

Continuous-Motion Text Input in Virtual Reality

Janis G. Jimenez, Jürgen P. Schulze; University of California San Diego; La Jolla, CA, USA

Abstract

Text input in virtual reality is a problem that does not currently have a widely accepted standard method. As VR headsets have become more commonplace, text input has also become more important. Using a physical keyboard is not possible with a head-mounted display that blocks the users visual field. The two most popular solutions for text input in VR today are a virtual keyboard interfacing with VR controllers and voice recognition. However, they either require a handheld controller or a quiet environment. 3D-tracked controllers with a virtual keyboard can simulate a real keyboard to an extent, but they suffer from a lack of tactile feedback that makes typing slow and unintuitive. A more intuitive solution is a Swype or SwiftKey-like algorithm, where the path the users finger travels is used as input, as opposed to individually pressing each key. We implemented a prototype for the Oculus Rift with a Leap Motion controller on it that combines a novel continuous-motion text input method with hand gestures to demonstrate an all-purpose, intuitive method of text input. We compare it to state-of-the-art VR keyboard input with a virtual keyboard, as well as a head-directed input method.

Introduction

The main barrier to finding a comfortable and reliable way of taking in text input is the fact that virtual reality headsets block the users vision. Some headsets are add-ons to existing hardware, so they have limited to no options for text input when they are used in VR mode. This means they might not even come with a controller, choosing instead to rely on gaze selection (looking at a specific spot long enough to select it) along with any buttons that are already on the headset itself. Examples of this are the Samsung Gear VR and the Google Cardboard.

Those who do have a pre-existing controller follow the Playstation VRs example. Their solution is to use either the controllers that are already being used for the Playstation 4 or their Playstation Move motion controllers. Their system for typing messages is therefore the same as it has been for most consoles: with a controller, the user uses the thumbstick (or touchpad) to slide over to the letter they want. With motion controls, the controller itself becomes a cursor in 2D space that the user moves over to the desired letter. Each letter must be individually confirmed with the push of a button.

Since these headsets are mainly for playing video games, there is no pressing demand to accommodate for text input, especially when the user has ways to input text on the existing system with their headset off. However, there is a growing demand when it comes to staying connected to ones phone while in VR. Some headsets allow the user to connect their phone through Bluetooth (or the headset display is their phone). The HTC Vive, for example, allows the user to see text messages sent by their contacts and call them with the built-in microphone. If its Vive for Android, the user can also send Quick Replies, which are pre-made mes-

sages [21]. Generally, the user is expected to take off their VR headset to type a custom message.

For headsets hooked up to a computer, the straightforward solution is to use the keyboard. But again, these headsets completely block the users vision. The Oculus Rift, HTC Vive, and Playstation VR all prefer to use their own controllers because they will always be in the users hand. This mostly removes the issue of having to search around in the dark for the right button once the user becomes familiar with the controllers button layout, and allows them to move around instead of having to sit at their desk. However, there are some independent projects that accepted this constraint and have made prototypes involving a physical keyboard.

Previous Work

Lachlan Sleight created a 3D replica of their keyboard and used the Leap Motion controller to track where their hands are, so the user can see both their hands and the keyboard through their VR headset [24]. It is an ideal solution when sitting in ones office and it is easy to learn, which will result in a fairly high word-per-minute performance compared to the one-letter-at-a-time input of most controllers, but not for general use.

In lieu of having to use a physical keyboard, some have opted to use a digital one. A big hurdle that this option has to overcome is the lack of tactile feedback, because there is no tangible controller to assure the user that their input was received. If there is no physical surface to press against, the user can't intuitively know when they are pressing something. For some, the solution is to create hardware that can simulate that feedback. A group of researchers from Rice University created special gloves that use bladders that expand and contract to simulate the sensation of touch in VR [30]. However, it is still in the prototype phase, so how well it performs with a 3D keyboard remains to be seen. Some have made prototypes that provide ample visual and audio feedback to try to make up for the lack of tactile feedback.

Researchers and VR enthusiasts have created a variety of virtual keyboard designs that can interact with hand-tracking sensors (such as the Leap Motion Controller or the Microsoft Kinect): One has the appearance of a regular QWERTY keyboard on a mostly flat surface, with keys that move down as the users fingers press them [5]. Another has the keyboard directly facing the user, with keys lighting up and making a clicking sound when the users finger is near enough to them [31]. Some may also add in a hover check, where the user keeps their finger pressed on a key for a short amount of time to confirm that it is the desired key [12]. These projects do free up the users hands by not requiring a controller, but the lack of tactile feedback and the imperfect tracking mechanisms of the sensors make typing slow and error-prone.

