previous research [42].

To enhance the acceptability and potential success of such interventions, designers should understand that the overall experience—especially the effects—must be of shared public interest. Experiences should also be designed so that they complement the surrounding environment, and do not simply overlay it.

Source: This suggestion is based on findings in Chapter 4, where we deployed systems in differing contexts. These contexts were, Swansea University (see Section 4.3) and the Hafod-Morfa Copperworks (see Section 4.4).

Stage-craft & improvisation

When designing a place-based experience, there is only so much that a visitor can be expected to do with a performative system, without further direction. Most visitors will not be used to public performances and will require further direction.

There are two ways in which designers can promote more performative behaviour; implicit and explicit guidance. Explicit guidance can be given through signage, prompts and further instruction, either digitally through a device or physically in the environment. Implicit guidance, is perhaps a cleverer, but more difficult solution, requiring very careful planning and curation of a public space. An example of implicit guidance would involve strategically placing exhibits to intercept the flow of visitors, or placing exhibits in locations that require visitors to be more pro-active and conspicuous when interacting.

However, designers should also consider the fact that performers and spectators may not be able to follow all guidance due to factors beyond their control. For example, there may be impassable obstacles or other spontaneous situations.

Source: This suggestion is based on observations in our experiment with projections (see Section 5.2). Improvisation appeared as an important element in our remote guidance experiments in public (see Section 6.3).

Variety

Although performative experiences can be playful and enjoyable, they can sometimes be lacking in hard information. Designers should remain mindful, that certain visitors will expect to see additional information that they can use to backup or validate the performance, allowing them to make sense of it.

It is possible—and deeply encouraged—to design experiences where informative and performative approaches compliment each other.

Source: This suggestion is based on our experiment in Section 5.2, where there was no clear preference between informative and performative approaches. It is also based on our prototype design in Section 4.4.

Ethics

Designers should be careful when designing performative place-based experiences, that what they are asking visitors to engage in is neither dangerous or unethical. Visitors should be able to interact with each other safely while unsupervised.

This point has great importance if the experience involves remote guidance, control or communication between visitors. Designers should ensure that the experience is safe, and that the area around the exhibit is suitable for the kind of interaction.

Source: This suggestion originates from our remote guidance experiment (see Section 6.3), where sending instructions between one another presented potential risks and accountability issues.

8.2.2 Point of contact

> *The physical point of contact with the experience. This concerns if and how visitors come into contact with the experience.*

Affordance

When a visitor arrives at a public space, most likely, they will be unaware of the place-based experience, or perhaps, not know how to interact with it. To make visitors aware of the experience, designers should ensure that the technology, access points and functionality are all *very* clearly advertised to visitors. This can be achieved through physical or digital markers, with at least one placed at each exhibit. Basic guidance should be explicit. Sheridan et al. [132] have previously document some of the difficulties of turning bystanders into performances.

Visitors to public spaces are used to being consumers, passively receiving information. For this reason, advertising performative functionality and interactivity is crucial to potential performers engaging with the experience.

Source: This suggestion is based on our only natural deployment (see Section 4.4), where it became very clear that visitors need to know how to access and control the experience. Our previous experiments had always primed participants.

Engagement

Amplified effects are very good at gaining bystanders' attention. Without the need for any complex design or placement, one can gain bystanders' attention and encourage them to talk amongst themselves merely by producing amplified audio and visual effects.

To encourage interaction between the performer and the spectator specifically, it is important to place the effect on or around the performer. Spectators are far more likely to approach or begin conversation with performers when the amplified effect is in the performers immediate vicinity. However, it must be noted that some performers enjoy anonymity, and would prefer not to be directly linked to causing the effect.

Source: This suggestions is based on all of our public deployments. In particular, in our projection (see Section 5.2) and remote guidance (see Section 6.3) experiments, we saw that producing an effect on or around the performer encouraged engagement.

8.2.3 Performer

> *A visitor who has chosen to actively participate by controlling and manipulating the experience, sustaining a performance.*

Control

For an experience to remain usable and enjoyable for a performer, it is imperative that the performer retains a high level of control at all times. Performers become less comfortable when they feel that they have lost control, or that their immediate manipulations are not directly affecting the performance in a linear and coherent way. Hornecker and Buur [70] have noted the importance of a clear link between manipulations and effects. Control should be kept simple and easy to understand for performers. In larger, more complex experiences, performers should be provided with a summarised overview of the current state of the performance. This can help performers situate themselves within the experience, both physically and mentally.

Designers should carefully consider the link between their manipulations and effects; and what kind of sensor, input, action or interface is best suited to the task.

Source: This suggestion is based on findings from all of our experiments. In particular, we saw that participants missed a high level of control in our expressive audio prototype (see Section 4.3). The simple control in our projection prototype was welcomed by participants (see Section 5.2). Our remote guidance experiment taught us that an overview was useful for larger, interconnected performances (see Section 6.3).

8.2.4 Performed

> *A visitor who has chosen to actively participate, though acts through instructions and effects caused by another performer's manipulations*

Split performing roles

Conventionally, in performative experiences, designers discuss the role of bystanders, performers and spectators [134]. Performers can manipulate these experiences, and are directly responsible for the effects caused.

Here, we present the idea of a new kind of performance, where the role of performing is split between multiple users. In this scenario, experienced performers receive a manipulation role, or an ability to control and orchestrate the performance. A secondary performer, known as the *performed*, passively receives instructions and effects on their mobile device. The secondary performer acts as a mobile automaton, where effects are performed through. This is similar to the theme of witting participation described by Sheridan et al. [132], though we have specifically designed a role which separates manipulations from effects, requiring little skill, yet remaining highly expressive and engaging for the *performed* within the performance frame. For this reason, we believe that the *performed* role lies between being a performer and being a spectator.

With performed users, performative experiences become much more scalable, and allow inexperienced performers to play a larger part in the experience.

Source: This suggestion is based on our work with remote guidance 6.3. We discovered the ability to have split manipulation and effect roles while designing our prototypes.

8.2.5 Spectator

> *A bystander who has acknowledged an ongoing performance, and chooses to engage in focused interaction, observing the experience.*

Visibility

When trying to gain the attention of bystanders, bystanders are much more susceptible to amplified effects rather than amplified manipulations. Amplified manipulations have a negligible effect in gaining attention.

However, when considering effects alone, there appears to be a large difference between the amount of attention that visual and audio effects garner. A property that makes projections less visible than sound is audible, is the fact that a bystander has to look around to notice a projection. Audio, however, is inherently easier to discover and experience, as it provides ambience through multi-directionality. Therefore, if increased visibility is desired, we suggest using audio as a primary effect.

Source: This suggestion is based on the difference in attention received by manipulations and effects in our performative audio experiments (see Section 4.3). The difference between attention received for audio and visual effects can be observed between Chapters 4 and 5.

8.3 Future work

The work presented here displays a solid foundation for performative place-based experiences. We have discussed what a performative place-based experience is, why one may want to design such an experience, and also how, with the aid of a framework. We have demonstrated fully functioning examples of audio and visual performances, and have also taken the framework a step further by trialling remote controllable performances. Because we had much greater success with gaining attention and facilitating a more socially active experience when using audio as an effect, we decided to focus more of our efforts on audio. Future work with visual performances could look at larger, more extravagant displays, and try to achieve a similar effect to that which was achieved with audio.

Naturally, the next step in this research would then involve experimenting with additional, new modalities, discovering other kinds of engaging experiences. Although we have used audio and visual as a baseline, there are a whole range of other modalities that one could utilise, including haptics and olfactory. Effects could even include physically manipulating the environment, using technology such as servos, motors and pulleys to move objects around. There is also then the possibility of combining effects to create more complex performances. In the context of the Hafod-Morfa Copperworks, this may include combining flashing lights, loud industrial noises and smoke, while at the same time experiencing vibrations through the floor. By combining a number of amplified effects, the performances would perhaps become even more engaging, expressive and extravagant. As well as exploring new effects, there is also the possibility of exploring additional mobile-based manipulations that could control these performances. In this research, we tested a select few manipulations, based mainly around gestures and movement through space. With the plethora of sensors in todays smartphone devices, and new ones being introduced all of the time, we imagine that there are many more—yet unimagined—possibilities for interacting with and controlling place-based experiences.

