Universitat Oberta de Catalunya

Máster Universitario Aplicaciones Multimedia

M1.925 - Área TFM de Investigación aula 1

Meta-review of text input approaches within VR - A study on the platform's viability as a productivity workspace -

Carlos Rodríguez Outerelo

Index of Contents

Index of Contents 2
Glossary of relevant acronyms 1
Abstract 1
[Spanish] Resumen documental 2
1. Introduction 3
1.1. Context – An overview of VR evolution 3
1.2. Today's VR, tomorrow's platform 7
2. State of Art 10
2.1. VR beyond entertainment 10
2.1.1. The neuroscience behind VR 11
2.1.2. VR as a workflow driver – VR interactivity 13
2.2. VR typewriting 16
2.2.1. Real physical keyboards within a virtual environment 16
2.2.2. Raycast-based virtual keyboards 20
2.2.3. Touch-based virtual keyboards 22
2.2.4. Hand tracking based virtual keyboards 23
2.2.5. Other novel approaches 24
3. Objectives 25
3.1. Summarizing the meta-review 25
3.2. Reinforcing the meta-review 27
4. Development of the prototype 28
4.1. Overview 28
4.2. Issues 31
5. Evaluation 33
5.1. The test 33
5.2. Results 35
6. Conclusions and future work 41
Bibliography 43

Glossary of relevant acronyms

- **6-DoF**: Six Degrees of Freedom
- **AR**: Augmented Reality
- **cpm**: characters per minute
- **FOV**: Field of View
- **HMD**: Head Mounted Display
- **MR**: Mixed Reality
- **VR**: Virtual Reality
- **VRE**: Virtual Reality Exposure
- **VRT**: Virtual Reality Treatment or Virtual Reality Therapy
- **WMR**: Windows Mixed Reality
- **wpm**: words per minute
- **XR**: Extended Reality

Abstract

VR technology has evolved significatively during the past decades both in terms of accessibility and technological feats, and this project reflects on its growth to then study the platform's viability as a productivity workspace. For that, the neuroscience behind VR is overviewed and VR interactivity extensively explored. Said interactivity is yet evolving and thereby so are the many text input approaches within VR. These are detailed throughout this document in a meta-review where their performance and usability are studied. Scientific literature suggests that the integration of real physical keyboards within virtual environments performs above virtual keyboards, with Logitech's approach being the most interesting take due to its accessibility and superior performance. However, strictly virtual keyboards fit better among other use case scenarios and, since there is not an alternative regarded as the defining standard among those covered within the meta-review, a prototype was developed to further evaluate their usability; the development of said prototype is covered within the pages of this document, and many were the text input approaches within its scope.

A study was performed with the prototype at issue, proving touch-based virtual keyboards to surpass raycast-based ones in typing speed and user experience. Such study also highlights which touch-based approaches performed best whilst also evaluating virtual size ownership sensations and the performance implications of the keyboards' size, suggesting that there is an ideal (intermediate) keyboard size for each approach and that, as opposed to what is implied in some of the works found in scientific literature, minimalistic representations of one's hand do not improve performance: given the mismatch between the visual and haptic experiences, these approaches negatively affect virtual hand ownership thus decreasing one's immersion and typing speed. The hand tracking support of the Oculus Quest HMD is also studied, proving its validity as text input approach despite not being robust enough to outperform controller-based techniques. All in all, it can be said that the goal of this project goes beyond the meta-review itself and aims to build upon it through the aforementioned study.

[Spanish] Resumen documental

La tecnología VR ha evolucionado significativamente durante las últimas décadas tanto en términos de accesibilidad como de hitos tecnológicos, y este proyecto se fundamenta en ese crecimiento en pos de estudiar la plataforma y su viabilidad como espacio de trabajo. Para ello, se supervisan y exploran la neurociencia e interactividad tras la VR. Esta interactividad está en constante evolución, y por lo tanto también lo están los muchos acercamientos a la tipografía en VR. Estos son detallados a lo largo del documento en una meta-revisión donde se estudia su rendimiento y usabilidad. La literatura científica sugiere que la integración de teclados reales en entornos virtuales supera en rendimiento a los teclados virtuales, siendo el acercamiento de Logitech el más interesante por su accesibilidad y superior rendimiento. Sin embargo, aquellos teclados estrictamente virtuales encajan mejor en determinadas circunstancias y, puesto que no hay una alternativa considerada como el estándar definitorio entre aquellas cubiertas por la meta-revisión, se ha desarrollado un prototipo para evaluar su usabilidad con mayor profundidad; el desarrollo de dicho prototipo es detallado en las páginas de este documento, y muchos son los acercamientos tipográficos bajo su alcance.

Se ha realizado pues un estudio con el prototipo en cuestión donde se demuestra que los teclados virtuales basados en el contacto superan a aquellos basados en un puntero láser (raycast) tanto en rendimiento como en experiencia de usuario, destacando qué acercamientos dentro del primer grupo funcionan/rinden mejor. También se evalúan la sensación de propiedad para con las manos y las implicaciones del tamaño del teclado, demostrando que existe un tamaño ideal (intermedio) para cada acercamiento y que, pese a ciertos alegatos encontrados en la literatura científica, las representaciones minimalistas de las manos no mejoran el rendimiento sino que, dada la discordancia entre las experiencias hápticas y visuales, este tipo de acercamientos afecta negativamente a la sensación de propiedad para con las manos disminuyendo la inmersión y velocidad de escritura. El seguimiento de manos del casco de realidad virtual Oculus Quest también ha sido estudiado, demostrando su validez como acercamiento tipográfico pese a no ser lo suficientemente robusto/estable como para superar aquellas técnicas fundamentadas en el uso de controladores. En resumidas cuentas, se puede decir que el objetivo de este proyecto va más allá de la meta-revisión, buscando construir sobre la misma a través del estudio en cuestión.

1. Introduction

1.1. Context – An overview of VR evolution

The ups and downs of virtual reality (VR) are a tale of its own.

While the first head-mounted displays (HMD) were developed in the 1960s, it was in the 1980s that this technology picked some momentum: VR was, at the time, being used as a flight simulator tool to train pilots; its interest was such that even NASA tried to build VR HMDs of their own. However, this momentum shifted during the 1990s as videogame companies tried to bring this technology to a broader audience and, in order to do so, two approaches were taken: they had to either cut costs down so that the average consumer could afford such items or they could try to manufacture a compelling product despite its high price. While concept-wise both approaches seem valid, given the technological limitations of the time none of them were successful: the released products were either too expensive or offered a subpar experience.

The perception of the masses shifted: VR was no longer a glimpse into the future, but a mere gimmick instead. Due to that, the early 2000s were stagnant in this regard.

Everything changed with the collaborative work of Palmer Luckey and John Carmack. Together, they founded Oculus VR and announced the most ambitious VR headset up to this point in time: the Oculus Rift. Through a tethered connection to their computers, the Oculus Rift userbase would be able to render any 3D visuals and environments inside their HMD and, with Oculus' promise of an incredibly wide field of view (FOV), a high-resolution display, extremely low-latency head tracking capabilities and a consumer-oriented price… Its success was just a matter of time.

The 1st of August of 2012 the Oculus Rift Kickstarter campaign was launched with an outstanding public reception: while the campaign goal was of $250,000, they raised $670,000 within the first 24 hours and after just three days, the campaign surpassed the $1,000,000 mark. By the time the Kickstarter ended, the amount crowdfunded had reached almost $2,500,000 (about 10 times their target goal).

After the success of the Rift's crowdfunding campaign, the number of companies involved in VR, the technology's popularity, the investments, the research… The VR market was revitalized as a whole, and the interest once lost was brought back by the Oculus team. At the time, Valve (the videogame company behind the Steam platform and the Half-Life videogame series) was also working on their own VR setups, with their main focus being the tracking technology that would allow HMDs to be rightfully positioned in three-dimensional environments (by the use of infrared lights and external sensors). That is why, instead of competing against each other, Oculus and Valve collaborated with one another (neither the date on which they began to work together nor the extent of such collaboration is certain).

Two development kits were launched for the Oculus Rift prior to its final consumer-oriented release in 2016: the first one, named Oculus Rift DK1, was launched in 2013 while the second one, the Oculus Rift DK2, was released a year later in 2014.

Both presented merits of their own, but only the latter saw the fruition of Valve's aid: by the use of an included external sensor meant to track the headset's position thanks to a few infrared LEDs built into the HMD itself, the users could move freely within a three-dimensional space. This made of the Oculus Rift DK2 the first major VR HMD to allow for 6 degrees of freedom (6-DoF). The term *room-scale VR* came into being, and the VR experience was elevated as a whole.

Figure 1. Graphical representation and effects/consequences of 6-DoF in VR
[Source: https://telecoms.com/484062/qualcomm-tries-to-muscle-in-on-vr-euphoria]

As the first evolutionary steps towards today's VR, these early development kits were not exempt of issues. While they presented what is known as *screen door effect*, a visual artefact of displays where the space between pixels becomes visible to the eye, the most concerning topic was the "motion sickness" some users were prone to. While the actual cause of cybersickness (as commonly referred to in scientific literature) is unknown, there are various theories regarding the issue. The most popular one is known as the *sensory conflict theory*, which claims this ill-feeling comes from the slight mismatch between the visual and vestibular systems; see Davis, Nesbitt and Nalivaiko (2014).
At the time this was, undoubtedly, VR's biggest drawback (and it arguably still is).

In between the Oculus Rift DK2 pre-orders and its release, Oculus was bought by Facebook (the social media giant) for $2 billion. The reception of this acquisition was controversial and, perhaps as a consequence, Oculus and Valve parted ways (the latter started collaborating with HTC instead). Besides that, 2014 was a bumper year for VR, also seeing the announcement of Sony's own VR headset for its PlayStation 4 console (codenamed Project Morpheus at the time), the launch of the Google Cardboard (the first major mobile VR experience) and the release of Leap Motion's first VR tracking module, providing hand tracking support within VR applications. While Leap Motion's offering did not become mainstream by any means, its undeniable potential captivated VR enthusiasts and researchers. It was a powerful tool, and it illustrated the potential of hand tracking support within VR environments; see Wozniak, Vauderwange, Mandal, Javahiraly and Curticapean (2016) along with Argelaguet, Hoyet, Trico and Lécuyer (2016).

After a relatively quiet 2015, 2016 saw the release of three major HMDs (all of which presented 6-DoF): the first consumer-aimed iteration of the Oculus Rift (referred to as the Oculus Rift CV1), the HTC Vive by HTC and Valve, and Sony's PlayStation VR (PSVR). The latter had a quiet release, but its relatively low price and its ever-growing catalogue of software ended up making it the best-selling VR headset of all time (at the time of writing) proving that both price and software are key selling points of VR. On the PC-based side of VR, the HTC Vive's success surpassed the Rift's and exceeded everyone's expectations: while there were not many differences hardware-wise, the Vive came bundled with a pair of wireless motion tracked controllers (one for each hand) that offered a much more immersive experience by featuring 6-DoF themselves, so their position could be tracked and reproduced inside virtual environments (meaning that the users' hands could be emulated within VR). These controllers were a leap forward in VR/XR interactivity and, as a consequence, the market share initially favored HTC and Valve's product. Oculus was known to be developing a similar kind of controllers which they finally released by the end of 2016: the Oculus Touch controllers improved upon Vive's offering by incorporating capacitive sensors within their buttons that detected whether or not some of the user's fingers were touching the controllers. This allowed for more complex and advanced hand emulation, improving virtual interactivity and greatly closing the initial gap between the Vive and the Rift; by the end of 2018 the Rift's market share ended up surpassing that of the Vive, mainly because of both its lower price point and Oculus' exclusive catalogue of games and software.

Microsoft HoloLens was also released in 2016, being the first major AR (augmented reality) headset, and along with it the Windows Mixed Reality (WMR) platform was launched, which was aimed towards HMDs of any XR (extended reality) nature (which encompasses VR, AR and MR). The HoloLens was not a consumer-oriented product, but it was an interesting one from a scientific standpoint as it brought costs down whilst also being state-of-the-art AR technology; see Coppens (2017) as well as Kalantari and Rauschnabel (2018). Based on the WMR platform, many companies such as Lenovo, Acer, Dell, HP, Samsung and Asus launched VR headsets of their own and, although those offerings where not as successful as neither the Vive nor the Rift, these HMDs were interesting from a technical standpoint as some of them used built-in cameras to track their position within three-dimensional spaces, thus allowing for 6-DoF without the need of any external sensors. This approach was denominated inside-out tracking, and while WMR headsets were not the first HMDs to make use of it, they contributed greatly to its popularization. One of the major complaints of VR at this point was its setup complexity: the amount of cables needed for both interconnectivity and power delivery, the sensors' installation and how their position constrained the space on which users could move… While the accuracy of inside-out tracking was inconsistent and its latency was higher than the one seen with outside-in approaches, it was an interesting take worth studying as it allowed the users to bypass the aforementioned inconveniences.

By the end of 2017 Oculus released the Oculus Go, an standalone VR experience that, unlike the Google Cardboard, did not need a smartphone to operate as, by having a built-in system-on-chip (SoC), no additional hardware was required to render VR environments and experiences.

However, this HMD presented a few downsides worth noting: not only was the headset limited to 3-DoF but so was the controller it came bundled with, which felt functionally lackluster when compared to the Oculus Touch. Furthermore, the integrated SoC (Qualcomm's Snapdragon 821) was not powerful enough to compete against tethered VR, unable to run demandant applications. Despite all of its flaws, the Go was (and is) a commercial success due to its price point and the support it received from both Oculus and developers.

The HTC Vive Pro and the HTC Wireless Adapter were released in 2018 (the latter being an accessory for both the original Vive and its Pro counterpart), although neither were adequately priced and thus their success was moderate. However, it is worth noting that HTC's wireless module and the Oculus Go itself meant that both companies had acknowledged the complexity of VR setups and were trying to offer ways to mitigate this issue. By the end of the year, this issue was no more than a choice: the Oculus Quest was announced, and with it came the very first standalone VR experience that could compete toe-to-toe with tethered VR.

The Quest was an all-in-one VR headset that aimed to fixed and/or improve the flaws and limitations of its predecessor, the Oculus Go. To achieve so, it offered 6-DoF by using inside-out tracking (Oculus' first attempt at it), a high-end SoC capable of running more hardware-dependent applications (Qualcomm's Snapdragon 835) and an across-the-board hardware overhaul. It also came bundled with an updated version of the Oculus Touch controllers and, even though its price was higher that the Oculus Go's, it was in line with the most inexpensive tethered VR headsets both pricewise (at $399) and performance-wise (but without any setup complexity whatsoever); although the most demandant applications could not run on the Quest natively, shortly after its launch an application by the name of Virtual Desktop was released on the Quest store allowing users to stream PC content to the headset through the network, granting the Quest the ability to run both native apps and PC-based ones.

Whereas the inside-out tracking on the Windows Mixed Reality headsets was hit-or-miss, the four cameras built within the Quest achieved an overall better reception, surpassing any previous implementation of this technology (Figure 2 illustrates the positional tracking techniques used among the Steam users who own a VR headset as of April 2020). This motivated its implementation in 2019's VR headsets, namely the Oculus Rift S (Oculus' latest PC-tethered experience) and the HTC Cosmos line of headsets; although the latter was poorly received, both the Quest and the Rift S are hard to find in stock due to their commercial success (more on the stock issues later on this document).

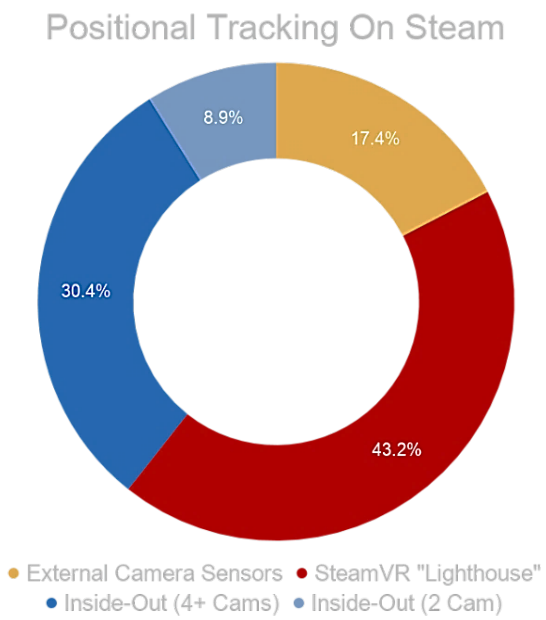

Figure 2. Positional tracking techniques used among Steam users [Source: https://www.superdataresearch.com/2019-year-in-review]

Amongst the HMDs released throughout 2019, the Valve Index (Valve's first VR headset of their own) is worth noting. While Oculus and HTC were trying to appeal the masses with lower prices and inside-out tracking to avoid setup complexity, Valve's approach was to target the most enthusiast VR users with a resolution as high as the one seen in the Vive Pro, the highest refresh rate and FOV up-to-date, external sensors to track more accurately the headset and its controllers… At the time of writing, the Index still represents state-of-the-art technology and so do its controllers, which further refined emulated hand controls by tracking each individual finger through additional capacitive sensors (an evolution on what was seen in the Oculus Touch).
On the other hand of the spectrum, both HTC and Oculus integrated (via software updates) hand tracking support within their headsets by using the integrated cameras built within them. As with Leap Motion's technology, it represented a leap forward in XR interactivity and, as such, upcoming HMDs are likely to implement the feature. Although it was an outstanding feat for both of them, Oculus' implementation of the feature within the Quest was the most successful take, practically eclipsing HTC's one.
The Oculus Quest hand tracking support was not the only major update it received during 2019: the Oculus Link was launched (although still in beta at the time of writing) allowing the HMD to function as a tethered VR headset through the use of a USB3 cable (recent updates allowed USB2 cables to work, which made of the included charging cable a perfectly viable one). While running PC-based applications was already possible through Virtual Desktop, Oculus Link provided lower latency and compression whilst also being a useful tool for developers, allowing them to test scenes without needing to build and sideload the application itself to test changes (there are workarounds to make this work with Virtual Desktop, but the Link made such an option more accessible).

Recent HMDs showcased many avantgarde features that are likely to be implemented in upcoming XR devices; not only has this technology evolved through the years but so have its controls, and with improved interactivity comes an improved experience. Now the question lies in where will said interactivity move towards.

1.2. Today's VR, tomorrow's platform

There is a reason for this whole context to be here, as to understand where the VR market is headed one has to first understand where it comes from. Not only has said market skyrocketed since the launch of the original Rift back in 2016, but sales of standalone VR headsets more than doubled in 2019 (due to the popularity of the Oculus Quest HMD) offsetting declining sales from older types of hardware. The Quest sold over 200,000 units in Q2 2019, illustrating its mainstream appeal; by eliminating the need of any additional hardware and reducing the setup complexity of the VR experience to the press of a button, it surpassed everyone's expectations and accounted for 49% of 2019's VR shipments.
Software revenue also rose by 41% during 2019. Consumers spent a total of $171 million in standalone VR experiences and around $86 million in PC-based ones within the same time frame, adding up to over $250 million in VR software. SuperData's Year In Review of 2019 showcases the aforementioned data and the overall performance of the VR industry proving that, by this point in time, VR's public perception had shifted away from the “gimmick” it was thought out to be after the 1990s flop.

Said data is shown in Figure 3 and Figure 4 which, within a margin of error, illustrate VR sales numbers and the marginal increase of standalone VR experiences.