There are other ways to type on a virtual keyboard, however. The difficulty in selecting keys partially stems from a depth-sensing issue, so a good way to circumvent this is to not rely on

depth at all. Some keyboards rely exclusively on the aforementioned hover check. This is how the Myo armband, a motion control device, allows the user to search the web [18]. But there is still a more efficient way that does not force the user to wait for every letter they want to type.

The Swype method of typing also does not need depth to work [25]. On a phone, typing with Swype works as follows: the user places their finger on the first letter of the word, then drags their finger to each required letter without stopping. They only lift up their finger once the word is finished. Swype will then guess what word they were typing. If it guesses wrong, multiple auto-complete options would be available to choose from. A phone knows when to start and stop tracking input by sensing when the users finger is touching the screen, but in VR, it could be controlled by a button on the headset or a hand gesture. There have been tests using the Leap Motion with a Swype keyboard [33], and there have been rumors of Swype looking into virtual reality [22], but to our knowledge nobody has implemented this yet. For our prototype, we chose to create a VR application with hand tracking and a 3D keyboard like the examples mentioned above, but we combined it with a Swype-like method of text input to demonstrate how this method works in virtual reality.

Hardware and Software

Our prototype was created in Unity version 5.5.1f3 and written in C#. When deciding on the developer environment, we chose Unity over Unreal Engine or Lumberyard for a number of reasons. A large majority of developers use Unity because of how quickly they can create a prototype [10]. Many of the projects mentioned in the Related Work section used Unity to develop their prototypes.

While there are many commercially available headsets, we focused on choosing a headset that is widely used and compatible with game-developing programs. The Oculus Rift and HTC Vive are two of the most popular PC-based headsets, and are compatible with a variety of devices. Unity is mainly compatible with the Oculus Rift, Gear VR, the HTC Vive, and the HoloLens, though it continues to add support for more types of headsets [26]. In this project we use an Oculus DK2 head-mounted display.

As for a hand-tracking device, the two main competitors in use are the Leap Motion Controller and the second-generation Kinect (because the first-generation Kinect cannot keep track of individual fingers). Both are compatible with the Oculus Rift, but the library the Leap Motion Controller comes with is much easier to work with when it comes to tracking individual fingers, while the Kinect V2 requires finding a third-party library or algorithm [Pte16]. The camera placement is also more convenient since the Leap can be placed directly onto the headset, so we ultimately chose the Leap Motion Controller, along with Leap Motion SDK version 2.3.1.

We developed on a PC with a Windows 10 operating system. It is the recommended OS by Oculus because they have stopped development on OS X and Linux for the time being [3]. The Leap Motion Controller SDK is also only available for Windows and Android [16].

Backend

For the Swype-like algorithm, Swype, Swiftkey, and Google Keyboard have no publicly available source code online, though

Android phones automatically have a Swype keyboard. However, the Android keyboard has a limited API to work with that is targeted towards touchscreens, and it is quite difficult to interact with through scripting, so we decided against directly using this keyboard and instead made our own keyboard with Unity. There are many alternative open source projects available online, but most if not all lack the complexity that the commercial apps have under the hood, and our own backend would be easier to debug and build upon. For our prototype, we focused on creating an intuitive, responsive keyboard and determining effective criteria for filtering user input in the front-end, and a fast, accurate way to compare a given word to a dictionary and auto-correct it on the back-end.

First, we needed a way to store words into a data structure that could be built and traversed in polynomial time. A hash table has quick $O(1)$ access, but searching for the right word could lead to $O(n)$ search time, which would be ill-advised for a large dictionary. Instead, we mainly looked at a ternary trie and a BK tree. Both essentially have $O(\log n)$ search time, so we looked at how the words were organized in each structure. A ternary trie has a structure similar to both Binary Search Trees and Tries. Like a trie, it would store words character-by-character, requiring w nodes for a word of length w . A trie, however, would have all 26 letters as its children at every letter, so it would be very space-inefficient with a large dictionary. Instead, a ternary trie is more like a binary tree, but with up to three children at every level instead of two. The left child has a value less than the current node, and the right child has a value greater than it. The middle child is for a node with a value that lies between the two other children.