In *Chapter 6*, we demonstrated a remote control mechanism that could be used to orchestrate a coherent performance over multiple devices. Our initial investigations focused on the usability of the approach, and the experience it offered. Further work is now required to understand how users feel about their personal device being commandeered and used as a peripheral in a public performance.

When conducting the focus groups in *Chapter 7*, it became apparent that when visiting the Hafod-Morfa Copperworks, sharing personal media was of collective interest. Upon making this discovery, we carried out a further community engagement, gauging visitors' thoughts on sharing their own media at the Hafod-Morfa Copperworks. This engagement proved equally encouraging, uncovering additional opinions and the kinds of media that visitors would share. Based on these findings, future work can now attempt to incorporate a sharing mechanism into a place-based experience, allowing visitors to share and experience their own media in a performative manner. By deploying and studying a fully functioning prototype in-situ, we would be able discover the true value of such an experience.

Finally, perhaps the most crucial area that requires future work is the further validation of the social benefits that performative place-based experiences offer. Although we observed a number of occasions where active engagements occurred between visitors, further tests are now required to confirm and validate the full effect. To achieve this kind of understanding, the interactions and experiences would need to be deployed in-the-wild for a much longer period of time, collecting data from performers, spectators and bystanders. Such a study would yield a detailed account of the effects, and also the sustainability of the approach.

Appendix A

Contributing Publications

Much of the research contained within this thesis has been published in co-author peer-reviewed international conferences and journal papers. A selection of the work has also been presented and discussed in international conference workshops. For each publication, a reference has been provided, along with an abstract that summarises the contents of the work. The author of this thesis is the first author in all contributing publications. This was due to the vast majority of the work being conducted solely by the author. All publications also feature Matt Jones as a co-author – an active voice who has been an invaluable inspiration in producing the work presented here. Whilst the vast majority of work here was conducted solely by the author, additional help was sometimes required for recruiting, observations and logistical reasons. In the following list of publications, where additional assistance was required, further details are provided.

As well as being given the opportunity to present and discuss this work with world leading researchers at numerous international conferences and workshops, these publications have also been distributed to, read by and inspired many of the stakeholders and other researchers involved in the *Cu@Swansea* project. Many of the ideas and concepts described in these publications are vastly different to current mainstream approaches to digital interpretation and social interaction in public spaces such as visitor centres and heritage sites. Of the project individuals that have read these publications, many have labelled the work as fresh and liberating, encouraging them to think more broadly about the kinds of visits and interactions that may be possible at the Hafod-Morfa Copperworks.

[P1] L. Betsworth, M. Jones Regenerative Social Engagement through Performative Technologies. In *CHI 2014 workshop on Socially Engaged Arts Practice in HCI*, Toronto, April 2014

Abstract. In this paper we discuss the use of performative technologies for regenerative purposes in a heritage context. Our motivation stems from a heritage site of world significance—The Hafod-Morfa Copperworks—being left abandoned for some time, leaving the area overgrown and unused. By re-opening the site and introducing performative technologies that encourage participation, we hope to encourage the local community and other visitors to engage with and re-use the site. Here, we introduce a number of prototype technologies that we have already developed and begun testing. Taking these experiences into account, we discuss the implications of visitors playing the role of a performer, and the trade-offs of using an experienced artist or performer.

[P2] L. Betsworth, N. Rajput, S. Srivastava, M. Jones Audvert: Using Spatial Audio to Gain a Sense of Place. In *Proceedings of INTERACT 2013*, Cape Town, September 2013. 455-462.

Abstract. We introduce Audvert – a system that facilitates serendipitous discovery and navigation through spatial audio; used to navigate and discover points of interest in large, unfamiliar indoor environments. Our main aim was to create a lightweight spatial audio display that can convey a sense of a place without complex point and select interactions. We conducted a preliminary study comparing two audio types to see which best suited sound localization and a study of Audvert used in a real world scenario. Our findings suggest that long continuous audio performs better than short intermittent audio for sound localisation. We also discover a change in behaviour when using the system, with a large percentage of users wanting to visit newly discovered shops after using the system. We discuss the findings and draw research conclusions.

Assistance. Two IBM colleagues—who are both co-authors on the paper—helped to recruit participants and translate during studies.

[P3] L. Betsworth, S. Robinson, M. Jones Pico Projection for Performative Place Based Services. In *CHI 2013 workshop on Experiencing Interactivity in Public Spaces (EIPS)*, Paris, April 2013

Abstract. In this paper we explore using handheld projectors in place of traditional location-based information services. We built a prototype system to compare performative projection of animations and images against conventional on-screen information. We conducted a user study to test the informative and the new performative design, gathering user feedback and reactions to the approach. Our findings highlighted design issues and the potential benefits of performative projection for prompting interaction with exhibits as part of the experience at a visitor attraction.

Assistance. Two University colleagues assisted in making observations of user behaviour – one of which is a co-author on the paper. This co-author contributed a figure and refined text.

[P4] L. Betsworth, H. Bowen, S. Robinson, M. Jones Performative Technologies for Heritage Site Regeneration. *Personal and Ubiquitous Computing* 18(7):16311650, 2014

Abstract. Heritage sites are an important part of understanding our role in history. They have the potential to teach us important lessons, such as where we came from and subsequently, the people it has made us today. As members of a large, heritage-led, regeneration project, we are working with the Hafod-Morfa Copperworks, a heritage site in the Lower Swansea Valley where there is not much to see or hear. The few ruins at the site make it difficult to imagine what the site would have been like back in its heyday. Our goal at the site is to draw people together, not to view a finished piece of curated heritage, but rather, to start conversations about their memories and the significance of the site to them and to discover what they would like to see at the site in the future. The technology we are producing is about engaging with the local community and stakeholders as groups to provoke discussion. This contrasts with previous uses of mobile guides which only attempt to be tourist aids. In this article, we report on two prototype technologies we have developed to help accomplish this task. Throughout

the article, we discuss how and why designing performative technologies could help encourage people to visit, socialise and communicate within the area. Our early results suggest that expressive performative technologies are good at gaining spectators attention and encouraging an active engagement between performer and spectator.

Assistance. In the initial performative projection experiments, two University colleagues assisted in making observations of user behaviour. The same two colleagues later helped to conduct a preliminary experiment with the situated loudspeaker system, leading a number of study sessions. One co-author contributed a figure and refined text. The other contributed factual text to the introduction.

[P5] L. Betsworth, M. Jones Director: A Remote Guidance Mechanism. In *Proceedings of CHI 2015: Extended Abstracts*, Seoul, April 2015

Abstract. When using a mobile device as a navigation aid, we are used to receiving computer-generated routes and directions. Remote guidance, however, remains an underexplored design space in mobile interaction design. In this paper, we introduce Director, a novel, remote guidance mechanism for the positioning of people in outdoor spaces using mobile devices. We conducted a study to test our novel positioning technique, testing its guiding accuracy and effect on Preferred Walking Speed (PWS). Our results suggest that Director offers users a fun and playful experience, and that our novel guidance technique is a very accurate remote mechanism.

Assistance. The experiment was conducted with the help of two University colleagues, both of which helped with setting up equipment and carrying out observations.