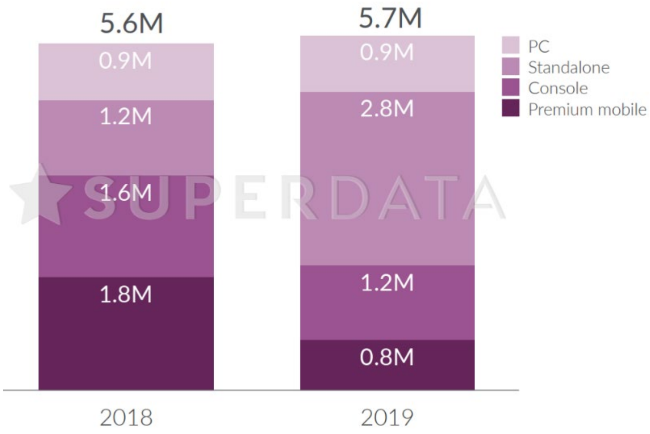

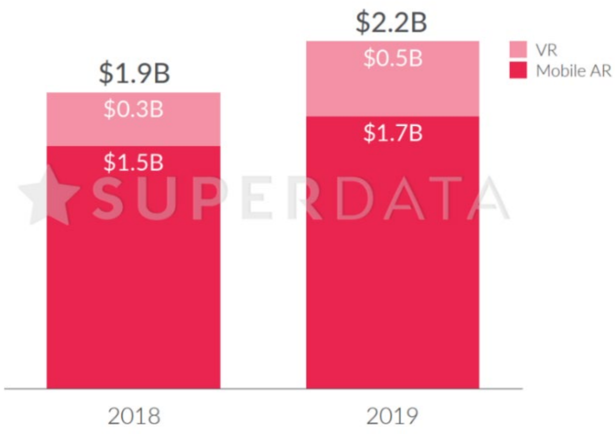

Figure 3 (Left). VR headset shipments (excluding Google Cardboard and the similar products)

Figure 4 (Right). Software revenue in both AR and VR markets [Source: https://www.superdataresearch.com/2019-year-in-review]

In fact, VR's success has been such that supply itself has been an issue during 2019 and even more so within its latest quarter and throughout Q1 2020. Although the lack of supply negatively affects sales numbers, COVID-19's situation and the confinement forced upon many has showcased VR's potential as a social platform and, due to the fewer out-of-home entertainment options, interest upon the platform has undoubtedly increased. As suggested in SuperData's overview of the first quarter of 2020, VR headsets are poised for a rebound as supply chains return to normal since its falling sales are due to a lack of supply, not demand.

On the software side of things, the release of Half-Life: Alyx (arguably the first AAA production for the platform) has been a commercial success. Steam Spy's data estimates around 1-2 million copies sold, which is outstanding for a VR-only release. Valve's game also caused a remarkable increase in the number of Steam users, which went from 1.29% to 1.92% of Steam's userbase (as of April 2020) and, since Steam has surpassed the mark of 90 million monthly active users, this translates in an increase of over half a million VR users for a total of roughly 2 million of them, surpassing the amount of users using a Linux-based operating system and closing in the number of users with a 4K main monitor.

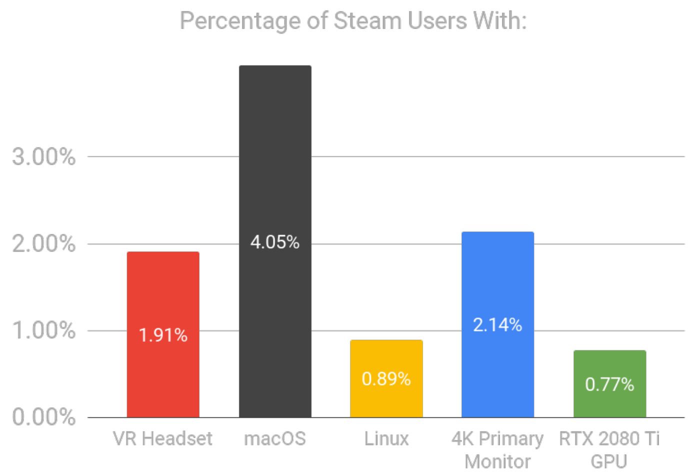

Figure 5. Percentage of Steam users with VR headsets [Source: https://uploadvr.com/steam-hardware-survey-2-million/]

Although COVID-19 has slowed down the sales numbers of the VR industry, the market growth of this technology is in continuous rise and the many possibilities it brings along are attracting developers, researchers and common users all the same. SuperData's overview of the virus outbreak and its effects on the industry lowers their previous 2020's prospect, but the market is expected to keep its steady growth on following years. Widespread adoption was unthinkable a decade ago but, at the moment, VR is on the verge of becoming commodity hardware; perhaps it won't be long before VR is but a platform of choice.

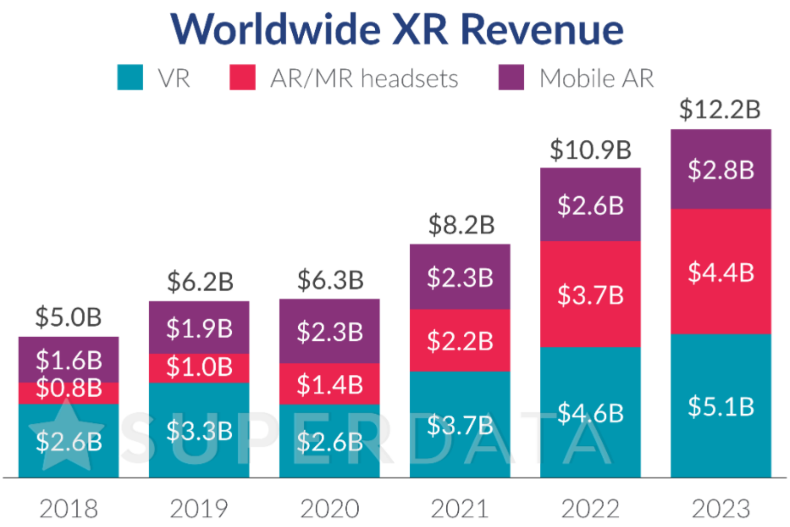

Figure 6. SuperData's past and future estimates for the XR market as a whole [Source: https://www.superdataresearch.com/blog/superdata-xr-update]

There is much upon which to improve today's VR (and XR as a whole). The screen door effect is still appreciable, cybersickness is still an issue and there are many aspects like weight and ergonomics, limited FOV and resolution/aliasing that have yet to be properly addressed. Interactivity within XR environments is yet evolving, but that's not a downside by any means: it is the process which motivates this very own project.

VR evolution is a tale of ups and downs: while there is much yet to be done, the many advances made during the last decade have made of VR a very powerful and versatile technology.

"I think I've seen five or six computer demos in my life that made me think the world was about to change. Apple II, Netscape, Google, iPhone … then Oculus. It was that kind of amazing."

Chris Dixon

2. State of Art

2.1. VR beyond entertainment

Even though the entertainment industries are the focus and drivers of today's VR, the use case scenarios of this technology are as many as imagination itself allows. The report from Grand View Research synthesized in Figure 7 illustrates the VR market share by industry/application (as of 2018) and, while entertainment accounts for over 50% of VR's market, fields like healthcare and education have a considerable presence.

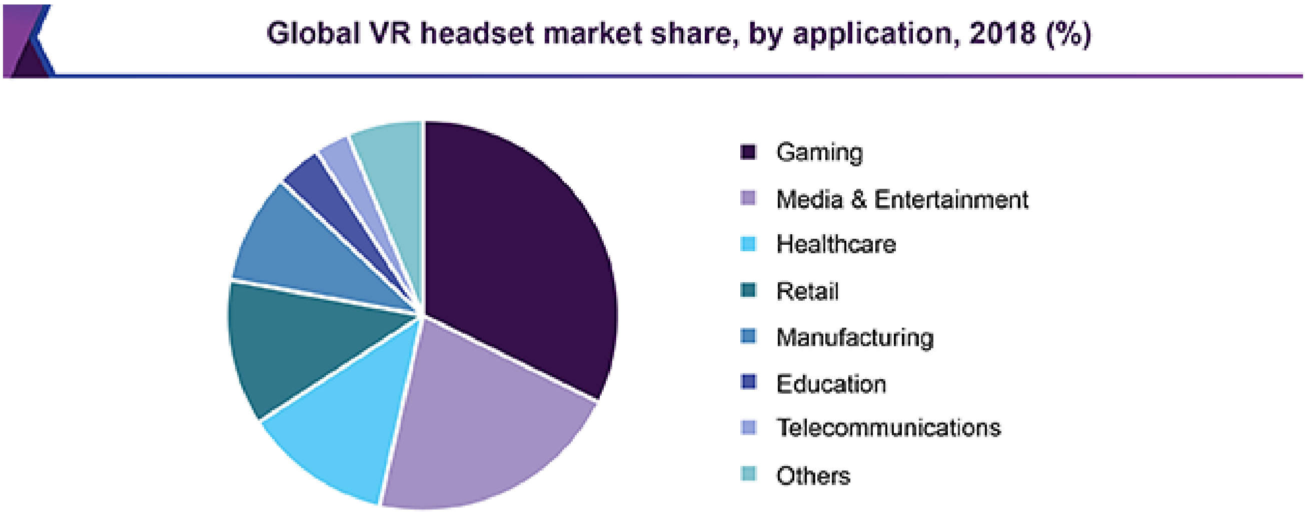

Figure 7. Global VR headset market share by industry (as of 2018)
[Source: https://www.grandviewresearch.com/industry-analysis/virtual-reality-vr-headset-market]

Within academics, researchers all over have been studying VR as a tool in numerous disciplines and many are the authors who have explored and approached this technology with scientific lenses. Nagata, Mikami, Miyashita, Wakayama and Takada (2017)'s study showcases some of the uses of VR in telecommunication services; in their work, they highlight VR technology's potential "as a tool for the three C's (creation, control and communication) and the three E's (elucidation, education and entertainment)". Education itself is one of the most researched topics with regards to VR as it is closely correlated to many other fields: the work of Martín Gutiérrez, Mora, Añorbe Díaz and González Marrero (2017) deepens into the educational possibilities and current trends of VR while Potkonjak et al. (2016) covers the use of VR within the education of science, technology and engineering. These fields are covered with great extent in scientific literature as VR can be used to remotely control and/or monitor related tasks: Berg and Vance (2017) evaluates VR in product design and manufacturing, while industries such as the automotive one have also made extensive use of this technology (see Lawson, Salanitri & Waterfield, 2015). VR has also been used as a training tool in the military industry (Lele, 2013) where its potential as a simulation technology has been exploited since VR's beginnings. This potential is key in behavioral health and in the understanding of the neuroscience behind this technology, which is discussed with great extent in upcoming pages.

2.1.1. The neuroscience behind VR

As shown in Figure 7, among the many uses and applications of VR technology the healthcare industry takes the lead in VR market share (entertainment aside). "Virtual Medicine" is the term some scientist use to tackle the many uses of VR technology within this field, and there are many researchers who have studied and explored with great extent said applications: upon the work of authors such as Székely and Satava (1999), Riva (2003) built an excellent overview which detailed the many medical areas, topics and tasks where this technology is no short of revolutionary. Riva's research has been updated and expanded in scientific literature, where it is worth highlighting the work of Kathuria, Gefen, Boulos, Claudio and Maddalena (2014) along with Li et al. (2017)'s, all of which offered an interesting perspective on VR's place within healthcare. The use case scenarios of this technology include medical imaging and diagnosis, preoperative planning, education and training programs, telemedicine/telesurgery… But among them, it is not farfetched to say that the most notable use is found within the behavioral health field, where VR has been used for over twenty years due to its effectiveness (Mazurek et al., 2019). In fact, virtual reality treatment (VRT) has been proven to surpass traditional cognitive behavioural therapies.
To further comprehend the scope on which VR is currently being applied in behavioural health, I would like to highlight the work done by Riva et al. (2018) and Riva, Wiederhold and Mantovani (2019). The latter conducted a well-documented meta-review evaluating VR applications within the aforementioned field, while both studies dive deeply into the understanding of the neuroscience behind VR to explore not only the technology itself but also what makes of it such a powerful tool in behavioural health and psychiatry as a whole. It is upon their work that I base most of the statements within this chapter.

Virtual reality exposure (VRE) is the most common form of VRT in behavioural health therapy. In traditional cognitive behavioural therapies, patients are progressively exposed to stimuli which trigger symptoms of the disorder, which involves the patients imagining themselves in certain difficult situations. However, this type of exposure is not always effective, mainly because many patients find it difficult to imagine an aversive situation. Therapy making use of the possibilities created by VR proves to be an effective approach, as seen in the work of Mazurek et al. (2019). Consequently, it has been used to treat different anxiety disorders: phobias, PTSD, panic disorder and agoraphobia, social anxiety disorders, psychological stress, eating and weight disorders and other generalized anxiety disorders; VR has also been explored for psychosis and pain management, proving VRE to be a versatile tool within psychotherapy. These many uses and applications are also covered in Riva (2005)'s research, although many of the limitations mentioned within the conclusions of said study are no more (as can be seen in Mazurek et al., 2019).

All in all, VR has been proven to be an effective clinical tool and an advanced imaginal system able to effectively induce experiences and emotions much like reality itself. As for the reason, different major discoveries in the field of neuroscience suggest that VR technology shares with our brain the same basic mechanism which is referred to in scientific literature as embodied simulations.

To regulate and control the body in the world effectively, the brain creates an embodied simulation of it which then uses to represent and predict sensory events: actions, concepts, emotions… VR works in a similar fashion: its hardware components track the motion of the user and transcribes it to his/her avatar (an embodied simulation of sorts), while the underlying software adjusts the imagery to be displayed. This simulation emulates the brain's expected sensory input, and the immersion itself is highly dependent of the correlation between the VR model and the brain one (the higher, the better). Consequently, VR provides a digital place where the user can be placed and live a synthetic but immersive experience. This leads to a feeling of "presence", which is extensively covered in scientific literature; the aforementioned works of Riva (2003) and Kathuria et al. (2014) explore VR's presence while other studies are centred upon it: Schuemie, Van Der Straaten, Krijn and Van der Mast (2001) conducted a survey on the subject and assessed their importance in VR, and said subject's study was later reinforced and researched by Lee (2004), Sánchez Vives and Slater (2005) along with Murray, Fox and Pettifer (2007). This presence is what makes VRE work so well (see Price & Anderson, 2007), and is the key behind the relationship between VR and a productive state of mind know as flow.

The aforementioned notion describes the mental state in which a person performing an activity is fully immersed and involved in the task (common literature refers to such state as "being in the zone"). When in said state, a person not only experiences an increase in his/her focus but also a boost in his/her productivity and performance levels. Flow's study began with the work of Csikszentmihalyi (1975) and since then, it has been researched and extensively explored by scientific literature in order to understand both its causes and its implications. Csikszentmihalyi's work on the subject has been an inspiration to many and while his first books are noteworthy, his later work in the psychology of optimal experience (Csikszentmihalyi, 1990) is arguably the most remarkable of his works. Kotler (2014) deconstructed this foreign yet familiar concept in order further understand its principles and, more importantly, its triggers. While there is no proven simple way to induce oneself into flow, it is commonly agreed that the more of these triggers can be hit, the easier it becomes to enter such a state. The neuroscience behind VR (detailed in the many works mentioned throughout this page) suggests that the user's subconscious is prone to thinking said user is indeed inside the virtual environment, due to this feeling referred to as "presence" in scientific literature (and already covered), VRE seems to be able to activate most (if not all) the triggers detailed in Kotler's studies meaning that, in other words, VR technology could increase the ease with which one enters the flow state. Consequently, using VR as a workspace could trigger flow activation (based on one's suspension of disbelief; see Duffy & Zawieska, 2012 as well as Brown, 2012) and thereby it could increase one's overall performance, potentially making VR headsets better productivity drivers than the ones being used in today's society.

However, besides the many uses of this technology in professional environments and its potential as a productivity-focused environment, most of the VR userbase is not yet using it as a productivity driver. Why so? To answer that question and in order to consider VR a viable media for productivity, one has to first understand and deconstruct the tools and devices used in today's society to drive most workflows.

2.1.2. VR as a workflow driver – VR interactivity

Today's main productivity driver is undoubtedly the personal computer (PC), with smart devices like tablets and smartphones being suitable alternatives. Most consumer PCs these days present either a desktop-based form factor or a "laptop" one, and there are also numerous operating systems (OS) used by people to drive their workflows among which to highlight Windows, MacOS and Linux-based ones. These PCs, regardless of their form factor and OS, have a number of elements that make them suitable tools for productivity workflows: the first one is the screen, a showcase of information which displays any document, program, image, video or electronic file the user desires; the second element involves the many hardware components of the unit (the CPU, the GPU, the memory…) as it regards the computing/processing power required to run whichever application or task the user wants to or needs to; finally, the last element to highlight is the one on which this research is focused: the interactivity itself. While most PCs use the keyboard and mouse to drive this human-computer interaction, many laptop PC users prefer the built-in trackpad of the unit to the mouse as hand-gestures are intuitive and said approach is proven to be almost as efficient as the alternative (General et al., 2015).

As previously stated, computers are not the only productivity drivers in today's society: some people use their smartphone or tablet to tackle their workflows, while others make use of them in combination with additional productivity tools (multi-device workflows are studied within the work of Santosa & Wigdor, 2013). By deconstructing said smart-devices one would find a similar scheme to the one present in PCs: a screen, to showcase the desired information and content, the computing/processing power, which is built inside the device, and the interactivity itself which, in the devices that are now being addressed, is the very same screen used as display (they are, in fact, touchscreen-based devices).

The introduction of this document illustrates how VR has been made more accessible with each passing year. Today, HMDs like the Oculus Quest present an all-in-one solution with a price-point comparable (if not lower) to the ones found in most PCs and smart-devices. Although their accessibility is a relevant matter to the subject, the key question lies in the viability of these HMDs as productivity tools: by deconstructing said headsets, one would see how their barebones are in fact quite similar to ones behind the previously covered devices.

First of all, VR headsets provide a pair of screens (one for each eye) to simulate 3D environments with real depth, and as such offer a superior presentation to those available in traditional productivity devices. However, while depth-based (3D) content exhibits the advantages of VR, the content most people use in their work environments is either 2D-based or can be projected as 2D content. This highlights the importance of traditional interface design even when working with XR technology. In fact, human brains have been trained to read words and text spread-out across 2D surfaces: paper made ones, the blackboards/whiteboards available in most classrooms, electronic displays as those seen in PCs and smart devices… Projected 2D content has been made a really efficient way of consuming information; interfaces built upon this concept (projected 2D content) are referred to and known as screen interfaces.

Screen interfaces' design is an intricate challenge that VR designers and developers have to face when working with this technology, as with the increasing amount of possibilities it opens comes an increasing amount of problematics that have to be properly addressed: the text has to be easily readable and buttons, triggers and whichever interaction tools are incorporated have to be easily understood and interpreted for the users to actually know how to interact with such elements; lists have to be easily scrollable, menus have to be easily understood so that the users can navigate through them as fluidly as possible… The list goes on.
Furthermore: within VR, due to its nature and limitless space, users can display as many screen interfaces of indefinite size as one wants; the works of Czerwinski et al. (2003) and Anderson et al. (2003) suggest that increasing the screen real estate improves one's productivity and VR platforms have no limits in this regard.

