Thresholds: Embedding Virtual Reality in the Museum

PAUL TENNENT, University of Nottingham, Nottingham, UK
SARAH MARTINDALE, University of Nottingham, Nottingham, UK
STEVE BENFORD, University of Nottingham, Nottingham, UK
DIMITRIOS DARZENTAS, University of Nottingham, Nottingham, UK
PAT BRUNDELL, University of Nottingham, Nottingham, UK
MATT COLLISHAW, Artist, UK

Abstract. We examine the experience of Thresholds, a virtual reality (VR) recreation of the world's first photographic exhibition, which has toured to multiple museums. Following the method of performance-led research in the wild, we provide an account of the artist's design rationale and the experiences of visitors as the work toured. We reveal how the overlaying and juxtaposing of virtual and physical spaces established a VR architecture that underpinned the extended user experience. Overlaying was used to layer a virtual model on to a corresponding physical set to deliver physical sensations of touch and movement alongside visual and audio stimuli. Juxtaposition was used to embed the VR installation within the surrounding gallery space at each host museum, dealing with the challenges of entering, exiting, spectating and invigilating the experience. We propose that museum designers can use these techniques to deliver VR installations that are compelling but also scalable and tourable.

CCS → Human-centered computing → Human computer interaction (HCI) → Interaction paradigms → **Mixed / augmented reality**

General terms: Design

Additional keywords and phrases: virtual reality, museum, passive haptics, augmented reality, tactile, touch, exhibition, interactive art, mixed reality boundaries.

INTRODUCTION

Virtual reality (VR) is a promising technology for museums, bringing the opportunity to immerse visitors in compelling and interactive interpretations of the past. The recent emergence of commodity headsets suggests that it may soon be practical and affordable for many museums to harness this potential. This raises the question of how museum professionals might set about integrating VR into their museums in a way that delivers compelling experiences while also being scalable in terms of visitor throughput and flexible to adapt to different settings.

We present a case study of designing and touring a large-scale public VR installation that tackled these challenges. Thresholds was a touring immersive artwork, created by a renowned professional artist, that recreated the world's first photographic exhibition as a room-sized installation that could be experienced by up to six participants at a time. Its most immediately striking feature was the use of passive haptics [1][19] to overlay the virtual recreation on a corresponding physical set, aligning the two to deliver physical sensations of touch and movement alongside visual and auditory stimuli. Moreover, Thresholds was also carefully designed to be a tourable experience, employing this virtual set as a 'box within a box' that could be readily installed in different galleries and museums. At the time of writing, the work had visited four galleries, been experienced by over 10,000 visitors, and had secured bookings for continued touring beyond that.

http://dx.doi.org/10.1145/3369394

We document and reflect on how the artist designed and implemented Thresholds and how visitors subsequently experienced the work in order to draw out key lessons for the design of VR user experiences in museums. Specifically, we show how the artist both overlaid and juxtaposed virtual and physical spaces to create an appropriate 'architecture for interaction' that met the challenges of delivering a compelling experience at scale with the flexibility to be installed in many museums.

RELATED WORK

VR has been widely applied to cultural heritage. To briefly summarise just a few examples (see [4] for a comprehensive survey), VR enables visitors to experience virtual exhibitions either off site or in a dedicated space in the museum as a form of virtual tourism [24][49][50][51], may improve accessibility [17], or enable communication between local and remote visitors [6]. Examples of virtual recreations include: an Etruscan town [13]; the Labyrinth of Versailles [15]; Nefertiti's tomb [52]; and a complete recreation of ancient Rome [12] amongst many others. Augmented reality has been used to overlaying digital content on existing exhibits, for example the Smithsonian's *Skin and Bones* [53], to create virtual guides [26], or introduce exhibits which are not physically present [40], [42] or [45], including reconstructing buildings outside of the museum as in *Dead men's eyes* [9] and the *Augurscope* [30].

One of the major challenges facing VR involves delivering physical sensations of touch and movement alongside visual and auditory stimuli. One way of addressing this is to *overlay* a virtual model onto a corresponding physical set, an approach known as passive haptics [18][19]. Previous research has explored how to extend this with dynamically reconfigurable props [1], props that are created 'on the fly' either as a single deformable object [11], assembled by robots [47], or underfoot to create representative floor surfaces [36], or that are configured in real time by human stagehands who are invisible to the user who is wearing a headmounted display [7].

A quite different approach to combining physical and virtual spaces is to *juxtapose* rather than overlay them. Koleva et al. introduced the concept of Mixed Reality Boundaries as points of connection between physical and virtual realities. As with everyday windows and doors, the spaces involved remain separate, but become connected, with their inhabitants being able to look or move between them. For example, the occupant of a physical room might peer into a connected virtual room through a 'window' projected onto the wall, while an occupant of that virtual room might look through a video texture. They presented a typology of mixed reality boundaries in which virtual inhabitants on one side and physical ones on the other could variously see and hear across permanent connections between their spaces [22] and considered how they might be made physically and virtually traversable [23].

Like many of the examples listed above, the work presented below engages visitors with a historic recreation. Its novelty lies in combining the two approaches of overlaying and juxtaposing so as to create an immersive tactile experience that also delivers sufficient throughput and flexibility that it can viably tour to many museums. Our contribution to practice is therefore to guide museum designers who wish to employ VR at scale as to how they might deliver similar experiences. Our technical contribution is to show how the combination of overlaying and juxtaposing real and virtual spaces can establish an appropriate architecture for extended VR user experiences.

APPROACH

We follow the approach of ‘performance-led research in the wild’ [5]. This involves collaborating with artists to develop new cultural works, touring these to venues such as galleries and museums so they can be experienced by public audiences, capturing rich documentation of the entire process, from the artist’s detailed rationale, including how they tackled key design challenges, to audience feedback. Reflection across all of this data then reveals wider design principles. As a form of Research Through Design [14,48] the approach is practice led, rather than hypothesis driven, with research findings emerging from the detailed and typically iterative process of making public experiences. There has been a particular emphasis on working with professional artists to create touring cultural products which requires refining designs and implementations to the point where they work robustly and at scale to the broad satisfaction of artists, venues, commissioners, audiences and perhaps even critics. The evaluation of such work is necessarily broad, engaging stakeholders throughout the design and touring process.

In this case, we have collaborated with the artist Matt Collishaw to create an installation called Thresholds, a touring virtual artwork that enables visitors to explore a historic recreation of one of the world’s first public exhibitions of photography. Collishaw is a renowned artist of international standing who first emerged as part of the Young British Artists movement of the late 1980s and 1990s [37] and who has made numerous photographic, video and interactive work over the decades since then. He approached us with the core concept for Thresholds with a view to us helping realise the work, a proposition to which we readily agreed.

INTRODUCING THRESHOLDS

Part history, part art, part science and part speculation, Thresholds is inspired by the ‘Model Room’, an exhibition that was staged at King Edward’s School (Birmingham, UK) in August 1839 at which photography pioneer Henry Fox Talbot presented a display of 93 ‘Photogenic Drawings’ (photographs).

Each visitor to Thresholds is equipped with a backpack pc wireless head-mounted display that enables them to explore a room-size VR recreation of the Model Room with up to five other visitors (Figure 1 top). They are guided into an all-white physical room which contains furniture in the form of model vitrines and whose walls feature blank outlines of windows, picture frames and other interior details (Figure 1, bottom). However, their headset ‘skins’ this blank physical canvas with the virtual world so that they can see and hear Collishaw’s recreation of the Model Room but also feel it whenever they reach out to touch a vitrine, lean against a wall or otherwise physically encounter the environment. Walking around, they can touch everything, from the vitrines, to the frames of paintings on the walls. They can peer out of windows into the mist to see and hear angry protesters outside. They can pull the photographs out of the vitrines for closer examination. A fire burns in the grate which feels warm and smells of wood smoke, while moths flit around the gaslights, and mice scuttle around the recesses of the room. Other visitors are represented as ghostlike auras, as are their own hands. A clock slowly ticks and, when six minutes have elapsed, chimes and they are asked to remove their headset, to find themselves once more in the bright white reality of the physical exhibit. In technical terms, Collishaw's vision is to immerse visitors in a room-sized, multi-sensory and multi-user VR.

Figure 1. Thresholds fuses virtual and physical realities. Top: the recreation of Fox Talbot's 1839 exhibition in the Model Room. Bottom: the physical 'canvas' on which this is overlaid.

The design process unfolded over 18 months and involved the artist, historical advisors, virtual modellers, physical builders and the research team. A shared CAD model was initially developed and used as a basis for development of both the physical and virtual spaces. From this, a series of iterations of both physical and virtual model were developed, assessed and refined. The physical/virtual alignment system was developed in parallel with these models, using stand-in objects of increasing scale and resolution. The complete system was first tested at a dedicated studio around six weeks before its premiere, at which time the historical advisors were able to further assess and refine it. Thresholds has been exhibited at four locations in the UK at the time of writing: Somerset House in London, Birmingham Museum and Art Gallery; Lacock Abbey; National Science and Media Museum, Bradford. It also visited Yapi Kredi Kültür Sanat Yayincilik A.S. in Istanbul, Turkey from May to July 2018. It has been experienced by more than ten thousand members of the public to date. In what follows, we explore Thresholds from three perspectives: its detailed design, its practical realisation as a touring work using today's VR technology, and the experiences of visitors.

DESIGNING THRESHOLDS

There were two distinct aspects to the artist's rationale: making an artistic interpretation of Fox Talbot's 1839 photographic exhibition; and ensuring that the work would be viable for touring. Artistically, Thresholds is not intended to be a straightforward historical re-creation; rather it is an artwork – a meditation on how technology changes our relationship with the world and a comment on how a technical innovation from 170 years ago has led to the deluge of visual information and simulated realities. Collishaw deliberately set out to juxtapose two technologies, using the currently radical technology of VR as a lens to interpret the previous radical technology of photography. It is perhaps difficult for a contemporary audience to appreciate the impact that photography would have had on an audience in 1839. The use of VR, especially with haptic and tactile extensions, is intended to thrill the modern audience as the artist imagines those at the original exhibition would have been thrilled. Practically, the artist was keen to ensure that Thresholds would be attractive to venues as a touring work, meaning that it would need to fit within their spaces, be easily relocatable, and support a sufficiently high throughput of visitors. These twin artistic and pragmatic concerns shaped the design of Thresholds as we now consider.

The virtual world. Designed by Augustus Pugin and Charles Barry, the original exhibition space is, or rather was (the school having been demolished in 1936), spectacular. Images of the original room do exist (Figure 2), and with support from Barry and Pugin scholars, an 'as faithful as possible' virtual recreation was developed by architectural CGI firm VMI Studios. In the virtual room, the scale, ambience and content of the original exhibition is recreated so that by donning VR headsets visitors are transported 'back in time' and find themselves apparently standing in the 1839 exhibition. Here glass vitrines containing Fox Talbot's prototype images sit alongside renderings of the philosophical instruments, scientific apparatus, inventions and manufacturing processes which shared the space. The choice of images was determined using the handbills from the original exhibition, though in some cases best guesses have been made with the support of photographic historians about which of a series of images might be correct for a given description. In a few cases, the original image has been lost or destroyed before a scan had ever been performed and has been substituted with an appropriate alternative from the existing collection.

Figure 2. Images of the original room. Left: P. Hamilton, Interior of The New Department, 1841, Engraving. Right: Photograph c.1935. Images courtesy of King Edwards Foundation Archive.

Prior to public deployment, the model was shown to architectural and photographic historians in order to test its accuracy and authenticity leading to several updates. Architectural historian David Blissett pointed out that the original rendering of the fireplace, a rather elaborate tiled affair was not consistent with the type of fireplace that Charles Barry would have installed in such a building, instead suggesting it be replaced with stone. Similarly, he had some comments about the initial version of the ceiling carvings, which were made more elaborate as a result. Fox Talbot scholar Larry Schaff queried the use of one particular image we had chosen to replace a missing original – pointing out that it was taken in 1841 – two years *after* the exhibition and so an alternative was found.

A letter exists written by Fox Talbot to J F W Herschel [43] citing his concern about running the exhibition in Birmingham with reports of Chartist demonstrations in the area. The artist wanted to represent this political tension, drawing a parallel with current tensions about the automation of jobs through digital technologies. We created an animated scene outside the windows of the hall, allowing visitors to look out and view Chartist protesters in the street below (Figure 3). The riot begins three minutes into each visitor's experience.

Figure 3. The Chartist riot taking place visible through the window of the Thresholds exhibition's virtual space.

The design includes minor effects (Figure 4) intended to breathe life into the model, direct the visitor's attention and encourage movement and exploration: mice scurry around near the skirting boards; moths flit around the gasoliers; a spider crawls across a painting; the flames in the fire dance; and the clouds and sun move in the sky. Similar attention was paid to the soundscape: a clock ticks loudly and approaches 10PM; a soundscape of a bustling (albeit quiet) large room was created, designed to imply a kind of abstract busyness; each of the avatar-ghosts emits a slightly disconcerting quiet hum, that becomes increasingly noticeable as you get close to them; and the mice's feet skitter slightly on the floor, though only loud enough to hear if you get down to their level. When the visitor's six-minute session elapses, the clock strikes 10.

Figure 4. The ambient life in the virtual room. Left: a spider crawls across a painting. Centre: moths flit about the gasoliers. Right: mice scuttle about the floor near the walls.

Room scale VR. Perhaps the most far reaching design decision was to use room scale VR to allow visitors to physically walk around the virtual world. However, the original, The New Department in King Edward's School, was enormous. The decision was therefore taken early on to allow visitors to physically traverse only part of the room. The artist chose to match one dimension (the width) of the original building with the physical set, leading to a movement area of 8.5m (length) by 6.5m (width) by 2.5m (height), located in the virtual building as shown in Figure 5. The length was determined as a factor of available space to deploy the physical exhibition at contributing venues and the reliable range of the HTC Vive tracking system at that time. The height was, for practical reasons, chosen to be higher than most people could reasonably reach, but low enough to keep down material costs and make access to above-ceiling cabling fairly easy. This reduced area was physically built exactly to scale as a 'room within a room' including all the necessary furniture and wall decorations (Figures 1 & 5).