[P6] L. Betsworth, M. Jones Designing a Performative Sharing Framework for Heritage Site Regeneration. In *IUI 2014 PATCH workshop: The Future of Experiencing Cultural Heritage*, Haifa, February 2014

Abstract. In this paper we introduce the concept of a performative sharing framework to promote the regeneration of heritage sites. We discuss our reasoning for the use of such technologies in this context, considering

the opportunities for exciting new interactions and experiences. Throughout the paper, we make reference to the Hafod-Morfa Copperworks – the heritage site where this framework is to be deployed. On-site, we imagine visitors being given the opportunity to play a larger participatory role, being able to craft and control not only their own experiences, but also others through a range of public control and feedback mechanisms. Taking into account some of the novel, public experiences that are possible through performative design in a heritage setting, we also foresee a new relationship between spectator and performer, with the introduction of a contributor. This contributor role has the potential to assist an underlying issue with a large number of heritage sites and other similar visitor attractions.

APPENDIX B

Project Roadmap

As our research method was exploratory based, all chapters and events do not appear in chronological order. This thesis was arranged in the current format to improve coherence between chapters, making clear how certain ideas and findings had influenced the direction of our research. The following roadmap in Figure B.1—along with descriptions—goes some way to explaining the thought process and sequence of events behind the main developments in this research.

[**B.1.1**] We initially began with the concept of performative projection – a more performative alternative to AR lens applications that merely augment the screen, rather than the physical environment. After conducting an experiment using a pico projector, it was clear that the effect was not gaining as much attention as we had hoped. Our thoughts were then based on where we should take the technology next, or if we were able to extend it in any way. We finally decided that the technology was perhaps less suitable for gaining lots of attention, and we placed the performative projection aside.

[**B.1.2**] Our second foray into performative technologies appeared in the form of audio. Building on our previous work with performative projection, we continued to use the mobile device as a control object, using similar pointing and gestural manipulations. We chose to experiment with audio as we believed that its multidirectional, omnipresent properties would make it a better suited effect for gaining spectator attention. Our early experiments with loudspeakers were quick to confirm this, with audio garnering far more attention than the small, dimly lit projections in our visual experiments.

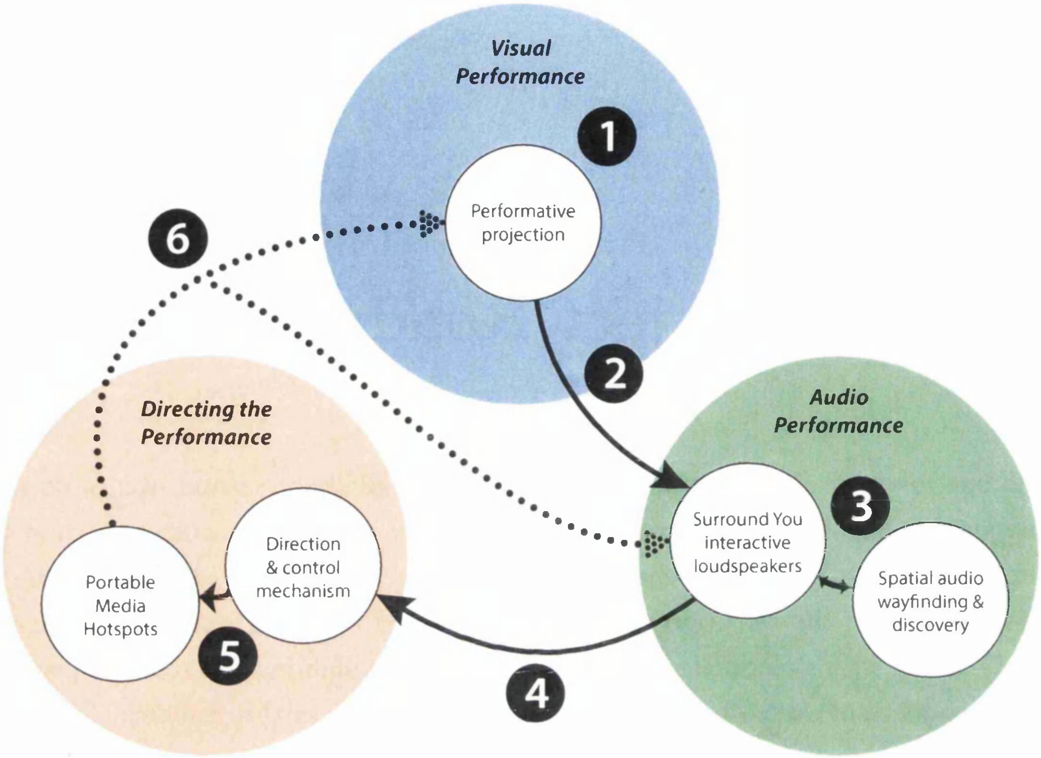

Figure B.1: A project roadmap with links between relevant decisions, concepts and developments (1–6).

In a later, naturalistic deployment at the Hafod-Morfa Copperworks, we took on board our findings from our initial projection experiment, providing the ability to perform, alongside useful information of the point of interest – allowing visitors to make sense of the performance whilst performing.

[**B.1.3**] At the same time as conducting our expressive audio experiments, we also began to experiment with navigation and discovery using spatial audio. This was an attempt to better understand the design space, seeking further opportunities for heads-up, face-on interaction with public spaces.

[**B.1.4**] After experimenting with a number of performative audio and visual experiences, we shifted our focus towards directing the performance. There were a number of reasons for this, including making the experience more mobile, and the ability to orchestrate larger collaborative performances.

Our novel guidance and control mechanism allowed us to remotely direct users on small-scale journeys very accurately. Early testing and a preliminary study allowed us to improve this mechanism further by adding additional options for directors, whilst keeping the interface minimal.

[**B.1.5**] Building on the previous direction mechanism, we developed a framework to allow visitors to act as dynamic performers, with their devices commandeered for use as media outlets. Initial experiments utilised the visitor's mobile device loudspeaker, with the director able to freely play sounds and send directions. By harnessing visitors' mobile devices, instead of being limited by a small number of curated speakers in the environment, performances could now become much larger and more versatile.

In this work, it became very obvious that—similar to in our performative projection experiments—placing the effect on or around the performer promoted more active social engagements between performers and spectators.

[**B.1.6**] The use of mobile devices for both manipulations and effects made it possible to create and orchestrate lightweight, scalable, impromptu performances. The work demonstrates that our original audio and visual experiences can easily be adapted to work in unison with this remote framework. In future work, additional modalities could be explored, including haptics, or perhaps a mix modalities to make the performance more immersive. Future work could also explore larger visual performative prototypes.

Unlike our previous designs, where the performer was responsible for both manipulations and effects, this new remote technique raises interesting questions between different performers being responsible for the manipulations and effects, with those responsibilities now split over different devices.

Bibliography

[1] Gregory D. Abowd, Christopher G. Atkeson, Jason Hong, Sue Long, Rob Kooper, and Mike Pinkerton. Cyberguide: A mobile context-aware tour guide. *Wirel. Netw.*, 3(5):421–433, October 1997.

[2] GregoryD. Abowd, AnindK. Dey, PeterJ. Brown, Nigel Davies, Mark Smith, and Pete Steggles. Towards a better understanding of context and context-awareness. In Hans-W. Gellersen, editor, *Handheld and Ubiquitous Computing*, volume 1707 of *Lecture Notes in Computer Science*, pages 304–307. Springer Berlin Heidelberg, 1999.

[3] C. Ardito, M. F. Costabile, and R. Lanzilotti. Gameplay on a multitouch screen to foster learning about historical sites. In *Proceedings of the International Conference on Advanced Visual Interfaces*, AVI '10, pages 75–78, New York, NY, USA, 2010. ACM.

[4] Carmelo Ardito, Maria Costabile, Antonella De Angeli, and Rosa Lanzilotti. Enriching archaeological parks with contextual sounds and mobile technology. *ACM Transactions on Computer-Human Interaction*, 19(4):29:1–29:30, 2012.

[5] Carmelo Ardito and Rosa Lanzilotti. "isn't this archaeological site exciting!": a mobile system enhancing school trips. In *Proceedings of the working conference on Advanced visual interfaces*, AVI '08, pages 488–489, New York, NY, USA, 2008. ACM.