Regarding the processing power of the unit: most HMDs base this fundamental requirement on an external unit such as one's PC, while some alternatives like the Oculus Quest use an integrated SoC built inside the HMD so that it can power itself. The interactivity, on the other hand, is the key point to be studied by this own project, as there is not an indisputable approach within VR environments. Whereas in PCs there is a prominent pointing tool (the mouse, seconded by the trackpad), there are many takes on pointing interactivity inside VR; currently, the most popular one is a raycast or laser-based approach, also known as controller pointing. As the name suggests, this technique is based on a straight beam/laser coming out of the front of one's controllers, granting the user full control of the casted ray so that he/she could point it towards an interactable of choice.
Closing the distance between the interactables and the users is also an option: most VR controllers can be emulated within the virtual environment and, by the use of the model's collider, users can actually touch whichever interactable is presented upon them. Due of its intuitiveness, as its behavior resemble reality itself, touch interactivity is the preferred approach in many use case scenarios. The use of emulated haptic feedback helps the users' immersion by mimicking real-world haptics through software-refined vibration from within the controllers. Furthermore, said controllers can also be used to emulate objects other than the controllers themselves: a showcase of this approach is found within the rendering of a pair of mallets which increase the users' reach by using the controllers as an additional joint; such an approach is actually quite common as a typewriting mechanism within VR applications, and is explored later on this document.
Virtual hand emulation is, as explained during the overview of VR evolution, an interesting interactivity approach which undoubtedly increases proprioception within VR. Advanced controllers such as the Oculus Touch and the Index controllers further refined this emulation, while real hand tracking like Leap Motion's or the one present on the Oculus Quest present a whole new horizon to explore in interactivity terms.
Which is the preferred approach? Different as they are, they have distinct use case scenarios and present their own set of advantages and drawbacks, although most hand tracking approaches are not mature enough to outperform controller-based ones: the studies of Caggianese, Gallo and Neroni (2018) and Gusai, Bassano, Solari and Cheesa (2017) prove Leap Motion's hand tracking implementation to be a slower pointing tool than the Vive controllers, and the work of General et al. (2015) illustrates how far Leap Motion's technology stands from the mouse and the trackpad both in terms of movement time (meaning speed) and error rate.

Trezl, Dittrich and Bullinger (n.d.)'s study compares the performance of Leap Motion's technology to both controller pointing and touch-based interactivity. Their performance was evaluated through the use of three measurable units: time (for the task to be completed), accuracy and user preference. The raycast approach proved to be most accurate but the slowest and least favorited approach, while the touch-capable controller was the fastest and the preferable alternative. On the other hand, Leap Motion's implementation performed in-between them both in terms of time and user preference, although it was the least accurate of them all. The hand tracking support of current VR headsets is not yet perfect: its latency could be improved and so could its accuracy, but it represents a new interaction paradigm to build upon. Most research regarding the tracking of real hands in VR is done using Leap Motion's technology, so the popularization of hand tracking support with headsets like the Quest is bound to improve the interactivity of this approach.
Contradicting information regarding the subject can be found in scientific literature: Speicher et al. (2018) overviews this comparison in typewriting tasks, which are covered later on this document, and their work suggests that the raycast approach outperforms the touch-based one, with Leap Motion's hand tracking support being surpassed by them both. This comparison is later reviewed with greater extent.

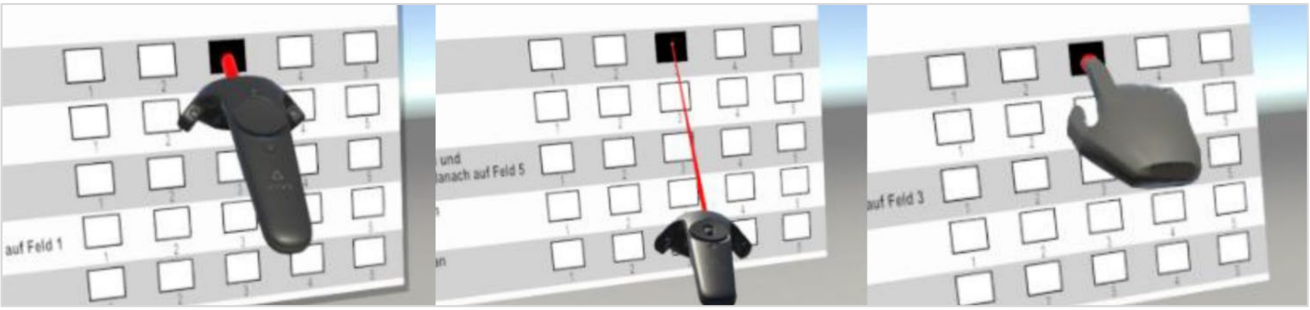

Figure 8. The approaches found in Trezl et al.'s work, side by side
[Extracted from Trezl et al. (n.d.)]

There are environments in which VR interactivity has been proven to be as much or even more effective than traditional tools due to the fact that VR technology allows for advanced depth simulation due to its dual-screen nature. The tasks most benefitted from the depth-sensitive capabilities of these virtual environments are those that best make use of this inherent quality: a clear example of this can be seen in 3D object manipulation, where XR has been proven to be more efficient than traditional 2D-based tools such as the mouse and its cursor: the work done by Krichenbauer et al. (2018) showcases how using 3D input devices (such as most VR controllers) yielded lower task completion times than the ones obtained through the use of a mouse.

3D object manipulation highlights how VR can outperform traditional tools, proving VR interactivity's potential. Speaking of which: while a variety of pointing approaches within VR have already been overviewed, typewriting by itself is an issue yet to be discussed and many productivity environments require text input of some sort. How can a keyboard and/or its function be used within virtual environments?

2.2. VR typewriting

As with pointing mechanism, there are many approaches to text input within VR. Half of these approaches base their functioning on the interaction techniques detailed during the previous chapter (such is the reason behind their elaborated commentary). As already stated, many productivity environments require text input of some sort, be it for document editing, social media tasks and publications, e-mailing, coding… However, text entry user interfaces have been a bottleneck of non-traditional computing devices. Smartphones, tablets and other smart devices have tackled this issue by using touchscreen-based virtual keyboards, but developing a text-input mechanism that works well within VR presents a challenge of its own: while there are advantages and capabilities exclusive of the VR platform, none of the interactivity approaches that had been covered so far can truthfully emulate the keyboard's sensations as the haptic feedback presented within VR controllers is not as responsive and "life-like" as touching a real and solid keyboard due to the volume it occupies.

Since the release of the first Rift iteration, research regarding VR has definitely increased, and thus so have the research done in VR interactivity and VR typewriting. There is no approach regarded as the defining standard and, given how quickly VR interactivity progresses, it is hard to determine if any of the current ones will preserve its viability in the future. However, the future is built upon the research done today, and that is why the following pages overview the many typewriting techniques being used and studied for the VR platform.

2.2.1. Real physical keyboards within a virtual environment

The VR-aimed keyboard technique on which most scientific literature is based on is none other than the least virtual one: researchers and developers have been meddling with the possibility of bringing real physical keyboards to virtual environments, with remarkable results performance-wise.
Many were the approaches taken, but all of them have something in common: they require additional and/or specialized hardware to track the keyboard in real-life so that it can be then emulated within VR. That is undoubtedly the biggest drawback of all of these approaches as, from an accessibility standpoint (which is of utmost importance if VR is to become a commodity platform), all of them require additional costs and a more complex setup (compromising mobility, among other factors). Whether the use of these sort of peripherals will become a standard is yet to be seen, but modern trends seem to suggest otherwise. Either way, since these typewriting approaches outperform alternative takes and it is upon them that most scientific literature is based on, the following paragraphs and pages are dedicated to detailing the tracking technology behind them.

During the introduction of this document I mentioned the potential behind Leap Motion's technology and commented upon the research done upon it. Depth perceptive cameras (Figure 11 showcases one mounted on a HMD) can track one's hands within the three-dimensional space, and the Leap Motion's device is not only remarkable performance-wise but also one of the most widely used. Besides its "age" (in VR terms), this specialized device and its usage are still being explored and studied in many different subjects and areas, one of them being typewriting within VR.

The study by Hoppe et al. (2018) is worth mention for its setup simplicity and overall performance. Their work combines Leap Motion's technology, which they used for the hand tracking, with HTC's Vive Trackers, which they used to track the keyboard itself.

A more advanced (and complex) version of Hoppe et al.'s approach was performed by Lin et al. (2017), whose approach was also based in the use of two devices, the first one being a high-cost specialized known as "Vicon Motion Capture System". Said system is an advanced motion tracking setup aimed at enterprises, universities and other organizations (expensive and not accessible to the average consumer). The second device used in the study at issue is a specialized external depth perceptive camera: the Intel RealSense F200. The work of Merriaux et al. (2017) over-validates Vicon's claims: their products are top notch when it comes to motion capture, and the virtual keyboard being reproduced was most likely indistinguishable from a real one from a tracking/mechanic standpoint.

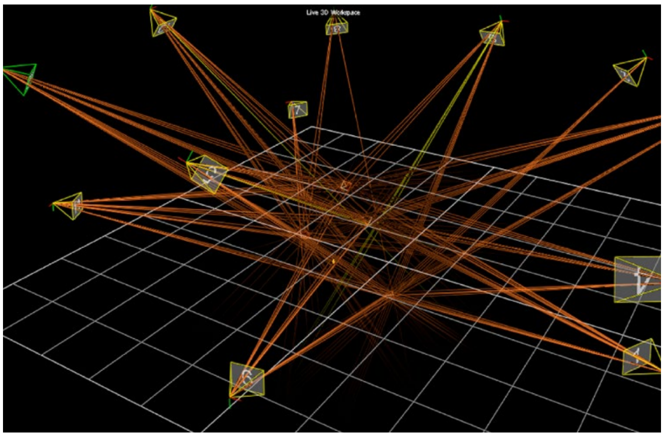

Figure 9. Graphical view of the Vicon Motion Capture System: twelve cameras are used to identify, locate and reconstruct unique markers in a 3D environment such as VR ones [Extracted from Wakkary et al. (2014)]

Regarding the camera used in their study, Intel RealSense products allow the computing of real-time depth perception data processing, so its use case scenarios are many: autonomous drones and vehicles, smart home devices, advanced face recognition and, expectedly, XR applications; it is, in fact, a line of products akin to Leap Motion's offering. Through the use of Intel's product, they integrated the user's hands within the virtual environment in a similar fashion to Leap Motion's technology and the hand tracking capabilities of the Oculus Quest HMD.

To achieve a better tracking of the hands, some researchers have used retroreflective markers placed on the users' fingers, namely Knierim et al. (2018) and Grubert et al. (2018a). Instead of Leap Motion's technology or Intel's RealSense, both studies use the OptiTrack SDK for streaming the positional data of bones, joints and the keyboard itself: the former uses eight OptiTrack 13W cameras, while the latter uses a single OptiTrack Flex 13.

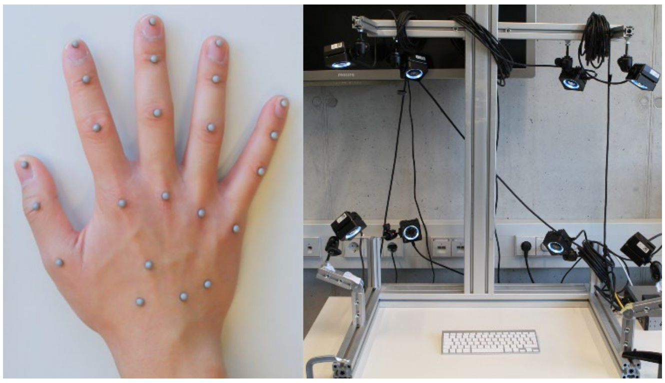

Figure 10. Hand with retroreflective markers on the left and OptiTrack's hardware setup for finger and keyboard tracking on the right [Extracted from Knierim et al. (2018)]

Alike Vicon's offerings, OptiTrack products and technology are top notch (see Furtado et al., 2019) and not consumer oriented but geared towards enterprises, universities and other organizations instead.

Returning to the approaches' overview, Grubert et al. (2018a)'s work also compares the performance difference between standard physical keyboards and touchscreen-based ones when integrated within a VR environment, to find out the former clearly outperforms the latter. Knierim et al. (2018)'s, on the other hand, compares the performance difference between different types of hand virtualization from a semi-realistic one to the most minimalistic one: just rendering one's fingertips. While they found no performance difference between the different hand visualizations, extensive research by Schwind et al. (2017a) and Schwind et al. (2018) suggest that any mismatch/inconsistency between the visual and haptic experiences affects virtual hand ownership, subconsciously disorienting the users and directly affecting the aforementioned "presence" that characterizes VR and thereby decreasing one's typing performance.
Furthermore, the research conducted by Grubert et al. (2018b) also studied minimalistic representations and within their work it is suggested that they achieve better typing results than more complex representations such as the hand tracking based ones: the usability of an approach on which only the fingertips were rendered was evaluated within their work, and their research suggests that the performance of such a novel (minimalistic) take was of great remark. This collides with the work of Knierim et al. (2018), whose study suggests that the approach on which only the fingertips were rendered performed below hand tracking based representations supposedly due to the lower virtual hand ownership levels of the approach at issue.

Another technique is found within the study performed by Walker et al. (2017), which is based on the fact that the physical keyboard is occluded from the users' view: instead of emulating the users' hand, a virtual assistant of sorts is used so that whenever a keystroke is detected upon the physical keyboard, a virtual keyboard (within the virtual environment) lights up the corresponding key thus granting the users the feedback they need to ubicate their hands' position and that of the keys. The performance achieved is on par with aforementioned approaches, partly due to the use of the VelociTap keyboard decoder (see Vertanen et al., 2015); this raises questions regarding the relevance of error measurement in these tests, which is assessed later on the document.

The possibility of integrating one's keyboard and hands into a virtual environment through a live video feed has also been studied. Lin et al. (2017) did so through rudimentary means (basically reproducing what was being recorded as a screen interface inside the virtual environment), but the works of McGuill et al. (2015) and Grubert et al. (2018b) used a green screen in their recordings so that they could effectively isolate both the hands and the keyboard itself through chroma key compositing to then reproduce them within the virtual environment (the latter's setup is illustrated in Figure 11, where the fiducials used to properly track one's hands can also be observed). Both works suggest that such an approach is perfectly functional (even outperforming common hand virtualization), but it heavily compromises immersion (whose importance in productivity tasks was already highlighted).

Figure 11. Grubert''s chroma key compositing setup
[Extracted from McGuill et al. (2015)]

Perhaps among all the approaches detailed so far the most interesting one was detailed by Logitech's team in the work of Bovet et al. (2018) as its simplicity, accessibility and performance are all noteworthy. In their study, they used HTC's Vive Tracker to capture the physical keyboard, a Logitech G810, and then used the keyboards' 3D model to properly reproduce it within the virtual environment. To track the user's hands, they used the Vive's integrated camera; it is an inexpensive approach when compared to aforementioned techniques, only needing one peripheral with an accessible price point.

The performance comparison done by Dube and Arif (2019) suggests that Logitech's approach performs better than most of the previously described ones; however, there's a considerable margin of error that has to be taken into account: on one hand, most studies measure typing performance in wpm (words per minute) so results depend on the text being written (and the average number of characters each word presents) and, on the other hand, each approach has its own userbase with each own level of typing expertise which, as seen in Knierim et al. (2018), is quite a relevant factor. While the latter of these issues is harder to resolve without a substantial amount of subjects, the former has been already addressed by some authors and works such as Hoppe et al. (2018)'s, where instead of measuring typing performance in wpm numbers they used cpm (characters per minute) to do so.

Figure 12. Logitech G810 recreated within a virtual environment with no more than the Vive's integrated camera and one Vive Tracker [Extracted from Bovet et al. (2018)]

2.2.2. Raycast-based virtual keyboards

Although real keyboards within VR have their own set of advantages, there are numerous use case scenarios where virtual keyboards might be the most convenient approach and thereby the preferred one. The most popular virtual keyboard within VR applications and environments bases its interactivity in the raycast approach that has already been discussed when overviewing VR interactivity. Despite it being the most commonly seen text input mechanism within virtual environments, there is limited amount of scientific literature on the subject (more so if compared to the integration of real physical keyboards in virtual environments). Whether these virtual keyboards are screen interfaces or not is a tagging matter but, for the best understanding of this technique, this document will be tackling them as such; concept-wise a keyboard is no more than an array of interactables (say, buttons) with a pre-established layout disposed over a 2D-like plane. In fact, most keyboards using this approach are usually displayed perpendicular to the user as a vertical and floating screen interface where, in order to type and input any text, the user has to point/control the casted ray or rays towards the keys to be pressed, clicking on a controller's button when the laser hovers over one of said keys thereby confirming the keypress itself.

If there's a study to be highlighted in this regard is the one done my Speicher et al. (2018), where this approach is studied through a typewriting exercise and proved to be the preferred approach both in terms of performance and user preference. However, there exists some contradiction on the subject (Trezl et al., n.d.) and there is no significative amount of scientific literature to properly explore and discuss the claims at issue. What seems more commonly accepted by researchers is the accuracy of this approach. It is also worth noting that Speicher et al. (2018)'s implementation of this technique allows both controllers to work at the same time, while most Oculus applications (namely: the Oculus environment itself, the Oculus Web Browser, YouTube, Amazon Prime Video, Netflix…) are limited to one casted ray at a time. The implications and performance differences between using a single ray and using both simultaneously is explored with great extent later on this document (Figure 13 showcases the comparison at issue).

While 6-DoF controllers are the preferable raycasters/pointers, it was quite uncommon to see them until the arrival of the HTC Vive. Before that time, the original raycaster/pointer was the headset itself, a technique that even today is still being used in some applications or headsets (in fact, the Microsoft HoloLens originally combined this approach with a gesture-based trigger mechanism to interact with the elements presented within its mixed reality interface). Speicher et al. (2018)'s research (which also covered such a technique) shows how the performance of this approach as a text input mechanism is lower than the one seen with common controller-based raycasting. However, the work of Yu et al. (2017) suggests that some head-based raycast approaches could yield remarkable performance results, although in Dube and Arif (2019)'s work it is suggested that Yu et al. (2017) might have augmented their system with predictive features.

Despite its performance and accuracy (whether high or low), this approach is to be avoided if possible due to it being an extremely tiresome technique which could potentially lead to fatigue. In long typing sessions, this could turn into significant (acute) drawback.

The Oculus Quest home environment and some native applications also allow for hand tracking based raycast/pointing interactivity using the hands as raycasters/pointers. So is the case with Microsoft's Mixed Reality Toolkit, built for the development of MR applications for Microsoft's own HoloLens and other HMDs of similar nature; it is worth noting that both frameworks use the pinching of the index finger and the thumb to confirm an interaction. An alternative take is explored in Ishii et al. (2017)'s work, where they use a smartphone head-mounted through the use of a Google Cardboard lookalike so that, by using the device's camera, a pointer is displayed on the screen based on the position of the user's hand, allowing him/her to interact with a target of choice by simply folding the thumb (within typewriting, said targets would be the keys themselves).

Lastly, another noteworthy approach to raycasting can be found within gaze typing, where one's eyes are used as pointing tools which the user can use to pinpoint the interactables of choice (the keys). Rajanna and Hansen (2018) evaluated two different approaches: a dwell-based one, where small pauses confirmed the keystrokes, and a click-based one, where the user's physical input/click is what dictates the key press. While functional, neither approach seemed to be remarkable performance-wise; the work of Ma et. al (2018) combined gaze typing with the brain electric signals, but it does not achieve better results and further complicates the typing experience not only by requiring an electroencephalogram but also by avoiding the common QWERY layout.
Since eye-tracking has already been implemented in some HMDs (an updated version of the HTC Vive Pro named the Vive Pro Eye is perhaps the best example of such an implementation) and given that foveated rendering eases the hardware load considerably, it is possible that gaze-typing finds its place with future HMDs.