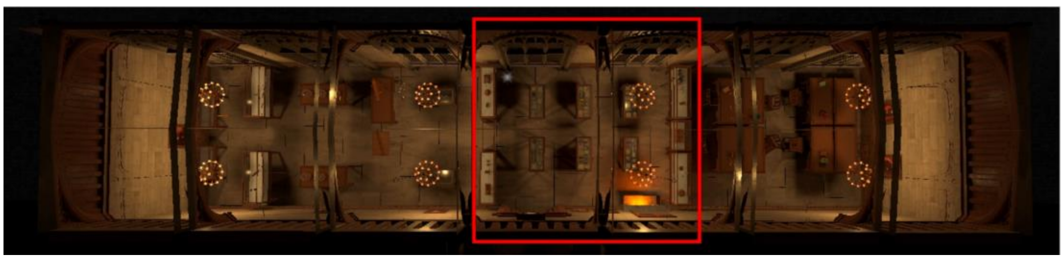

Figure 5. The complete virtual room. Only the area shown in red was constructed as a physical set, so only this area is directly accessible by walking.

Visitors therefore actually walk around a relatively small area within a much larger virtual room (and indeed wider virtual world if one considers the street outside). In order to preserve the illusion of being within a far larger room and discourage inappropriate movement, large virtual cabinets were placed as barriers along the ends of the space. These were transparent so that visitors could see into the virtual space beyond but large enough to apparently block their way. Six vitrines made of wood and glass, each around 1m tall in two rows of three hold photographs (Figure 6), while five further vitrines of double height, display a collection of scientific curiosities (for example flywheels, a microscope and a mirror-based optical illusion).

A concession to the practicalities of delivery was the addition of bars on the windows. While these were not present in the original room, the alcove nature of the windows tended to cause the VR tracking system to fail when users put their heads inside them. The addition of the bars prevents the users from placing their heads in this area, while remaining consistent with the 'feel' of the architecture. A transition happens again as users remove the headset: a sharp return to reality from the vibrant warm colours of the virtual space to the harsh white of the real physical space, this return is intended by the artist to be jarring.

Figure 6. A wood and glass vitrine holds several of the photographs. Other vitrines displaying scientific equipment are visible in the background.

Physical set. A second critical aspect of Thresholds was that the environment that a visitor sees in the virtual world should, as far as possible, be touchable. Aesthetically the artist wished to create a room that looked like a 'canvas', taking cues from films such as 2001: A Space Odyssey and the Matrix. The sci-fi-like white space, particularly when accented with the look of the headsets and computer backpacks, contrasts dramatically with the rich visual nature of the virtual space. Generally, the physical models are quite spartan, with much of the detail added visually through textures in the CGI. This leads to some haptic inconsistencies – for example, the model is largely built of wood, while the CGI version often shows stone. These inconsistencies tend to be borne of practicality – for example a stone floor and fireplace would be difficult to deploy as a touring exhibition. To enhance the virtual fire, a ceramic heater was installed in the physical space, providing a strong heat source, and a wood smoke scent, delivered though a motion activated 'air freshener' was used to further enhance this illusion (Figure 7).

Figure 7. The virtual fire (top) is augmented with a ceramic heater and wood smoke scent (bottom).

Ghostly hands. The artist was keen that the visitors shouldn't have to hold controllers, but rather should interact using their hands. To achieve this we used a head mounted Leap Motion (2017) as we discuss below. Testing revealed a tension between displaying a representation of the user's hands so as to increase immersion while not making this so realistic that it would reveal any misalignments between the virtual world and physical set. Our solution involved making the representations of the visitors' hands 'ghostly' (Figure 8).

Figure 8. Ghostly hands.

Picking up photos. Despite significant improvements, the resolution of today's VR headsets remains limited. For an exhibition of something as visual as photographs it seemed churlish to present them at low resolutions, particularly when we were working with extremely high-resolution scans of the images. After much discussion, the artist elected to compromise the "if you can reach it, you can touch it' rule so as to enable people to look more closely at the images. We therefore introduced a 'summon' gesture: the visitor holds their hand palm down over the image they want to pick up, then turns their hand face up and the photograph flies up into their hand (Figure 9). They can then carry the image around and use a pinch-to-zoom gesture to increase or decrease its size. Turning their hand palm-down again 'drops' the image back to its original location. On balance, it was felt that the benefits of allowing visitors to examine the photographs in detail was a fair exchange for the loss of tactility.

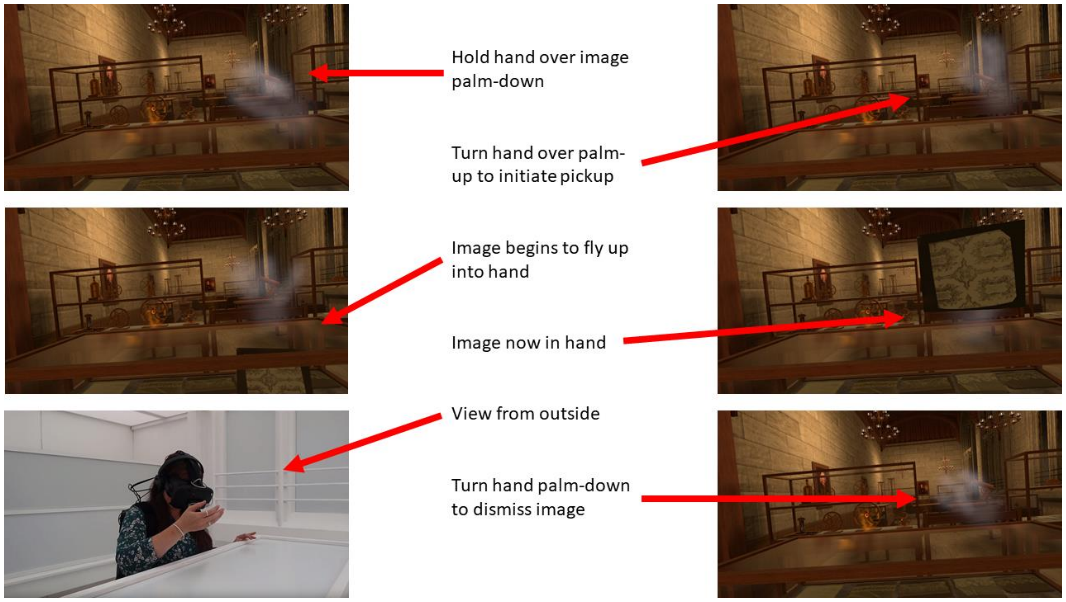

Figure 9. Picking up a photograph.

Multiple users. In the interest of achieving sufficient visitor throughput to make Thresholds an attractive proposition for museums and galleries, and enabled by the wireless capability of the HTC Vive VR system, we created a multi-user experience for up to six simultaneous visitors. To make it safe to walk around the space, bearing in mind that in VR a visitor is to all intents and purposes blindfolded, it was necessary to represent the other visitors in some way. We therefore extended the ghostly representation of hands described above to apply to each visitor's entire avatar (Figure 10). These visible ghostly shapes were intended to help visitors avoid colliding with one another. They also imply one's own ghostly presence as a time-traveler, with the ghosts serving to highlight the slightly unreal nature of the experience – you are 'there, but also not there'. From a technical perspective, the ghosts simplify aspects of the development. Because we only track visitors' headsets, we have no reasonable way to fully animate their bodies for humanoid avatars. Using the ghostly 'pillars' we inhabit a space roughly equivalent to a human with a relatively simple yet also artistically appropriate representation.

Figure 10. Two ghostly avatars showing the positions of other visitors in the space.

Connecting to the wider museum environment. We created a waiting area outside of the main set where visitors queued at busy times and donned the VR interface before being led – effectively blindfolded – into the set through a single doorway. A decision was made to allow visitors to see into the set through a viewing window before entering (Figure 11), partly to stimulate them while queuing, partly to prepare them for how to behave in the experience, but also because the sight of participants wandering around, feeling the interior of the room like blind insects, provides a powerful spectacle. From inside the virtual model, the viewing window appears as a painting of Edward VI, the patron of the school (Figure 4, Left). The artist embellished this with a small animated spider with the intention of drawing the viewer's attention and ideally causing them to stop and stare closely, being posed directly in front of spectators who are watching them from outside. Artistically, this unusual one-way window was intended to create a moment in which participants inside the virtual model are looking back in time several hundred years by inspecting an image from the 16th Century, while viewers outside are looking back at them, bedecked in modern technology in a 'futuristic' space.

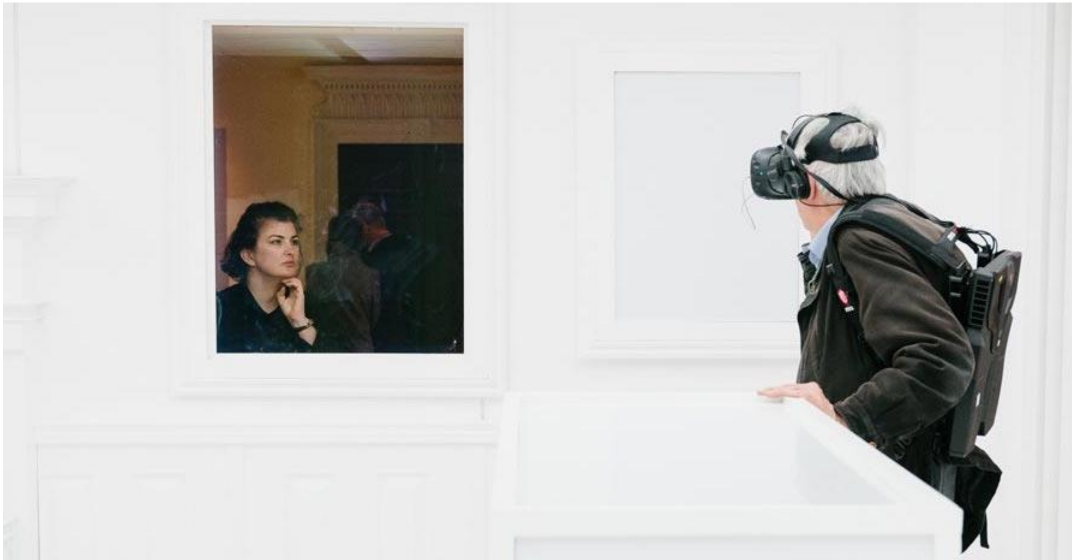

Figure 11. Spectators view the space through a window which appears as a painting in the virtual world, creating an interesting view into the experience from outside.

IMPLEMENTING THRESHOLDS

Several technologies were considered to deliver the VR experience. To create a VR space in which visitors could walk about and explore necessitated six-axis freedom (translation and rotation), and the high-quality graphics required significant rendering capabilities. The 8.5x6.5x2.5m space is relatively large by current VR standards. We considered the use of an UWB positioning system like pozyx [20], but found the tracking quality insufficient for our needs. We ran some initial tests with pozyx, and were impressed by the range, but we found the moment to moment positioning, especially the height axis to be too jittery for comfortable use with a head mounted display – and the use of smoothing, such as a low pass filter made the positioning too laggy, and thus quite motion-sickness inducing. We then made a prototype using Google's Tango augmented reality tablet [16], but found the rendering capabilities insufficient to deliver the graphical quality we intended. Ultimately, despite the size of the tracking space, we kept returning to the HTC Vive, the technology that first popularised room scale VR. The Vive uses "outside-in" tracking, that is, it uses a pair of "base stations" that emit infrared pulses at a rate of 60Hz and the headset (and controllers) have a large number of receivers for these pulses positioned on them which are then used to triangulate the position of the headset in relation to the base stations. At the time of development "inside-out" tracking, where the headset uses cameras to build a model of its location and thus track without the need for external base-stations, was limited to the Microsoft HoloLens which didn't offer full VR capabilities, and Google's Tango, which we had already established lacked rendering capability. At the time of writing there are several inside-out solutions on the market, including the Vive Cosmos and Oculus Quest, however, as mentioned, when Thresholds was developed, outside-in tracking was the de facto standard.

HTC claim sub-millimetre accuracy for their tracking, though our anecdotal experience of building Thresholds suggests that this varies by up to a few centimetres depending on the range and visibility of the base stations. However, we concluded that the Vive tracking system offered the best balance of cost, reliability and tracking quality for our needs.

It was also necessary to deliver the system "wirelessly", by which we mean not having the trailing wires often associated with (high-end) VR headsets. Six free roaming visitors dragging long cables around would be impractical at best. At the time of development the current TPCast and HTC Wireless systems, both based on Intel's WiGig technology were unavailable – and indeed even now would be unable to support six simultaneous users – thus it was necessary to use portable PCs to run the system.

Consequently, Thresholds runs on six MSI VR One backpacks PCs, each with an Intel i7-6820HK CPU and an nVidia GTX1070 GPU. The headsets are HTC Vives, connected via short power, USB and HDMI cables. A Leap Motion (IR camera) using the Orion SDK to detect the visitor's hands, is mounted on the front of the Vive, just above the main camera and connected to the Vive's onboard USB port. A pair of headphones provides personalized spatial audio, and finally a mouse is affixed to the back of the backpack to give the invigilators a simple control interface for resetting the system. A single pair of Vive base stations mounted in diagonally opposite corners of the physical set (connected with a sync cable as they are out of range for optical sync) are sufficient to track six visitors at a time. Each backpack has a Wi-Fi connection to a local network to talk to a server that runs on a similarly equipped desktop PC. The server's screen displays an image of the virtual room, showing the locations of each of the visitors, and is mirrored to a 60" Screen displayed in the set-up area and used by invigilators for briefing the visitors. A diagram of the physical system architecture can be seen in Figure 12, and images of the visitors' equipment from different angles can be seen in Figure 13.

