

Project anywhere: An interface for virtual architecture

International Journal of
Architectural Computing
1–12

© The Author(s) 2016

Reprints and permissions:

sagepub.co.uk/journalsPermissions.nav

DOI: 10.1177/1478077116670746

jac.sagepub.com



Constantinos Miltiadis

Abstract

Virtual and augmented realities open a new world of great potential for spatial research and experimentation by allowing new forms of unbuilt sensible architectural space. This article starts with a sketch of the current context in virtual reality and continues by outlining the development and structure of the research ‘project Anywhere’. The project is an easily deployable, wireless, multi-user, augmented reality app system that offers full body immersion through body, head and hands tracking. It can host multiple concurrent users, able to move freely in the virtual space, by moving in the real and also perform actions through a gesture interface to affect their shared environment. In conclusion, we describe the inherent properties of such a space, which we propose as a novel spatio-temporal medium for architecture that suggests an enriched notion of space for exploration and experimentation, through an example of a potential application.

Keywords

Virtual reality, augmented reality, interactive environments, virtual space

Introduction

With 2016 seen in the public eye either sceptically¹ or concretely² as the year of virtual reality (VR), one cannot help to acknowledge some given facts. After the revival of the concept of VR triggered by the Oculus Kickstarter³ campaign in 2012, the development of various developer-targeted products in the next 2 years, followed by the acquisition of Oculus by Facebook for USD2 billion in 2014,⁴ the term ‘virtual reality’ has established itself in mainstream culture and has created a significant market demand from gamers and non-gamers alike. Today, we reached the point of having consumer VR devices^{5,6} readily available – for either consumption or production – while more tech-companies⁷ are engaging the field by announcing products and content to be available later this year.

What we are going to argue about is that the experience provided by VR devices is mainly spatial; replacing one’s visual and auditory physical surroundings with a believable and responsive virtual alternative is inherently a question of space that requires a new design paradigm. At this stage, we are still in the formative years of the adoption of such technologies, which is still problematic in its conception. The best example is

Institute of Architecture and Media, Graz University of Technology, Graz, Austria

Corresponding author:

Constantinos Miltiadis, Institute of Architecture and Media, Kronesgasse 5/3, A-8010, Graz, Austria.

Email: c.miltiadis@gmail.com

the two main market rivals offering vastly different types of experience: Oculus at the moment suggests a sitting pose and interfacing through a gamepad, while the HTC Vive that enables head and hands tracking in a 4.5 m × 4.5 m – approximately – area allows a bodily immersive interaction in space. The democratisation of the development tools for these technologies, also allow us, architects, to join this quest.

The credibility of VR

The field that this article is concerned about is VR through the implementation of head-mounted displays (HMDs), like the ones mentioned above. Their function is to render two digital images at a time, using two cameras arranged in a distance of human eyes (interpupillary distance (IPD)), which with the display worn in front of the users' eyes create the illusion of depth and hence stereoscopic view, in an interactive manner in which any tilting of the head corresponds and affects the virtual point of view in real time. Paramount to the application of HMD VR was computer performance.⁸ Because the premise and basis of HMD VR technology are essentially tricking the brain to experience or perceive something virtual (not real), the major obstacles that delayed its development were maintaining a high refresh rate, at least 60 Hz, and imperceptible low latency – a combination of hardware and software that renders graphics in a very high frequency and that responds very quickly to the user's head rotation. Implementations lacking or underperforming in any of these parameters result in a form of nausea, which is in this case called simulator sickness.⁹ Since the 1980s, multiple efforts were made in order to develop and bring VR technology to the masses. However, only today, one can truly admit that the technology enabling VR has been advanced to such an extent that renders its implementation both viable and worthwhile.