[6] Gaynor Bagnall. Performance and performativity at heritage sites. *Museum and society*, 1(2), 2003.

[7] Paramvir Bahl and Venkata N Padmanabhan. Radar: An in-building rf-based user location and tracking system. In *INFOCOM 2000. Nineteenth Annual Joint Conference of the IEEE Computer and Communications Societies. Proceedings. IEEE*, volume 2, pages 775–784. Ieee, 2000.

[8] P. Baranski, M. Polanczyk, and P. Strumillo. A remote guidance system for the blind. In *e-Health Networking Applications and Services (Healthcom), 2010 12th IEEE International Conference on*, pages 386–390, July 2010.

[9] Benjamin B. Bederson. Audio augmented reality: A prototype automated tour guide. In *Conference Companion on Human Factors in Computing Systems*, CHI '95, pages 210–211, New York, NY, USA, 1995. ACM.

[10] S. Benford and G. Giannachi. *Performing mixed reality.* MIT Press, 2011.

[11] Steve Benford. Future location-based experiences. *JISC Technology and Standards Watch.*, 2005.

[12] Steve Benford, Andy Crabtree, Martin Flintham, Adam Drozd, Rob Anastasi, Mark Paxton, Nick Tandavanitj, Matt Adams, and Ju Row-Farr. Can you see me now? *ACM Trans. Comput.-Hum. Interact.*, 13(1):100–133, March 2006.

[13] Steve Benford, Andy Crabtree, Martin Flintham, Chris Greenhalgh, Boriana Koleva, Matt Adams, Nick Tandavanitj, Ju Row Farr, Gabriella Giannachi, and Irma Lindt. Creating the spectacle: Designing interactional trajectories through spectator interfaces. *ACM Transactions on Computer-Human Interaction*, 18(3):11:1–11:28, 2011.

[14] Steve Benford, Andy Crabtree, Stuart Reeves, Jennifer Sheridan, Alan Dix, Martin Flintham, and Adam Drozd. The frame of the game: Blurring the boundary between fiction and reality in mobile experiences. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems*, CHI '06, pages 427–436, New York, NY, USA, 2006. ACM.

[15] Steve Benford, Chris Greenhalgh, Andy Crabtree, Martin Flintham, Brendan Walker, Joe Marshall, Boriana Koleva, Stefan Rennick Egglestone,

Gabriella Giannachi, Matt Adams, et al. Performance-led research in the wild. *ACM Transactions on Computer-Human Interaction (TOCHI)*, 20(3):14, 2013.

[16] Steve Benford, Chris Greenhalgh, Gabriella Giannachi, Brendan Walker, Joe Marshall, and Tom Rodden. Uncomfortable interactions. In *Proc. CHI '12*, pages 2005–2014. ACM, 2012.

[17] Steve Benford, Holger Schnädelbach, Boriana Koleva, Rob Anastasi, Chris Greenhalgh, Tom Rodden, Jonathan Green, Ahmed Ghali, Tony Pridmore, Bill Gaver, Andy Boucher, Brendan Walker, Sarah Pennington, Albrecht Schmidt, Hans Gellersen, and Anthony Steed. Expected, sensed, and desired: A framework for designing sensing-based interaction. *ACM Trans. Comput.-Hum. Interact.*, 12(1):3–30, March 2005.

[18] Steve Benford, Will Seager, Martin Flintham, Rob Anastasi, Duncan Rowland, Jan Humble, Dana Stanton, John Bowers, Nick Tandavanitj, Matt Adams, JuRow Farr, Amanda Oldroyd, and Jon Sutton. The error of our ways: The experience of self-reported position in a location-based game. In Nigel Davies, ElizabethD. Mynatt, and Itiro Siio, editors, *UbiComp 2004: Ubiquitous Computing*, volume 3205 of *Lecture Notes in Computer Science*, pages 70–87. Springer Berlin Heidelberg, 2004.

[19] Bert Bongers. The projector as instrument. *Personal Ubiquitous Computing*, 16(1):65–75, 2012.

[20] Virginia Braun and Victoria Clarke. Using thematic analysis in psychology. *Qualitative Research in Psychology*, 3(2):77–101, 2006.

[21] Stephen Brewster. Head-up interaction: Can we break our addiction to the screen and keyboard? In *Proceedings of the 2009 International Conference on Multimodal Interfaces*, ICMI-MLMI '09, pages 151–152, New York, NY, USA, 2009. ACM.

[22] Stephen Brewster, Joanna Lumsden, Marek Bell, Malcolm Hall, and Stuart Tasker. Multimodal 'eyes-free' interaction techniques for wearable devices.

In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems*, CHI '03, pages 473–480, New York, NY, USA, 2003. ACM.

[23] Gregor Broll, John Hamard, Massimo Paolucci, Markus Haarländer, Matthias Wagner, Sven Siorpaes, Enrico Rukzio, Albrecht Schmidt, and Kevin Wiesner. Mobile interaction with web services through associated real world objects. In *Proceedings of the 9th international conference on Human computer interaction with mobile devices and services*, MobileHCI '07, pages 319–321, New York, NY, USA, 2007. ACM.

[24] Gregor Broll, Susanne Keck, Paul Holleis, and Andreas Butz. Improving the accessibility of nfc/rfid-based mobile interaction through learnability and guidance. In *Proceedings of the 11th International Conference on Human-Computer Interaction with Mobile Devices and Services*, MobileHCI '09, pages 37:1–37:10, New York, NY, USA, 2009. ACM.

[25] Barry Brown and Matthew Chalmers. Tourism and mobile technology. In *Proceedings of the Eighth Conference on European Conference on Computer Supported Cooperative Work*, ECSCW'03, pages 335–354, Norwell, MA, USA, 2003. Kluwer Academic Publishers.

[26] Barry Brown, Ian MacColl, Matthew Chalmers, Areti Galani, Cliff Randell, and Anthony Steed. Lessons from the lighthouse: Collaboration in a shared mixed reality system. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems*, CHI '03, pages 577–584, New York, NY, USA, 2003. ACM.

[27] Jorge Simarro Cabrera, Henar Muñoz Frutos, Adrian G. Stoica, Nikolaos Avouris, Yannis Dimitriadis, Georgios Fiotakis, and Katerina Demeti Liveri. Mystery in the museum: collaborative learning activities using handheld devices. In *Proceedings of the 7th international conference on Human computer interaction with mobile devices & services*, MobileHCI '05, pages 315–318, New York, NY, USA, 2005. ACM.

[28] Francesco Cafaro, Leilah Lyons, Joshua Radinsky, and Jessica Roberts. Rfid localization for tangible and embodied multi-user interaction with museum

exhibits. In *Proceedings of the 12th ACM international conference adjunct papers on Ubiquitous computing*, Ubicomp '10, pages 397–398, New York, NY, USA, 2010. ACM.

[29] Heloisa Candello. Developing principles for outdoor mobile multimedia guides in cultural heritage settings. In *Proceedings of the 11th International Conference on Human-Computer Interaction with Mobile Devices and Services*, MobileHCI '09, pages 96:1–96:1, New York, NY, USA, 2009. ACM.

[30] Xiang Cao and Ravin Balakrishnan. Visionwand: interaction techniques for large displays using a passive wand tracked in 3d. In *Proceedings of the 16th annual ACM symposium on User interface software and technology*, UIST '03, pages 173–182, New York, NY, USA, 2003. ACM.

[31] Janet Cardiff. Her long black hair. `http://www.cardiffmiller.com/artworks/walks/longhair.html`. Accessed: 2013-08-20.

[32] Ugo B. Ceipidor, Carlo M. Medaglia, Amedeo Perrone, Maria De Marsico, and Giorgia Di Romano. A museum mobile game for children using qr-codes. In *Proceedings of the 8th International Conference on Interaction Design and Children*, IDC '09, pages 282–283, New York, NY, USA, 2009. ACM.