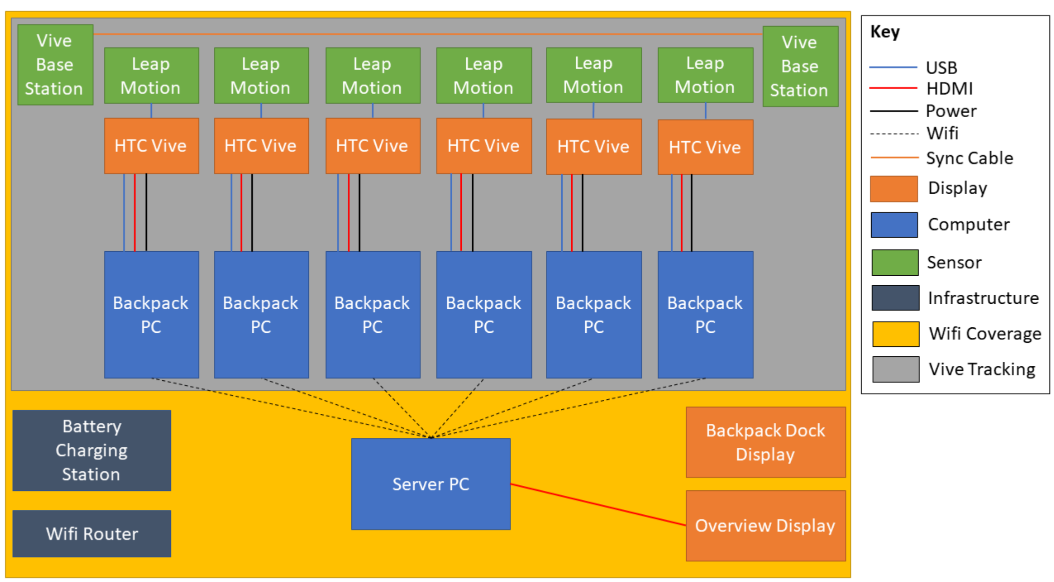

Figure 12. Thresholds Physical System Architecture.

Our physical room size was well beyond the (then) supported bounds of the Vive (a maximum of 5m between diagonally opposite base stations). After some experimentation, we were able to achieve reliable tracking of the headset in our space;

though notably not of the controllers which have fewer receivers than the headset. Exceeding the maximum bounds to this degree cost us some tracking resolution however, and introduced more scope for significant tracking error to occur. Experimentation suggested a maximum achievable tracking space of 8.5x6.5m, leading us to slightly reduce the physical length of the set from a planned 10m to 8.5m. The artist deemed this reduction acceptable for the practicality of working with the Vive tracking system.

The Thresholds application itself is built in Unity3D. The objects in the scene were modelled in Autodesk 3D Studio Max[1], then imported into a Unity3D scene. A client-server architecture was applied, with both using the same graphical assets. The server has a software camera placed above the room giving an overview of the locations of each visitor. This server listens for connections from each client, and after registering them listens for position updates from each. It constantly broadcasts the positions of all known clients. The client is the VR application running on each of the backpack PCs, which, when started, connects to the server and begins sending its position in the room and listening for the broadcasts of everybody else's position. When it receives the broadcast message it draws the "ghostly figures" that represent the other visitors into the scene. A timer on the client is used to limit the experience to six minutes and is reset by pressing the mouse button. This reset also includes the timed elements such as the riot – which was implemented as a chroma-keyed video.

Figure 13. Thresholds visitor equipment form multiple angles. Backpack PC, HTC Vive, Leap Motion, Headphones.

Aligning the Physical and Virtual Spaces

To create properly aligned physical and virtual rooms it was necessary to match the dimensions of the two spaces and their contents precisely. A CAD model of the

[1] The CGI for Thresholds was created by VMI Studio (vmistudio.com/).

physical room was built to carefully measured dimensions, then every object to be physically built was measured to millimetre accuracy. Physical materials and paints were selected for both their haptic and non-reflective properties. Virtual wood was matched with real wood, metal with metal, and the glass of the vitrines with a powder-coated aluminium that was carefully chosen to be a good analogue for glass but without the reflections. The artist wanted to keep the room white, so an initial test used a sheet of glass with white wood behind it, however the glass tended to scatter the IR used by the tracking system making it unreliable around the vitrines, so the matt effect of power-coated aluminium was selected as a good analogue, with its cool-to-the-touch finish. The spectator window was one of the most reflective surfaces in the space, and as such tended to cause the most instances of tracking error.

Correctly aligning the physical and virtual spaces proved to be a significant challenge that required configuring the relationships between: 1) The physical room, 2) The virtual model of the room, 3) The tracking space of the VR headsets, 4) The tracking space of their associated Leap Motions and 5) The driver-level tracking configuration for each PC.

In the VR engine used to develop the system, in this case Unity, positions of both the tracking space and virtual room are available as software objects that can be moved and rotated relative to each other. The manipulation of these objects is a core part of establishing a correct alignment of the virtual world with the physical space. Figure 17 shows how the various spaces involved interact. Each object has a root point in the software through which the rest of it can be translated and rotated. If we assume (for now) that the tracking systems deliver the accuracy they claim and that driver-level tracking configuration is consistent across all the PCs, then we need only manipulate the positions and rotations of these root objects to align our physical and virtual spaces. However, it transpired that slight variations in the manufacture of the Vive headsets (which become more pronounced when working at the extended scale of Thresholds) combined with small variations in how the Leap Motions were attached to the headsets, also required us to carefully tune the spatial relationship between each grouping of Headset, PC and Leap, specifying individual device mappings in a series of configuration files.

Once the Headset/Leap Motion pairings were configured we then had to align the physical world and virtual set. This needed to be done afresh each time the physical set was reconstructed – i.e. each time Thresholds was installed in a new location and separately for each PC. To help perform this alignment, we built a real-time 'configurator' tool using the Vive controllers as an interface to manipulate the pitch roll and yaw, as well as the x, y and z positioning of the root of the tracking space. The rotation and position of the tracking space is determined by the positioning of the light boxes and the room setup configuration file in the Vive's settings.

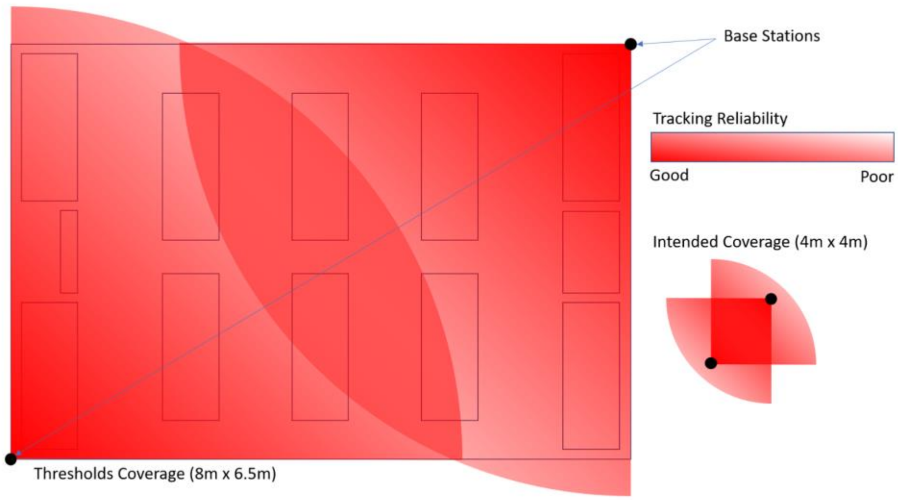

Figure 14. Tracking coverage of the thresholds space compared to the intended coverage. Lighter areas show less reliable tracking at the limits of the base stations' range.

The scale of the space makes setting this up additionally difficult for two reasons: first, because as the tracking reaches the limits of the base station's range (see Figure 14), tracking consequently becomes less reliable. In general, for reliable tracking a sensor needs to receive data from both base stations, but in spaces as large as Thresholds, we are often relying on a single base station. This appears to result in "stretching" of the tracking space towards the limits, which can mean that perfect tracking in one part of the room is not matched elsewhere, even though the dimensions are identical. Figure 15 shows a somewhat exaggerated (and only two dimensional) representation of this mismatch between the positions reported by the tracking and the real physical positions.

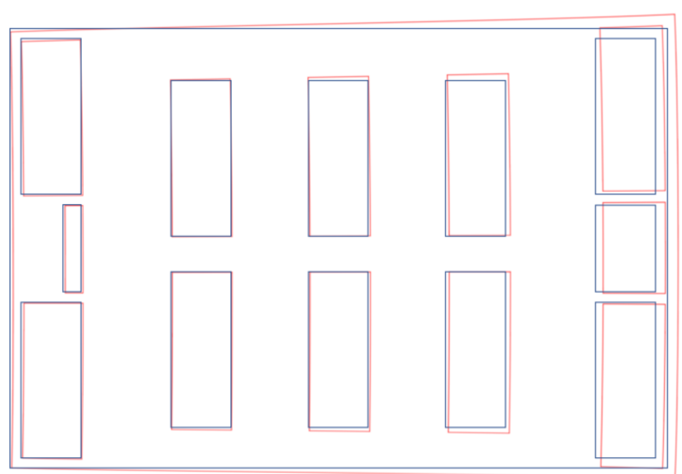

Figure 15. Distortion of the tracked space from one lightbox. Tracked space is shown in red, real space shown in blue. This effect is exaggerated here for clarity.

We found it possible to mostly correct for this distortion by careful rotation of the root of the tracking space. Certain objects, particularly at the limits of both trackers, may have a slightly less than perfect mapping, and part of the skill of deploying the work is to find the best compromise. Second, the tracking space has a root that is a single point. This point is defined as the centre of the tracked space. However because our room is significantly larger than the 'officially' tracked space, this leads to a second challenge: the root of our tracking space – and subsequently the origin of rotation is not actually in the centre of the real room (see Figure 16) which makes performing these rotations in situ challenging to understand – for example a small rotation in one part of the room can have a dramatic effect at the other end of the room. A significant practicality of deployment was teaching museum staff how to perform this calibration.

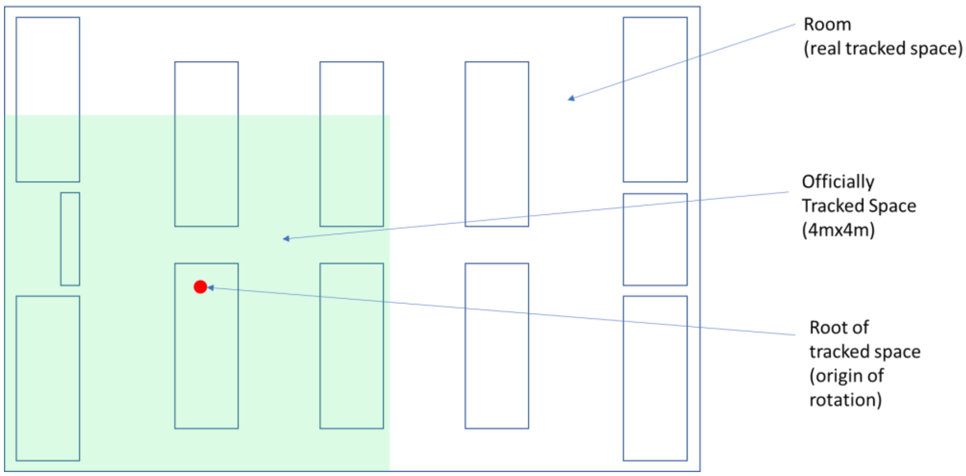

Figure 16. The origin of rotation for the tracked space is not at the centre of the room. This makes rotation difficult to understand in practice.

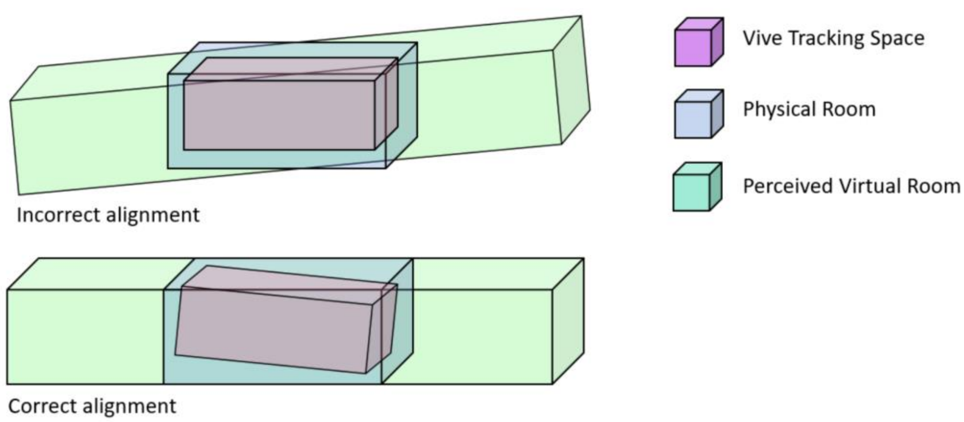

Figure 17. Moving or rotating the root of the tracking space changes perception of the position of the virtual room with respect to the physical room.

Given these practical challenges with aligning the various spaces involved, our calibration process unfolded as follows. To begin the alignment process we walk to a known location – say the corner of a vitrine – and position our head carefully with respect to this known reference point. We then use the controllers to move the root of the tracking space (which has the effect of moving our own apparent position and rotation) until we appear to be in the correct position and rotation. It is necessary to repeat this in several positions around the room to confirm that the overall position and rotation settings are correct. These values can then be written to a configuration file (by clicking the grip button), so the next time the system is started (without paired controllers) the alignment is retained. This process requires being able to think about rotations in 3D, but once trained and with a little practice, we found that it can be performed in a few minutes. Figure 17 shows how rotating the tracking space has the effect of rotating the perceived virtual room to make the alignment correct. In the top image the Vive tracking space and physical set appear to be aligned, but due to the various distortions involved, the perceived effect is misalignment. In the bottom, the Viva tracking space has been misaligned using the procedure described above to correct for this.