At this point, regarding the scientific credibility of VR, it is worth and suffices to mention CAVE labs,¹⁰ which were also the first substantial reproducible and functional VR setups. CAVEs are relatively small rooms with computer-generated stereoscopic projections on most of their surfaces, in which through head-tracking and polarised glasses, a subject can move and experience an interactive stereoscopic environment. They were first put into practice in the early 1990s and are since producing qualitative scientific results in cognitive and psychological studies, as human behaviour in virtual environments has proven to stem from the same processes as in real environments and produce similar results.¹¹ More recently, with the advent of HMDs, VR experiments became even more widespread for their ability to consistently simulate controlled environments and situations, having also proven to some extent their linguistic credibility.¹²

We, therefore, can empirically maintain that VR technologies can produce functional experiences of environments and spaces that are both cognitively engaging and affective. These characteristics of such a spatial configuration we intend to transfer to the domain of architecture as tools in order to experiment with qualities and notions of unbuilt space.

VR and architecture

VR in its current form was invented by Ivan E. Sutherland^{13,14} in the late 1960s. The interesting fact, though, is that the same person is single-handedly responsible also for 'Sketchpad',¹⁵ just 5 years earlier. The question arises: How come the same individual who invented the precursor to the most commonly used architectural software go on to invent a device that enabled alternate realities? Are virtual design and experiencing virtual designs such closely related fields? Our reply is that the answer depends on one's standpoint. If digital design software is taken as means of replacing an equivalent manual task, then the answer is negative; if, on the other hand, design software is paired with a means of sensible manifestation to an audience, then the result is unprecedented, as design is then decoupled from an end purpose of materialisation in the physical world. In the latter case, design can begin and end in the digital realm and can be experienced and celebrated as such.

Going further back in time to explore the relationship of architecture and the concept of the virtual as well as means and ends, we observe two main examples. In the first, Brunelleschi, credited for the modern discovery of linear perspective, devised an interesting experiment to prove his discovery in 1420.¹⁶ In his setup, the subject would look at a building through a hole from the backside of a perspective painting of the same building. Then by placing a mirror in front of the painting, the subject would see the painting instead of the building while maintaining its surroundings and therefore comparing his perspective construction to the actual subject's view of the real building and creating what we could today consider as the first augmented reality (AR) experience. The second example supporting the fascination of architecture with the virtual or unreal is trompe-l'œil murals, which, using an established technique and utilising a forced perspective, intended to extend the real, either as a cognitive play on space or as a conceptual extension to a virtual dimension, such as the examples of Andrea Pozzo. (From Andrea Pozzo's work, we can examine the case of the dome of the Jesuit Church in Vienna (~1703) that creates the illusion of a deeper space and the ceiling of Saint Ignatius in Rome (1685–1694) that expands space by placing the spectator under a deep perspective divine scenery.)

Our claim is that these former examples are no longer experiments or additive elements to architecture, respectively, but because of their spatial effect, they are potentially architecture themselves.

The same viewpoint we would like to take for VR and AR technology, which presently, in the field of architecture, has been adopted as means to visualise designs intended to be built.¹⁷ Although this application is effective for the purpose, our experience shows that VR is much more capable than that. Our proposal is to surpass the intention of previewing, and utilising the fact that VR can produce believable environments, we suggest that there can be another use of VR, as an architectural choro-poietic (space-creating) medium to design and experience spaces solely intended for that purpose. As demonstrated by previous works, like the proposal for a virtual museum and 'project Anywhere' – that will be elaborated in what follows – creating rich virtual environments in which one can move, experience, act and interact is within our reach. VR can today serve as a new architectural platform, with significant emergent properties to classify it as a distinct form of space and where real and virtual can be mixed in various degrees.¹⁸ We not only can design spaces otherwise unthought of, overlay virtual on top of real spaces and expand notions of space, but can furthermore design in four dimensions – not just spaces, but also eventful situations. Lending the cognitive credibility of VR, we want to proclaim and prove virtual space (VS) as a new spatial platform for novel design and experimentation, where the limits are not any more what is possible in our physical world, but what is imaginable.