The main issue with this structure was that it was suited more for spell-checking and auto-complete, not auto-correct. The first two rely on a sequential procedure, while auto-correct has to cover many more cases. When a user types a word, they could accidentally transpose two letters, or forget a letter in the middle of the word. They could even do this multiple times within the same word. A ternary trie orders the letters of words sequentially, though it can store a distance that would allow it to compare just how different two words are from each other. With some tolerance to check against, it could still traverse the data structure and find all words that fit within the tolerance, but this means that it would have to compare every substring it finds in its path, not just every word.

We ultimately created a BK-tree [11, 32], a tree structure proposed by Burkhard and Keller that relies on sorting words by the distance between them. This means that each node is a word,



Figure 1. The Oculus Rift DK2 headset with the Leap Motion Controller attached. [17]

not a letter. First, it picks an arbitrary node as its root. Then, to add a node, it finds the distance between new node and the root of the tree. It then traverses its children until it finds another node that has the same distance. If it does, then it traverses that nodes children, and continues doing so recursively until it cannot find a node with the same distance. When it reaches this point, the new node is added as a child to the node it stopped on. Building the tree takes about $O(n \log n)$: it takes $O(\log n)$ to place a single node into the tree, and it places $O(n)$ nodes.




Figure 2. An example of a BK-tree sorted by the Levenshtein Distance between the words.

The distance used by the BK-tree is the Levenshtein Distance, [29] a number representing how many changes need to be made to a word (deletions, insertions, or mismatched letters) in order to turn it into another word. We relied on this metric because it is a simple, discrete way to keep track of differences between strings that is commonly utilized for a BK-tree (and is sometimes used in ternary tries as well). For example, to change the word cat to gate, the c must be changed to a g and an e must be added to the end. Two changes must be made, so the distance between cat and gate is 2. The Levenshtein Distance code works by maintaining a matrix of integers holding the distance between substrings of the two words. It iterates through both words, letter by letter, and adds 1 to the count in the current index. Once it finishes iterating, it returns the value at the last index.

The original Levenshtein Distance equation does not take transpositions into account (cat and chat would have a distance of 3 without transpositions, and 1 with transpositions), so we coded in the transposition check made by the Damerau-Levenshtein Distance version [28]. In addition to this, we added an extra check: a keyboard neighbors check. When typing, it is not uncommon for the user to accidentally choose a letter neighboring the letter they actually wanted. To account for this, we have a dictionary storing the surrounding neighbors for every letter and symbol on the keyboard. Each time it makes a comparison between the two words, it will check if the two letters are neighbors of each other. If so, it will add 1 to the distance; otherwise, it will add 2. To balance this change, every increment that is not between two neighboring letters will be doubled. This means that each deletion, insertion, and non-neighbor mismatch will add 2 to the distance instead of 1.

The next step is using the tree to come up with auto-correct candidates. Once the user has finished typing a word, it traverses the tree to find candidates within a certain tolerance threshold. This $O(\log n)$ search process is as follows: it starts with a queue containing the root of the tree and finds the distance between the new word and the root word. If it falls within the tolerance threshold, it is added to a list of possible auto-correct candidates. Then it looks through its children and adds them to the queue if they

fall between the range of the previously calculated distance plus or minus the tolerance. The current node is then dequeued and the process loops until nothing is left in the queue.

		k	i	t	t	e	n
	0	1	2	3	4	5	6
s	1	<u>1</u>	2	3	4	5	6
i	2	2	<u>1</u>	2	3	4	5
t	3	3	2	<u>1</u>	2	3	4
t	4	4	3	2	<u>1</u>	2	3
i	5	5	4	3	2	<u>2</u>	3
n	6	6	5	4	3	3	<u>2</u>
g	7	7	6	5	4	4	<u>3</u>

		S	a	t	u	r	d	a	y
	0	1	2	3	4	5	6	7	8
S	1	<u>0</u>	1	2	3	4	5	6	7
u	2	1	<u>1</u>	2	2	3	4	5	6
n	3	2	2	2	<u>3</u>	4	5	6	
d	4	3	3	3	3	<u>3</u>	4	5	
a	5	4	3	4	4	4	<u>3</u>	4	
y	6	5	4	4	5	5	5	<u>3</u>	

Figure 3. Two matrices showing the Levenshtein Distances between the substrings of kitten and sitting on the left, and Saturday and Sunday on the right.[29]

The list of candidates is sorted by lowest to highest distance to the new word. The first three words in the list are then displayed in the scene as possible candidates that the user can select to correct the word they just typed. The results are then passed on to the front-end logic, which will then display the results to the user.