In a smaller space, where the Vive controllers are tracked reliably, and the above distortion does not occur, it is possible to simply place the two controllers in known locations then translate and rotate the root of the tracking space until the controllers appear to be in the correct positions. This process can even be trivially automated as long as the controller's physical and intended positions are consistently used. The recently released Vive 'trackers' provide more suitably shaped objects to use for this – including the ability to track several of them at once – and may further simplify the process. HTC have also recently released their Tracking 2.0 which tracks a space of up to $10m^2$, however the same challenges would occur if we wished to build something exceeding that range. In other words, the techniques we report can help designers cope with using Vive and related technologies at or even beyond the edges of their tracking ranges.

Mapping of the hand tracking space works similarly to the headset tracking. After having correctly set the headset alignment, we place our hand in the corner of a vitrine and use the controller in the other hand to tweak the position and rotation of the hand-tracking space to ensure they are correctly aligned. These settings are then also written to the configuration file.

Using the method described above, we were generally able to create a reliable physical-virtual mapping by eye correct across the space varying by no more than a few centimetres at worst, and within one centimetre at best. We address this remaining error by design as noted above. By making the outline of the visitor's hand 'fuzzy', the error may be hidden. The 3D model of the hands is actually made out of a series of 'blobs of white light' each attached to one of the bones of the hand, which move appropriately as tracked by the Leap. The result is a rather oversized hand shape that fades out around the edges. It looks sufficiently like a hand to correctly represent one, but is sufficiently vague about where its edges are as to fool most users into ignoring minor alignment error.

Thresholds in Practice

The 93 photographs were created from high quality scans of the originals, gathered from a range of sources. However, there were some cases where the originals have been lost, it was unclear which of several in a set were displayed in the original exhibition, or they are too light-sensitive to be scanned. In these cases, images were substituted with others that were from the correct time period with appropriate subject matter under the guidance of Fox Talbot scholars.

We introduced accessibility for wheelchair users. The vitrines are quite high (around a metre) and it is difficult for wheelchair users to look into them. By moving the root of the tracking system up by around 50cm, we raise the visitor's apparent position by the same amount. Of course, this breaks the alignment of the props in one dimension, but does allow such users to enjoy the visual aspect of the experience.

The physical world and virtual set needed to be aligned afresh each time the physical set was reconstructed, i.e. each time Thresholds was installed in a new location. Moreover, one further tracking challenge emerged. In practice, with an installation running for a period of weeks or months, the alignment configuration tended to drift as a result of settling of the installation, lightboxes being shaken, Leap Motions being knocked out of position and so forth. We therefore instigated a weekly 'health check', where a technician would check the alignments of each of the headsets, make necessary adjustments and save configuration files.

We supported invigilators by introducing a separate Monitor interface to show the virtual space with the ghosts of visitors moving around, so as to allow invigilators to monitor that everything was running smoothly. This was displayed on a large screen near the viewing window providing spectators with a tantalising glimpse of what awaited them in the virtual space. As the system can only track the presence of visitors wearing the equipment it is necessary for any additional people who enter the physical set, most notably invigilators, to be careful when moving around as they are not visible to the visitors.

The system has several safety-nets in place. In order to support the physical presence of invigilators inside the set, we placed a 'fake' avatar in one corner of the room where an invigilator would normally be stationed, creating a kind of 'safe space' where they could stand without having to constantly move out of the way of visitors as they walk around the space. This fake person brought the additional benefit of allowing us to 'block off' one particular space in the extreme corner of the room that had poor tracking coverage. If the client has not received a message from the server for five seconds, a warning is shown asking the visitor to remove the headset. This is necessary, because if there is no connection to the server, then there is not correct information about the location of other visitors, which could lead to accidents. Similarly, if the reported position is outside the expected space, the system shows a message asking the user to remove the headset as this implies the tracking has failed.

In most cases when the tracking does fail, either by being in a blind spot or by visitors covering the IR receivers on their headset, this occurs as a sudden jump or drifting of the viewpoint. This can be quite disorienting, and as such we build a warning about this and the instruction to close one's eyes and/or turn around (usually enough to re-establish tracking) into the instructions given to visitors by the invigilators, along with fire-safety instructions etc.

EXPERIENCING THRESHOLDS

We draw on multiple data sources including interviews, observations and system logs to validate that Thresholds was successful as a touring cultural product.

Thresholds as a successful touring experience

Our first observation is that Thresholds proved successful as a touring installation, including at: Somerset House, London (18th May to 11th June 2017, ~5000 visitors); Birmingham Museum and Art Gallery (24th June 2017 - 6th August 2017, ~1200 visitors); Lacock Abbey, Wiltshire, UK (16th September - 29th October 2017, ~800 visitors); and the National Science and Media Museum Bradford, UK (2nd March - 7th May 2018, ~2600 visitors). Thresholds achieved an average throughput of 54 visitors per day across these deployments, peaking at an average of 200 at Somerset House. Media and social media coverage reveals that the work was well received critically. Perhaps the most significant cultural review to date was an extended discussion by four art critics on the BBC Radio 4's Saturday Review who debated the artist's core concept of using the emerging technology of VR to reinterpret the earlier radical technology of photography. Other media coverage, including press reports (e.g. [3] [8] [44]), BBC television (Click) and radio (Saturday Review) coverage, and specialist art and culture blogs [2] [29] [41], presented a positive critical reception to the work. It was also shortlisted for the prestigious South Bank Sky Award for Visual Arts [33].

How visitors behaved

We now inspect our systems logs to help contribute to our overall picture of how visitors behaved in terms of where they went, where they looked and how they interacted in the virtual world of Thresholds. Figure 18 shows a heat map of the horizontal positions of the headsets of our 2,349 logged visitors as seen from above and set against the backdrop of virtual model. In other words, it summarises the time spent by visitors at different locations in the virtual model. Red shows the most popular locations, orange and yellow the next, green lees so, while areas that are not coloured were not visited at all.

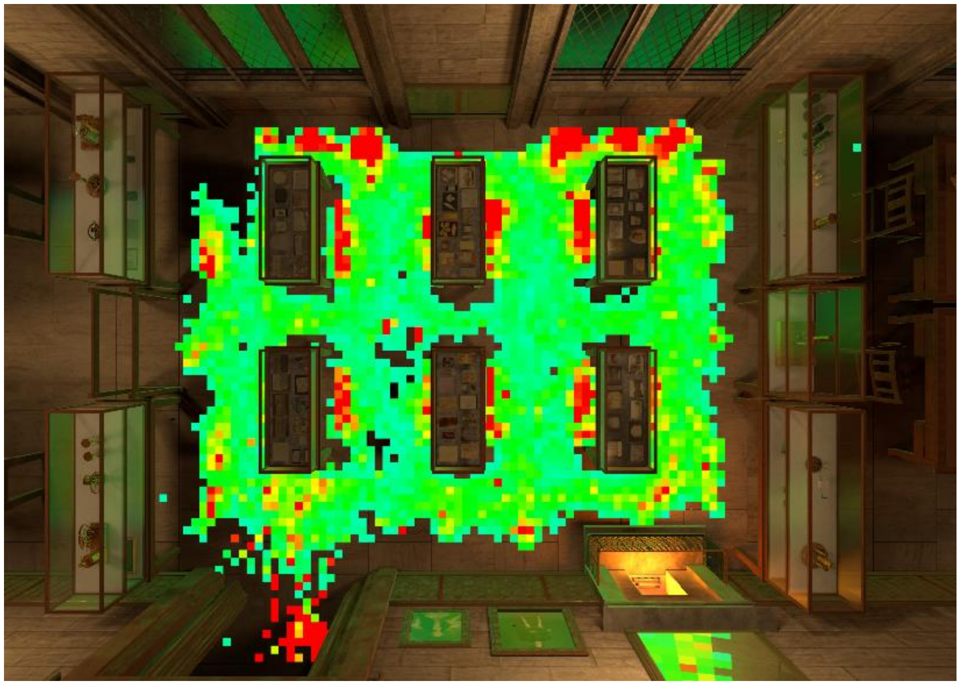

Figure 18. Aerial view of headset positions.

There is substantial clustering around the door as we might expect as this is the entry and exit point for all visitors. The vitrines were evidently popular locations and visitors tended to stand at their sides rather than their ends, reflecting the orientation of the photographs. They also appeared to stand more at their inner sides, perhaps avoiding the relatively busy corridors around the outside of the room and horizontally through its centre. The windows were also popular locations with many visitors pausing to look out at the riot. Figure 19 visualises the before and after effects of the riot happening at 3 minutes in to each visitor's experience. The heatmap on the left shows the first 180 seconds during this period, and the one on the right, the second 180 seconds. There appears to be more traffic around the windows in the second image, suggesting that the riot animation did its job of attracting visitors to the windows. The painting with the spider (the spectator window in the physical set) was more popular than the one next to it. The notable gap at the top left is where the invigilator ghost avatar was placed, showing how visitors actively avoided this area. They also avoided the narrow gap between the protruding wall and central vitrines (top middle) and the area directly to the left of the door, perhaps tending to head into room on first arrival.

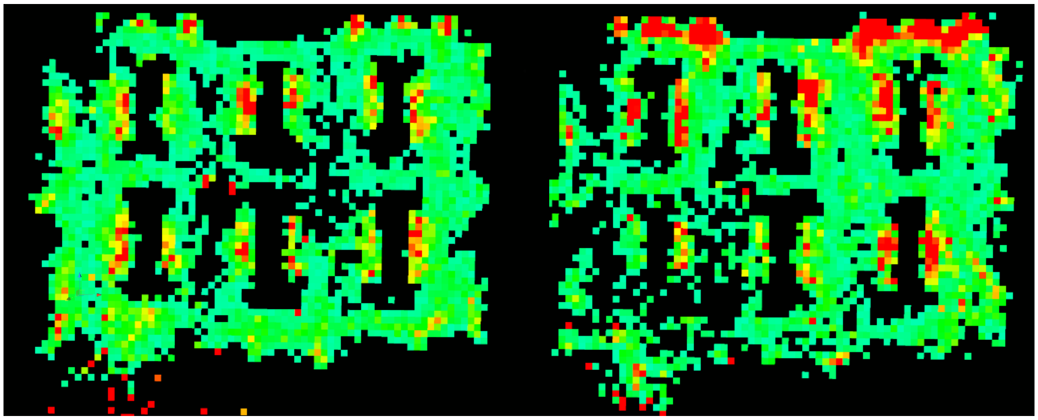

Figure 19. Left the first 18 seconds of the riot and right the second 180 seconds showing how the animation attracted visitors to the window.

Figure 20 shows a corresponding heat map of the vertical positions of visitors' headsets. Crouching tends to occur on the left side of the room (the entrance side) which is consistent with the distribution of the mice, which were concentrated on this side.

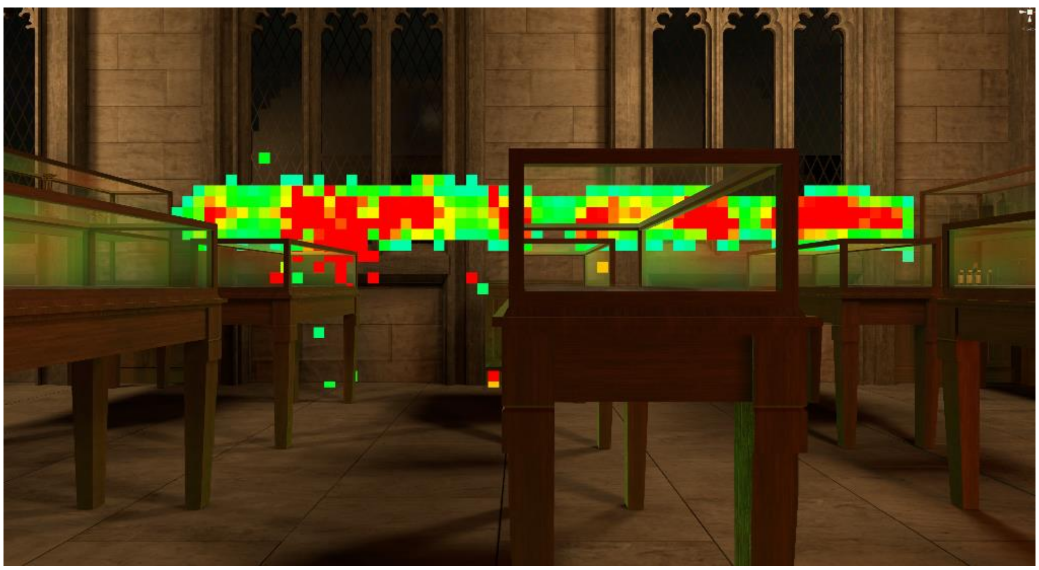

Figure 20. Vertical positions of the headsets.

In terms of path movement around the space an observation of visitors' paths shows a tendency to move straight to the nearest vitrine, then move around the room distributed fairly evenly between clockwise and anticlockwise, pausing at each of the vitrines. Most users also move up to the fire as they reach it and stop to look at the paintings. Once a complete circuit has been made, movement around the space becomes less predictable, with visitors returning to images and objects in which they were interested or pausing next to another avatar (which may suggest conversation, despite the fact that earphones form a part of the system which would seemingly preclude this). Of perhaps more interest (and as covered in Figure 19) is that this is interrupted (either the initial circuit or the later wandering depending how long the initial circuit takes) by the arrival of the riot outside the window at 180 seconds, after which point many (but not all) visitors tend to move over to the nearest windows to observe the riot before either returning to their previous path or moving to the vitrine nearest the window.

To complement knowledge of where visitors stood, Figure 21 summarises the top 15 most significant objects that they looked at. This takes into account all *significant* objects viewed by the visitors (i.e. ignores the walls, ceiling and floor). The y-axis in the Figure shows the total number of seconds on average each visitor spent examining each of the top 15 items.