Project Anywhere

Research goals

Project Anywhere was initiated in May 2014 and was developed by the author, at the Chair of CAAD ETH Zurich. Experiencing an immersive environment with a contemporary HMD motivated us to imagine scenarios of VR applications in architecture, beyond the one of visualisation. Our research goal was to attempt a 'viable' VS, by enhancing both the immersive experience and the nature of its space, in an easily deployable and reproducible setup. Consequently, the three main pillars we set were to include more properties from the human spatial presence in the virtual environment, rather than only vision and head movement, in order to create a greater degree of immersion; to invent intuitive ways of interfacing and performing actions in the VS; and eventually enable multiple users to coexist simultaneously in a common, shared VS, which is not allowed by CAVEs. In other words, to increase the degree of cognitive self-awareness in the virtual world, allow the user to affect space and add a social aspect to it. Eventually, our aim was to explore the possibility and potential of experiencing VS under the proposed circumstances and what it could offer to architecture.

Setup outline

The next section will describe the setup that was necessary to develop in order to realise the project that consists of commercial devices (smartphones, infrared marker-less motion trackers) as well as hardware and software prototypes (passive head mount, data gloves, tracking software and synchronised desktop and mobile apps). While we are not advocating that our prototypes are individually exemplary or original inventions (e.g. the data gloves), we do hold that the overall setup configuration is for the aforementioned implementation purposes, as we intend to prove, optimal.

To mention as an introduction sketch, the setup includes a smartphone functioning as an HMD and also sensors for wireless motion and gesture tracking. The tracking data are recorded on a stationary computer and then sent through a web cloud to the smartphone application in real time. All components are modular and can be easily interchanged or upgraded, therefore making the whole system flexible to adapt to future technological advancements.

HMD platform

The major property that we wanted to include and create a correspondence of, from the real to the VS, was physical presence and movement. In order to have the user freely moving in a given space where the project would be installed, we had to minimise cable dependencies and therefore abandoned the idea of using an HMD device. Instead, we chose to implement a wireless mobile setup using a smartphone as an HMD platform. At the time, the concept of using smartphones as VR HMDs was at its beginnings, and commercial products were still not available. (Google Cardboard, which was the first to be released, was announced at the Google I/O conference in June 2014.)¹⁹ Eventually, we made use of the Durovis Dive SDK²⁰ that provided a framework for adequately utilising a smartphone as an HMD. For a smartphone mount, Durovis²¹ offered the ‘Open Dive’, an open computer-aided design (CAD) drawing for a three-dimensional (3D) printable smartphone head mount which we used while prototyping consecutive versions that solved various problems of the original IPD adjustment, light proofing, rigidity, ergonomics and so on. The final version labelled ‘omni’ was optimised for an iPhone 5s and can be fabricated on a regular 3D printer at a minimal cost, requiring as extra parts lenses (50 mm \varnothing 25 mm), padding and an elastic strap (Figure 1).

In terms of performance, while active HMDs have a clear advantage in both performance and quality, using smartphones is an easily deployable – also in larger quantities – cheaper solution, that offers a decent, satisfactory performance, under certain load thresholds.

Real-time tracking

Tracking devices. Already having the head movement of the user captured by the IMU (Inertia Measurement Unit) of the phone, we furthermore needed to find solutions for how to perform body and gesture tracking. For the former, we decided to implement Microsoft Kinect cameras, which operate using infrared light and can track movement in space by calculating 15 points of a person’s skeleton in a frequency of about 30 Hz. Each camera covers a frustum of 58 \times 45 in a depth of roughly 6 m and can track up to six skeletons without the need of tracker markers for full body interaction.²²

Regarding gesture tracking, we initially tested two commercial devices: the Leap Motion Controller,²³ which is optical, and Thalmic Labs Myo,²⁴ which tracks gestures from the arm muscle’s electromagnetic behaviour – both of which we rejected for our purpose, the first due to its cable dependencies, range and orientation limitations, and the second because of limited gesture pool and tracking imprecision. Without any other commercial solutions left, we proceeded by prototyping a pair of interactive data gloves. The ‘inteliglove’ system much like other similar devices²⁵ uses flex sensors for capturing finger movements and a 9-degrees-of-freedom IMU for calculating orientation. It was programmed using Arduino micro-controllers, and the IMU sensor fusion was based on the ‘Razor 9DOF AHRS’ open-source algorithm.²⁶



Figure 1. Omni mask prototype.