Continuous-Motion Text Input

After looking through some demos online, we decided to build off of the Unity's UI System in VR demo and tutorial provided online from the Oculus site [4]. The main component that was crucial to this prototype was their gaze pointing code, which has a cursor directly in the middle of the screen that follows the users head motions. It provided a reliable backend for interacting with Unitys user interface (UI) with a cursor. We combined this with the Leap Motion Controller in two different ways: camera gaze mode and finger pointer mode. In gaze mode, the cursor is still tied to head movement, while in pointer mode, the cursor is controlled by the pointer finger of the right hand. Both modes are augmented by hand gestures, which control keyboard interactions. There is also a third mode called one-letter, which will turn off the Swype aspect of either mode and only take in one letter at a time. Note that here, when a button is pressed, it means the user made the appropriate gesture while the cursor is above that button.

The Leap Motion Controller SDK provides multiple default hand models. We chose the wire hand models because the colors make it easy to know which hand the computer considers left or right at any point. The wire model also blocks less of the users view, which makes it easier to see the rest of the scene.

The scene is a flat 2D keyboard panel in world space made with Unitys UI that has the same basic layout as a typical phone keyboard, mainly so that most users will already be familiar with the layout. It has the default ABC mode, and a SYM mode with some basic symbols. To the left of the keyboard is the text box panel where the users text is displayed. The backspace key will delete the rightmost letter that has been typed into the text box if pressed. Note that if the shift key is pressed, then all subsequent words typed by the user (if not in one-letter mode) will be capitalized until it is pressed again. If the user is in one-letter mode, then

every subsequent letter typed by the user will be capitalized until it is pressed again. Above these two panels is the auto-correct panel, which will update with the recommended auto-correct candidates. It also has buttons that the user can press to switch between the different modes, and a Clear button that will erase all the text in the text box panel when pressed.

We put the keyboard in World Space mainly due to the resulting size of keyboard and text box. If we shrank the keyboard more, it became difficult to see the individual keys, and putting it in Screen Space would have forced it to this size so that the whole UI component would be visible. Although the text box panel would technically be on top of the keyboard in a typical Android keyboard layout, we placed it to the side of the keyboard to give the keyboard center stage in the scene, rather than push it towards the bottom. The spacing between the keys also imitates the Samsung Galaxy S7 phone layout we used as a reference, but it was also so that it would be harder for the user to swipe over a letter they didnt intend to include, due to the large size of the cursor.




Figure 4. A screenshot of gaze mode showing the letters keyboard.




Figure 5. A screenshot of gaze mode showing the symbols keyboard.

The cursor itself was taken from the Oculus demo and repainted to match the colors of the scene. The size of the cursor is different depending on the mode the user is in. In gaze mode, the cursor is the same size as what the demo provided because it doesnt get in the way of the background but is also big enough to see. The circle shape with a hole in the center is also useful because the letter will be visible in the center when the user is hovering over it instead of blocking it. Pointer mode, however, slightly increases the size of the cursor, as seen in image 4.3. This is because pointer mode is harder to control and jitters a bit, so the larger cursor size helps to stabilize it.

The Oculus demo handles the logic behind how the cursor

knows what button it is hovering over. Originally we had tried using 2D collisions between the colliders of the cursor and the keyboard buttons, but there was a peculiar invisible offset that the collider boundaries would not show, which skewed the input the keyboard received. The demos logic instead relies on casting a ray away from the camera towards the scene and keeping track of what it collides with.

Below the keyboard is some text kept for the purpose of debugging. The top left text states whether it correctly knows that the user is pinching or not. If the user is pinching, it will say Debug: pinching!, otherwise it will say Debug: not pinching. The text on the right keeps track of the users speed. The threshold for what is considered fast or slow is currently determined by a hard-coded number in the code, so it will say Speed: SLOW if at or below that speed, and Speed: FAST otherwise. The bottom left text object states what the current mode is. At startup, the default mode is gaze, so the text will say Current Mode: GAZE.

The main gesture used in this project is the pinch gesture. This gesture was chosen because most of the other gestures were removed by the Orion update to Leap Motion to be reworked. The pinch gesture is actually a short method that checks if each finger is bent or not. If all five fingers are bent, then the hand is pinching, otherwise it is not. When the left hand does a pinch gesture, the cursor will start reading in the letters it is directly hovering over. It will keep doing this until the left hand stops pinching. At that point, it will add the word to the text box, then calculate and display the possible auto-correct candidates. If it is in one-letter mode for either of these options, then instead the user will type one letter at a time. They can go to the desired letter, make the pinching gesture with their left hand, then stop pinching and it will add that letter to the text box panel.