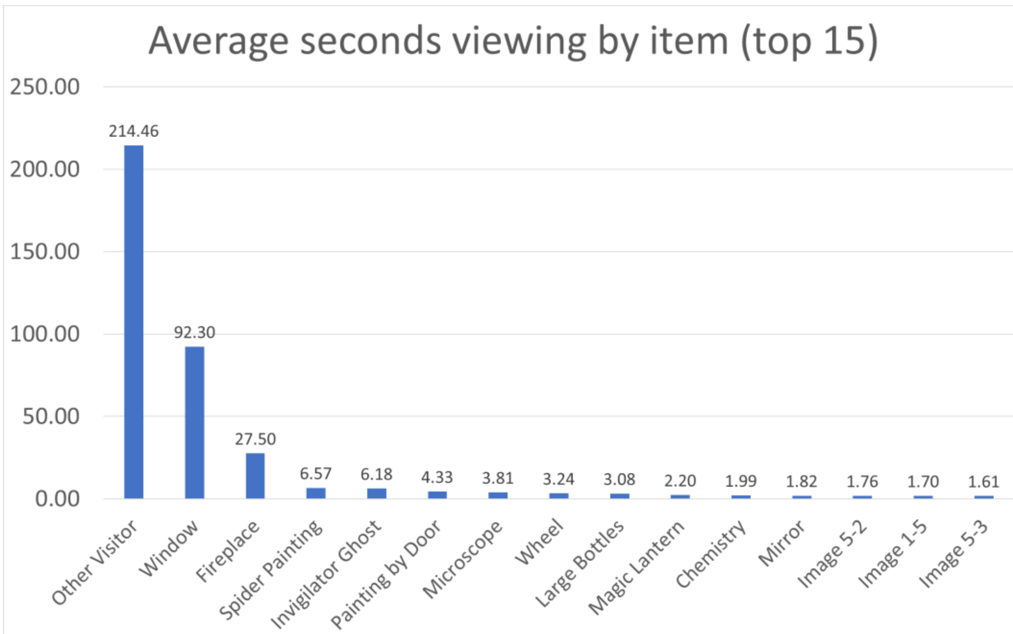

Figure 21. Time spent looking at different objects in the world.

Object views are calculated as described above. This takes into account *significant* objects viewed by the visitors (ignoring walls etc.). We can see that by far and away the most viewed thing was other visitors. This isn't that surprising as the surface areas of the avatars were relatively large compared to other objects and up to six of them were present in the world. However, this does provide evidence to suggest that visitors were aware of each other through their avatars. Figure 21 also suggests that objects naturally visible at eye-level (ghosts, windows, paintings and fireplace) were looked at more than those that required looking down (photos and other objects in the vitrines).

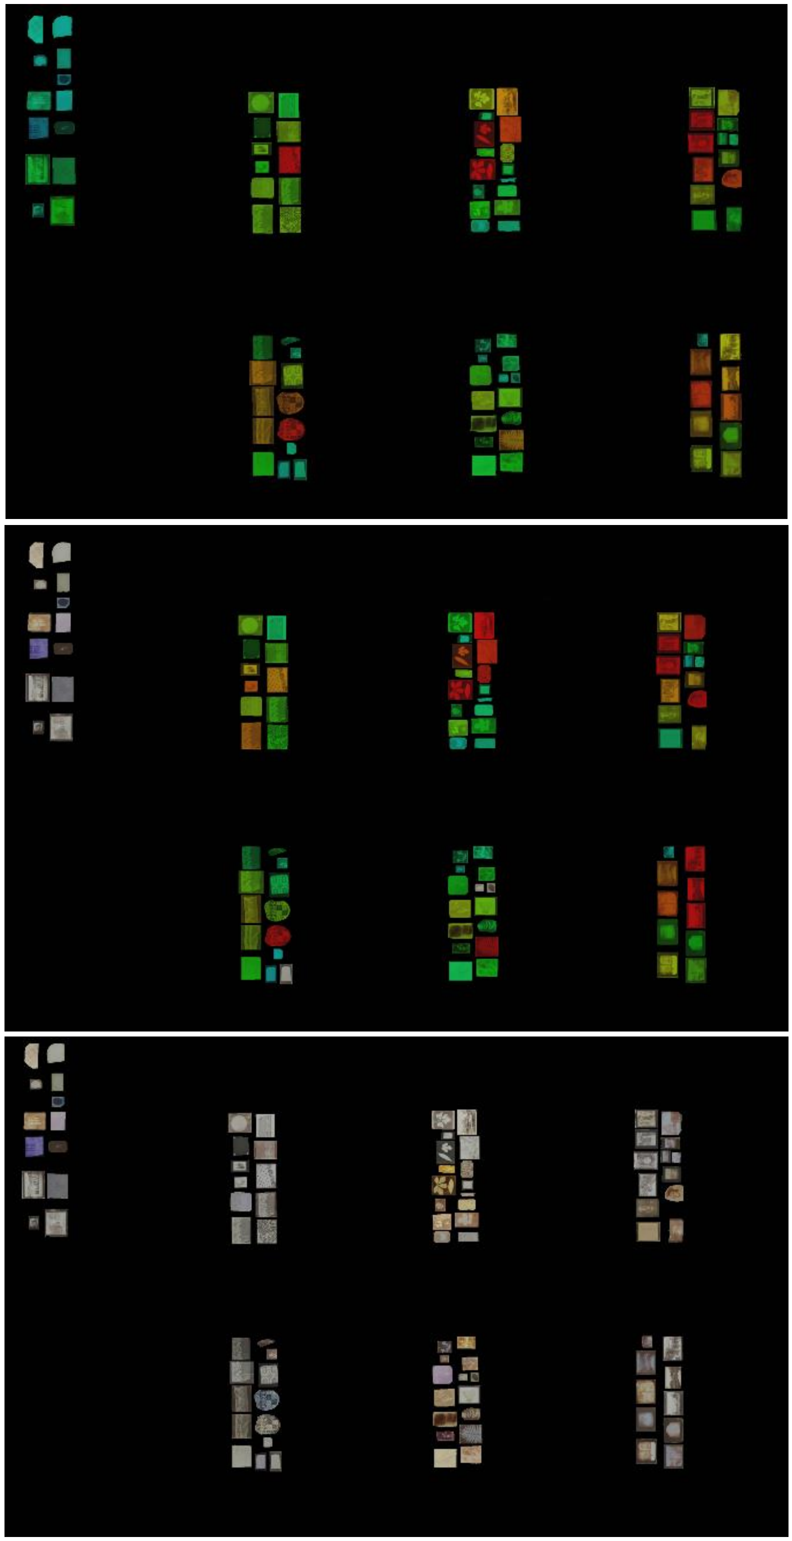

Figure 22. Viewing (top) and picking up (middle) the photographs with reference image (bottom).

In short, as one might expect, being large and at eye level seems to lead to being looked at more. This doesn't always hold true though: size doesn't seem to matter with the paintings so much as other details such as the spider as noted above (the spider painting gained ~1.5X more views than the one next to it, despite being only about ~1.2X the size).

The two heatmaps in Figure 22 reveal how much time was spent looking at and picking up the various reproductions of the photographs in the vitrines (the reference image below shows their baseline colours for comparison). Viewing is calculated by casting a ray directly forward from the headset, so doesn't take account of peripheral vision. However, in our experience, VR users do tend to look directly forward, partly due to the limited field of view, and partly due to parallaxing of the image when viewed side on. The images in the top left vitrine have comparatively few views, but it should be noted that this is not exactly a like-for-like comparison. That particular vitrine had two levels, with a row of scientific equipment on the level above the images. Additionally, the invigilator ghost is situated next to this vitrine, making it less accessible than the others.

The middle image shows those photographs that were picked up the most (note that those in the top left vitrine could not be picked up). These are broadly consistent with the viewing data shown top and suggest that larger images are more likely to be picked up, perhaps because they are easier to grasp. This is notable as the reasoning behind including the grasping option was to enable smaller images to be viewed more easily 'up close'. Additionally, we note that it is the images furthest from the entrance (top right in Figure 22) that are more commonly picked up, suggesting this happens more later in a visit, perhaps once they are more confident moving around and more used to their surroundings.

Finally, with reference to the log data, Figure 23 summarises our estimates of the spatial distribution of tracking errors. These are defined as being reported positions that were either outside the physical constraints of the space or more than 50cm away from the previous recorded point (unlikely to occur with logging at 90 Hz). The visualisation shows the last reported 'good' position just before the tracking error occurred. We see the most errors around the entrance – this is to be expected as this was a popular location, is effectively outside the tracking space and additionally many users were observed holding onto the headset as they began which might cover its sensors and so compound tracking errors. We see several errors around the windows, another popular locations, but this is also consistent with users banging the headset against the bars as discussed further below, or with the mouldings around the windows obscuring the sensors. There are also some towards the centre of the room which is consistent with being on the edge of the maximum range of the sensors which we physically placed in the corners.

Figure 23. Visualisation of estimated tracking errors.

What visitors thought

We also gathered qualitative data to help us understand visitors' thoughts and feelings about the experience. This included visitor feedback collected by the venues themselves using their own established mechanisms, including comment cards, and an in-depth observational study of 12 volunteers (9 female and 3 male visitors; 3 <35 years, 9 >35 years) when Thresholds was exhibited at Birmingham Museum and Art Gallery. During a dedicated study session, when the exhibition would not otherwise have been open, study participants experienced the work and various types of data were collected about their experiences by the research team. Participants were recruited through the networks of local contacts engaged with artistic practice, using a snowballing method. Participants completed pre- and post-experience surveys, which mixed closed quantitative questions using Likert scales with open qualitative questions where free text responses were recorded. After undertaking the experience and surveys, participants also participated in short interviews. Venue comment cards were combined with the survey results and interview transcripts into a single dataset, which was thematically coded in order to investigate the following issues: user confidence, embodied engagement, the experience of room scale VR and its perceived value.

Our participants arrived at the exhibition with varied experience and expectations of VR. By and large, they enjoyed the experience, with 11/12 reporting in the short survey that it was better than expected. Moreover, everyone agreed that it was an unusual experience. A common concern recorded on feedback cards was the need for more time to explore the exhibit and also to *"adjust to the peculiarities of VR"* and "*practise touching the photos"*. Many visitors appeared to appreciate the artist's historical interpretation: *"looking at the photographs and realising the development that had taken place imagining that VR will be like this in 90 years"* (feedback card).

However, some did note the danger that the VR experience might overwhelm the original photographs: "*the new VR experience, kind of, overshadowed the fact that I was seeing these photographs, but I think once you get used to VR experience what they were achieving in photography will come through a lot more"* (P4).

Tactility. 11/12 agreed that tactility created a feeling of immersion (6/12 strongly agreed), especially touching objects in the physical space, such as vitrines and window bars: *"It became a curious thing then to run my hands across edges of things to see how that sort of lined up with what I was seeing, which all became really interesting and fun to do because it was a very successful alignment. I feel it could've quite easily hampered the experience slightly if things weren't happening in that way. But then it was, then that made me want to almost touch everything and sort of scrutinise everything in there"* (P1). The moment of looking out of the window while leaning against or holding onto the bars (Figure 24, left) was especially powerful: P6 highlighted "*the one bit where I was looking out of the window, the bars on the windows, they were exactly where I was expecting them to be" … "I just stood on my tip-toes and did that to look better through the window, and as soon as I did it I realised that shows how I didn't even think about doing it"*.

Confidence. This sense of a good alignment between virtual and physical contributed to visitors' confidence in moving and acting. Prior to the experience, some users expressed concern about how it might feel 'disorienting' or be perceived as 'weird' or 'surreal'. But once inside Thresholds, visitors moved freely around the space, looked around the room and into the vitrines, and reached out to touch things: "*because the alignment was all correct, I felt pretty confident about walking around. When I first walked in, my initial reaction was wanting to put my hands out to touch things. But very quickly I became comfortable in acknowledging where things were and what space I had to move around*" (P1). Only one of our interviewees commented on feeling disoriented and our general observations of many visitors (backed up by the visualisation above and their self-reported activities) were that, by and large, they explored freely and widely. Visitors took pleasure in the realisation that they acted naturally within the virtual environment: *"I was prompted to change what I was doing by the sounds I suppose. So, went over to the window and I could hear something was going on out there and then did something else then went back to the window again"* (P7). This confidence manifested in various observable 'relaxed' behaviours (Figure 24), such as standing on tiptoes to see something better or leaning against a cabinet or the wall.

Constraints. The large vitrines and barriers placed at the ends of the virtual model were largely effective in dissuading people from trying to move through them (and hence entering the non-aligned physical world beyond), and any frustration that arose due to not being able to explore more widely was understood in context: "*I had that similar sense of impatience that you get in real galleries when you've got something barriered, and you're, like, I want to go in that bit"* (P11). However, there was some confusion about their contents not being interactive: "*I think I realised the stuff in the cabinets you probably couldn't interact with, the side cabinets rather than the ones in the middle*" (P10) and "*I was also trying to pick up the pictures in the cabinets on the* *ends, and I couldn't do it with those because my hands were bashing against the cabinet"* (P7).

Figure 24. Many visitors appeared to be physically comfortable in the space.

The devil is in the detail. 11/12 of our interviewees agreed that the sounds, riot, and creatures added greatly to the sense of immersion the experience: "*it's those extra touches that make you feel that you are back in 1839, with the mice and people outside"* (P8). The multisensory experience of the fireplace appears to have been particularly compelling, with a combination of sensory stimuli contributing to an aligned moment of experience (Figure 7): *"It didn't look like a real fire. But, you were interacting with it in a different way, it wasn't just visual. When you can hear the crackling of the wood in your ears, like the normal fire, that sort of like made it a bit more immersive"* (P9) and *"I think going nearer to the fire and feeling it and also hearing it was really immersive, and that felt really realistic"* (P4).