The circuit also features an XBee radio module and is powered by a Lithium Polymer battery which renders the device wireless and autonomous for more than 6 h. The micro-controller allowed for a calculation frequency of 50 Hz, which was handled adequately by the communication protocol on the condition that each glove had a dedicated receiver on the other end; the sending rate was very high for one receiver to collect data from more than one transmitters (Figure 2). Although our data glove prototype performed adequately, it is worth mentioning that it has significant disadvantages mainly due to the nature of the flex sensors, such as the flex-sensor fragility, the difficulty to adapt and fit to various hand sizes and the tracking limitations of the multi-axis finger movements. Additionally, contemporary mobile egocentric optical gesture tracking technology is for this application a more promising approach.²⁷

Tracking software

Regarding the sensor data, we developed ‘Omnitracker’ – a Java desktop software using the Processing and SimpleOpenNI libraries. Its main functionality is to capture and process data from the Kinect cameras and the intelliglove devices. It allows for a multiple Kinect camera setup, each of which can capture up to six skeleton sets at a time. Eventually, a complete tracking dataset for a whole body includes a 15-vector set for a skeleton’s position and for each hand, a pair of five finger flex values and a palm orientation vector, adding up eventually to 73 degrees of freedom.

General setup and node roles

The project is described as an app system; since, though, it mainly consists of a mobile app, it also requires, as described before, other hardware and software, controlled by one or multiple computers. Therefore, its implementation depends on two nodes: the user or subject node (SN) on a mobile device and the data server node (DN), on a computer which handles the tracking data. Both are needed to run simultaneously and exchange data in real time. The way we proceeded to achieve this was to develop a single application with an embedded networking solution that can run both on mobile and desktop devices, having, though, distinct roles for each node type. While both nodes participate in the network by exchanging data, the SN would be an active agent, whereas the DN is just a passive element of the VS.

Application development

The project’s main application was developed with the game engine Unity 3D, which was selected because it provides a framework capable of real-time graphics processing unit (GPU) rendering, compatibility with




Figure 2. Inteliglove, data glove prototype.

various VR devices and an array of plug-ins (‘Assets’) that can extend its functionality. Additionally, besides graphics and geometry, it facilitates designing a complete environment including interaction and audio. Finally, utilising the Xamarin Mono framework, Unity 3D provides C Sharp scripting capabilities, along with deployment across all major desktop and mobile platforms.

Networking and multiplayer functionality

The two networking issues we faced were, first, the effective sending or receiving of sensor data from the DN to the SN and, second, allowing the system to host multiple concurrent users in the same VS. These issues required a solution for transferring data both on request and on regular intervals. We chose to use the Exit Games Photon PUN²⁸ cloud service, which was specifically designed for real-time multiplayer games and could perform the two requirements by ‘Remote Procedure Calls’ and ‘Serialisation’ update cycles, respectively, between all concurrent users of the application.

The first data cycle regarding sensor datasets starts by DN capturing and processing the sensor data through the Omnitracker software. The data are then forwarded via the Open Sound Control (OSC) protocol to the Unity 3D application that is running in parallel and from there are uploaded to the cloud. Instantaneously, the SN related to the specific dataset downloads and applies it to update its state. This process is asynchronous because while the Omnitracker can secure a static frame rate, its frequency is not necessarily matching the frame rate which the application is running (Figure 3).