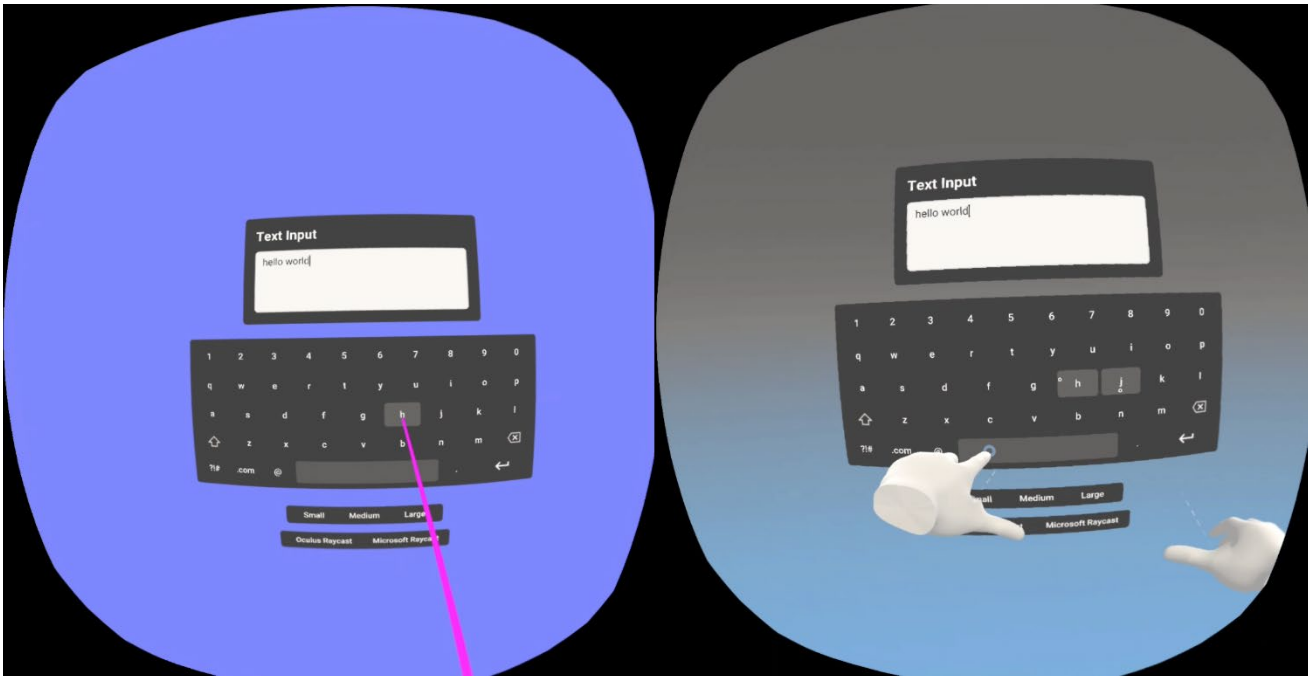

Figure 13. Single raycast approach (left) and dual raycast approach (right)
[Extracted from the prototype discussed later on this document]

2.2.3. *Touch-based virtual keyboards*

As with the raycast approach, touch interactivity was briefly discussed in previous chapters when discussing VR interactivity as a whole and, as was the case with said approach, there is not much scientific literature evaluating this interactivity's viability as text input mechanism. It is also noteworthy that, as mentioned earlier in this document, contradictions on its performance and usability can be found within scientific literature: while the work of Speicher et al. (2018) suggests that controller-based touch interactivity is surpassed by the raycast-based approach both in terms of performance and user preference (although ranking second in their performance comparison, which is remarkable per se), Trezl et al. (n.d.)'s work suggests otherwise, claiming touch interactivity to be both the fastest and preferred approach among their userbase. The former seems to be more adequate for this study as it is focused on text input as well, but fact is that there is not much scientific literature to back the claims and results of the work at issue.

Figure 14 showcases an example of touch-based interaction as typewriting mechanism. Said figure uses the 6-DoF controllers, their gyroscope and accelerometer's data and their built-in capacitive sensors to remarkably emulate one's hands (virtual hand emulation has been touched upon throughout the introductory context) which the user can use to interact with the keyboard's keys. However, as detailed during the touch-based interactivity commentary, the controllers at issue can be used to emulate other objects which could lead to more novel approaches to the typewriting exercise. Among those takes, the most remarkable one lies within the emulation of a pair of mallets which increase the users' reach by using the controllers as an additional joint, making of the keyboard a drum or xylophone of sorts with noteworthy performance and accuracy as typewriting mechanism.

While current 6-DoF controllers are usually the preferred input devices, there are a few other (novel) approaches to touch-based interactivity. The work of Gugenheimer et al. (2016) uses the backside of the HMD as a touch-sensitive surface (through the use of a smartphone attached to said side) mimicking the trackpad's interactivity previously detailed within the document. A similar take was studied in Kim and Kim (2017)'s work, but in this case the smartphone was not attached to the backside of the HMD, being held by the user at all times. More avant-garde approaches have also been explored, like using hand tracking based interactivity to touch and thus type upon a contact-based virtual keyboard.

Figure 14. Touch-based virtual keyboard example using emulated hand controls based on the Oculus Touch controllers' data
[Extracted from the prototype discussed later on this document]

2.2.4. Hand tracking based virtual keyboards

As detailed throughout the introductory context, hand tracking support within virtual environments have been researched for years now. Leap Motion's technology caused a spiked in the interest behind hand-based interactivity, and hand tracking itself is now built-in some HMDs (the most notorious one to do so being the Oculus Quest). The most notorious hand-based approach has already been mentioned within the touch-based virtual keyboards' commentary, as the user's tracked hands can be intuitively used alongside contact-based interactivity. The works of Trezl et al. (n.d.) and Speicher et al. (2018) evaluated this technique through the use of Leap Motion's technology, although only the latter studies it from a typewriting perspective. However, there is some contradicting information on its performance: the former of the works suggests it outperforms controller-based raycasting/pointing whereas the latter suggests hand tracking based touch interactivity is surpassed by both controller-based touch interactivity and controller-based raycasting/pointing.

Hand tracking based controller pointing has also been explored inside the Oculus development teams, as the Oculus Quest presents such interactivity within some spaces (namely the home one, where it can be used to navigate between the available menus) and applications are starting to make use of this technology (it was added to Virtual Desktop in June 2020). A similar take is found within Microsoft's Mixed Reality Toolkit, as it was already stated within the raycast text input overview. Scientific literature has not yet researched this approach, although gesture-based interactivity as a typewriting mechanism has been explored in the works of Mehring, Kuester, Singh and Chen (2004), where they use a keyboard-independent approach which maps the QWERTY layout to one's fingers instead (a similar approach is seen in Mourouzis, Kilintzis, Chouvarda & Maglaveras, 2014), whilst the work of Markussen, Jakobsen and Hornbæk (2013) evaluated and compared a variety of hand-based approaches (although none achieved remarkable results). What seems accepted within literature is that haptic feedback significantly affects the user experience and his/her performance, and its absence reduces immersion and typing speed considerably.

Current interactivity advances, like the native hand tracking support in the Oculus Quest, could also improve the efficiency of hand tracking based typewriting which, as commented throughout the introductory context, might be the next milestone in VR interactivity. In fact, Oculus itself is known to be studying novel hand tracking text input techniques (see Figure 16) for the Quest and future HMDs.

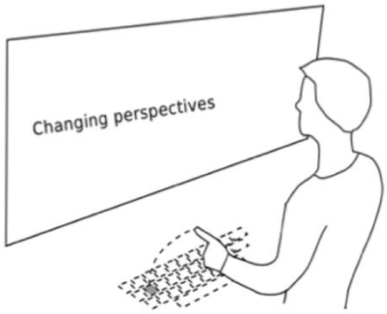

Figure 15 (Top). Touch-based keyboard concept used alongside hand tracking technology [Extracted from Markussen et al. (2013)]

Figure 16 (Bottom). Oculus' text input study [Extracted from https://uploadvr.com/frl-pinchtype-ar-vr-keyboard/]

2.2.5. Other novel approaches

There are many other text input mechanisms contemplated within scientific literature. Some novel approaches were already highlighted throughout the document, but the work of Dube and Arif (2019) covers many of these more novel approaches with greater extent and is the recommended lecture for doing so.

Despite the lower performance of these takes, some use case scenarios may benefit from them (that is why these approaches are included within the scope of this document). Amongst them, it is worth mentioning the use of specialized hardware (besides the keyboard itself, as such an approach was already covered with great extent) as text-input mechanisms for, precisely, the novelty of such a take and the outstanding performance of the real physical keyboard even within virtual environments.

Arif and Stuerzlinger (2009) suggest that the Twiddler, a specialized text-input device, is the fastest approach amongst their evaluated non-QWERTY typewriting methods. Also known as chord keyboard, its form factor makes it, theoretically, a great candidate for VR usage (said device is found on the right-hand side of Figure 17). However, the studies done by Bowman, Rhoton and Pinho (2002) and González et al. (2009) disregard the performance of this technique. Both of these works studied and researched a considerable amount of hardware and their viability as text input approaches within VR (see Figure 17), although none was remarkable enough to compete against the approaches already discussed throughout the chapter.

Figure 17. Specialized hardware for text input within VR
[Extracted from Bowman et al. (2002)]

Videogame controllers, also known as gamepads, have already been used as text input for years both in videogames/applications and within each console's OS. These controllers have also found a place within VR (most notably within VR-based games) and scientific literature have explored these peripherals' viability as novel approaches to typewriting within VR (see Yu et al., 2018). However, as is the case with all the takes found within this chapter, the performance levels obtained through said approaches is not significant enough to compete against the most relevant techniques highlighted throughout this meta-review.

3. Objectives

3.1. Summarizing the meta-review

The meta-review nature of this project makes its objectives very clear: to study and evaluate the many text-input approaches within VR. Such task has already been performed throughout the previous chapter but, although many are the typewriting mechanisms detailed, many are yet to be explored and properly studied. Due to that, following pages will try to reinforce the meta-review with a prototype which aims to build upon the many works highlighted throughout this document; said prototype is covered later on with great extent.

First of all, and summing up the meta-review at issue, literature suggests that the best approach among those covered is found within the real physical keyboard's integration inside virtual environments. There are many takes on said integration with many different instruments and setups at use and, among them, Logitech's approach (seen in Bovet et al., 2018) not only seems to outperform the alternatives but it is also the most accessible one.

Other techniques do not require specialized hardware or peripherals (an advantage from an accessibility standpoint), basing its functionality in the use of 6-DoF controllers (a pair of which is bundled with most notable HMDs). The most commonly used approaches are found within the raycast-based virtual keyboards and the touch-based ones. Although there are use cases for each of them, fact is that it is yet unclear which performs better overall as there is some misleading information on the subject. On one hand, the work of Speicher et al. (2018), which was built around this comparison, suggests that the raycast-based technique outperforms touch-based virtual keyboards in typewriting exercises although the performance difference between them is not significant enough to conclude the discussion (and there is not much research to further validate their claims). On the other hand, there is some literature which actually contradicts the study at issue: Trezl et al. (n.d.)'s work compares the performance of Leap Motion's technology to both the raycasting approach and the touch-based one, and their research suggests that the latter outperforms them all in terms of performance and user preference. However, it is worth noting that this study is not focused on typewriting: the task from which their results were obtained consists on a questionnaire, which has fewer interactions than typewriting exercises where many keystrokes (interactions) are required along an extended period of time. The difference between their tests opens up the possibility of their results not actually contradicting each other: perhaps touch-based interactivity works best in short periods of time while controller pointing outperforms it in prolonged periods of time. The study of Dube and Arif (2019) also compares the approaches at issue, among others, but they do not explore this difference any further. Their work, while remarkable, presents a noteworthy limitation: there is an inherent margin of error in their comparison as each approach has its own userbase with each own level of typing expertise which, as seen in Knierim et al. (2018), is quite a relevant factor. Furthermore, each approach/study uses a different text so the average number of characters per word is distinct, meaning that the use of wpm as measurement unit slightly increases this margin of error; the comparison at issue is remarkable indeed, but the lack of reproducibility and repeatability means that only significant differences should be considered.

This issue has already been addressed by some authors and works such as Hoppe et al. (2018)'s, where instead of measuring typing performance in wpm numbers they used cpm (characters per minute) to do so.

There is another noteworthy contradiction to be highlighted, as it arises another question to which find an answer: the research conducted on Grubert et al. (2018b) studied minimalistic representations and within their findings it is suggested that those simplistic takes achieve better typing results than more complex representations such as the hand-tracking based ones. Within their work, they evaluated the usability of an approach on which only the fingertips were rendered, and concluded that the results of such a novel take were as performant as the ones obtained through a live video feed of one's hands and the keyboard (which is akin to typing in a real-world non-virtual scenario, something agreed upon among scientific literature). This collides with the work of Knierim et al. (2018), whose research suggests that the approach on which only the fingertips were rendered performed below hand tracking based representations. Whichever the answer, the comparison itself is deeply tied to virtual hand ownership, a factor of utmost importance in VR text input (more so in long typing sessions); the work of Schwind et. al (2018) suggests that inconsistencies between the visual and haptic experiences distract the users, decreasing their virtual hand ownership and thereby negatively affecting immersion and presence, of utmost importance to neuropsychology (Price & Anderson, 2007) and to VR's viability as a workspace due to their effects in one's productivity levels (Csikszentmihalyi, 1975; Csikszentmihalyi, 1990; Kotler, 2014).

Going back to the interactivity-based comparison, hand tracking support offers an interesting take on VR interactivity and, if the arrival of Leap Motion's technology spiked its interest (Wozniak et al., 2016; Argelaguet et al., 2016), the Oculus Quest has simplified and widespread its use. There is not much scientific research regarding its usability as a typewriting mechanism, and the existing one suggests it is best used as a touch-based text input approach. The work of Speicher et al. (2018) suggests that such approach (by the use of Leap Motion's technology) is not remarkable performance-wise; Trezl et al. (n.d.) claim otherwise, as within their study hand-based touch interactivity outperforms the controller-based raycasting one. However, both works coincide in that, accuracy-wise, it is less reliable than other approaches due to current technology's tracking limitations.

There exist a variety of novel approaches within raycasting and touch-based interaction and many more beyond their scope. Among them, the most remarkable one lies in gaze-typing as some consumer-oriented headsets have already incorporated such technology within their capabilities (an updated version of the HTC Vive Pro named the Vive Pro Eye is perhaps the best example of such an implementation). The work of Rajanna and Hansen (2018), among others, dwells on the interactivity at issue and, although the performance showcased within their study does not seem to be that outstanding, fact is that foveated rendering (which renders one's peripherical vision at lower resolution to save on hardware resources) reduces the performance hit (upon the computing hardware) so much that it makes eye tracking support remarkable for performance-dependent applications, making gaze-based interactivity likely to be implemented in upcoming HMD. If so, eye-tracking support and thus gaze-based interactivity are bound to improve in the future.

3.2. Reinforcing the meta-review

As seen along the meta-review and its summary, there are many questions and contradictions left to unravel. To tackle those uncertainties, a prototype was developed with a variety of operating modes that have been designed and developed to evaluate the performance and usability of the approaches within its scope (by contrasting the results obtained through empiric experimentation with a set of subjects of significant size).
It can be said that the goal of this project goes beyond the meta-review itself and aims to build upon it (science itself is a pyramid of information which grows taller and broader as new pieces are incorporated).

The first technique to be covered by the prototype is the raycast-based one. As it has been already stated during the overview of this approach, the work of Speicher et al. (2018) uses both controllers simultaneously whereas most Oculus applications are limited to one casted ray at a time; both approaches are included within the scope of the prototype at issue to evaluate their performance difference and their overall validity within typewriting tasks. The hand tracking support of the Oculus Quest HMD and its viability as a typewriting tool also falls within the prototype's scope, and a hand tracking based raycasting approach is also present. Another goal behind the protype is to compare, through results of my own, raycast-based virtual keyboards and touch-based ones. To address the latter, a mode with emulated hand controllers was developed along with a mode based on the hand tracking support of the Quest (hand tracking based performance is unclear due to contradicting information on the subject: see Speicher et al., 2018 and Trezl et al., n.d.). Another touch-based technique within the prototype's scope is regarded as "the drum-like approach" throughout the document and was previously commented upon: its functionality is based on the controllers and their built-in gyroscope and accelerometer's data, which is used to emulate a pair of mallets by using the controllers as an additional joint making of the keyboard a drum or xylophone of sorts with noteworthy performance and accuracy as a typewriting mechanism.
Furthermore, the prototype also aims to evaluate the importance of virtual hand virtualization and, based on the approaches found within the works of Grubert et al. (2018b), Knierim et al. (2018), and Schwind et. al (2018), a typing mode was developed where only the fingertips themselves were visible as opposed to rendering the whole hand (see Figure 19).

To sum it up, the approaches being studied with this prototype are: **[1]** single raycast, **[2]** dual raycast, **[3]** dual raycast using one's hands as raycasters/pointers, **[4]** touch-based virtual keyboard with emulated hands, **[5]** touch-based virtual keyboard with emulated mallets, **[6]** touch-based virtual keyboard with the users' own hands and **[7]** touch-based virtual keyboard with the users' own hands but only rendering the fingertips.
At the time of writing, Microsoft Mixed Reality Toolkit (covered with more extent later on) limits the hand's interactivity with buttons to a single finger: the index. It has been pointed out to me that the reason behind such limitation is that Microsoft had conducted research suggesting that users end up accidentally pressing buttons more often than it adds value. To confirm those claims, two additional approaches were developed: **[8]** a touch-based virtual keyboard with which the users can interact through their index fingers and **[9]** the very same exercise although this time only rendering the index fingertips (to evaluate virtual hand ownership).

4. Development of the prototype

4.1. Overview

There are many frameworks upon which to build VR, AR and MR applications. The concept of XR came into being to embrace all of the previously mentioned “realities” under one mantle, and the work of MacIntyre and Smith (2018) reflects on the potential behind WebXR and its development. Building on the web provides wide accessibility whilst also allowing web developers to easily tackle this kind of projects with ease, using the tools with which they are already familiarized. With WebXR, multiplatform development takes XR into account.
This prototype was initially conceived as a WebXR project using A-Frame’s framework (see Dibbern et al., 2018), which uses a declarative entity-component-system (ECS) to grant web developers control over objects through HTML (using the custom tags of said framework) and JavaScript whilst also presenting a visual inspector to manipulate and tweak VR scenes as one would do in more object-oriented frameworks like Unity and Unreal Engine. However, at the time of writing, there is no way to work and exploit the Oculus Quest hand tracking capabilities with this framework and, as a result, the prototype’s development was moved over to Unity.

The development of raycast-based virtual keyboards for the Oculus Quest has served as means to understand the current limitations of the Oculus framework regarding this kind of interactivity. As previously stated, Speicher et al. (2018) use both Vive controllers to simultaneously operate two rays; however, Oculus own apps are limited to one casted ray at a time. To compare the performance difference between the two approaches both were to be developed within Unity’s engine, and although the development of the single raycast technique was quite intuitive through the use of the Oculus Integration package, the Oculus input script doesn't support two raycasts at the same time (at the time of writing): a button input is required to switch from one controller to the other.