Picking up photographs. Although visitors navigated the exhibition confidently and appeared to inhabit the space naturalistically, they found it noticeably harder to interact with the virtual photographs. The inclusion of this mechanism had been a matter of some debate during the design stage and so it proved with visitors too. There was evidently a knack to be learned (Figure 9) and the success with which visitors mastered this had an impact on their judgement of the experience: *"I'd figured out there was a certain knack to do that. I thought that was really good to be able to do that because they're quite small in the cabinets and looking down you can't really see them in their full glory. But when you can expand them up it's really good"* (P6). However, some commented that this virtual interaction broke the perceived reality of the experience: *"Although, it sort of breaks the reality of the situation, of course you can take a thing out of case. So, you know, constantly battling, trying to allow your mind to be there and then, you know, taking a thing out of a case, although it's part of the game, it sort of breaks the fiction"* (P2). Others anticipated that the technique would apply to the other objects that were available within the environment and were confused when this was not possible: *"There was model as you went in on the left-hand side on the approach, did that do anything? Could you tell me? Because I was just trying to reach for those, and I thought I could turn a wheel and it didn't work"* (P11).

Avatars. The ghostlike figures (Figure 10) appear to have been broadly successful at conveying the presence of other visitors: "*I thought it might feel quite, kind of,* *crowded because I was conscious that there was lots of people and I thought that would be a problem. But actually that was fine. You could, sort of, see some hazy figures kind of, but it wasn't, like, oh, they're in my way"* (P11). The introduction of these more abstract representations into what was otherwise a naturalistically represented world established a sense of creepiness for some: *"Is it meant to be a slightly creepy experience? I mean, if it is then they're adding to that"* (P2). The representations of visitors' own hands (Figure 8) were generally accepted and occasioned less comment, which tended to focus on practical matters, for example tracking issues: *"I couldn't always see my hands. So, they were white at some point and then sometimes when I lifted my hands up higher, they were there and then they'd disappear"* (P10).

The spectator window. This provided waiting visitors with a first impression of the physical space. On the positive side, it could help establish a sense of confidence in moving: "*You might be more apprehensive and slightly less confident if you hadn't seen it."* (P6). On the negative side, prior glimpses may reduce the sense of mystery: *"So, I was wondering whether if there had been some veil before going in or not having seen it, whether I'd have equated the visual cue of this kind of wood or something, feeling a certain way"* (P1). The window revealed the dimensions of the set, which appeared to have an impact on one visitor*: "I think you're a bit influenced by looking at the environment inside before you go in, and I think the one thing that I realised I didn't do as a result was I didn't look up. And I think that's a shame actually because I missed that and the chandeliers on the ceiling and everything."* (P7). The presence of spectators didn't appear to detract from the experience once inside. P7 was initially aware of the possibility of being watched, but soon forgot: *"I was just completely engrossed. I remember when I was watching people before I went in I thought, whatever's in that window try not to look at it because everyone's going to be staring at you. And I remember I found myself watching a spider crawl across the picture, looking gormless probably"*.

Glitches. While the experience mostly worked well, there were occasional glitches including the tracking errors noted above; occasional hardware failures such as batteries suddenly dying or the Leap Motions overheating; occasional lag either in frame rate, because some process on the PC interrupted the experience, or due to network delay when communicating with the server; occasional operator errors such as leaving the battery to drop below 30% charge, which initiates a battery saving mode on the backpack PC that is no longer sufficient to render the experience; or leaving the backpack on 'wheelchair mode' in which the viewpoint is raised by nearly a metre; or once memorably knocking the plug out of the server. Glitches that caused misalignment between the virtual and physical were especially troublesome for visitors: *"For me, my physical space was misaligned by about 15 centimetres, so I would try to put my hand on the case and it'd already be there and that kept jarring me"* (P2) and "*there was a notable lag and when I was standing still looking out the window, the image I was seeing actually moved around even though I was stood still. But that corrected itself and so, there was only those instances really"* (P1). This important of correct technical setup for the overall quality of experience is highlighted by a visitor for whom errors had to be corrected: *"I had a false start because the equipment wasn't* *working as it should. But it was half-working. So, I was able to go around the exhibition, but when the chap made it work in full technicolour, it was amazing, the room just came alive"* (P8). One notable glitch that emerged during the experience was that the headset extends several centimetres beyond the visitor's face. While perhaps not an issue for many VR experiences, this caused problems with the kinds of naturalistic and close-up interactions that we saw in Thresholds in which visitors peered closely through windows and into vitrines. Specifically, visitors reported knocking their goggles against physical boundaries as they (and we) had forgotten that the VR hardware physically extends the head: *"It was easy not to bump into the furnishings, but I did keep leaning in too far and hitting things with the glasses"* (P6).

DISCUSSION

Our account of the design and experience of Thresholds reveals how it is possible to deliver a compelling historical recreation in VR with sufficient throughput and flexibility to successfully tour to multiple museums. We propose that a key factor underlying this success is that Thresholds is built on an appropriate *VR architecture* that embeds its virtual model into various physical museum galleries. We use the term 'architecture' here very much in the conventional architect's sense of a spatial structure that affords and constrains movements, encounters and interactions (rather than in the sense of software or interface architecture). We call it a mixed reality architecture because this spatial structure combines both virtual and physical spaces into a complex whole. Our core argument is that, by establishing such an architecture in which virtual and real are both overlaid and juxtaposed in multiple ways, the artist was able to shape an overall user experience that provided a compelling moment of immersion while also dealing with the pragmatic challenges of embedding this into a museum and managing visitor throughput, entry, exit, spectating and ease of installation.

Overlaying virtual and physical spaces

The metaphor of overlaying in which virtual models appear to become layered on top of and aligned with physical sets and props lies at the heart of Thresholds. It is also central to various forms of immersive experience including substitutional reality [32] and augmented reality [25] as reviewed earlier. Overlaying involves using a suitably accurate and reliable tracking system combined with a headset display (opaque for substitutional reality and transparent for augmented reality) to register the virtual model with a physical environment. This can be seen in Thresholds' use of the Vive and Leap Motion tracking systems to register the virtual Model Room onto the physical set, combined with the careful design of the virtual world, choice of physical materials and addition of multisensory stimuli such as heat. However, overlaying in Thresholds is a more complex proposition than simply registering a virtual world with a physical one. As shown in Figure 25, Thresholds requires that several different spaces – or layers – become overlaid:

- The *original building* (as far as is known through plans and blue prints);
- The *virtual model* that recreates aspects of this;
- The *physical set* that then reproduces key elements of the virtual world;
- The *gallery space* into which this physical set is ultimately installed.

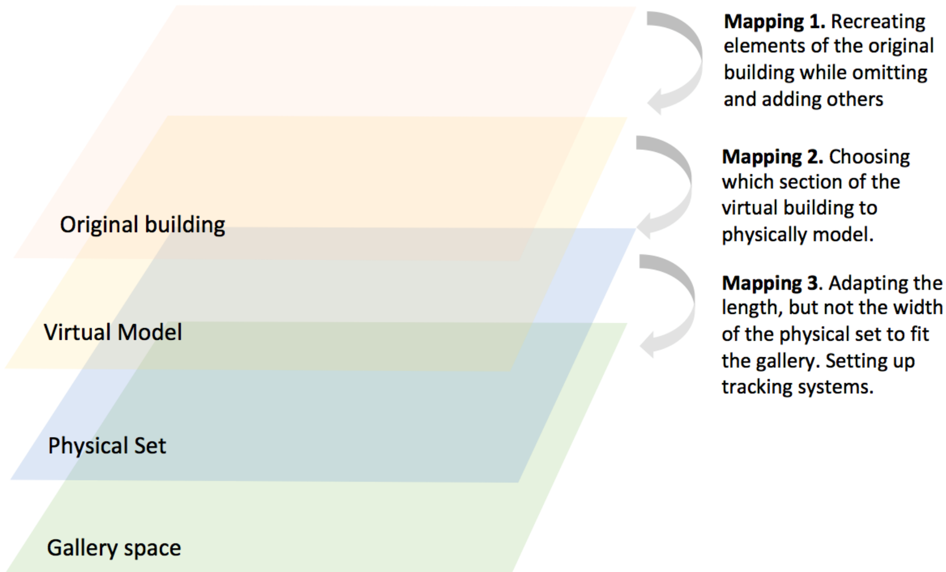

Figure 25. Mapping between four spatial layers in Thresholds.

Overlaying virtual upon real involves recognising the presence of all four of these layers and creating appropriate mappings between them. Each such mapping recreates some details of one layer in another but importantly, also omits some details while adding further ones. The first mapping (Mapping 1 in Figure 25) is between the original building – or at least the plans and photographs that are available – and the virtual model. Some elements of the original building such as its dimensions and architectural style are recreated while many others are undoubtedly omitted (not least because they are not known) and many further ones are added (the riot, fire, mice, moths and spider). The resulting virtual model is therefore an artistically made interpretation of the original building. All such virtual models are of course interpretations, even where they appear to strive for great realism, as some details will necessarily always be omitted and others added in the process.

The physical set is then a separate layer in Thresholds. Rather than directly mapping the virtual model onto each gallery space to which the work tours, it is instead mapped to a single physical set with the idea that this can then be installed in multiple galleries in order to facilitate touring (Mapping 2 in Figure 25). Again, this involves reproducing some details of the virtual model (the panelled walls, window bars and of course, the vitrines) while omitting others (the chandeliers for example) and introducing yet others (the spectator window). Adding the end walls (as the set is considerably smaller than the virtual model) required introducing the end cabinets as barriers into the virtual model in order to hide these from and/or explain them to visitors in a consistent way. Thus, we see that mappings work in two directions – introducing elements into one 'lower' layer (the physical set) may cause equivalents to be added back into the layer above (the virtual model).

The third mapping (Mapping 3 in Figure 25) is between the physical set and each gallery space in which it is installed. This involves finding a sufficiently large gallery to house the set but also determining the exact position and orientation of the set within this and addressing any local idiosyncrasies such as the positions of power supplies and ramps.

Thus, while overlaying might at first sight appear to be a straightforward metaphor for aligning virtual with physical, our experience of Thresholds reveals how it spans multiple layers, involves designing multiple mappings between these, how each such mapping involves reproducing but also omitting and adding elements, how this may work in both directions and how specialised tools may be needed to manage this. This observation reflects Steed et al.'s [38] exposition of mixed reality systems involving relating multiple geometric and symbolic spatial models including: '3D' models, room coordinate models and sensor models. Whereas their discussion largely focuses on the various representations of space that enable the mapping of virtual models to physical spaces, ours has extended this to also recognise the presence and importance of other layers and the questions of omitting and adding key content.

While presented here as a technical approach to embedding a virtual world within a physical museum, this notion of layering and also overlaying resonates with ideas from the world of art, most notably with the concept of mise en abyme, which describes the creative technique of placing a copy of an artwork within itself, e.g. a play within a play, in order to suggest an infinitely recurring sequence and to encourage reflection on different meanings [10]. In Thresholds, the reproductions of the original photographs are placed in a virtual reconstruction of the original building. This virtual reconstruction is situated within a wider historical setting of political and technical upheaval, represented by the Chartist riots outside the windows. In reality though, the physical exhibition space is contained within a gallery space, and visitors can gaze in at the physical set, intrigued, before entering. Moreover, in one case, when the exhibition was hosted at Lacock Abbey, the set was installed into Fox Talbot's historical home: the site where the technology to produce the original photographs was developed in the first place; and the location where several of the original images were taken (indeed the abbey is shown in several of those photos). Presenting the photographs in multiple levels of time and space may serve to superimpose their meaning then and now.

Juxtaposing virtual and physical spaces

Overlaying isn't the only spatial relationship at play in Thresholds. A second key one involves juxtaposing virtual and physical spaces, by which we mean appearing to arrange them side-by-side rather than on top of one another. Various forms of spatial juxtaposition are a commonplace feature of everyday built environments whose architecture relies on many techniques for separating and/or connecting different spaces including walls, windows and doors and other forms of internal barrier. Unsurprisingly, equivalent methods for juxtaposing spaces appear in virtual worlds, with virtual walls, windows and doors, though their solidity may sometimes differ from their physical counterparts. As discussed earlier, Koleva at al. [22] have previously extended juxtaposition to also cover mixed virtual and physical spaces, arguing that various forms of *mixed reality boundary* can be used to connect physical spaces to virtual spaces in order to establish opportunities for both seeing and moving between them.

We revisit and extend this concept of mixed reality boundaries to explain some of the important features of Thresholds. The architectural structures in each of the four layers described above have their own boundaries. The virtual Model Room is bounded, as is the street outside where the riot takes place, so that visitors cannot see beyond these to the empty space that lies beyond (i.e. the wider virtual world beyond is not modelled). The physical set is a large box bounded by wooden walls and ceiling. Each hosting gallery may also contain many further boundaries that define its rooms, corridors, exterior walls and so forth.

However, the various layers and their boundaries are in fact only partially aligned when they are overlaid, resulting in a mixed reality architecture that supports four different combinations of partially overlapping physical and virtual space.

- A central area in which the virtual model, physical set and gallery space are overlaid and aligned in the sense that participants inhabit all three simultaneously and experience congruent content. Participants in this central space can see, hear, move and touch within aligned virtual and physical worlds.
- The vitrines within this central area operate differently, defining a space into which the participants can see but into which they cannot physically reach and touch (instead having to virtually extract photographs using the summoning gesture).
- At the far ends of the virtual room and beyond the virtual windows lie parts of the virtual word into which visitors can see but cannot move or touch.
- Outside of the physical set is the antechamber, the queue and preparation area within the gallery space, where visitors prepare for the experience and cannot see the virtual world and from which invigilators may watch though the monitor screen.