The second data cycle concerns the horizontal updates between concurrent nodes running the application. In this scheme, while each device running the application controls locally its own avatar, it also hosts passive instances of the avatars controlled by remote nodes. Each local instance, besides performing the data cycle described before, is responsible for uploading its state changes to the cloud, at regular intervals, so as to inform its instances hosted on remote nodes (Figure 4).

Additionally, actions performed by any SN which affect the shared VS are asynchronously sent through the cloud to all nodes to apply to their version of the VS.

Digital avatar and virtual kinaesthetics

Of topmost importance for creating the immersive experience we set for was to achieve virtual kinaesthetics – proprioception, to create a virtual body that would follow the user’s real body and to contribute to enhance the sense of presence and self-awareness in the VS.²⁹ We therefore used the sensor dataset downloaded from the cloud, responsible for the user’s body and hands, along with a rotation vector produced




Figure 3. Data stream from data node to subject node.




Figure 4. Horizontal data synchronisation between concurrent users.

by the smartphone's IMU, responsible for tracking head orientation, fused together to produce a multidimensional vector describing the user's state in space. To formalise the user's presence, we therefore implemented as a digital avatar a humanoid 3D model capable of skeletal animation to which we applied the aforementioned quantisation vector in order to create real-time animation, which proved sufficient for this purpose. Eventually, the users' body, hands and head rotations and movements in real space were aligned and correlated with the ones of the digital avatar in a 1:1 relationship (Figure 5).

It is important to note that since the virtual body of the user is entangled to the real, and therefore, the virtual viewpoint moves with the user's body, simulator sickness, which describes the illusory perception of self-motion and is one of the main causes of disorientation in VR applications, in our case was minimal.³⁰

Additionally, the digital avatar featured physics behaviour, meaning it could physically interact – collide – when in contact with other virtual objects or avatars of the VS.

Active presence

Having solved the previous, we wanted to add some actions in order to display and test the interactive nature of the project and, therefore, prove the active presence of the user in the VS. As a proof-of-concept scenario, we programmed hand gestures to perform actions on geometrical objects. These were object creation, scaling, colouring, rotation and movement, which were assigned to unique gestures. We could therefore show




Figure 5. Users can see their virtual hands animating in real time.




Figure 6. User creating and scaling a box using gestures.

that individual users can control and affect their shared VS, given the actions are synchronised over the network (Figure 6).

Virtual environment

As a virtual environment, the appearance of the VS overlapping the real, the game engine we used to develop the application allowed the use of 3D digital geometry. We were therefore able to import 3D CAD models with textures which successfully dressed the VS and were rendered in real time by the application. With the only constrain of using the available space of the installation which was effectively tracked by the motion capture sensors, we could overlay any context, similar or dissimilar to the real. Not needing to play any functional role, since it is immaterial, this environment simulating a built object or even landscape does not additionally need to either follow the way we design to build what our real environment – urban or natural – looks like or its physical laws (Sutherland, 1965). Furthermore, surpassing the limitations of what we conceive as built environment, the virtual environment can also be programmed to animate and alter in time.

The characteristics of a new spatial medium

Putting everything together, the project creates a first of its kind experience in a novel spatial form. To justify our claim, we intend to offer a list of characteristics or properties that define the singularity of this type of

VS. To render, however, these properties easier to grasp, we will provide an example of a potential application of the VR setup described through the project in the form of a virtual museum.

The general concept for this application supposes the hosting of a virtual exhibition in a physical building. The setup requires for the space to be covered by a layer of motion capture sensors, which are responsible for quantising the visitors' body position and movements. Furthermore, each visitor is handed a mobile HMD – smartphone with head mount – connected to the Internet and running the application. A data centre controlling the motion capture sensors, scans the space and sends each skeletal dataset to the device of the visitor it corresponds for it to update their position and avatar state in real-time. The VS where the exhibits are hosted is overlapping the real. The visitors, consequently, wandering in a seemingly empty space, just by wearing their HMD, are exposed to VS overlapping and enhancing the real, by creating a mixed spatial experience full of exhibits. Eventually, the museum emerges from the augmentation of the real space where the exhibition is hosted and the exhibition content which is found in a digital form in the smartphone application. The characteristics of this augmentation are as follows:

- The exhibition space itself does not necessarily need to follow spatial requirements typical for museums. Since only the floor plan is of relevance – as the effective space that can be used as the ground of the VS overlay – it can as well be implemented in an elaborate museum space, as in an underground parking garage with a ceiling at 2.5 m.
- The content of an exhibition, in this case, is not limited to a predefined form or collection. It can host any type of medium of digital 'reproduction',³¹ which can furthermore occupy the four dimensions: image, audio, video, 3D geometry, 3D animation. For example, one can experience and walk inside a virtual Parthenon while the ancient Greeks are having a religious ceremony, regardless of their physical location.
- Since the virtual environment does not require a Virtruvian 'firmitas', which is borrowed from the physical space on which it overlays, it also does not need to resemble physical or built space. A physical building resemblance is only required as a way-finding and cognitive-semiotic³² element. The VS can be designed to resemble a physical structure, but as well be abstract, non-volumetric, as it is in fact a spatial visualisation of data – geometric or not.³³
- Another emergent characteristic of this space is parallel heterotopy.³⁴ Since each visitor is experiencing the VS from their individual viewport – HMD device – each VS can be different. Each visitor can concurrently experience different contents and different surroundings. While one has the Parthenon in front of them, another can have the Great Pyramid of Giza and yet another can be in the sculpture section of a virtual Louvre, and all of them can interact with each other. The museum can overlay multiple exhibition floors in one.
- Subsequently and in relation to the dimension of time, the space is heterochronic. Since each user's experience is generated locally, each one can have a different localised dimension of time. For example, exhibits with a timeline such as audio, 3D animation and film allow for them to be in a different state for each visitor: the film screening begins at the moment when the visitor enters the room for each one individually. Therefore, the museum experience can be furthermore personalised and function in multiple local dimensions of time concurrently.
- Since the visitors' motion is synchronised over the Internet, the data representing their presence are non-contextual. Overlaid heterotopy is the inherent property, in which multiple virtual museums of this kind can overlay their visitors in the same VS, regardless of where they are located. A visitor in a Japan museum can be found in the same VS next to a visitor in a London one. The museum is a multiplicity.³⁵
- Eventually, bringing people together in a shared spatial experience, who can originate either from the same physical space or from vastly remote geographical locations, opens another great potential for the social aspect of VS. With an added functionality of voice messaging – like in computer games – a virtual environment can expand to spatial social network.

Further research and application in architecture education

Since 2014, the project setup has been optimised and improved, although the general configuration is still the same. We are still using contemporary smartphone devices such as HMDs as they are easier to deploy while offering a comparatively good-quality VR experience, although requiring more restricted and careful graphics optimisations in comparison with their counterparts. The motion tracking was upgraded with the more precise Kinect v2 sensors (that offer tracking of 25 skeleton joints as opposed to 15 offered by the previous model). At this time, the setup has proven to be stable and very easy to deploy through various implementations we did for exhibitions and educational workshops.

More recently, to further explore the potential of the fusion of VR and architectural design, we offered a Master Design Studio titled 'Virtual Spaces', which at the time of writing was under publication. With the goal of creating interactive spatio-temporal narratives, 13 students engaged in developing individual VR experiences. In the educational process, besides teaching the technical aspects of working in such a setup, we emphasised on building a glossary and sensibility for interactive and narrative space, which proved very valuable since the notion of time is a significant part missing and usually underdeveloped in classical architectural education. For that purpose, we drew from a wide range of media and sources, such as media theory, architecture theory, cinema, video games, novels and also ludology.³⁶ Additionally, to further develop skills of communicating space, students were assigned to experiment on conveying their spatial concepts by alternative media such as text and video. Eventually, the 13 students yielded a wide range of projects that were publicly exhibited on two occasions. The two most prevalent categories were 'space-cognitive' and 'abstract-narrative' focussed projects. The former category was dedicated on 'games' that require a high level of attentiveness on space. Often utilising simple tricks that alter the behaviour of the VS like teleportation, changing the direction of gravity or the orientation of the user to the world, these applications engaged their subjects on a creative spatial narrative or game by questioning their conception of space-time. The second category of 'abstract-narrative' applications was concerned by storytelling that was articulated through the spatial experience itself. By curating the nature of space (size, aesthetics, events and sounds) as well the nature of the participant and their correlation to the VS time, these examples focussed on creating an affective experience.