A few approaches were contemplated to achieve dual raycast support: modifying Oculus’ own code was a possibility, its development ended up being based upon Microsoft’s Mixed Reality Toolkit (MRTK) instead. While this toolkit was designed for another platform (the Windows Mixed Reality one), a customized version of it allowed the Quest to make use of most of its functionality; it was through the use of said toolkit that the evaluation and comparison between the single raycast technique (from the Oculus input script) and the dual raycast one (from Microsoft’s tools) was possible.
The Oculus Quest hand tracking feature can also be used as raycaster/pointer and, since no research has yet evaluated the performance of such an approach, the prototype at issue incorporates said technique in order to study its usability and compare its performance to that of controller-based raycasting. In this case, the raycast direction is a result of the palms’ orientation, and the key press is performed by pinching the thumb and the index fingers together. The development of such approach was based on Microsoft’s Mixed Reality Toolkit as it simplified the task and allowed for dual raycast support (meaning that both hands could be used simultaneously).

The prototype also incorporates three sizes in order to evaluate the relevance of such factor in VR typewriting: large, medium and small. This is the case with all the approaches covered within the prototype’s scope.

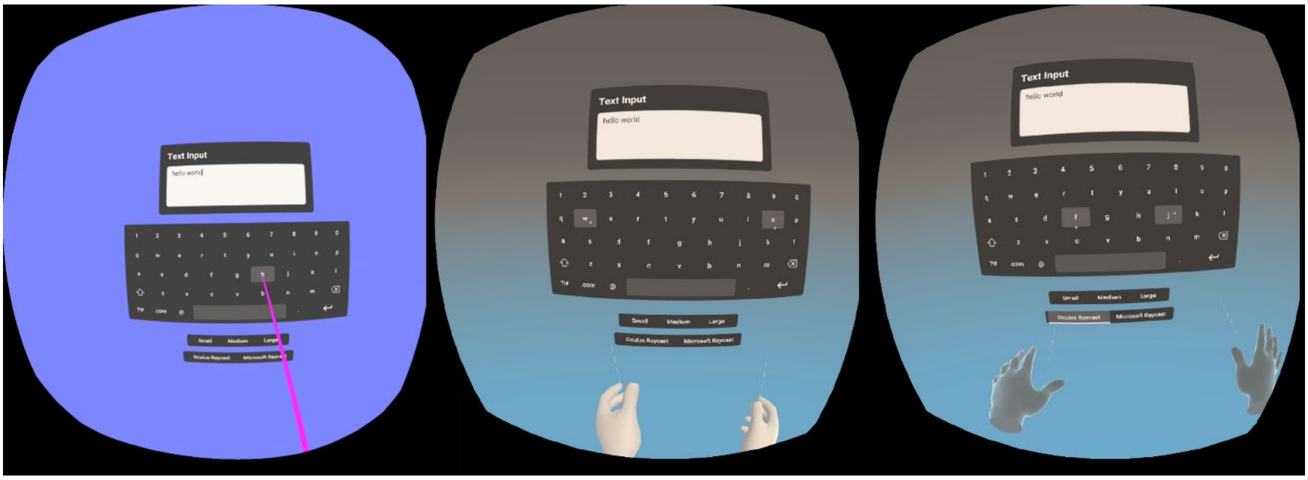

Figure 18. Oculus' single raycast approach on the left, Microsoft's dual raycast approach on the right and hand-tracking based raycasting on the right

The development of touch-based virtual keyboards with which the users could interact with their own hands was meant to be based on Oculus' interactable tools. However, as detailed later on when describing the major development issues encountered throughout this development, said tools presented a series of problematics which led to the study of other approaches. Using the hands' meshes as colliders found no avail, and it is worth mention that mesh colliders are far too demandant for most applications and should be avoided if possible. In the end, placing small spheres in each fingertip's location and then use those spheres as colliders made this mode function as intended: this is inspired by Oculus' own interactable tools, but the fingertips being tracked have colliders of their own. Such approach eased the development of the modes on which only the fingertips are meant to be rendered (by not rendering the hands' meshes and rendering the spheres' ones instead).

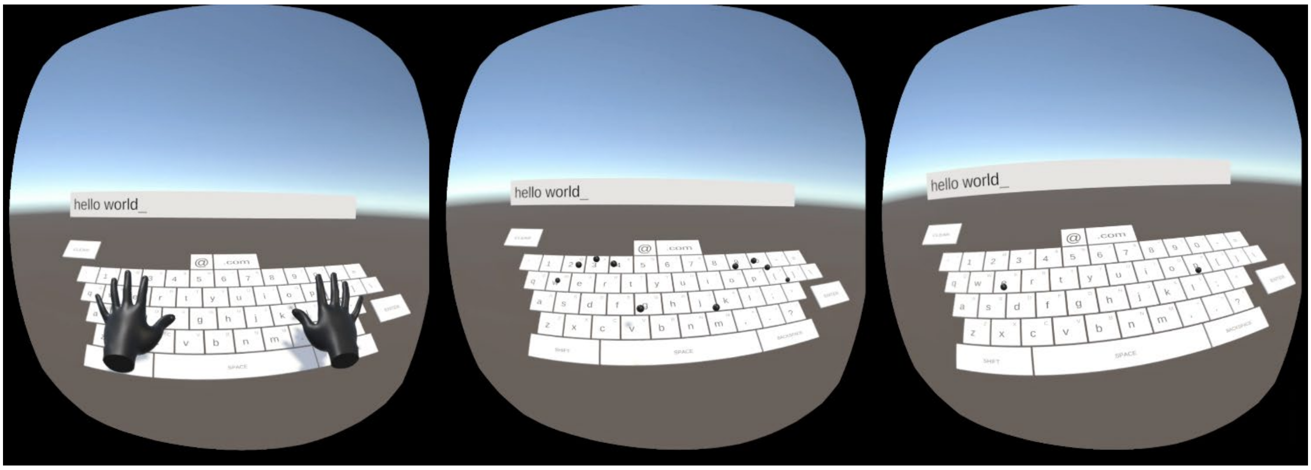

Figure 19. Touch-based virtual keyboards using the hand tracking capabities of the Quest

On the left: the hand is fully rendered (it is not possible to visually distinguish if all fingers are functional just the index is)
On the centre: the five fingertips' colliders are the rendered while the hands' meshes are now occluded
On the right: the index fingertip is now the only one whose collider is being rendered

The very same keyboard used in the hand tracking test is recycled throughout all touch-based keyboards to ensure consistency within the obtained results. It is worth noting that the keyboard at issue is based on an asset from a github repository (also available within the Unity asset store) and has been properly adapted and modified to best fit the scenes developed for the prototype. To best fit the many approaches within its scope, many design choices were considered, explored and/or applied to favour both their functionality and overall comfort.

Two are the touch-based text input techniques left to discuss: the emulated hands' approach and the drum-like one. Both takes require the use of controllers in order to reproduce, based on the gyroscopes' position and rotation, whichever model is needed (the former needs the hands, which can be found within Oculus' framework, while the latter emulates a pair of mallets). The emulated hands' approach presents a collider placed on each index, which required unpacking the prefab used for the hands in order to find the index tip position for each individual hand. By applying the colliders' properties used in the hand tracking implementation scene upon the empty object found and placed in the tip of each index, the intended behaviour was achieved: the left side of Figure 20 shows the emulated hand controls, where the user is isolating both index fingers through the use of the Oculus Touch capacitive sensors to then use said fingers as the interactables with which to write.

Regarding the drum-like approach: for the mallets to interact with the keyboard, a sphere with a collider was placed at the end of each one so that, by manipulating the mallets through the use of the controllers, users can trigger/activate the keyboard's keys at will.

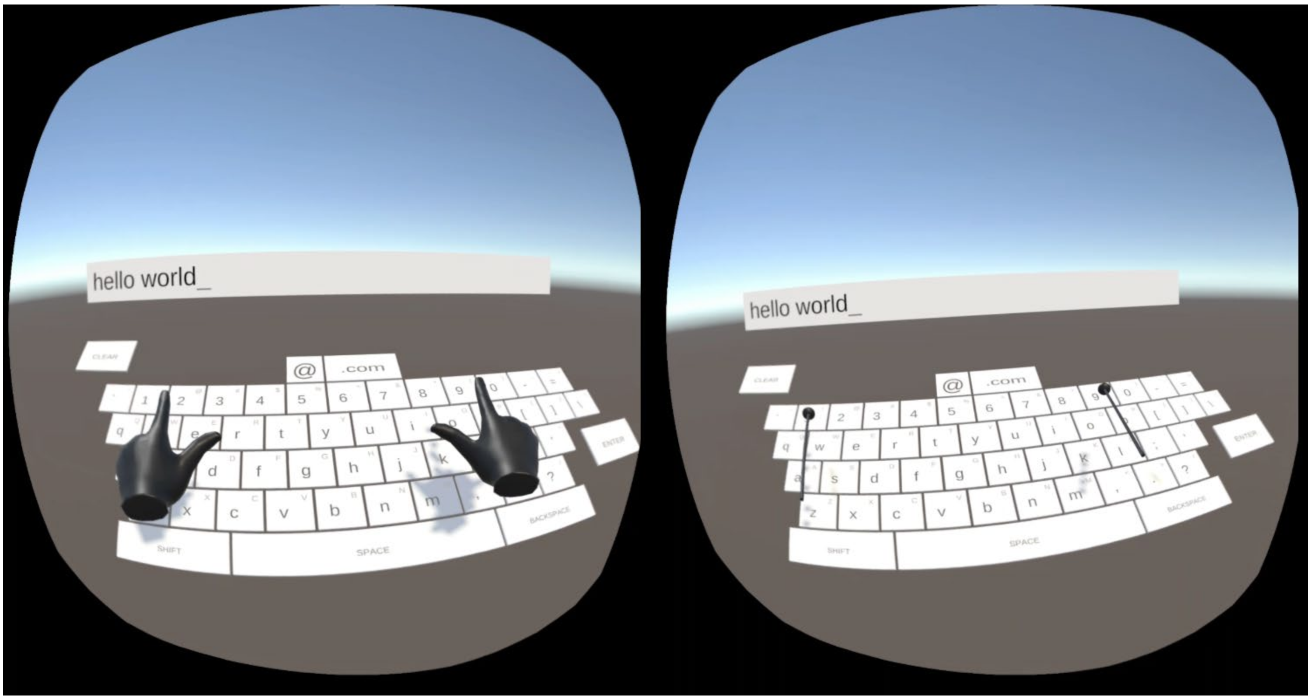

Figure 20. Touch-based virtual keyboards using the Oculus Touch controllers

On the left: emulated hands' controls (based on the capacitive sensors built-in the Oculus Touch)
On the right: the drum-like approach using emulated mallets

4.2. Issues

It is said that "all roads lead to Rome", and there are indeed many paths one can take to tackle hand tracking support within Unity. One would think that the most accessible and approachable of them all is the one found within Oculus' documentation, and while there are many remarkable tools within the framework itself, the hand tracking elements of it are lacking in many ways. For one, the last few months have seen quite a few iterations of the framework itself. While this might seem as something positive, there have been quite a few changes as to how the hands are implemented inside the scene (now requiring the HandsManager and the InteractableToolsSDKDriver prefabs to work as intended), which is confusing from a developer standpoint.

The next issue to be highlighted takes place when using the aforementioned elements: the hand tracking itself does not function correctly with the default time settings of Unity's player. To fix this issue, one has to significantly reduce both the "Fixed Timestep" and the "Maximum Allowed Timestep"; Figure 21 shows the values that worked best for the prototype, although it is worth noting that demandant applications may find a more significant performance impact with these values due to the higher frequency at which the hands are being tracked.

Figure 21. Timestep values that worked best with the intended approach (based on the InteractableToolsSDKDriver)

Another problematic took place when trying to track one's thumbs. The tracking of each individual finger is performed through the use of the InteractableToolsSDKDriver element and although there is a tracking prefab for the index finger, the middle finger, the ring finger and the pinky finger… There is none for the thumb (see Figure 22). Despite existing a prefab for all fingers but the thumbs, each of said prefabs gives developers the choice of which finger to track (meaning that, even with one finger's name, these elements can actually follow a different one). Said choice is found within the dropdown selection shown in Figure 23 where, despite the thumb's presence within the options, its selection does not yield the intended behavior.

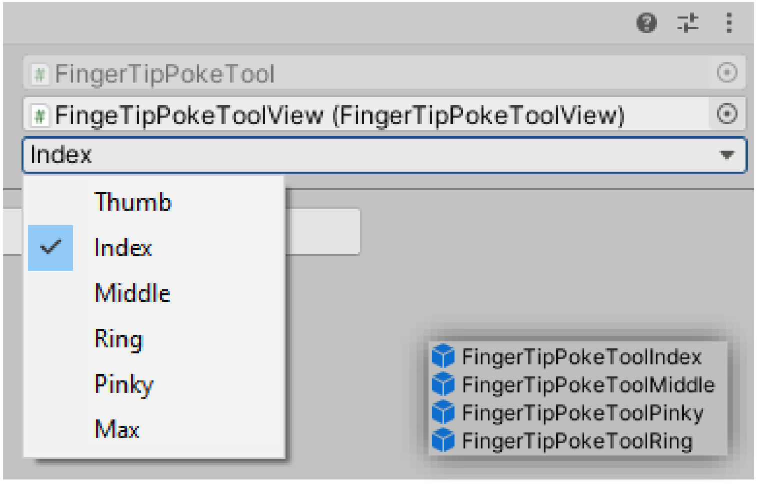

Figure 22 (Bottom right). The thumb-based FingerTipPokeTool prefab is missing

Figure 23. Thumb tracking is within the dropdown menu, yet it does not work

The reason behind this odd behavior is found in the underlying code which has some instructions incorrectly written, and selecting the thumb on the aforementioned dropdown menu makes the tracking element follow the index instead. The study found in Schwind et al. (2017b) specifically tackles the issues that a missing finger brings along (both performance-wise and with regards to hand ownership), so this was definitely an issue and Oculus have not fixed it yet at the time of writing (and this prototype development has seen every Oculus Framework update from the 13th to the 16th).

It is not a major problematic by any means, but it should have been already been addressed.

As previously stated, there are many paths one can take to tackle hand tracking support within Unity. The one being overviewed is the one found in Oculus documentation, although developing interactables for the InteractableToolsSDKDriver was not as intuitive as it should, so in the end the prototype was developed with a physics-based approach which, as detailed during the overview of this prototype's development, uses colliders placed at each fingertip.

While the many updates to the Oculus Framework brought some useful functionalities, such as being able to run the scene on the Quest via USB cable (using the Oculus Link functionality detailed within this document's introduction) without having to build the APK and sideload it every time (which was a major inconvenience). However, running the scene with the Oculus Quest's mixed reality active (named/referred to as the guardian system by Oculus) makes Unity crash and close instantly losing any unsaved changes, so one has to work quite meticulously to successfully avoid the aforementioned issues and crash-inducing problematics.

Lastly, it is worth mention that although the Microsoft Mixed Reality Toolkit allowed the prototype to include dual raycast support (even allowing the use of one's own hands as raycasters/pointers), said toolkit did not work with Unity 2019 and was prone to errors within Unity 2018, where every time one accesses a scene (specifically through Unity's own editor) on which the toolkit was being used, every single MRTK canvas had to be converted to the defaulted one by Unity and then converted back to MRTK canvas in order for the scene to work at all.

5. Evaluation

5.1. The test

As stated earlier in the document, the approaches being studied with this prototype are: **[1]** single raycast, **[2]** dual raycast, **[3]** dual raycast using one's hands as raycasters/pointers, **[4]** touch-based virtual keyboard with emulated hands, **[5]** touch-based virtual keyboard with emulated mallets, **[6]** touch-based virtual keyboard with the users' own hands and **[7]** touch-based virtual keyboard with which the users can interact through their index fingers. To evaluate hand ownership and its performance and usability implications another pair of approaches is included within the prototype's scope: **[8]** a touch-based virtual keyboard with the users' own hands but only rendering the fingertips and **[9]** the very same exercise although this time only rendering the index fingertips.

While most researchers have been using wpm as the performance measurement unit, the work of Hoppe et al. (2018) uses cpm disregarding the drawbacks of using words instead of characters: using wpm leads to inconsistencies when comparing different studies done by different authors, as each work uses a different text as input and, consequently, the average number of characters within each word is different and the margin of error is thereby increased. Thereby, the study at issue uses cpm. Regarding error measurement: the approach found in Walker et al. (2017)'s study performs as it does partly due to the use of the VelociTap keyboard decoder (see Vertanen et al., 2015) meaning that their work is based on the use of a decoder, which is arguably ideal from a practical standpoint as what matters is the end result. Dube & Arif (2009) also claims some of the approaches found in scientific literature might have augmented their systems with predictive features. Research such Ghost and Kristensson (2017)'s one showcase the potential of said decoders (whose convolution now may involve underlying neural network algorithms) and consequently their measurement is of secondary importance to this study, which values legibility instead (the text being legible imply that the amount of typos is not significant enough to hinder the decoding and interpretation of the intended input).

The prototype on which the test is based also aims to review the implication of the keyboards' size, which could be directly correlated with typing performance. In order to do so, each of the aforementioned modes is presented in three different sizes: small, medium and large. That, along with the need to repeat each task multiple times to further validates the results of this study, should minimize the learning process (although at the cost of being over-repetitive). As can be observed in Figures 24 and 25, there are two different keyboard designs; each design was chosen and developed to adequate the most each of the approaches within the prototype's scope: the raycast-based model is seen in Figure 24 while Figure 25 illustrates the design used in the touch-based virtual keyboards. Within said figures are showcased the keyboards' sizes, which were developed based on the FOV covered by each: raycast-based keyboards occupy around an 80%, 60% and 40% of the users' FOV in their large, medium and small sizes (respectively), whereas touch-based keyboards occupy around a 100%-110%, 80-90% and 60-70% of the users' FOV (two values are provided due to the mallet-based approach taking place at a greater distance, thereby making the keyboard occupy less of the users' FOV).

The difference in FOV values between the raycast-based and the touch-based virtual keyboards was in order to best accommodate each approach.