[33] Alan Chamberlain, Leif Oppermann, Martin Flintham, Steve Benford, Peter Tolmie, Matt Adams, Ju Row Farr, Nick Tandavanitj, Joe Marshall, and Tom Rodden. Locating experience: touring a pervasive performance. *Personal Ubiquitous Computing*, 15(7):717–730, 2011.

[34] Angie Chandler, Joe Finney, Carl Lewis, and Alan Dix. Toward emergent technology for blended public displays. In *Proceedings of the 11th international conference on Ubiquitous computing*, Ubicomp '09, pages 101–104, New York, NY, USA, 2009. ACM.

[35] David J. Chatting and Jon S. Sutton. Visible imagination: projected play. In *Proceedings of the 23rd British HCI Group Annual Conference on People*

and Computers: Celebrating People and Technology, BCS-HCI '09, pages 405–409, Swinton, UK, UK, 2009. British Computer Society.

[36] Keith Cheverst, Nigel Davies, Keith Mitchell, and Christos Efstratiou. Using context as a crystal ball: Rewards and pitfalls. *Personal and Ubiquitous Computing*, 5(1):8–11, 2001.

[37] Luca Chittaro and Stefano Burigat. Augmenting audio messages with visual directions in mobile guides: An evaluation of three approaches. In *Proc. MobileHCI '05*, pages 107–114. ACM, 2005.

[38] Omar Choudary, Vincent Charvillat, Romulus Grigoras, and Pierre Gurdjos. March: mobile augmented reality for cultural heritage. In *Proceedings of the 17th ACM international conference on Multimedia*, MM '09, pages 1023–1024, New York, NY, USA, 2009. ACM.

[39] Kyo-il Chung. Mobile rfid technologies in korea. http://www.gs1tw.org/twct/gs1w/download/GS1_Kyoil_Chung_Korea.pdf. [Online; accessed 21-October-2011].

[40] Maria Francesca Costabile, Carmelo Ardito, and Rosa Lanzilotti. Enjoying cultural heritage thanks to mobile technology. *interactions*, 17:30–33, May 2010.

[41] Murray Crease and Stephen Brewster. Making progress with sounds - the design & evaluation of an audio progress bar. In *Proceedings of the 1998 International Conference on Auditory Display*, ICAD'98, pages 9–9, Swinton, UK, UK, 1998. British Computer Society.

[42] Peter Dalsgaard and Lone Koefoed Hansen. Performing perception–staging aesthetics of interaction. *ACM Trans. Comput.-Hum. Interact.*, 15(3):13:1–13:33, 2008.

[43] Areti Damala, Pierre Cubaud, Anne Bationo, Pascal Houlier, and Isabelle Marchal. Bridging the gap between the digital and the physical: design and evaluation of a mobile augmented reality guide for the museum visit. In *Proceedings of the 3rd international conference on Digital Interactive*

Media in Entertainment and Arts, DIMEA '08, pages 120–127, New York, NY, USA, 2008. ACM.

[44] Nigel Davies, Keith Cheverst, Alan Dix, and Andre Hesse. Understanding the role of image recognition in mobile tour guides. In *Proceedings of the 7th international conference on Human computer interaction with mobile devices & services*, MobileHCI '05, pages 191–198, New York, NY, USA, 2005. ACM.

[45] Norman K. Denzin. *Sociological Methods: A Sourcebook.* Methodological perspectives. Aldine Transaction, 2006.

[46] Rogerio DePaula. A new era in human computer interaction: the challenges of technology as a social proxy. In *Proceedings of the Latin American conference on Human-computer interaction*, CLIHC '03, pages 219–222, New York, NY, USA, 2003. ACM.

[47] Christina Dicke, Shaleen Deo, Mark Billinghurst, Nathan Adams, and Juha Lehikoinen. Experiments in mobile spatial audio-conferencing: Key-based and gesture-based interaction. In *Proceedings of the 10th International Conference on Human Computer Interaction with Mobile Devices and Services*, MobileHCI '08, pages 91–100, New York, NY, USA, 2008. ACM.

[48] Alan Dix, Jennifer G Sheridan, Stuart Reeves, Steve Benford, and Claire OMalley. Formalising performative interaction. In *Interactive Systems. Design, Specification, and Verification*, pages 15–25. Springer, 2006.

[49] Paul Dourish. *Where the action is: the foundations of embodied interaction.* MIT press, 2004.

[50] Arvid Engström, Mark Perry, and Oskar Juhlin. Amateur vision and recreational orientation:: Creating live video together. In *Proc. CSCW '12*, pages 651–660. ACM, 2012.

[51] Silke Feldmann, Kyandoghere Kyamakya, Ana Zapater, and Zighuo Lue. An indoor bluetooth-based positioning system: Concept, implementation

and experimental evaluation. In *International Conference on Wireless Networks*, pages 109–113, 2003.

[52] Brian Ferris, Dirk Hhnel, and Dieter Fox. Gaussian processes for signal strength-based location estimation. In *In Proc. of Robotics Science and Systems*, 2006.

[53] Lesley Fosh, Steve Benford, Stuart Reeves, and Boriana Koleva. Gifting personal interpretations in galleries. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems*, CHI '14, pages 625–634, New York, NY, USA, 2014. ACM.

[54] Lesley Fosh, Steve Benford, Stuart Reeves, Boriana Koleva, and Patrick Brundell. See me, feel me, touch me, hear me: trajectories and interpretation in a sculpture garden. In *Proc. CHI '13*, pages 149–158. ACM, 2013.

[55] Mike Fraser, Danaë Stanton, Kher Hui Ng, Steve Benford, Claire O'Malley, John Bowers, Gustav Taxén, Kieran Ferris, and Jon Hindmarsh. Assembling history: achieving coherent experiences with diverse technologies. In *Proceedings of the eighth conference on European Conference on Computer Supported Cooperative Work*, pages 179–198, Norwell, MA, USA, 2003. Kluwer Academic Publishers.

[56] Anar Gasimov, Fabio Magagna, and Juliana Sutanto. Camb: Context-aware mobile browser. In *Proceedings of the 9th International Conference on Mobile and Ubiquitous Multimedia*, MUM '10, pages 22:1–22:5, New York, NY, USA, 2010. ACM.

[57] E. Goffman. *The presentation of self in everyday life*. Doubleday anchor books. Doubleday, 1959.

[58] E. Goffman. *Behavior in Public Places; Notes on the Social Organization of Gatherings*. Free Press of Glencoe, 1966.

[59] Joy Goodman, Stephen Brewster, and Philip Gray. Using field experiments to evaluate mobile guides. In *Proc. HCI in Mobile Guides, workshop at Mobile HCI '04*, 2004.

[60] Google. Google glass. `http://www.google.com/glass/start/`. Accessed: 2014-12-14.

[61] Andrew Greaves and Enrico Rukzio. View & share: supporting co-present viewing and sharing of media using personal projection. In *Proc. MobileHCI '09*, pages 44:1–44:4, 2009.

[62] Stefan Greuter and David J. Roberts. Spacewalk: Movement and interaction in virtual space with commodity hardware. In *Proceedings of the 2014 Conference on Interactive Entertainment*, IE2014, pages 1:1–1:7, New York, NY, USA, 2014. ACM.

[63] Nancy Grudens-Schuck, Beverlyn Lundy Allen, and Kathlene Larson. Methodology brief: Focus group fundamentals. *Iowa State University Extension and Outreach*, 2004.

[64] Tony Hall, Luigina Ciolfi, Liam Bannon, Mike Fraser, Steve Benford, John Bowers, Chris Greenhalgh, Sten-Olof Hellström, Shahram Izadi, Holger Schnädelbach, and Martin Flintham. The visitor as virtual archaeologist: explorations in mixed reality technology to enhance educational and social interaction in the museum. In *Proceedings of the 2001 conference on Virtual reality, archeology, and cultural heritage*, VAST '01, pages 91–96, New York, NY, USA, 2001. ACM.