Providing a coherent user experience of this complex and partial alignment of layers then relies on establishing various kinds of mixed reality boundary that afford different possibilities for seeing and moving between the constituent layers as shown in Figure 26 (which for simplicity of visualisation collapses the layers onto physical and virtual). Specifically, we propose that the designers of experiences similar to Thresholds need to consider how to employ five distinct types of mixed reality boundary so as to create an appropriate architecture for interaction:

- **Cordons** – these allow visitors to see into a wider virtual world but without being able to physically move there. Examples include the tall virtual cabinets at the end of the set and the virtual window that looks out into the street where the riot passes by. These two examples work subtly differently. The visual design of the cabinets discourages users from physical trying to step through them (if they did they would soon encounter the solid wall of the set just behind them), whereas the virtual window encourages physical engagement with an apparently solid boundary (e.g. leaning against or holding on to the window bars). In each case, the effect is to enhance the illusion that the visitor is in a much larger virtual space than the physical set that is actually constraining their ability to move.
- **Portals** – these are physical doorways that enable entry and exit to and from the central virtual part of the experience. Designers need to consider whether there is one portal for both entry and exit or multiple ones (possibly allowing for a one-

way flow through the experience) and also on which side of each such boundary visitors put on or take off the VR equipment. Thresholds, for example, uses a single portal for both entry and exit between the virtual experience and the physical antechamber outside, with the experience of passing through being consciously designed to be asymmetric: visitors initially don their headsets in the antechamber to be led blindfold into the set before the virtual world is introduced, whereas at the end of the experience, they remove their headsets while still inside the set. Knibbe et al. [21] have discussed the challenge of designing the exit from VR, arguing that the moment of exit is an opportunity for designing enhanced engagement. Our artist deliberately designed this moment of reveal at the end of the experience in which the headset was removed while the participant was still inside the physical set (even though they had been led in blindfolded by it) so as to create a shocking contrast between the visual richness of the virtual model and blank sparseness of the set.

- **Spectator portholes** – these allow spectators who await their turn to view aspects of the experience in advance so as to be prepared (as well as passing time while waiting) or those who have completed it to watch others and reflect back on their experience. A key design tension concerns how much of the experience to reveal. Reeves et al. [27] presented a taxonomy of spectator interfaces according to the extent that they reveal or hide the primary user's physical manipulations of an interface versus the consequent effects of these (e.g. digital content), which in turn revealed four design strategies: expressive (manipulations and effects both revealed); secretive (manipulations and effects both hidden); magical (manipulations hidden and effects revealed); and suspenseful (manipulations revealed but effects hidden). The spectator window in Thresholds adopts the latter strategy as spectators can see how visitors are moving about the blank set but cannot see the content that appears in their headsets. Moreover, this boundary is unidirectional – spectators can see the visitors but not vice versa.
- **Vitrines** – vitrines are mixed reality boundaries behind which virtual artefacts can be seen but cannot be touched. Whereas traditional physical vitrines lock physical artefacts behind glass to prevent them being touched for their own security and safety, our mixed reality vitrines lock away objects that cannot be physically felt because there is no mapping of virtual content onto a corresponding physical prop. They also signal that an additional interaction mechanism now has to be employed instead, in the case of Thresholds holding one's hand above the vitrine to lift up a virtual photograph for closer inspection.
- **Monitors** – are boundaries that enable invigilators, technicians and other operators to monitor what is happening within the experience in case they should need to intervene in some way. Thresholds employs two of these. The first is the *monitor screen* positioned in the antechamber that provides a window into the virtual world through which invigilators can view its virtual inhabitants. The second is the ghost-like avatar in the corner of the model that masks an area where invigilator can physically stand in the set. Both of these are one-way boundaries with those inside not being able to look back out to see those outside who are watching them.

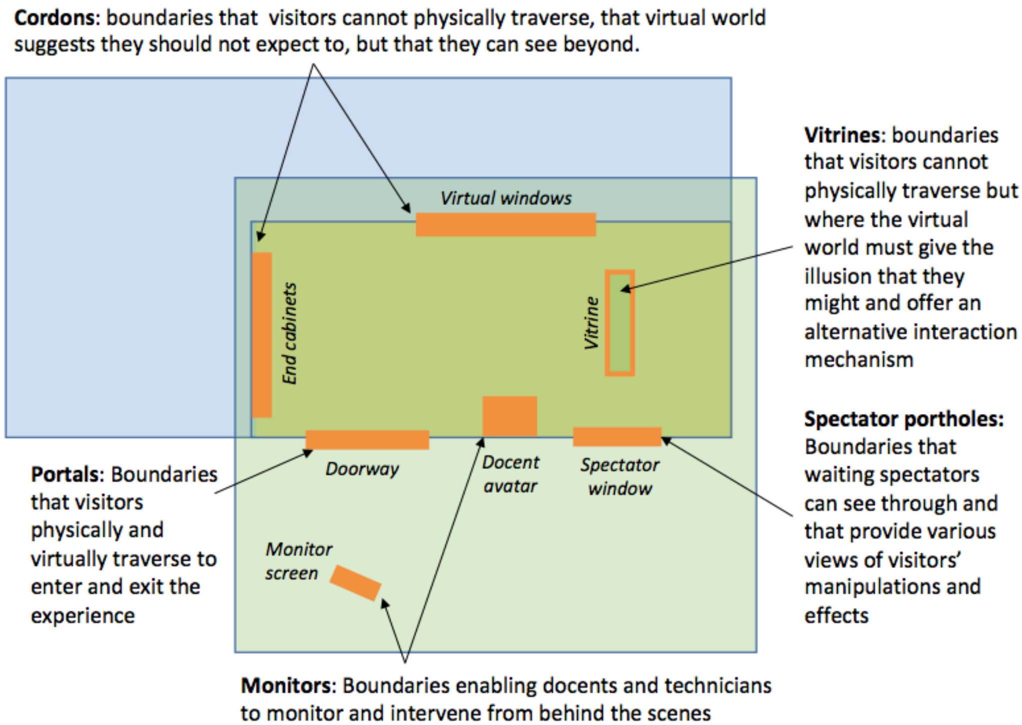

Figure 26. Mixed Reality Boundaries in Thresholds.

This analysis of how the various layers and spaces in Thresholds are connected by different kinds of mixed reality boundaries serves to illustrate the complexities of juxtaposing physical and virtual spaces. We propose that the designers of museum VR experiences similarly need to employ cordons, portals, spectator portholes, vitrines and monitors to divide and reconnect virtual and physical spaces in a such a way as to deliver a compelling illusion of being immersed in a virtual world in which many things can be touched, while also addressing the pragmatic challenges of entry, exit, spectating and monitoring that arise when delivering high-throughput public experiences. Some boundaries may be introduced as a matter of necessity due to the constraints of the situation (e.g. to cover up the limited space of physical movement), while others may be carefully designed as a matter of choice so as to enhance the experience in unusual ways (e.g. providing unusual ways of spectating). Some serve to partition spaces while others establish connections between layers.

In the case of Thresholds, the resulting *partial-overlay* of the layers ensures the provision of an antechamber and physically inaccessible areas of the virtual world arranged around the central area. Indeed, this central area was the only part the whole experience where the conventional overlaying of virtual onto physical was attempted. All of the other surrounding spaces were *connected but not overlaid*, but nevertheless remained vital parts of the experience.

Other factors in the user experience

While the focus of this paper is on the unusual mixed reality architecture of Thresholds, we recognise that many other key decision decisions also contributed to the overall user experience and briefly note several of these here for completeness.

One key decision was the design of ghostlike fuzzy avatars which served to (i) show users their own hands (showing people their own virtual bodies has be shown to lead to a stronger sense of presence in VR [28,34,35]) while managing to successfully mask minor misalignments between the virtual model and physical set; and (ii) convey the presence of others in the space so as to help avoid physical collisions while retaining anonymity and not inviting social interaction so as to deliver a largely solitary but shared experience. We also note a key challenge for embodiment in that headsets protruded several centimetres beyond visitors' faces, which could lead to physical collisions if they peered very close to vitrines.

Perhaps the most taxing decision was the choice to break the metaphor of physical alignment by enabling visitors to pick up the virtual photographs. While broadly accepted, we note the possibility of drawing on emerging techniques for passive-haptic props to address this challenge in the future such as: the work of Stoakley et al. [39] that used clipboards as props to provide passive haptic feedback; the concept of a tangible hologram system proposed by Signer and Curtin [31]; and the weight shifting physical proxy object developed by Zenner and Kruger [46].

Finally, we note the importance of the additional sensory stimulus of temperature to help deliver a compelling moment of immersive experience at the fireplace. Heat is interesting in this regard, being both subtly evocative and yet also tolerant of far coarser alignment.

KEY PERFORMANCE INDICATORS AND GUIDELINES

Here we provide some key performance indicators (KPIs) that we have developed through our analysis of Thresholds. We believe these could reasonably be applied to any VR exhibition experience. These measures may mostly be assessed by keeping logs of headset positions and directions, and further supplemented with gaze (or even eye tracking) data and qualitative follow ups.

- **All items are being looked at**. We can see from the log data (Figure 22) that while distribution of time spent looking at the images is not necessarily even, all of the images are getting some attention.
- **Visitors explore the full space**. We can see from the spatial distribution of tracking (Figure 18) that visitors fully explored the available space.
- **Event based items are visited at the appropriate time.** We can see from the log files (Figure 19) that the riot – the only timed event in the experience – does indeed encourage visitors to attend to it at the appropriate time.
- **Noticing key features.** We can see from the spatial distribution (Figure 18) that visitors did indeed spend time in front of the fire, and the spectator window as our design intended. This is further backed up by the visitor accounts from the study.
- **Avoiding areas specifically blocked off.** We can see from the spatial distribution (Figure 18) that visitors did not step into the area we 'protected' with an avatar – the location where we intended our invigilator

to stand. Similarly, visitors did not attempt to cross the virtual boundaries between the real edge of the room and the extended virtual room – our use of both physical and digital barriers here was effective.

- **Tracking error density is low**. While tracking errors do occur in Thresholds, we can see from the log recordings that they are not too common, and tend to occur in known places and with known reasons (Figure 23). We can further mitigate these errors through scripting the experience – that is telling visitors before they begin not to hold the headset (covering the receivers), and what to do when such errors occur.
- **Alignment error is tolerable**. The use of fuzzy edges for tracked hands helps disguise minor alignment errors. As long as significant tracking errors do not occur, this strategy has been sufficient to provide a compelling haptic experience.

We now provide some guidelines for designers looking to create similar types of mixed reality experiences.

- **Picking the right tracking system.** There are many different tracking systems to choose from, e.g. "inside out", "outside-in", etc. For any given system it is necessary to balance reliability and flexibility with cost. For VR, a tracking system needs to be stable (i.e. not jittery), fast (i.e. not laggy) and provide sufficiently few errors not to break the trust of the users. In instances such as Thresholds where there is a physical/digital alignment such as the use of props, this trust is doubly important.
- **Creating sufficiently correct alignment.** Regardless of the in-practice (or in-lab) quality of the tracking system, the alignment between physical and digital assets is necessarily subject to human error as it is calibrated by humans – even in automated systems such as "place the controller on the edge of the table". When multiple tracking systems are in use (such as the addition of Leap Motions to track hands) these errors can be multiplied or magnified. Mitigation strategies should be used to reduce the recognisability of such errors. This may include not showing hands/bodies if appropriate or using fuzzy edges, as in Thresholds, and relying on proprioception error margins to disguise errors in alignment.
- **No inclusion of physical items without digital counterparts.** Similarly to the showing of other visitors' positions, it is necessary to recall that the visitors in VR are effectively blindfolded. As such, physical objects (including people) need to be represented in some way in the digital space to prevent accidents.
- **Clear barriers between accessible and non-accessible spaces.** These barriers need not always be physical, however they must be represented digitally (as above). Such digital barriers need to be clear and understandable by users. While systems such as chaperones (e.g. in Vive and Oculus Rift) do not attempt to "fit in" with the experience, simply appearing as a saftey warning, it is possible to achieve the same effect using "in world" objects, as in Thresholds.
- **Understanding mixed reality boundaries.** It is necessary to design spaces, both physical and digital, which correctly reflect the different boundaries between, and characteristics of those spaces – i.e. which are passable, which are visual, what purpose each serves, and to whom they are relevant.
- **Handling multiple users.** If the system needs to handle multiple users simultaneously then they need either be physically separated or made

aware of each other's positions. This position information needs to be delivered in a way that makes sense within the context (e.g. Thresholds' ghosts). Similarly, if having multiple roaming users, then it's necessary to provide wireless client systems for safety.

- **Carefully scripted support.** It is necessary to provide appropriate training for visitors and to set their expectations for the experience. This may include safety briefings, as well as usage of the system (e.g. telling users not to hold the headset).

CONCLUSION

The design and experience of Thresholds reveals how to deliver compelling multi-sensory VR experiences in museums at scale. Its key innovation has been to establish an 'architecture for interaction', a multi-layered arrangement of virtual and physical spaces that are both overlaid and juxtaposed to embed a virtual recreation of a historical scene into the surrounding museum space. At its core, it employs passive haptics – a form of overlaying virtual and physical and virtual spaces – to deliver a powerful tactile and kinaesthetic VR experience. However, it also uses a 'box in a box' approach – that involves various juxtapositions – to carefully situate the experience within the wider museum spaces, relating it to the challenges of existing and entering the space, enabling spectating, delivering throughput and being able to fit a large virtual model into the constraints of far smaller museum galleries.

We suggest that museum designers who wish to employ VR can employ the same techniques to deliver compelling and scalable experiences. Turning to the VR community, we suggest that it is important to consider virtual spaces as potentially being both overlaid and also juxtaposed with physical ones as part of a wider notion of VR architecture.

ACKNOWLEDGEMENTS

This research has been supported by the Engineering and Physical Sciences Research Council (EPSRC) Horizon: My Life in Data Centre for Doctoral Training (EP/L015463/1) and by the European Commission through the GIFT project (grant 727040). The authors would like to express special gratitude to photographic historian Pete James, a major contributor to the project who passed away in early 2018; Pete Gilbert for running the touring show; and all the venues where 'Thresholds' has been exhibited.