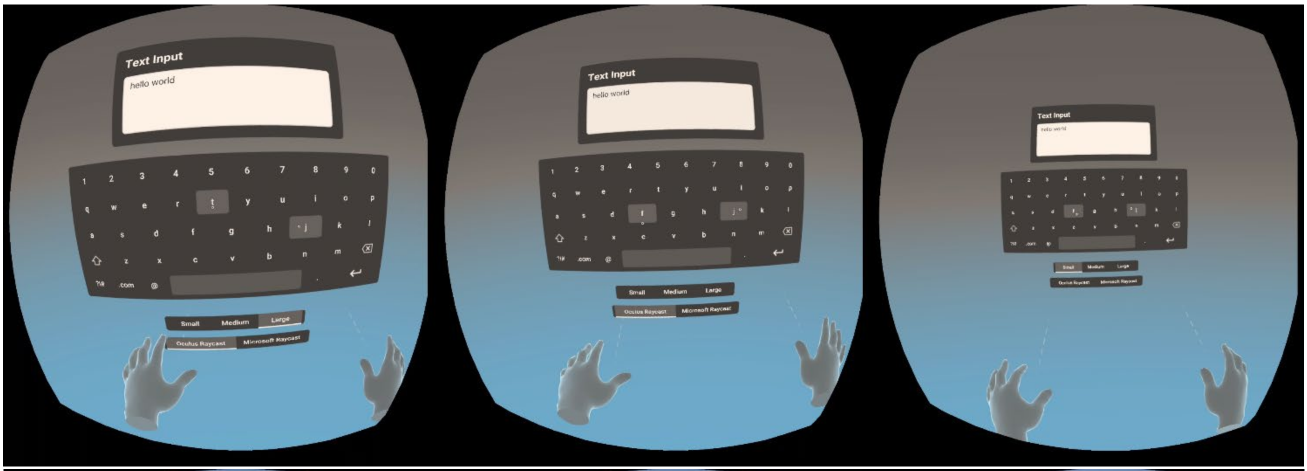

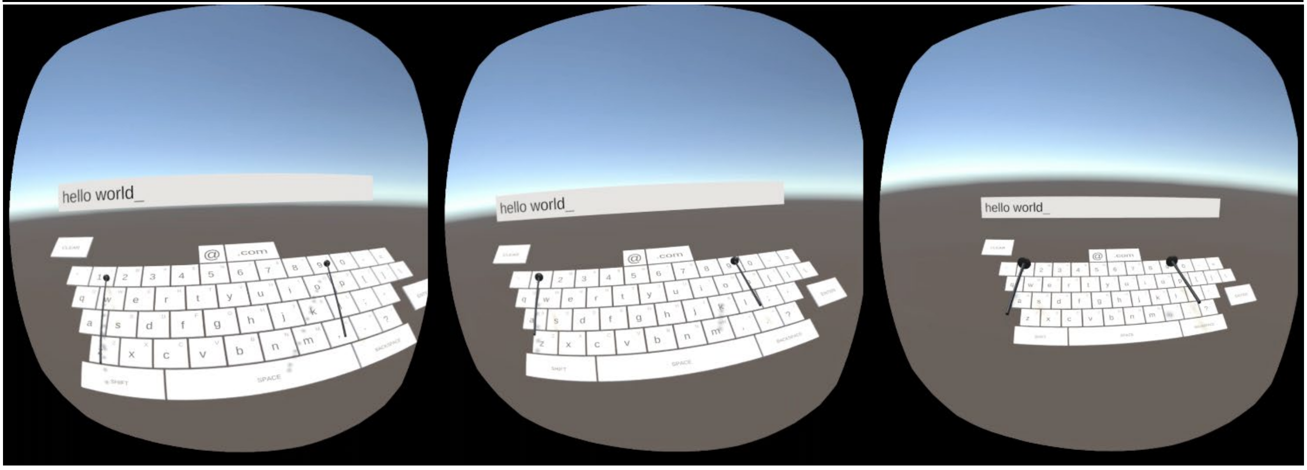

Figure 24 (Top). Size comparison in the raycast-based approaches

Figure 25 (Bottom). Size comparison in the touch-based approaches. From left to right: large, medium and small sizes

The test itself is meant to be supervised to control the environment and ensure the userbase is evaluated under the same conditions, as well as to observe their behavior with the different text input approaches and to listen to their feedback as they face the test. Due to the amount of approaches considered and the need to repeat each one to validate results and evaluate the performance implication of the keyboards' size along with the appropriate feedback acquisition, the test's duration ended up being around (if not over) an hour long.
The test was designed to be faced with two distinct mentalities: a calligraphic one, where users should try to input the text with little to none mistakes, and a speed-typing one, where the goal was to see how fast the users could theoretically type on each keyboard (based on the time in between keys) and the fluidity of each approach.

The following results were obtained with the test at issue and 10 participants ranging from 20 to 24 years of age, encompassing experienced and inexperienced typists as well as experienced and inexperienced VR users.

5.2. Results

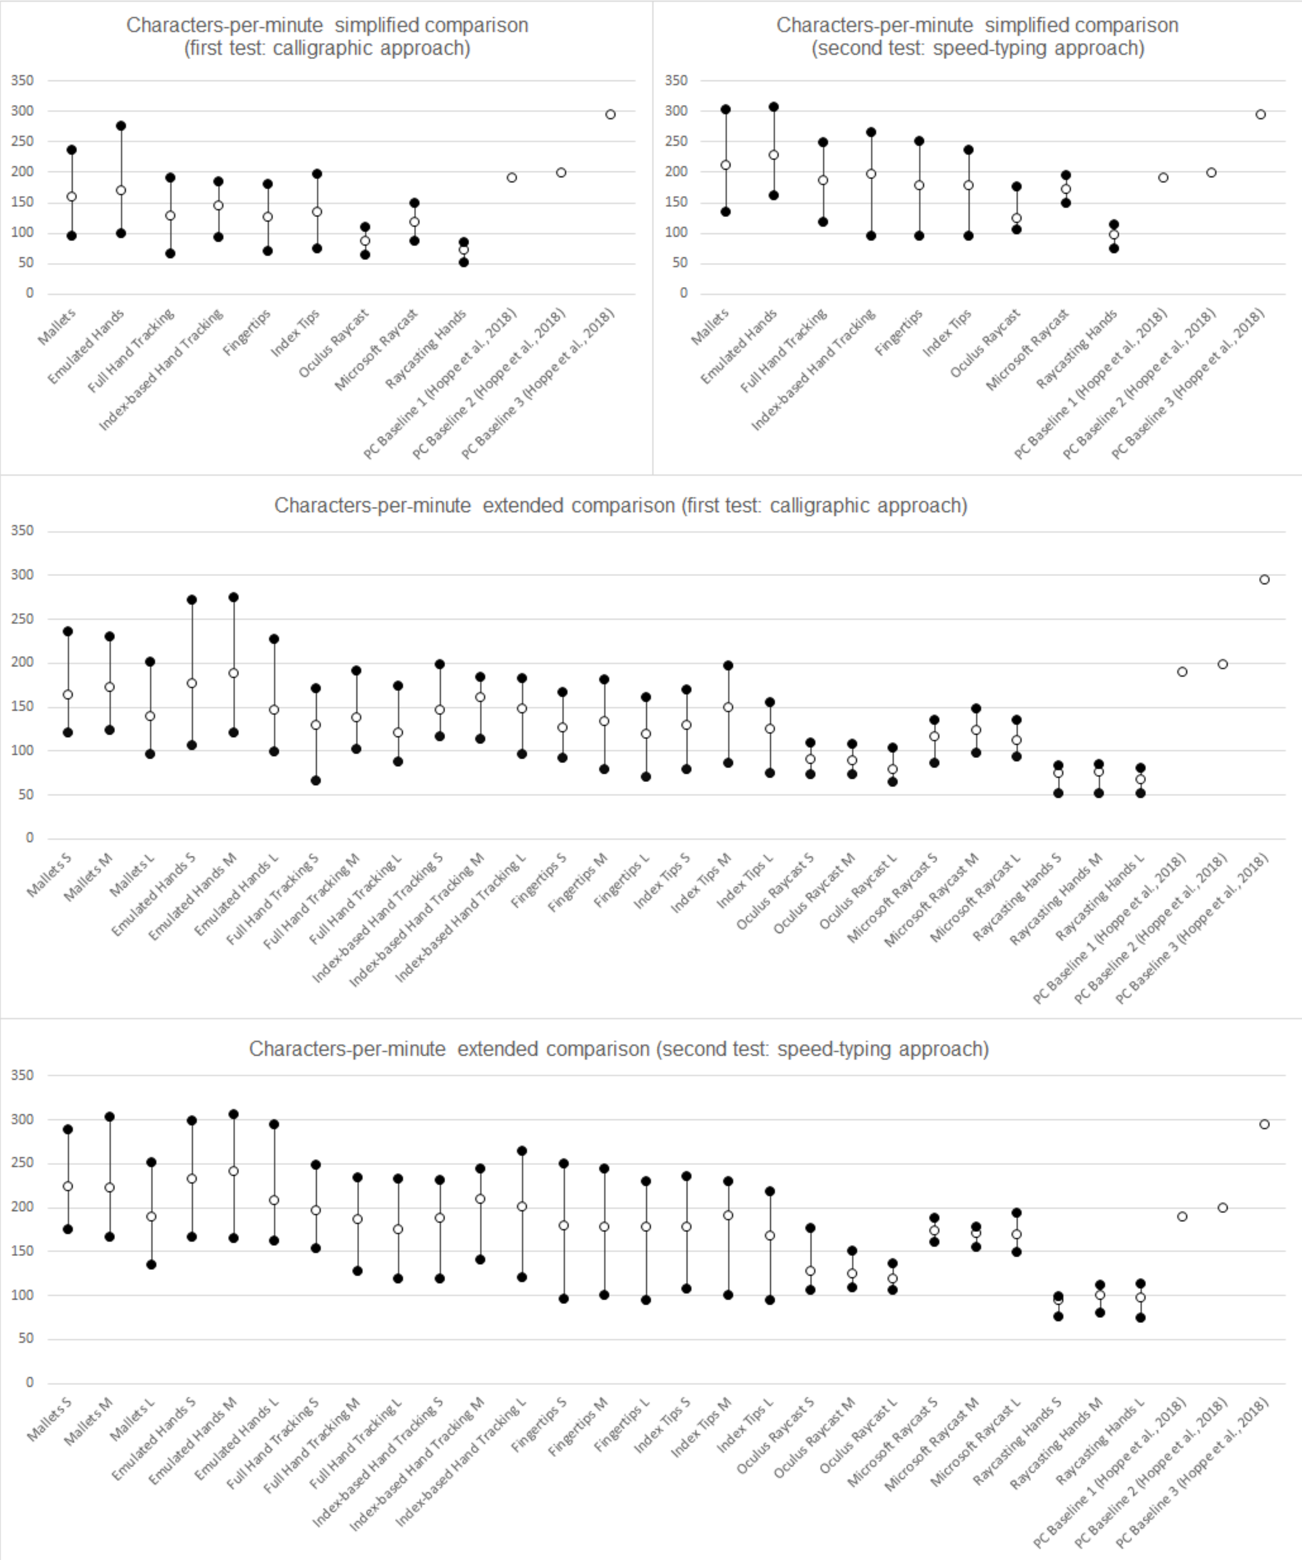

Figure 26. Performance comparison in characters-per-minute

As is shown in Figure 26, the raycast-based text input approaches consistently performed below the touch-based ones, and feedback regarding the matter suggest users prefer the latter to the former. Touch-based virtual keyboards performing above raycast-based ones collide with the results found within the work of Speicher et al. (2018), who suggested otherwise. There are many reasons which could justify this discrepancy: the sizes and models designed for their study are unknown and perhaps inadequate for some of the approaches evaluated, the proximity of the screen interface on which the keyboard is projected and which could be relevant in raycast approaches, the angle at which the touch-based keyboards are placed (the prototype presents them with an inclination of 30 degrees while the figures found within their research show a vertical keyboard being used, which is less ergonomic thus hindering the results; see Figure 27), etc.

However, the obtained results coincide with those found within the work of Trezl et al. (n.d.) despite their work not being based on a typewriting exercise. Contrasting their study with Speicher et al. (2018)'s and with this very own is proof of one of the following: either Speicher et al. (2018)'s claims are misleading, perhaps due to one of the aforementioned factors, or neither of these works contradict each other being slightly different use case scenarios. Among them, the former seems most likely due to the similarities between this study and Speicher et al. (2018)'s.

Figure 27. Speicher et al.'s work showcase a vertical keyboard being used with touch-based interaction [Extracted from Speicher et al. (2018)]

Microsoft's toolkit, which allowed for the dual raycast approach on which the simultaneous use of both controllers was possible, clearly outperformed Oculus' raycast implementation proving the controller-based dual raycast approach to work better than the controller-based single raycast one. Using one's hands as raycasters/pointers ended up being not only less precise and harder to control than both controller-based approaches but it was also the overall worst performer among the many techniques contemplated within the prototype (also being the less preferred approach between the test participants).

Besides their low typing speed, controller-based raycast techniques are undoubtedly viable text input mechanisms as the amount of typos committed through their use is not significant enough to invalidate its functionality, so it might find some use in certain applications and use case scenarios. However, it is worth mentioning that although these typos are not substantial during the first half of the test (the calligraphic approach), they are frequent during the second one (the speed-typing approach) using either of the frameworks at issue. The reason lies in how the keypress is interpreted: if the user presses the interaction button when moving the raycast, then the keypress might not register and, when speed-typing, taking the time to steady one's grip breaks the typewriting flow. This is a limitation of the frameworks at use, although the magnitude of the issue's impact upon typing performance is uncertain.

Among the touch-based virtual keyboards, the best performers were the ones using the 6-DoF controllers: the emulated hands approach and the drum-like one. Typing mistakes were uncommon with both approaches, proving their validity as text input techniques; with regards to their typing speed, the emulated hands controls outperformed the mallets, although the recollected feedback suggests that most users prefer the latter as it is less tiresome than the former, only requiring precise wrist movement to control the mallets and make them interact with the desired keys (instead of having to move both hands/arms around to type on each key).
Among the other reasons behind this preference, the “fun” of it seems to be quite predominant; that is not to say users did not enjoy the emulated hands typing experience, as it ranks seconds in terms of preference and seems to be (again, based on users’ feedback) slightly more intuitive than the mallets.

Hand tracking support proved to be more valid with touch-based keyboards than with raycast-based ones. Performance-wise they do not fall that far from the drum-like approach, and they have the advantage of not requiring any controllers whatsoever which could be useful in some applications and use case scenarios. Among its many other advantages, it also presents the user a pair of hands so alike his/her own that the gap between the virtual embodied simulation and the real one is reduced, theoretically improving immersion. However, based on the users’ feedback and the legibility of the text being input, although the tracking itself works wonders with simple interactables and gestures it is neither robust nor stable/consistent enough to avoid breaking said immersion when performing a task as precise as typewriting. These inconsistencies not only break the immersion, which is deeply tied with one’s productivity, but also increase one’s typing mistakes.
Regarding these typos, it is worth noting that only the index-based approaches allow the user to write a discernible text, proving Microsoft’s claims to be right. Such claims, which were previously detailed, suggested that the use of more than one finger makes the users prone to accidentally press non-intended buttons/keys and, as this study illustrates, that is indeed the case: sometimes, when trying to press a given key with a given finger, users ended up mistyping due to the rest of the fingers’ colliders making contact with other keys, usually placed rows below the intended one. While a physical keyboard requires a given force to properly press any of its keys, virtual ones do not; the lack of any haptic resistance is the cause of these typos, and without the use of specialized hardware that is an issue hard to address. Most users noticed that and tried to adjust the hand accordingly, trying to find some way to avoid the mistakes at issue: isolating the index to use said finger as the sole interactable, isolating the thumb with the same purpose, making a fist (or a “karate chop” pose) so the colliders’ separation was shortened and thus the hand could be used as a single collider of sorts…
All in all, it has been proven that feedback-less hand tracking support with all of its fingers is hardly a valid typewriting approach. While there was a participant which managed to use said approach with some expertise through exaggerated hand movements, the rest of the participants could not input a recognizable text. Index-based approaches, on the other hand, yielded better results overall not only regarding these typos but also performance-wise (as shown in Figure 26 and Figure 28, specially within the top graphs). Consequently, limiting the interactables to a single finger (the index) seems to be the recommended course of action (at least with regards to hand-based typewriting exercises).

As previously commented, the prototype at issue also aims to evaluate the importance of the hands' presence when typewriting within virtual environments. Figure 28 illustrates how its absence affects the typing performance, with the approaches on which only the colliders were rendered performing worse than those approaches on which the whole hand was present. On one hand, this collides with the claims found in the work of Grubert et al. (2018b) who suggested a user preference towards the more minimalistic approaches but, on the other hand, it coincides with the results found within the studies of Knierim et al. (2018).

It is worth noting that users have reported that rendering the fingertips/colliders instead of the whole hand reduces hand ownership significantly: instead of sensing the fingertips as such, they felt as if the fingers were acting as controllers/joysticks through which they control the spheres' location.

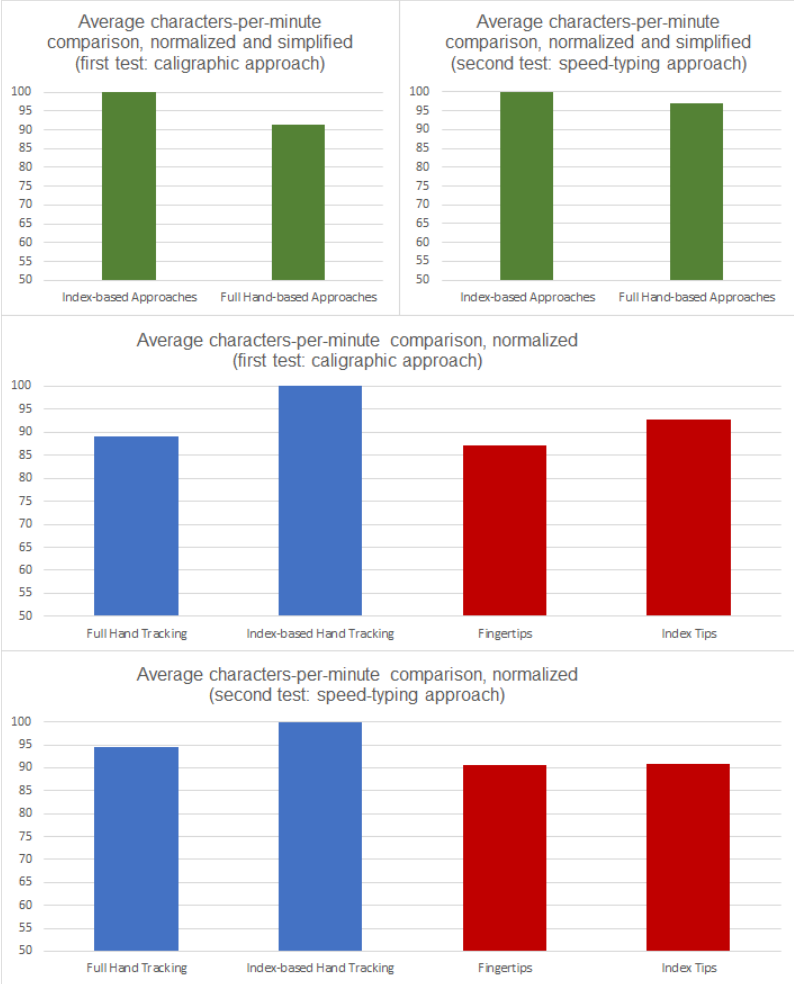

Figure 28. Percentual performance comparison between hand-tracking based approaches

This reinforces the claims found in the work of Schwind et. al (2018), where feedback from their subject pool showed that inconsistencies between the visual and haptic experiences distracted the users decreasing their virtual hand ownership, thus leading to lower performance levels on tasks such as the typewriting exercise at issue. However, their work also evaluated a fingertip-only rendering approach and they claim it achieved high limb ownership among their subject pool, whereas the results of this study and the acquired feedback suggest otherwise.

Amongst all the hand tracking based approaches explored throughout this study, the recommended one is undoubtedly that on which the whole hands are rendered but only the index fingers are functionally enabled to input text. However, it is worth mentioning that the minimalistic take on which only the index fingertips were rendered yielded acceptable results and the typewritten text was legible enough to validate its viability as a text input mechanism (despite its viability, it is still hard to recommend over an alternative).

Regarding the different sizes: Figure 29 showcases how the intermediate size outperforms the rest of them in almost all scenarios and is in fact the preferred size as suggested by the users' feedback. Nonetheless, some of the participants have shown predilection towards certain small sizes. Fact is that if the keys are closer to one another, the users have to perform a less pronounced motion to move from one key to the following, thus decreasing the time needed to move from one key to the following. This event is clearly noticeable within the speed-typing approach comparison found in Figure 29 (where its performance is on par with the intermediate size), as such approach was precisely built to study how fast the users could theoretically type on each keyboard and smaller sizes benefit the most from this take. However, in a normal typing experience (calligraphic approach) the smallest size requires precision and focus to aim, negatively affecting the typing's flow and thereby decreasing its performance, and due to that the intermediate size is the preferred one among users and the fastest within the calligraphic part of the test.

On the other side of the spectrum, the larger sizes have their keys too far apart from each other, making typos slightly less common at the cost of being slower than the intermediate size. These results prove not only that keyboards' size matter quite significantly but also that there are sizes which are either too large or too small to perform acceptably (while functional, their performance and typing experience is hindered).