[65] Robert Hardy and Enrico Rukzio. Touch & interact: touch-based interaction with a tourist application. In *Proceedings of the 10th international conference on Human computer interaction with mobile devices and services*, MobileHCI '08, pages 531–534, New York, NY, USA, 2008. ACM.

[66] Graham Healy and Alan F. Smeaton. An outdoor spatially-aware audio playback platform exemplified by a virtual zoo. In *Proceedings of the 17th ACM International Conference on Multimedia*, MM '09, pages 837–840, New York, NY, USA, 2009. ACM.

[67] Florian Heller, Thomas Knott, Malte Weiss, and Jan Borchers. Multi-user interaction in virtual audio spaces. In *Proc. CHI EA '09*, pages 4489–4494. ACM, 2009.

[68] Florian Heller, Aaron Krämer, and Jan Borchers. Simplifying orientation measurement for mobile audio augmented reality applications. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems*, CHI '14, pages 615–624, New York, NY, USA, 2014. ACM.

[69] Simon Holland, David R. Morse, and Henrik Gedenryd. Audiogps: Spatial audio navigation with a minimal attention interface. *Personal Ubiquitous Comput.*, 6(4):253–259, January 2002.

[70] Eva Hornecker and Jacob Buur. Getting a grip on tangible interaction: A framework on physical space and social interaction. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems*, CHI '06, pages 437–446, New York, NY, USA, 2006. ACM.

[71] Hiroshi Ishii and Brygg Ullmer. Tangible bits: Towards seamless interfaces between people, bits and atoms. In *Proceedings of the ACM SIGCHI Conference on Human Factors in Computing Systems*, CHI '97, pages 234–241, New York, NY, USA, 1997. ACM.

[72] Christian Jacquemin, Wai Kit Chan, and Mathieu Courgeon. Bateau ivre: an artistic markerless outdoor mobile augmented reality installation on a riverboat. In *Proc. MM '10*, pages 1353–1362. ACM, 2010.

[73] Carlo Jacucci, Giulio Jacucci, Ina Wagner, and Thomas Psik. A manifesto for the performative development of ubiquitous media. In *Proc. CC '05*, pages 19–28. ACM, 2005.

[74] Mohit Jain, Jeremy Birnholtz, Edward Cutrell, and Ravin Balakrishnan. Exploring display techniques for mobile collaborative learning in developing regions. In *Proc. MobileHCI '11*, pages 81–90, 2011.

[75] Rui Jesus, Ricardo Dias, Rute Frias, Arnaldo J. Abrantes, and Nuno Correia. Memoria mobile: Sharing pictures of a point of interest. In *Proceedings of the Working Conference on Advanced Visual Interfaces*, AVI '08, pages 412–415, New York, NY, USA, 2008. ACM.

[76] Matt Jones. Journeying toward extravagant, expressive, place-based computing. *interactions*, 18(1):26–31, 2011.

[77] Matt Jones and Steve Jones. The music is the message. *interactions*, 13:24–27, July 2006.

[78] Matt Jones, Steve Jones, Gareth Bradley, Nigel Warren, David Bainbridge, and Geoff Holmes. Ontrack: Dynamically adapting music playback to support navigation. *Personal Ubiquitous Comput.*, 12(7):513–525, October 2008.

[79] Maurits Clemens Kaptein, Clifford Nass, and Panos Markopoulos. Powerful and consistent analysis of likert-type ratingscales. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems*, CHI '10, pages 2391–2394, New York, NY, USA, 2010. ACM.

[80] Fahim Kawsar, Enrico Rukzio, and Gerd Kortuem. An explorative comparison of magic lens and personal projection for interacting with smart objects. In *Proc. MobileHCI '10*, pages 157–160, 2010.

[81] Michael Kenteris, Damianos Gavalas, and Daphne Economou. Electronic mobile guides: a survey. *Personal Ubiquitous Computing*, 15(1):97–111, 2011.

[82] KewGardens. Kew gardens mobile application. `http://www.kew.org/visit-kew-gardens/visit-information/garden-guides/mobile-app/`. Accessed: 2013-05-15.

[83] Karen Johanne Kortbek and Kaj Grønbæk. Communicating art through interactive technology: new approaches for interaction design in art museums. In *Proceedings of the 5th Nordic conference on Human-computer interaction: building bridges*, NordiCHI '08, pages 229–238, New York, NY, USA, 2008. ACM.

[84] Graffiti Research Lab. L.A.S.E.R. tag. `http://www.graffitiresearchlab.com/blog/projects/laser-tag/`. Accessed: 2013-08-15.

[85] Rock Leung, Karon MacLean, Martin Bue Bertelsen, and Mayukh Saubhasik. Evaluation of haptically augmented touchscreen gui elements under cognitive load. In *Proceedings of the 9th International Conference on Multimodal Interfaces*, ICMI '07, pages 374–381, New York, NY, USA, 2007. ACM.

[86] Shuai Li, Yuesheng Lou, and Bo Liu. Bluetooth aided mobile phone localization: A nonlinear neural circuit approach. *ACM Trans. Embed. Comput. Syst.*, 13(4):78:1–78:15, March 2014.

[87] Rensis Likert. A technique for the measurement of attitudes. *Archives of psychology*, 1932.

[88] Markus Löchtefeld. Advanced interaction with mobile projection interfaces. In *Proceedings of the 24th annual ACM symposium adjunct on User interface software and technology*, UIST '11 Adjunct, pages 43–46, New York, NY, USA, 2011. ACM.

[89] Vince Mancuso. Take me home: Designing safer in-vehicle navigation devices. In *CHI '09 Extended Abstracts on Human Factors in Computing Systems*, CHI EA '09, pages 4591–4596, New York, NY, USA, 2009. ACM.

[90] Jani Mantyjarvi, Fabio Paternò, Zigor Salvador, and Carmen Santoro. Scan and tilt: towards natural interaction for mobile museum guides. In *Proceedings of the 8th conference on Human-computer interaction with mobile devices and services*, MobileHCI '06, pages 191–194, New York, NY, USA, 2006. ACM.

[91] Georgios N. Marentakis and Stephen A. Brewster. Effects of feedback, mobility and index of difficulty on deictic spatial audio target acquisition in the horizontal plane. In *Proc. CHI '06*, pages 359–368. ACM, 2006.

[92] Paul Marshall, Richard Morris, Yvonne Rogers, Stefan Kreitmayer, and Matt Davies. Rethinking 'multi-user': an in-the-wild study of how groups approach a walk-up-and-use tabletop interface. In *Proceedings of the 2011 annual conference on Human factors in computing systems*, CHI '11, pages 3033–3042, New York, NY, USA, 2011. ACM.

[93] David McGookin and Stephen Brewster. Pulse: the design and evaluation of an auditory display to provide a social vibe. In *Proceedings of the 2012 ACM annual conference on Human Factors in Computing Systems*, CHI '12, pages 1263–1272, New York, NY, USA, 2012. ACM.

[94] David Mcgookin, Stephen Brewster, and Pablo Priego. Audio bubbles: Employing non-speech audio to support tourist wayfinding. In *Proceedings of the 4th International Conference on Haptic and Audio Interaction Design*, HAID '09, pages 41–50, Berlin, Heidelberg, 2009. Springer-Verlag.

[95] David McGookin, Yolanda Vazquez-Alvarez, Stephen Brewster, and Joanna Bergstrom-Lehtovirta. Shaking the dead: multimodal location based experiences for un-stewarded archaeological sites. In *Proc. NordiCHI '12*, pages 199–208. ACM, 2012.

[96] Microsoft. Microsoft hololens. `http://www.microsoft.com/microsoft-hololens/`. Accessed: 2015-01-23.

[97] Nigel Newton. The use of semi-structured interviews in qualitative research: strengths and weaknesses. `http://www.academia.edu/1561689/The_use_of_semi-structured_interviews_in_qualitative_research_strengths_and_weaknesses`, 2010. Accessed: 2014-10-27.