REFERENCES

1. Laurent Aguerreche, Thierry Duval, and Anatole Lécuyer. 2010. Reconfigurable Tangible Devices for 3D Virtual Object Manipulation by Single or Multiple Users. In *Proceedings of the 17th ACM Symposium on Virtual Reality Software and Technology* (VRST '10), 227–230. https://doi.org/10.1145/1889863.1889913
2. Eli Anapur. Mat Collishaw's Thresholds Explore Photography through VR at Somerset House. *WideWalls*. Retrieved September 18, 2017 from http://www.widewalls.ch/mat-collishaw-thresholds-somerset-house/
3. Brian Appleyard. 2017. VIRTUALLY BRILLIANT; Once a gritty YBA, Mat Collishaw is now one of our best grown-up artists. And his new VR show is genius. *The Sunday Times (London)*, CULTURE;FEATURES; Pg. 10,11.
4. Mafkereseb Bekele, Roberto Pierdicca, Emanuele Frontoni, Eva Malinverni, and James Gain. 2018. A Survey of Augmented, Virtual, and Mixed Reality for

Cultural Heritage. *Journal on Computing and Cultural Heritage* 11: 1–36. https://doi.org/10.1145/3145534
5. Steve Benford, Chris Greenhalgh, Andy Crabtree, Martin Flintham, Brendan Walker, Joe Marshall, Boriana Koleva, Stefan Rennick Egglestone, Gabriella Giannachi, Matt Adams, Nick Tandavanitj, and Ju Row Farr. 2013. Performance-Led Research in the Wild. *ACM Trans. Comput.-Hum. Interact.* 20, 3: 14:1–14:22. https://doi.org/10.1145/2491500.2491502
6. Barry Brown, Ian MacColl, Matthew Chalmers, Areti Galani, Cliff Randell, and Anthony Steed. 2003. 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), 577–584. https://doi.org/10.1145/642611.642711
7. Lung-Pan Cheng, Li Chang, Sebastian Marwecki, and Patrick Baudisch. 2018. iTurk: Turning Passive Haptics into Active Haptics by Making Users Reconfigure Props in Virtual Reality. In *Proceedings of the 2018 CHI Conference on Human Factors in Computing Systems* (CHI '18), 89:1–89:10. https://doi.org/10.1145/3173574.3173663
8. Hannah Ellis-Petersen. 2017. Mat Collishaw restages 1839 photography show in virtual reality; Artist says VR will change our outlook as he prepares Somerset House display based on Henry Fox Talbot's seminal exhibition. *The Guardian(London)*, TECHNOLOGY; Version:1.
9. SJ Eve. 2014. Dead men's eyes: embodied GIS, mixed reality and landscape archaeology. UCL (University College London).
10. Gerhard Fischer and Bernhard Greiner. 2007. *The Play Within the Play: The Performance of Meta-theatre and Self-reflection*. Rodopi.
11. Daniel Fitzgerald and Hiroshi Ishii. 2018. Mediate: A Spatial Tangible Interface for Mixed Reality. In *Extended Abstracts of the 2018 CHI Conference on Human Factors in Computing Systems* (CHI EA '18), LBW625:1–LBW625:6. https://doi.org/10.1145/3170427.3188472
12. Bernard Frischer, Dean Abernathy, Gabriele Guidi, Joel Myers, Cassie Thibodeau, Antonio Salvemini, Pascal Müller, Peter Hofstee, and Barry Minor. 2008. Rome Reborn. In *ACM SIGGRAPH 2008 New Tech Demos* (SIGGRAPH '08), 34:1–34:1. https://doi.org/10.1145/1401615.1401649
13. A. Gaucci, S. Garagnani, and A. M. Manferdini. 2015. Reconstructing the lost reality archaeological analysis and Transmedial Technologies for a perspective of Virtual Reality in the Etruscan city of Kainua. In *2015 Digital Heritage*, 227–234. https://doi.org/10.1109/DigitalHeritage.2015.7419502
14. William Gaver. 2012. What Should We Expect from Research Through Design? In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems* (CHI '12), 937–946. https://doi.org/10.1145/2207676.2208538
15. C. F. Giloth and J. Tanant. 2015. User experiences in three approaches to a visit to a 3D Labyrinthe of Versailles. In *2015 Digital Heritage*, 403–404. https://doi.org/10.1109/DigitalHeritage.2015.7413914
16. Google. 2017. Tango. Retrieved September 19, 2017 from https://get.google.com/tango/
17. Lilian de Greef, Meredith Morris, and Kori Inkpen. 2016. TeleTourist: Immersive Telepresence Tourism for Mobility-Restricted Participants. In *Proceedings of the 19th ACM Conference on Computer Supported Cooperative Work and Social Computing Companion* (CSCW '16 Companion), 273–276. https://doi.org/10.1145/2818052.2869082
18. H. G. Hoffman. 1998. Physically touching virtual objects using tactile augmentation enhances the realism of virtual environments. In *Proceedings.*

IEEE 1998 Virtual Reality Annual International Symposium (Cat. No.98CB36180), 59–63. https://doi.org/10.1109/VRAIS.1998.658423
19. Brent Edward Insko. 2001. Passive Haptics Significantly Enhances Virtual Environments. The University of North Carolina at Chapel Hill.
20. Phisan Kaewprapha, Nattakan Puttarak, and Thaewa Tansarn. 2016. Multi-hop network localization in unit disk graph model under noisy measurement using tree-search algorithm with graph-properties-assist traversing selection. *Engineering and Applied Science Research* 43: 114–117.
21. Jarrod Knibbe, Jonas Schjerlund, Mathias Petraeus, and Kasper Hornb\a ek. 2018. The Dream is Collapsing: The Experience of Exiting VR. In *Proceedings of the 2018 CHI Conference on Human Factors in Computing Systems* (CHI '18), 483:1–483:13. https://doi.org/10.1145/3173574.3174057
22. Boriana Koleva, Steve Benford, and Chris Greenhalgh. 1999. The Properties of Mixed Reality Boundaries. In *ECSCW '99*. Springer, Dordrecht, 119–137. https://doi.org/10.1007/978-94-011-4441-4_7
23. Boriana Koleva, Holger Schnädelbach, Steve Benford, and Chris Greenhalgh. 2000. Developing Mixed Reality Boundaries. In *Proceedings of DARE 2000 on Designing Augmented Reality Environments* (DARE '00), 155–156. https://doi.org/10.1145/354666.354690
24. G. Lepouras, A. Katifori, C. Vassilakis, and D. Charitos. 2004. Real exhibitions in a virtual museum. *Virtual Reality* 7, 2: 120–128. https://doi.org/10.1007/s10055-004-0121-5
25. Paul Milgram, Haruo Takemura, Akira Utsumi, and Fumio Kishino. 1995. Augmented reality: a class of displays on the reality-virtuality continuum. In *Telemanipulator and Telepresence Technologies*, 282–293. https://doi.org/10.1117/12.197321
26. T. Miyashita, P. Meier, T. Tachikawa, S. Orlic, T. Eble, V. Scholz, A. Gapel, O. Gerl, S. Arnaudov, and S. Lieberknecht. 2008. An Augmented Reality museum guide. In *2008 7th IEEE/ACM International Symposium on Mixed and Augmented Reality*, 103–106. https://doi.org/10.1109/ISMAR.2008.4637334
27. Stuart Reeves, Mike Fraser, Holger Schnadelbach, Claire O'Malley, and Steve Benford. 2005. Engaging Augmented Reality in Public Places. In *Adjunct proceedings of SIGCHI Conference on Human Factors in Computing Systems*, 11.
28. Maria V. Sanchez-Vives and Mel Slater. 2005. From presence to consciousness through virtual reality. *Nature Reviews Neuroscience* 6: 332.
29. Larry J. Schaaf. Perception of Realities, No. 1. Retrieved September 18, 2017 from http://foxtalbot.bodleian.ox.ac.uk/2017/05/19/perception-of-realities-no-1/
30. Holger Schnädelbach, Boriana Koleva, Martin Flintham, Mike Fraser, Shahram Izadi, Paul Chandler, Malcolm Foster, Steve Benford, Chris Greenhalgh, and Tom Rodden. 2002. The Augurscope: A Mixed Reality Interface for Outdoors. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems* (CHI '02), 9–16. https://doi.org/10.1145/503376.503379
31. Beat Signer and Timothy J. Curtin. 2017. Tangible Holograms: Towards Mobile Physical Augmentation of Virtual Objects. *arXiv:1703.08288 [cs]*. Retrieved from http://arxiv.org/abs/1703.08288
32. Adalberto L. Simeone, Eduardo Velloso, and Hans Gellersen. 2015. Substitutional Reality: Using the Physical Environment to Design Virtual Reality Experiences. In *Proceedings of the 33rd Annual ACM Conference on Human Factors in Computing Systems* (CHI '15), 3307–3316. https://doi.org/10.1145/2702123.2702389
33. Sky. Nominations announced for The South Bank Sky Arts Awards 2018. *Sky Corporate*. Retrieved February 22, 2019 from

https://www.skygroup.sky/corporate/media-centre/articles/en-gb/nominations-announced-for-the-south-bank-sky-arts-awards-2018

34. Mel Slater and Martin Usoh. 1993. Representations Systems, Perceptual Position, and Presence in Immersive Virtual Environments. *Presence: Teleoperators and Virtual Environments* 2, 3: 221–233. https://doi.org/10.1162/pres.1993.2.3.221
35. Mel Slater, Martin Usoh, and Anthony Steed. 1994. Depth of Presence in Virtual Environments. *Presence: Teleoperators and Virtual Environments* 3, 2: 130–144. https://doi.org/10.1162/pres.1994.3.2.130
36. Hyungki Son, Hyunjae Gil, Sangkyu Byeon, Sang-Youn Kim, and Jin Ryong Kim. 2018. RealWalk: Feeling Ground Surfaces While Walking in Virtual Reality. In *Extended Abstracts of the 2018 CHI Conference on Human Factors in Computing Systems* (CHI EA '18), D400:1–D400:4. https://doi.org/10.1145/3170427.3186474
37. Julian Stallabrass. 1999. *High Art Lite: British Art in the 1990s*. Verso.
38. A. Steed, J. MacColl, C. Randell, B. Brown, M. Chalmers, and C. Greenhalgh. 2004. Models of space in a mixed-reality system. In *Proceedings. Eighth International Conference on Information Visualisation, 2004. IV 2004.*, 768–777. https://doi.org/10.1109/IV.2004.1320228
39. Richard Stoakley, Matthew J. Conway, and Randy Pausch. 1995. Virtual Reality on a WIM: Interactive Worlds in Miniature. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems* (CHI '95), 265–272. https://doi.org/10.1145/223904.223938
40. Stella Sylaiou, Katerina Mania, Athanasis Karoulis, and Martin White. 2010. Exploring the relationship between presence and enjoyment in a virtual museum. *International Journal of Human-Computer Studies* 68, 5: 243–253. https://doi.org/10.1016/j.ijhcs.2009.11.002
41. Laurie Taylor. It's a Kind of Magic. Retrieved September 18, 2017 from https://frieze.com/article/its-kind-magic
42. M. White, N. Mourkoussis, J. Darcy, P. Petridis, F. Liarokapis, P. Lister, K. Walczak, R. Wojciechowski, W. Cellary, J. Chmielewski, M. Stawniak, W. Wiza, M. Patel, J. Stevenson, J. Manley, F. Giorgini, P. Sayd, and F. Gaspard. 2004. ARCO - an architecture for digitization, management and presentation of virtual exhibitions. In *Proceedings Computer Graphics International, 2004.*, 622–625. https://doi.org/10.1109/CGI.2004.1309277
43. William Henry Fox Talbot. 1839. Talbot Correspondence Project: TALBOT William Henry Fox to HERSCHEL John Frederick William. Retrieved June 13, 2018 from http://foxtalbot.dmu.ac.uk/letters/transcriptDocnum.php?docnum=3908
44. Gaby Wood. 2017. Somewhere between art and illusion; Virtual reality meets Victorian fakery in Mat Collishaw's eerie new works, finds Gaby Wood. *The Daily Telegraph (London)*, REVIEW;NEWS; Pg. 6,7.
45. Eric Woods, Mark Billinghurst, Julian Looser, Graham Aldridge, Deidre Brown, Barbara Garrie, and Claudia Nelles. 2004. Augmenting the Science Centre and Museum Experience. In *Proceedings of the 2Nd International Conference on Computer Graphics and Interactive Techniques in Australasia and South East Asia* (GRAPHITE '04), 230–236. https://doi.org/10.1145/988834.988873
46. A. Zenner and A. Krüger. 2017. Shifty: A Weight-Shifting Dynamic Passive Haptic Proxy to Enhance Object Perception in Virtual Reality. *IEEE Transactions on Visualization and Computer Graphics* 23, 4: 1285–1294. https://doi.org/10.1109/TVCG.2017.2656978
47. Yiwei Zhao, Lawrence H. Kim, Ye Wang, Mathieu Le Goc, and Sean Follmer. 2017. Robotic Assembly of Haptic Proxy Objects for Tangible Interaction and Virtual Reality. In *Proceedings of the 2017 ACM International Conference on*

Interactive Surfaces and Spaces (ISS '17), 82–91. https://doi.org/10.1145/3132272.3134143
48. John Zimmerman, Jodi Forlizzi, and Shelley Evenson. 2007. Research Through Design As a Method for Interaction Design Research in HCI. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems* (CHI '07), 493–502. https://doi.org/10.1145/1240624.1240704
49. Take a virtual tour of the National Gallery | Visiting | National Gallery, London. Retrieved July 4, 2018 from https://www.nationalgallery.org.uk/visiting/virtual-tours
50. Louvre Museum. Retrieved July 4, 2018 from http://www.youvisit.com/tour/louvremuseum
51. factory-42 | HOLD THE WORLD. *factory-42*. Retrieved July 4, 2018 from https://www.factory42.uk/holdtheworld
52. Nefertari: Journey to Eternity on Steam. Retrieved July 4, 2018 from https://store.steampowered.com/app/861400/Nefertari_Journey_to_Eternity/
53. Skin and Bones - Mobile Augmented Reality App for The National Museum of Natural History's Hall of Bones. Retrieved July 4, 2018 from https://naturalhistory.si.edu/exhibits/bone-hall/index.cfm