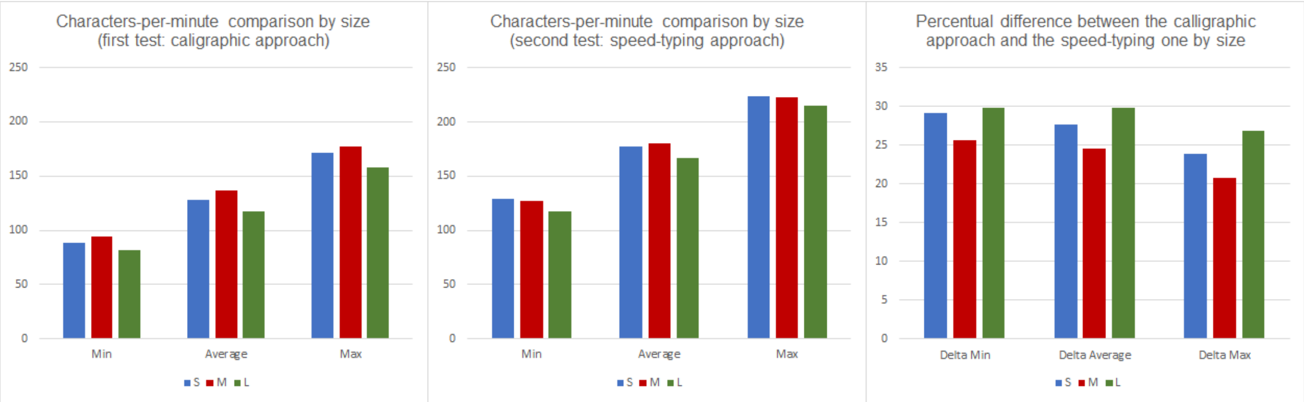

Figure 29. Performance comparison by size

The right-hand side of Figure 29 further reinforces the prevalence of the medium size when compared to the smaller and larger ones: it is, in fact, the size which presents the less difference between the calligraphic and the speed-typing approaches, meaning that typewriting correctly is easy enough for the subject pool so that users were already speed-typing to some degree. To further reinforce this, and as was already stated, feedback on the subject indicates an overall preference towards the intermediate size, which seems to be the closest one to being ideal/optimal for the typewriting experience. Despite that, it is also suggested that a middle point in between the small and medium sizes could be preferred, and further research is needed to validate such claims.

All in all, this study suggests that each approach has an optimal (intermediate) keyboard size, but whether or not there is an "universally optimal size" is yet to be studied.

Figure 30 showcases the percentual difference between the calligraphic and speed-typing approaches for each of the text input mechanisms/techniques within the prototype's scope. While the right-hand side of Figure 29 served as a means to reinforce how the less pronounced difference of the intermediate size played in its favor, Figure 30 does the same with the techniques at issue: those less affected by the change in approach are those on which the userbase is more comfortable typing. Although this does not mean the less affected techniques are the preferred ones (as using the hands as raycasters/pointers is considerably underperforming and prone to typos), it illustrates with which approaches the user is closer to speed-typing even within normal typing conditions.

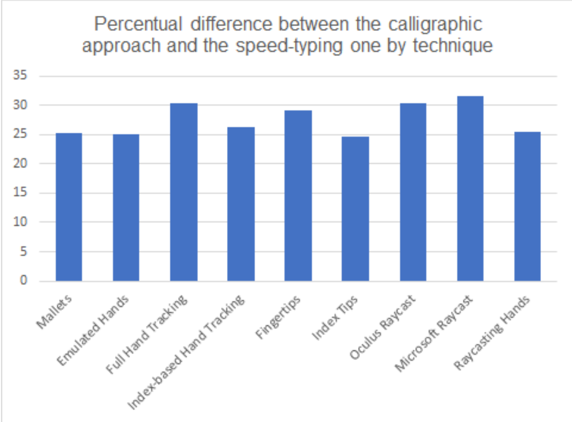

Figure 30. Percentual difference between the calligraphic and speed-typing approaches by technique

Figure 31 further explores this comparison, showcasing the percentual difference between the calligraphic and speed-typing approaches for every size and technique, and it can be observed how in almost all cases the preferred approaches (those being the mallets-based one, the emulated hands controls and the index-based touch interactivity) along with the intermediate sizes are the takes presenting a less pronounced change.
This reinforces the results obtained through Figure 26, the axis around which this chapter is built upon.

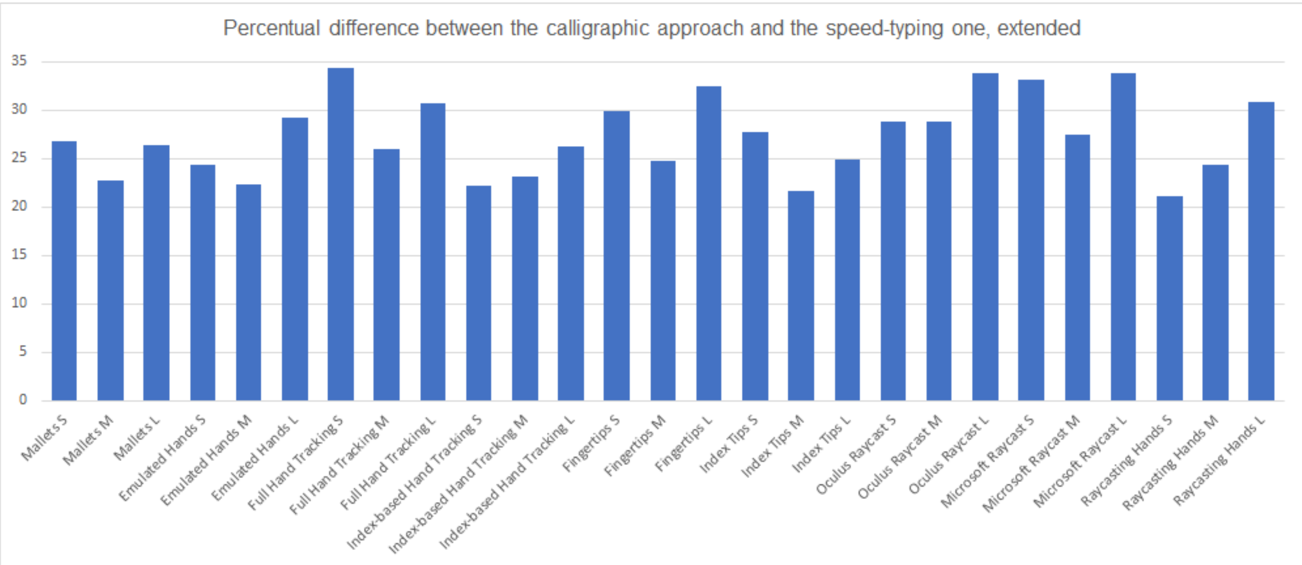

Figure 31. Extended percentual difference between the calligraphic and the speed-typing approaches

6. Conclusions and future work

VR technology has evolved significatively during the past decades both in terms of accessibility and technological feats, and this project reflects on its growth to then study the platform's viability as a productivity-focused workspace. For that, the neuroscience behind VR is overviewed: to regulate and control the body in the world effectively, the brain creates an embodied simulation of it which then uses to represent and predict sensory events: actions, concepts, emotions… VR works in a similar fashion, as its hardware components track the motion of the user and transcribes it to his/her avatar (an embodied simulation of sorts), while the underlying software adjusts the imagery to be displayed. This simulation emulates the brain's expected sensory input, and the immersion itself is highly dependent of the correlation between the VR model and the brain one (the higher, the better). Consequently, VR provides a digital place where the user can be placed and live a synthetic but immersive experience. This remarkable degree of immersion and presence has made of VR an effective clinical tool and an advanced imaginal system able to effectively induce experiences and emotions much like reality itself, which has been exploited for decades in VRE.

Scientific literature also suggests this presence is deeply tied to one's productivity, and although VR interactivity is yet evolving, the higher degree of immersion and focus could make of VR headsets better productivity drivers than the ones being used in today's society.

Indeed, VR interactivity is yet evolving and consequently so are the many text input approaches within VR. These are detailed throughout this document in the form of a meta-review where their performance and usability are studied. From the meta-review itself it is extracted that the best approach among those covered is found within the real physical keyboard's integration inside virtual environments, particularly within Logitech's approach (seen in Bovet et al., 2018) which not only seems to outperform the alternative takes but it is also the most accessible one. However, strictly virtual keyboards fit better among other use case scenarios and cannot be casted aside as the variety of use case scenarios sometimes favors one typewriting technique over other. There is not an alternative regarded as the defining standard among those covered within the meta-review: each of them has its own set of advantages and drawbacks. However, among the many virtual keyboard approaches detailed throughout the document, the most relevant ones are found within raycast-based keyboards and touch-based ones. However, the difference in their performance was uncertain due to some contradictions within the scientific literature upon which this research is built upon, so a prototype was developed so that an answer could be found.

A study was performed with the prototype at issue, proving touch-based virtual keyboards to surpass raycast-based ones in typing speed and user experience. The emulated hands controls and the drum-like mallet-based approach are the preferred techniques, and recommending one over the other highly depends on the use case scenario and user preference. Among the raycast-based approaches, the conducted research proves the dual raycast one (based on Microsoft's Mixed Reality Toolkit) to outperform the single raycast/pointing technique.

The Oculus Quest hand tracking support was also assessed, proving to be a subpar alternative in raycast-based virtual keyboards although a valid alternative within touch-based ones. Amongst all the hand tracking based approaches explored throughout this study, the recommended one is undoubtedly that on which the whole hands are rendered but only the index fingers are functionally enabled to input text.
However, hand tracking inconsistencies makes of said approach a worse alternative than those previously detailed, although its uniqueness and that of its use case scenarios could be of great remark if the tracking itself is improved upon with future HMDs.

The study at issue also evaluated the virtual size ownership levels of the approaches covered within the prototype's scope proving that, that, as opposed to what is implied in some of the works found in scientific literature, minimalistic representations of one's hand do not improve typing performance in the slightest: given the mismatch between the visual and haptic experiences, these approaches negatively affect virtual hand ownership thereby decreasing one's immersion and typing speed. However, it is worth mentioning that the minimalistic take on which only the index fingertips were rendered yielded acceptable results and the typewritten text was legible enough to validate its viability as a text input mechanism (despite its viability, it is still hard to recommend over an alternative).
The performance implications of the keyboards' size are also assessed within this study. The obtained results and recollected feedback on the subject indicate an overall preference towards a given size suggesting that there is an ideal (intermediate) one for each approach. Said feedback also suggested that such an optimal size could be found in between the small and medium sizes presented within the prototype, although further research is needed to validate such claims (whether or not there is an "universally optimal size" is yet to be researched).

A series of conditioning factors could have affected the results and further research could back up and expand this project's work. Among those conditioning factors, it is worth highlighting the number of participants evaluated (a subject pool of 10 provided this research with hours of test data, but may not be representative of a larger scope) and the habituation of said participants: although the Quest has spiked VR's popularity, all the habituation and the results were obtained through the use of a single unit.
Future work could build upon this research in many ways: evaluating more people (with a significant size, it would be ideal to study the different populations), controlling and estimating their habituation, etc.

There are many more approaches that have not been covered within this study, and future work could tackle such techniques or instigate on more novel approaches that not yet explored within scientific literature. Not only has VR evolved through the years but so have its controls: advanced emulated hand controls, real hand tracking support, eye tracking support (gaze-based typewriting was also covered by the meta-review at issue)… With improved interactivity comes an improved experience and, as such, there are many takes not yet developed nor contemplated, and given how quickly VR interactivity progresses it is hard to determine if any of the current ones will preserve its viability in the future. However, the future is built upon the research done today: now the question lies in where will said interactivity move towards.

Bibliography

- *History Of Virtual Reality* (2020, January 2). https://www.vrs.org.uk/virtual-reality/history.html
- Barnard, D. (2019, August 6). *History of VR - Timeline of Events and Tech Development.* https://virtualspeech.com/blog/history-of-vr
- Rubin, P. (2014, May 20). *The Inside Story of Oculus Rift and How Virtual Reality Became Reality.* https://www.wired.com/2014/05/oculus-rift-4/
- Kumparak, G. (2014, March 26). *A Brief History Of Oculus.* https://techcrunch.com/2014/03/26/a-brief-history-of-oculus/
- Clark, T. (2014, October 16). *How Palmer Luckey Created Oculus Rift.* https://www.smithsonianmag.com/innovation/how-palmer-luckey-created-oculus-rift-180953049/
- Purchese, R. (2013, July 11). *Happy Go Luckey: Meet the 20-year-old creator of Oculus Rift.* https://www.eurogamer.net/articles/2013-07-11-happy-go-luckey-meet-the-20-year-old-creator-of-oculus-rift
- Marlett, D. (2015, September). *The Virtual Reality of John Carmack.* https://www.dmagazine.com/publications/d-ceo/2015/september/virtual-reality-of-john-carmack/
- Oculus. (2012). *Oculus Rift: Step Into the Game by Oculus*. https://www.kickstarter.com/projects/1523379957/oculus-rift-step-into-the-game
- Dingman, H. (2014, January 16). *Valve and Oculus team up for Virtual Reality supergroup.* https://www.pcworld.com/article/2088747/valve-and-oculus-team-up-for-virtual-reality-supergroup.html
- Hamilton, I. (2016, May 24). *Inside The Growing Rift Between Valve And Oculus*. https://uploadvr.com/valve-shared-vr-oculus/
- Rajesh Desai, P., Nikhil Desai, P., Deepak Ajmera, K., & Mehta, K. (2014). *A Review Paper on Oculus Rift-A Virtual Reality Headset. International Journal of Engineering Trends and Technology, 13(*4), 175–179. https://doi.org/10.14445/22315381/ijett-v13p237
- Davies, J. (2017, August 15). *Qualcomm tries to muscle in on VR euphoria*. https://telecoms.com/484062/qualcomm-tries-to-muscle-in-on-vr-euphoria/
- Davis, S., Nesbitt, K., & Nalivaiko, E. (2014). Systematic review of cybersickness. *ACM International Conference Proceeding Series*, *02-03-December-2014*, 1–9. https://doi.org/10.1145/2677758.2677780
- Hamilton, K. (2014, March 26). *What The Heck Is Oculus Rift? A Guide To Facebook's $2 Billion Deal*. https://kotaku.com/what-the-heck-is-oculus-rift-a-guide-to-facebooks-2-1552186794
- Wozniak, P., Vauderwange, O., Mandal, A., Javahiraly, N., & Curticapean, D. (2016). Possible applications of the LEAP motion controller for more interactive simulated experiments in augmented or virtual reality. *Optics Education and Outreach IV*, *9946*(September 2016), 99460P. https://doi.org/10.1117/12.2237673
- Argelaguet, F., Hoyet, L., Trico, M., & Lécuyer, A. (2016). The role of interaction in virtual embodiment: Effects of the virtual hand representation. *Proceedings - IEEE Virtual Reality*, *2016-July*, 3–10. https://doi.org/10.1109/VR.2016.7504682
- Brown, M. (2016, July 28). *Exploring the magic behind the HTC Vive controller*. https://www.vrheads.com/exposing-magic-behind-htc-vive-controller
- Durbin, J. (2016, December 5). *Oculus Touch Review: The World's Best VR Controller*. https://uploadvr.com/oculus-touch-controllers-review/
- Coppens, A. (2017). *Merging real and virtual worlds: An analysis of the state of the art and practical evaluation of Microsoft Hololens*. http://arxiv.org/abs/1706.08096
- Kalantari, M., & Rauschnabel, P. (2018). *Exploring the Early Adopters of Augmented Reality Smart Glasses: The Case of Microsoft HoloLens*. *December*, 229–245. https://doi.org/10.1007/978-3-319-64027-3_16
- Langley, H. (2017, May 3). *Inside-out v Outside-in: How VR tracking works, and how it's going to change.* https://www.wareable.com/vr/inside-out-vs-outside-in-vr-tracking-343

- Hunt, C. (2018, July 30). *How to fix motion controller tracking loss in Windows Mixed Reality*. https://www.windowscentral.com/how-fix-motion-controller-tracking-loss-windows-mixed-reality
- Painter, L. (2020, January 17). *Oculus Go Review: The PC-free VR experience is finally here.* https://www.techadvisor.co.uk/review/gaming/oculus-go-3676570/
- Hamilton, I. (2018, April 3). *Vive Pro Hardware Review: For VR Buyers With Deep Pockets.* https://uploadvr.com/vive-pro-hardware-review/
- Hamilton, I. (2018, October 19). *Vive Wireless Adapter Review: The Best VR Of 2018 Is Too Expensive.* https://uploadvr.com/vive-wireless-adapter-pro-review/
- Stein, S. (2019, June 28). *Oculus Quest review: Facebook's newest VR headset is the best thing I've tried this year*. https://www.cnet.com/reviews/oculus-quest-review/
- Jones, R. (2019, May 21). *Oculus Quest Review: A VR game changer.* https://www.trustedreviews.com/reviews/oculus-quest
- Hesch, J., Kozminksi, A., & Linde, O. (2019, August 22). *Powered by AI: Oculus Insight*. https://ai.facebook.com/blog/powered-by-ai-oculus-insight/
- *Qualcomm Snapdragon Mobile VR Platform Powers Oculus Quest*. (2019, May 21). https://www.qualcomm.com/news/releases/2019/05/21/qualcomm-snapdragon-mobile-vr-platform-powers-oculus-quest
- Hamilton, I. (2020, February 26). *Virtual Desktop Developer Rewrites Wireless PC VR Oculus Quest Feature*. https://uploadvr.com/virtual-desktop-quest-streaming/
- Lang, B. (2019, April 3). *Oculus Quest Tracking Handles "Beat Saber" on Expert+ Difficulty*. https://www.roadtovr.com/beat-saber-oculus-quest-expert-plus-tracking-performance/
- Lang, B. (2019, May 21). *Oculus Rift S Review – A Good Choice for VR Newcomers, a Difficult Choice for VR Vets*. https://www.roadtovr.com/oculus-rift-s-review-good-choice-for-newcomers-difficult-choice-for-vr-vets/
- Hardawa, D. (2019, May 4). *Oculus Rift S review: Just another tethered VR headset*. https://www.engadget.com/2019-05-04-oculus-rift-s-review-vr.html
- Jagneaux, D. (2019, October 3). *HTC Vive Cosmos Review: A Strong Debut, But Still Not Enough*. https://uploadvr.com/htc-vive-cosmos-review/
- Pino, N. (2020, March 16). *HTC Vive Cosmos and Cosmos Elite review*. https://www.techradar.com/reviews/htc-vive-cosmos-and-cosmos-elite
- Dingman, H. (2019, June 28). *Valve Index review: The new bar for VR headsets*. https://www.pcworld.com/article/3405558/valve-index-vr-review.html
- Dingman. (2019, June 28). *Valve Index Controllers review: These revolutionary VR controllers need patience to hit full potential*. https://www.pcworld.com/article/3405879/valve-index-controllers-review.html
- Carbotte, K. (2020, February 25). *Valve Index VR Headset and Controllers Review: A New Champion*. https://www.tomshardware.com/reviews/valve-index-vr-headset-controllers,6205.html
- Heaney, D. (2019, March 18). *HTC Brings Finger Tracking To Vive and Vive Pro*. https://uploadvr.com/htc-vive-finger-tracking/
- Hardawar, D. (2019, September 27). *Oculus Quest's hand tracking is a new level of VR immersion | Engadget*. https://www.engadget.com/2019-09-27-oculus-quest-hand-tracking-hands-on.html
- Carlton, B. (2019, October 1). *Hands-On: Oculus Quest Hand Tracking Feels Great, But It's Not Perfect - VRScout*. https://vrscout.com/news/hands-on-oculus-quest-hand-tracking/
- Baker, H. (2019, November 18). *Oculus Link Beta: Guide To The Oculus Quest's PC VR Headset Makeover*. https://uploadvr.com/oculus-link-beta/
- Nagata, H., Mikami, D., Miyashita, H., Wakayama, K., & Takada, H. (2017). Virtual reality technologies in telecommunication services. *Journal of Information Processing*, *25*, 142–152. https://doi.org/10.2197/ipsjjip.25.142