[98] Jakob Nielsen. Iterative user-interface design. *Computer*, 26(11):32–41, 1993.

[99] Tsouknidas Nikolaos and Tomimatsu Kiyoshi. Qr-code calibration for mobile augmented reality applications: Linking a unique physical location to the digital world. In *ACM SIGGRAPH 2010 Posters*, SIGGRAPH '10, pages 144:1–144:1, New York, NY, USA, 2010. ACM.

[100] Ian Oakley and Jun-Seok Park. Designing eyes-free interaction. In Ian Oakley and Stephen Brewster, editors, *Haptic and Audio Interaction Design*, volume 4813 of *Lecture Notes in Computer Science*, pages 121–132. Springer Berlin Heidelberg, 2007.

[101] Kenton O'Hara. Understanding geocaching practices and motivations. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems*, CHI '08, pages 1177–1186, New York, NY, USA, 2008. ACM.

[102] Antti Oulasvirta. The fragmentation of attention in mobile interaction, and what to do with it. *interactions*, 12(6):16–18, November 2005.

[103] Antti Oulasvirta. Field experiments in hci: Promises and challenges. In Hannakaisa Isomki and Pertti Saariluoma, editors, *Future Interaction Design II*, pages 87–116. Springer London, 2009.

[104] Helen Petrie, Steven Furner, and Thomas Strothotte. Design lifecycles and wearable computers for users with disabilities. In *First Workshop on Human-Computer Interaction with Mobile Devices, Glasgow, UK*, 1998.

[105] Antti Pirhonen, Stephen Brewster, and Christopher Holguin. Gestural and audio metaphors as a means of control for mobile devices. In *Proc. CHI '02*, pages 291–298. ACM, 2002.

[106] Omer Rashid, Will Bamford, Paul Coulton, Reuben Edwards, and Jurgen Scheible. Pac-lan: mixed-reality gaming with rfid-enabled mobile phones. *Comput. Entertain.*, 4, October 2006.

[107] Fiona Redhead and Margot Brereton. Iterative design within a local community communication fabric. In *Proceedings of the 22nd Conference of the Computer-Human Interaction Special Interest Group of Australia on Computer-Human Interaction*, OZCHI '10, pages 388–391, New York, NY, USA, 2010. ACM.

[108] Stuart Reeves, Steve Benford, Claire O'Malley, and Mike Fraser. Designing the spectator experience. In *Proc. CHI '05*, pages 741–750. ACM, 2005.

[109] Josephine Reid, Erik Geelhoed, Richard Hull, Kirsten Cater, and Ben Clayton. Parallel worlds: immersion in location-based experiences. In *Proc. CHI EA '05*, pages 1733–1736. ACM, 2005.

[110] Ingrid Richardson, Amanda Third, and Ian MacColl. Moblogging and belonging: new mobile phone practices and young people's sense of social

inclusion. In *Proceedings of the 2nd international conference on Digital interactive media in entertainment and arts*, DIMEA '07, pages 73–78, New York, NY, USA, 2007. ACM.

[111] Julie Rico and Stephen Brewster. Usable gestures for mobile interfaces: evaluating social acceptability. In *Proc. CHI '10*, pages 887–896. ACM, 2010.

[112] Jane Ritchie, Jane Lewis, Carol McNaughton Nicholls, Rachel Ormston, et al. *Qualitative research practice: A guide for social science students and researchers.* Sage, 2013.

[113] S. Robinson. *Eyes-Off Physically Grounded Mobile Interaction. PhD thesis.* PhD thesis, Swansea University, 2012.

[114] Simon Robinson, Parisa Eslambolchilar, and Matt Jones. Evaluating haptics for information discovery while walking. In *Proceedings of the 23rd British HCI Group Annual Conference on People and Computers: Celebrating People and Technology*, BCS-HCI '09, pages 93–102, Swinton, UK, UK, 2009. British Computer Society.

[115] Simon Robinson, Matt Jones, Elina Vartiainen, and Gary Marsden. Picotales: Collaborative authoring of animated stories using handheld projectors. In *Proceedings of the ACM 2012 Conference on Computer Supported Cooperative Work*, CSCW '12, pages 671–680, New York, NY, USA, 2012. ACM.

[116] Simon Robinson, Matt Jones, John Williamson, Roderick Murray-Smith, Parisa Eslambolchilar, and Mads Lindborg. Navigation your way: From spontaneous independent exploration to dynamic social journeys. *Personal and Ubiquitous Computing*, 16:973–985, 2012.

[117] Simon Robinson, Matt Jones, John Williamson, Roderick Murray-Smith, Parisa Eslambolchilar, and Mads Lindborg. Navigation your way: From spontaneous independent exploration to dynamic social journeys. *Personal and Ubiquitous Computing*, Published online 20 September 2011.

[118] Simon Robinson, Elina Vartiainen, Matt Jones, and Gary Marsden. Picotales: Collaborative authoring of animated stories using handheld projectors. In *CSCW '12: Proceedings of the ACM 2012 Conference on Computer Supported Cooperative Work*, pages 671–680, New York, NY, USA, 2012. ACM.

[119] Y. Rogers, S. Price, G. Fitzpatrick, R. Fleck, E. Harris, H. Smith, C. Randell, H. Muller, C. O'Malley, D. Stanton, M. Thompson, and M. Weal. Ambient wood: designing new forms of digital augmentation for learning outdoors. In *Proc. IDC '04*, pages 3–10. ACM, 2004.

[120] Yvonne Rogers. Interaction design gone wild: Striving for wild theory. *interactions*, 18(4):58–62, July 2011.

[121] Virpi Roto and Antti Oulasvirta. Need for non-visual feedback with long response times in mobile hci. In *Special Interest Tracks and Posters of the 14th International Conference on World Wide Web*, WWW '05, pages 775–781, New York, NY, USA, 2005. ACM.

[122] José Rouillard and Mona Laroussi. PerZoovasive: contextual pervasive qr codes as tool to provide an adaptive learning support. In *Proc. CSTST '08*, pages 542–548, 2008.

[123] Duncan Rowland, Martin Flintham, Leif Oppermann, Joe Marshall, Alan Chamberlain, Boriana Koleva, Steve Benford, and Citlali Perez. Ubikequitous computing: designing interactive experiences for cyclists. In *Proceedings of the 11th International Conference on Human-Computer Interaction with Mobile Devices and Services*, MobileHCI '09, pages 21:1–21:11, New York, NY, USA, 2009. ACM.

[124] Teri Rueb. Core sample. `http://www.terirueb.net/core_sample/index.html`. Accessed: 2013-08-20.

[125] Daniel Sauter. Light attack. In *ACM SIGGRAPH 2006 Art gallery*. ACM, 2006.

[126] Jürgen Scheible and Timo Ojala. Mobispray: mobile phone as virtual spray can for painting big anytime anywhere on anything. In *ACM SIGGRAPH 2009 Art Gallery*, pages 5:1–5:10. ACM, 2009.

[127] B. Schilit, N. Adams, and R. Want. Context-aware computing applications. In *Mobile Computing Systems and Applications, 1994. WMCSA 1994. First Workshop on*, pages 85–90, Dec 1994.

[128] W. Schilit. *A system architecture for context-aware mobile computing. PhD thesis.* PhD thesis, Columbia University, 1995.

[129] Jochen Schiller and Agnès Voisard. *Location-based services.* Elsevier, 2004.

[130] Holger Schndelbach, Boriana Koleva, Martin Flintham, Mike Fraser, Shahram Izadi, Paul Chandler, Malcolm Foster, Steve Benford, Chris Greenhalgh, and Tom Rodden. Augurscope: a mixed reality interface for outdoors. In *Proc. CHI '02*, pages 9–16, 2002.

[131] Julia Schwarz, David Klionsky, Chris Harrison, Paul Dietz, and Andrew Wilson. Phone as a pixel: Enabling ad-hoc, large-scale displays using mobile devices. In *Proc. CHI '12*, pages 2235–2238. ACM, 2012.