Overall, the studio managed to engage students and cultivate their approach to designing space through a holistic approach on space and time while further supporting that architecture as a discipline has the latent sensibilities to come up with and develop cognitively and emotionally engaging experiences by designing space itself in a VR context.

Conclusions

VR is already transforming architectural visualisation by providing means to preview and evaluate designs before they are materialised. However, there is a much greater potential for architecture, we want to argue, in a speculative use of these technologies for designing virtual environments. The experience of VS in a full body immersion is one that does not easily accept classification under the deterministic dipole of real and unreal. While not material, VS is definitely sensible and affective to the extent that it requires for itself a new spatial paradigm. As we have demonstrated, the characteristics and circumstances that collective, fully immersive, AR space allows open new unimaginable horizons and can serve as a platform for developing and exploring novel, previously inconceivable spatio-temporal configurations. Furthermore, in these formative years of the adoption of VR, we maintain that architecture as a discipline can and should be in the forefront of the development of VS.

Acknowledgement

The author wants to thank the *Hutchison Drei Austria GmbH* for their support with mobile VR hardware for the Virtual Spaces Master Studio, of the Institute of Architecture and Media, TU Graz taught in the summer semester of 2016.

Declaration of conflicting interests

The author(s) declared no potential conflicts of interest with respect to the research, authorship and/or publication of this article.

Funding

The author(s) received no financial support for the research, authorship and/or publication of this article.