- Martín-Gutiérrez, J., Mora, C. E., Añorbe-Díaz, B., & González-Marrero, A. (2017). Virtual technologies trends in education. *Eurasia Journal of Mathematics, Science and Technology Education*, *13*(2), 469–486. https://doi.org/10.12973/eurasia.2017.00626a
- Potkonjak, V., Gardner, M., Callaghan, V., Mattila, P., Guetl, C., Petrović, V. M., & Jovanović, K. (2016). Virtual laboratories for education in science, technology, and engineering: A review. In *Computers and Education* (Vol. 95). https://doi.org/10.1016/j.compedu.2016.02.002
- Berg, L. P., & Vance, J. M. (2017). Industry use of virtual reality in product design and manufacturing: a survey. *Virtual Reality*, *21*(1). https://doi.org/10.1007/s10055-016-0293-9
- Lawson, G., Salanitri, D., & Waterfield, B. (2015). VR processes in the automotive industry. *Lecture Notes in Computer Science (Including Subseries Lecture Notes in Artificial Intelligence and Lecture Notes in Bioinformatics)*, *9171*, 208–217. https://doi.org/10.1007/978-3-319-21006-3_21
- Lele, A. (2013). Virtual reality and its military utility. *Journal of Ambient Intelligence and Humanized Computing*, *4*(1), 17–26. https://doi.org/10.1007/s12652-011-0052-4
- Székely, G., & Satava, R. M. (1999). Virtual reality in medicine. Interview by Judy Jones. In *BMJ (Clinical research ed.)* (Vol. 319, Issue 7220). https://doi.org/10.1136/bmj.319.7220.1305
- Riva, G. (2003). *Applications of Virtual Environments in Medicine * 2 . The Role of VR in Health Care*. 524–534. https://doi.org/10.1267/METH03050524
- Kathuria, A., Gefen, D., Boulos, M. K., Claudio, P., & Maddalena, P. (2014). Number 1 Overview. In *Virtual Reality in Medicine 1 Lantern* (Vol. 7, Issue 1). http://jvwresearch.org
- Li, L., Yu, F., Shi, D., Shi, J., Tian, Z., Yang, J., Wang, X., & Jiang, Q. (2017). Application of virtual reality technology in clinical medicine. In *American Journal of Translational Research* (Vol. 9, Issue 9). www.ajtr.org/ISSN:1943-8141/AJTR0055713
- Mazurek, J., Kiper, P., Cieślik, B., Rutkowski, S., Mehlich, K., Turolla, A., & Szczepańska-Gieracha, J. (2019). Virtual reality in medicine: A brief overview and future research directions. *Human Movement*, *20*(3), 16–22. https://doi.org/10.5114/hm.2019.83529
- Riva, G., Wiederhold, B. K., Chirico, A., Di Lernia, D., Mantovani, F., & Gaggioli, A. (2018). Brain and virtual reality: What do they have in common and how to exploit their potential. In *Annual Review of CyberTherapy and Telemedicine* (Vol. 2018, Issue 16). https://www.researchgate.net/publication/332370861
- Riva, G., Wiederhold, B. K., & Mantovani, F. (2019). Neuroscience of Virtual Reality: From Virtual Exposure to Embodied Medicine. *Cyberpsychology, Behavior, and Social Networking*, *22*(1), 82–96. https://doi.org/10.1089/cyber.2017.29099.gri
- Riva, G. (2005). Virtual reality in psychotherapy: Review. *Cyberpsychology and Behavior*, *8*(3), 220–230. https://doi.org/10.1089/cpb.2005.8.220
- Schuemie, M. J., Van der Straaten, P., Krijn, M., & Van der Mast, C. A. P. G. (2001). Research on presence in virtual reality: A survey. *Cyberpsychology and Behavior*, *4*(2), 183–201. https://doi.org/10.1089/109493101300117884
- Lee, K. M. (2004). Presence, explicated. *Communication Theory*, *14*(1), 27–50. https://doi.org/10.1111/j.1468-2885.2004.tb00302.x
- Murray, C. D., Fox, J., & Pettifer, S. (2007). Absorption, dissociation, locus of control and presence in virtual reality. *Computers in Human Behavior*, *23*(3), 1347–1354. https://doi.org/10.1016/j.chb.2004.12.010
- Sanchez-vives, M. V, & Slater, M. (2005). From Presence Towards Consciousness. *Nature Reviews Neuroscience*, *6*(10), 332. http://citeseerx.ist.psu.edu/viewdoc/download?doi=10.1.1.99.5596&rep=rep1&type=pdf
- Price, M., & Anderson, P. (2007). The role of presence in virtual reality exposure therapy. *Journal of Anxiety Disorders*, *21*(5), 742–751. https://doi.org/10.1016/j.janxdis.2006.11.002
- Syverson, C. (2011). What determines productivity. *Journal of Economic Literature*, *49*(2), 326–365. https://doi.org/10.1257/jel.49.2.326
- Csikszentmihalyi, M. (1975). *Beyond boredom and anxiety*. Jossey-Bass.

- Csikszentmihalyi, M. (1990). *Optimal experience: Psychological studies of flow in consciousness*. Cambridge University Press.
- Kotler, S. (2014). *The rise of superman: Decoding the science of ultimate human performance*. Houghton Mifflin Harcourt.
- Schaper, E. (1978). Fiction and the suspension of disbelief. *The British Journal of Aesthetics*, *18*(1), 31–44. https://doi.org/10.1093/bjaesthetics/18.1.31
- Duffy, B. R., & Zawieska, K. (2012). Suspension of disbelief in social robotics. *Proceedings - IEEE International Workshop on Robot and Human Interactive Communication*, *May*, 484–489. https://doi.org/10.1109/ROMAN.2012.6343798
- Brown, D. W. (2012). *The suspension of disbelief in videogames*. *August*, 1–251. http://bura.brunel.ac.uk/handle/2438/7457
- General, A., Silva, B. Da, Esteves, D., Halleran, M., & Liut, M. (2015). A Comparative Analysis Between the Mouse , Trackpad and the Leap Motion. In *HCI International Conference*.
- Santosa, S., & Wigdor, D. (2013). A field study of multi-device workflows in distributed workspaces. *UbiComp 2013 - Proceedings of the 2013 ACM International Joint Conference on Pervasive and Ubiquitous Computing*, 63–72. https://doi.org/10.1145/2493432.2493476
- Czerwinski, M., Smith, G., Regan, T., Meyers, B., Robertson, G., & Starkweather, G. (2003). Toward Characterizing the Productivity Benefits of Very Large Displays. In *Interact* (Vol. 3). http://research.microsoft.com/en-us/um/redmond/groups/cue/publications/interact2003-productivitylargedisplays.pdf
- Anderson, J. A., Colvin, J., Tobler, N., Rennert, N., Mulcahy, R., Connery, C., & Mitsubishi, N. /. (2003). *Productivity and Multi-Screen Displays A Comparison of Single Monitor, Multiple Monitor, and Multiple Monitor with Hydravision ® Computer Displays over Simulated Office Tasks across Performance and Usability Productivity and Multi-Screen Displays*. http://cdmbuntu.lib.utah.edu/utils/getfile/collection/ir-eua/id/3137/filename/3138.pdf
- Trezl, J., Dittrich, F., & Bullinger, A. C. (n.d.). *Capturing quantitative user feedback using virtual questionnaires in virtual reality*. 37878.
- Speicher, M., Feit, A. M., Ziegler, P., & Krüger, A. (2018). Selection-based text entry In Virtual Reality. *Conference on Human Factors in Computing Systems - Proceedings*, *2018-April*, 1–13. https://doi.org/10.1145/3173574.3174221
- Krichenbauer, M., Yamamoto, G., Taketom, T., Sandor, C., & Kato, H. (2018). Augmented Reality versus Virtual Reality for 3D Object Manipulation. *IEEE Transactions on Visualization and Computer Graphics*, *24*(2), 1038–1048. https://doi.org/10.1109/TVCG.2017.2658570
- Hoppe, A. H., Otto, L., van de Camp, F., Stiefelhagen, R., & Unmüßig, G. (2018). qVRty: Virtual keyboard with a haptic, real-world representation. In *Communications in Computer and Information Science* (Vol. 851). https://doi.org/10.1007/978-3-319-92279-9_36
- Lin, J. W., Han, P. H., Lee, J. Y., Chen, Y. S., Chang, T. W., Chen, K. W., & Hung, Y. P. (2017). Visualizing the keyboard in virtual reality for enhancing immersive experience. *ACM SIGGRAPH 2017 Posters, SIGGRAPH 2017*, *1*. https://doi.org/10.1145/3102163.3102175
- Merriaux, P., Dupuis, Y., Boutteau, R., Vasseur, P., & Savatier, X. (2017). A study of vicon system positioning performance. *Sensors (Switzerland)*, *17*(7), 1–18. https://doi.org/10.3390/s17071591
- Wakkary, R., Hatala, M., Lovell, R., & Droumeva, M. (2005). An ambient intelligence platform for physical play. *Proceedings of the 13th ACM International Conference on Multimedia, MM 2005*, *January*, 764–773. https://doi.org/10.1145/1101149.1101313
- Knierim, P., Schwind, V., Feit, A. M., Nieuwenhuizen, F., & Henze, N. (2018). Physical keyboards in Virtual reality: Analysis of typing performance and effects of avatar hands. *Conference on Human Factors in Computing Systems - Proceedings*, *2018-April*, 1–9. https://doi.org/10.1145/3173574.3173919
- Grubert, J., Witzani, L., Ofek, E., Pahud, M., Kranz, M., & Kristensson, P. O. (2018a). Text Entry in Immersive Head-Mounted Display-Based Virtual Reality Using Standard Keyboards. In *25th IEEE Conference on Virtual Reality and 3D User Interfaces, VR 2018 - Proceedings*. https://doi.org/10.1109/VR.2018.8446059
- Furtado, J. S., Liu, H. H. T., Lai, G., Lacheray, H., & Desouza-Coelho, J. (2019). Comparative analysis of OptiTrack motion capture systems. *Lecture Notes in Mechanical Engineering*, 15–31. https://doi.org/10.1007/978-3-030-17369-2_2

- Schwind, V., Knierim, P., Tasci, C., Franczak, P., Haas, N., & Henze, N. (2017a). “These are not my hands!”: Effect of gender on the perception of avatar hands in virtual reality. *Conference on Human Factors in Computing Systems - Proceedings*, *2017-May*, 1577–1582. https://doi.org/10.1145/3025453.3025602
- Schwind, V., Lin, L., Di Luca, M., Jörg, S., & Hillis, J. (2018). Touch with foreign hands: The effect of virtual hand appearance on visual-haptic integration. *Proceedings - SAP 2018: ACM Symposium on Applied Perception*. https://doi.org/10.1145/3225153.3225158
- Grubert, J., Witzani, L., Ofek, E., Pahud, M., Kranz, M., & Kristensson, P. O. (2018b). Effects of Hand Representations for Typing in Virtual Reality. In *25th IEEE Conference on Virtual Reality and 3D User Interfaces, VR 2018 - Proceedings*. https://doi.org/10.1109/VR.2018.8446250
- Walker, J., Li, B., Vertanen, K., & Kuhl, S. (2017). Efficient typing on a visually occluded physical keyboard. *Conference on Human Factors in Computing Systems - Proceedings*, *2017-May*, 5457–5461. https://doi.org/10.1145/3025453.3025783
- Mcgill, M., Boland, D., Murray-Smith, R., & Brewster, S. (2015). A dose of reality: Overcoming usability challenges in VR head-mounted displays. *Conference on Human Factors in Computing Systems - Proceedings*, *2015-April*, 2143–2152. https://doi.org/10.1145/2702123.2702382
- Vertanen, K., Memmi, H., Emge, J., Reyal, S., & Kristensson, P. O. (2015). Velocitap: Investigating fast mobile text entry using sentence-based decoding of touchscreen keyboard input. *Conference on Human Factors in Computing Systems - Proceedings*, *2015-April*, 659–668. https://doi.org/10.1145/2702123.2702135
- Bovet, S., Kehoe, A., Crowley, K., Curran, N., Gutierrez, M., Meisser, M., Sullivan, D. O., & Rouvinez, T. (2018). Using Traditional Keyboards in VR: SteamVR Developer Kit and Pilot Game User Study. *2018 IEEE Games, Entertainment, Media Conference, GEM 2018*, 132–135. https://doi.org/10.1109/GEM.2018.8516449
- Dube, T. J., & Arif, A. S. (2019). Text Entry in Virtual Reality: A Comprehensive Review of the Literature. *Lecture Notes in Computer Science (Including Subseries Lecture Notes in Artificial Intelligence and Lecture Notes in Bioinformatics)*, *11567 LNCS*, 419–437. https://doi.org/10.1007/978-3-030-22643-5_33
- Yu, C., Gu, Y., Yang, Z., Yi, X., Luo, H., & Shi, Y. (2017). Tap, dwell or gesture?: Exploring head-based text entry techniques for HMDs. *Conference on Human Factors in Computing Systems - Proceedings*, *2017-May*, 4479–4488. https://doi.org/10.1145/3025453.3025964
- Ishii, A., Adachi, T., Shima, K., Nakamae, S., Shizuki, B., & Takahashi, S. (2017). FistPointer. *Proceedings of the 2017 CHI Conference Extended Abstracts on Human Factors in Computing Systems*, 474–474. https://doi.org/10.1145/3027063.3049795
- Rajanna, V., & Hansen, J. P. (2018). Gaze typing in virtual reality: Impact of keyboard design, selection method, and motion. *Eye Tracking Research and Applications Symposium (ETRA)*. https://doi.org/10.1145/3204493.3204541
- Ma, X., Yao, Z., Wang, Y., Pei, W., & Chen, H. (2018). Combining brain-computer interface and eye tracking for high-speed text entry in virtual reality. *International Conference on Intelligent User Interfaces, Proceedings IUI*, 263–267. https://doi.org/10.1145/3172944.3172988
- Gugenheimer, J., Dobbelstein, D., Winkler, C., Haas, G., & Rukzio, E. (2016). FaceTouch: Touch interaction for mobile virtual reality. *Conference on Human Factors in Computing Systems - Proceedings*, *07-12-May-2016*, 3679–3682. https://doi.org/10.1145/2851581.2890242
- Kim, Y. R., & Kim, G. J. (2017). HoVR-Type: Smartphone as a typing interface in VR using hovering. *2017 IEEE International Conference on Consumer Electronics, ICCE 2017*, *February*, 200–203. https://doi.org/10.1109/ICCE.2017.7889285
- Mehring, C., Kuester, F., Singh, K. D., & Chen, M. (2004). KITTY: Keyboard independent touch typing in VR. *Proceedings - Virtual Reality Annual International Symposium*, *April*, 243–244. https://doi.org/10.1109/VR.2004.1310090
- Mourouzis, A., Kilintzis, V., Chouvarda, I., & Maglaveras, N. (2014). Using the leap 3D motion controller for implementing text entry for All. *Proceedings of the European, Mediterranean and Middle Eastern Conference on Information Systems, EMCIS 2014*, *2014*, 1–18.

- Markussen, A., Jakobsen, M. R., & Hornbæk, K. (2013). Selection-Based Mid-Air Text Entry on Large Displays Design Space for Mid-Air Text Entry. *Interact*, 401–418.
- Arif, A. S., & Stuerzlinger, W. (2009). Analysis of text entry performance metrics. *TIC-STH'09: 2009 IEEE Toronto International Conference - Science and Technology for Humanity*, 100–105. https://doi.org/10.1109/TIC-STH.2009.5444533
- Bowman, D. A., Rhoton, C. J., & Pinho, M. S. (2002). Text Input Techniques for Immersive Virtual Environments: An Empirical Comparison. *Proceedings of the Human Factors and Ergonomics Society Annual Meeting*, *46*(26), 2154–2158. https://doi.org/10.1177/154193120204602611
- González, G., Molina, J. P., García, A. S., Martínez, D., & González, P. (2009). Evaluation of text input techniques in immersive virtual environments. In *New Trends on Human-Computer Interaction: Research, Development, New Tools and Methods* (pp. 109–118). Springer London. https://doi.org/10.1007/978-1-84882-352-5_11
- Yu, D., Fan, K., Zhang, H., Monteiro, D., Xu, W., & Liang, H. N. (2018). PizzaText: Text entry for virtual reality systems using dual thumbsticks. *IEEE Transactions on Visualization and Computer Graphics*, *24*(11), 2927–2935. https://doi.org/10.1109/TVCG.2018.2868581
- Microsoft (2020). MixedRealityToolkit-Unity. *GitHub repository*. https://github.com/microsoft/MixedRealityToolkit-Unity
- Eric Provencher (2020). MRTK-Quest. *GitHub repository*. https://github.com/provencher/MRTK-Quest
- *What is the Mixed Reality Toolkit | Mixed Reality Toolkit Documentation*. (n.d.). https://microsoft.github.io/MixedRealityToolkit-Unity/README.html
- MacIntyre, B., & Smith, T. F. (2018). Thoughts on the Future of WebXR and the Immersive Web. *Adjunct Proceedings - 2018 IEEE International Symposium on Mixed and Augmented Reality, ISMAR-Adjunct 2018*, 338–342. https://doi.org/10.1109/ISMAR-Adjunct.2018.00099
- Dibbern, C., Uhr, M., Krupke, D., & Steinicke, F. (2018). Can WebVR further the adoption of Virtual Reality? *Mensch Und Computer 2018 - Usability Professionals*, *September 2018*, 377–384. https://doi.org/10.18420/muc2018-up-0249
- *Hand Tracking*. (n.d.). https://developer.oculus.com/documentation/unity/unity-handtracking
- *Blog - Oculus Touch and Finger Stuff Part 1*. (n.d.). https://www.rgbschemes.com/blog/oculus-touch-and-finger-stuff-part-1/
- *Blog - Oculus Touch and Finger Stuff Part 2*. (n.d.). https://www.rgbschemes.com/blog/oculus-touch-and-finger-stuff-part-2/
- Dilmer Valecillos (2020). OculusQuestHandTackingPhysics. *Github repository.* https://github.com/dilmerv/OculusQuestHandTrackingPhysics
- The Campfire Union (2019). VRKeys. *GitHub repository.* https://github.com/campfireunion/VRKeys
- Vosinakis, S., & Koutsabasis, P. (2018). Evaluation of visual feedback techniques for virtual grasping with bare hands using Leap Motion and Oculus Rift. *Virtual Reality, 22*(1), 47–62. https://doi.org/10.1007/s10055-017-0313-4
- Schwind, V., Knierim, P., Chuang, L., & Henze, N. (2017b). “Where's pinky?”: The effects of a reduced number of fingers in virtual reality. *CHI PLAY 2017 - Proceedings of the Annual Symposium on Computer-Human Interaction in Play*, 507–515. https://doi.org/10.1145/3116595.3116596
- Ghosh, S., & Kristensson, P. O. (2017). *Neural Networks for Text Correction and Completion in Keyboard Decoding*. *14*(8), 1–14. http://arxiv.org/abs/1709.06429