[132] Jennifer G. Sheridan, Nick Bryan-Kinns, and Alice Bayliss. Encouraging witting participation and performance in digital live art. In *Proceedings of the 21st British HCI Group Annual Conference on People and Computers: HCI...But Not As We Know It - Volume 1*, BCS-HCI '07, pages 13–23, Swinton, UK, UK, 2007. British Computer Society.

[133] Todd Simcock, Stephen Peter Hillenbrand, and Bruce H. Thomas. Developing a location based tourist guide application. In *Proceedings of the Australasian information security workshop conference on ACSW frontiers 2003 - Volume 21*, ACSW Frontiers '03, pages 177–183, Darlinghurst, Australia, Australia, 2003. Australian Computer Society, Inc.

[134] Jocelyn Spence, David M. Frohlich, and Stuart Andrews. Performative experience design. In *CHI '13 Extended Abstracts on Human Factors in*

Computing Systems, CHI EA '13, pages 2049–2058, New York, NY, USA, 2013. ACM.

[135] Christoph Stahl. The roaring navigator: a group guide for the zoo with shared auditory landmark display. In *Proc. MobileHCI '07*, pages 383–386. ACM, 2007.

[136] Steven Strachan, Parisa Eslambolchilar, Roderick Murray-Smith, Stephen Hughes, and Sile O'Modhrain. Gpstunes: controlling navigation via audio feedback. In *Proceedings of the 7th international conference on Human computer interaction with mobile devices & services*, MobileHCI '05, pages 275–278, New York, NY, USA, 2005. ACM.

[137] Ewan Sutherland. Counting customers, subscribers and mobile phone numbers. *info*, 11(2):6–23, 2009.

[138] Colin Swindells and Kori M. Inkpen. What is that?: gesturing to determine device identity. In *CHI '02 extended abstracts on Human factors in computing systems*, CHI EA '02, pages 588–589, New York, NY, USA, 2002. ACM.

[139] Colin Swindells, Kori M. Inkpen, John C. Dill, and Melanie Tory. That one there! pointing to establish device identity. In *Proceedings of the 15th annual ACM symposium on User interface software and technology*, UIST '02, pages 151–160, New York, NY, USA, 2002. ACM.

[140] Robyn Taylor, Guy Schofield, John Shearer, Jayne Wallace, Peter Wright, Pierre Boulanger, and Patrick Olivier. Designing from within: humanaquarium. In *Proc. CHI '11*, pages 1855–1864. ACM, 2011.

[141] Nava Tintarev, Ana Flores, and Xavier Amatriain. Off the beaten track: a mobile field study exploring the long tail of tourist recommendations. In *Proceedings of the 12th international conference on Human computer interaction with mobile devices and services*, MobileHCI '10, pages 209–218, New York, NY, USA, 2010. ACM.

[142] TomTom. Drive (satellite navigation). `http://www.tomtom.com/en_gb/drive/car/`. Accessed: 2014-12-11.

[143] Sherry Turkle. *Alone Together: Why We Expect More from Technology and Less from Each Other.* Basic Books, Inc., New York, NY, USA, 2011.

[144] Dieter Uckelmann, Mark Harrison, and Florian Michahelles. *Architecting the Internet of Things.* Springer Publishing Company, Incorporated, 1st edition, 2011.

[145] Yolanda Vazquez-Alvarez and Stephen Brewster. Investigating background & foreground interactions using spatial audio cues. In *Proceedings of the 27th international conference extended abstracts on Human factors in computing systems*, CHI EA '09, pages 3823–3828, New York, NY, USA, 2009. ACM.

[146] Yolanda Vazquez-Alvarez and Stephen A. Brewster. Eyes-free multitasking: The effect of cognitive load on mobile spatial audio interfaces. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems*, CHI '11, pages 2173–2176, New York, NY, USA, 2011. ACM.

[147] Yolanda Vazquez-Alvarez, Ian Oakley, and Stephen A. Brewster. Auditory display design for exploration in mobile audio-augmented reality. *Personal Ubiquitous Comput.*, 16(8):987–999, December 2012.

[148] Vassilios Vlahakis, John Karigiannis, Manolis Tsotros, Michael Gounaris, Luis Almeida, Didier Stricker, Tim Gleue, Ioannis T. Christou, Renzo Carlucci, and Nikos Ioannidis. Archeoguide: first results of an augmented reality, mobile computing system in cultural heritage sites. In *Proc. VAST '01*, pages 131–140. ACM, 2001.

[149] Ge Wang, Georg Essl, Deutsche Telekom, and Henri Penttinen. Mopho: Do mobile phones dream of electric orchestras. In *Proc. ICMC 2008*, page 29, 2008.

[150] Roy Want, Andy Hopper, Veronica Falcão, and Jonathan Gibbons. The active badge location system. *ACM Trans. Inf. Syst.*, 10(1):91–102, January 1992.

[151] Roy Want, BillN. Schilit, NormanI. Adams, Rich Gold, Karin Petersen, David Goldberg, JohnR. Ellis, and Mark Weiser. The parctab ubiquitous computing experiment. In Tomasz Imielinski and HenryF. Korth, editors, *Mobile Computing*, volume 353 of *The Kluwer International Series in Engineering and Computer Science*, pages 45–101. Springer US, 1996.

[152] Alan J. Wecker, Joel Lanir, Tsvi Kuflik, and Oliviero Stock. Pathlight: supporting navigation of small groups in the museum context. In *Proc. MobileHCI '11*, pages 569–574. ACM, 2011.

[153] Leonard Wein. Visual recognition in museum guide apps: Do visitors want it? In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems*, CHI '14, pages 635–638, New York, NY, USA, 2014. ACM.

[154] D.J. Weintraub, M. Ensing, and Crew System Ergonomics Information Analysis Center. *Human Factors Issues in Head-up Display Design: the Book of HUD.* CSERIAC SOAR. CSERIAC, 1992.

[155] Mark Weiser. The computer for the 21st century. *SIGMOBILE Mob. Comput. Commun. Rev.*, 3(3):3–11, July 1999.

[156] John Williamson, Simon Robinson, Craig Stewart, Roderick Murray-Smith, Matt Jones, and Stephen Brewster. Social gravity: A virtual elastic tether for casual, privacy-preserving pedestrian rendezvous. In *CHI '10: Proceedings of the 28th International Conference on Human Factors in Computing Systems*, pages 1485–1494, New York, NY, USA, 2010. ACM.

[157] Julie Rico Williamson. Send me bubbles: multimodal performance and social acceptability. In *Proc. CHI EA '11*, pages 899–904. ACM, 2011.

[158] Karl Willis, Ivan Poupyrev, and Takaaki Shiratori. Motionbeam: a metaphor for character interaction with handheld projectors. In *Proc. CHI '11*, pages 1031–1040, 2011.

[159] Max L. Wilson, Simon Robinson, Dan Craggs, Kristian Brimble, and Matt Jones. Pico-ing into the future of mobile projector phones. In *Proceedings of the 28th of the international conference extended abstracts on Human factors in computing systems*, CHI EA '10, pages 3997–4002, New York, NY, USA, 2010. ACM.

[160] Toshiya Yamada, Jun Shingu, Elizabeth Churchill, Les Nelson, Jonathan Helfman, and Paul Murphy. Who cares?: reflecting who is reading what on distributed community bulletin boards. In *Proceedings of the 17th annual ACM symposium on User interface software and technology*, UIST '04, pages 109–118, New York, NY, USA, 2004. ACM.

[161] Nikoleta Yiannoutsou, Ioanna Papadimitriou, Vassilis Komis, and Nikolaos Avouris. "playing with" museum exhibits: designing educational games mediated by mobile technology. In *Proceedings of the 8th International Conference on Interaction Design and Children*, IDC '09, pages 230–233, New York, NY, USA, 2009. ACM.