References

1. Fortune. Is 2016 the year of virtual reality? <http://fortune.com/2015/12/04/2016-the-year-of-virtual-reality/>, 2015.
2. BBC. 2016: the year when VR goes from virtual to reality, <http://www.bbc.com/news/technology-35205783>, 2016.
3. Kickstarter. Oculus Rift: step into the game, <https://www.kickstarter.com/projects/1523379957/oculus-rift-step-into-the-game>, 2012.
4. *The Guardian*. Facebook closes its \$2bn Oculus Rift acquisition. What next? <https://www.theguardian.com/technology/2014/jul/22/facebook-oculus-rift-acquisition-virtual-reality>, 2014.
5. Oculus. Oculus Rift, <https://www.oculus.com/en-us/rift/>
6. HTC. HTC Vive, <http://www.htcvive.com/eu/product/>
7. Sony-Playstation. Playstation VR, <https://www.playstation.com/en-us/explore/playstation-vr/>
8. Brooks FP. What's real about virtual reality? *IEEE Comput Graph* 1999; 19: 16–27.
9. Moss JD and Muth ER. Characteristics of head-mounted displays and their effects on simulator sickness. *Hum Factors* 2011; 53: 308–319.
10. Cruz-Neira C, Sandin DJ, DeFanti TA, et al. The CAVE: audio visual experience automatic virtual environment. *Commun ACM* 1992; 35: 64–73.
11. Kozlov MD and Johansen MK. Real behavior in virtual environments: psychology experiments in a simple virtual-reality paradigm using video games. *Cyberpsychology, Behav Soc Netw* 2010; 13: 711–714.
12. Heyselaar E, Hagoort P and Segaert K. In dialogue with an avatar, language behavior is identical to dialogue with a human partner. *Behav Res Methods*. Epub ahead of print 16 December 2015. DOI: 10.3758/s13428-015-0688-7.
13. Sutherland IE. The ultimate display. In: *Proceedings of the IFIPS Congress*. New York: IFIP, 1965, pp. 506–508.
14. Sutherland IE. A head-mounted three dimensional display. In: *Proceedings of the December 9–11, 1968, fall joint computer conference, part I* (ed Blue R and Rosenberg A), San Francisco, CA, 9–11 December 1968, pp. 757–764. New York: ACM.
15. Sutherland IE. *SKETCHPAD-a man-machine graphical interface*. PhD thesis, MIT, Cambridge, MA, 1963.
16. Edgerton SY. *The mirror, the window, and the telescope: how Renaissance linear perspective changed our vision of the universe*. New York: Cornell University Press, 2009.
17. Devaki M and Mondal J. *Construction Tech Report*. Tracxn Technologies, 2016. <http://blog.tracxn.com/2016/02/10/tracxn-report-construction-tech/> (web report)
18. Milgram P, Takemura H, Utsumi A, et al. Augmented reality: a class of displays on the reality-virtuality continuum. In: *Photonics for industrial applications. International Society for Optics and Photonics*, 21 December 1995, pp. 282–292.
19. Google. Google cardboard, <https://vr.google.com/cardboard/get-cardboard/>, 2014.
20. Durovis. Durovis SDK, <https://www.durovis.com/en/sdk.html>, 2014.
21. Durovis. OpenDive, <https://www.durovis.com/opendive.html>
22. Lange B, Rizzo A, Chang C-Y, et al. Markerless full body tracking: depth-sensing technology within virtual environments. In: *Interservice/industry training, simulation, and education conference (IITSEC)*, Orlando, FL, 28 November 2011–1 December 2011. National Training and Simulation Association.
23. Leap Motion. Leap motion controller, <https://www.leapmotion.com/>, 2014.
24. Thalmic Labs. Myo, <https://www.myo.com/>
25. Benali-Khoudja M, Hafez M, Alexandre J-M, et al. Tactile interfaces: a state-of-the-art survey. In: *ISR 35th symposium on robotics*, 23–26 March 2004, pp. 23–26. New York: ACM.
26. Github/ptrbrtz. Razor AHRS v1.4.2, <https://github.com/ptrbrtz/razor-9dof-ahrs>, 2014.

27. Liang H, Yuan J, Thalmann D, et al. AR in hand: egocentric palm pose tracking and gesture recognition for augmented reality applications. In: *Proceedings of the 23rd ACM international conference on Multimedia*, Brisbane, QLD, Australia, 26–30 October 2015, pp. 743–744. New York: ACM.
28. Games E and Photon PUN. <https://www.photonengine.com/en/PUN>, 2016.
29. Slater M and Usoh M. Body centred interaction in immersive virtual environments. *Artificial Life Virtual Real* 1994; 1: 125–148.
30. Hettinger LJ, Berbaum KS, Kennedy RS, et al. Vection and simulator sickness. *Military Psychology* 1990; 2: 171–181.
31. Benjamin W. *The work of art in the age of mechanical reproduction*. London: Penguin Books, 2008.
32. Eco U. Function and sign: the semiotics of architecture. In: Broadbent G, Bunt R and Jencks C (eds) *Signs, symbols and architecture*. New York: John Wiley & Sons, 1980, pp. 11–69.
33. Novak M. Liquid architectures in cyberspace. In: Benedickt M (ed.) *Cyberspace: first steps*. Cambridge, MA: The MIT Press, 1991, pp. 225–254.
34. Foucault M and Miskowiec J. Of other spaces. *Diacritics* 1986; 16: 22–27.
35. Deleuze G and Guattari F. Introduction: rhizome. In: *A thousand plateaus: capitalism and schizophrenia*. New York: Columbia University Press, 1988, pp. 3–25.
36. Aarseth E. Computer game studies, year one. *Game Stud* 2001; 1: 1–15.