As for the amount of auto-correct candidates, the Android phone also has three candidates. We attempted adding more, but that meant shrinking the buttons. Sometimes the length of the word would extend past the button, and shrinking the font made it difficult to read the word. It would be possible to comfortably add more if the UI panels were stretched out more to the sides, but we kept it at three to maintain the current layout style. If the cursor just took in every letter that the user hovered over, it would be difficult to make out any comprehensible word, so there had to be a restriction on what letters were included in the final word. We had at first tried to weight the letters using a combination of velocity checking and the change in direction. Basically, as the cursor is moving, it would constantly update the current velocity of the cursor. Every few frames, it would create a vector in the direction its currently facing. When it reached a letter, if the velocity was below a certain number, it would compare the current vector to the previous vector it calculated to find the change in angle. This angle would be used to give the letter a weight, and only letters above a certain weight would be included in the final word. The rest would be handled by the auto-correct.

In practice, however, the change in angle did not seem to follow the desired behavior, and we instead decided to simplify so it would only do a velocity check. When the user is going across the letters, all of them are added to the word, but they have a negative weight. If the user slows down on a letter below some threshold, the letters weight is updated to a positive value. Once the user stops making a pinching gesture, all the letters that have a non-negative weight are added to the final word and returned.

Discussion

Since this is a simpler prototype, there are several areas where it could benefit from further improvements. One major issue is how unstable and shaky the cursor is when in the second mode (using the right hand). This is due to the cursors dependency on the movement of the head, and the shorter distance between the pointer finger and cursor (compared to the camera and cursor). Smoothing out the movement would likely improve the speed and accuracy of this mode.

Concerning the code, the most troublesome part would have to be the auto-correct. While it does contain the desired word most of the time, our UI only displays the first three results, weighted by distance. If there are more than three words that have the same distance, there's a chance the desired word will not be displayed as one of the three options. For earlier tests, we used a small dictionary of a few hundred words to catch any larger-scale bugs that could appear. Then we switched over to a much larger dictionary, which is directly from Linux 14.04's dictionary file. Suddenly, the auto-corrects accuracy dropped significantly because there were many more words that fell into the same Levenshtein Distance tolerance. We attempted to fix this using two different approaches, and neither one passed what we called a hello test. This simply means that the user types out the word hello using the Swype-like method and looks at the three candidates that show up. We chose the word hello because it is a fairly common word that has double letters in it. Duplicate letters cannot currently be typed out with this prototype. We considered adding in a gesture that would indicate that a letter would be counted twice, but the more natural solution is that the autocorrect should be smart enough to account for it. So instead, we have a test for local repeats that decreases the overall distance by one for every immediate double found in the word (i.e., the two *ts* in *potted count*, but not the ones in *potato*). This could have unintended consequences on the calculated Levenshtein distances, however.

The first approach we wrote was to sort the candidates by their Levenshtein distance to the users typed word. Those with the lowest Levenshtein distance would be at the front of the list. However, other words like *hell*, *well*, and *bell* would show up first simply because they were found as candidates before *hello* was found, and they have the same Levenshtein distance. Then we instead used a 5,000 word frequency dictionary from a word frequency site [8] and sorted the words by frequency. Even then, words like *well*, *tell*, and *help* would come first because they were considered more frequent than *hello*. Even sorting by both methods combined would not be enough. Note that in both of these

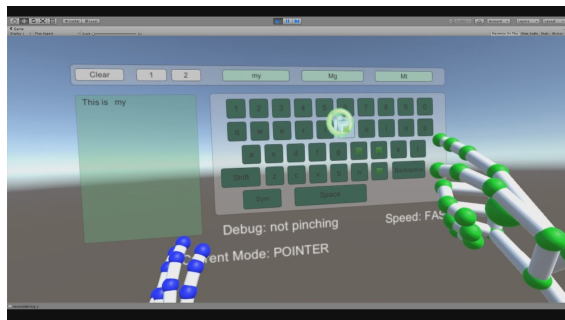


Figure 6. A screenshot of pointer mode.

cases, *hello* is correctly found as a candidate, but it does not reach the criteria to make it to the front of the list. The ideal solution would be to train it on user input so that it would learn what words are more commonly typed and offer more accurate suggestions, but I did not have the time to implement this.

In the back-end chapter, we mentioned that when calculating the Levenshtein Distance, it would check if the two compared letters were neighbors on the keyboard. This keyboard neighbors dictionary stores the neighbors of letters as they appear in a typical Android keyboard, but any other keyboard would have to manually change this because the entries are all hard-coded strings of characters. It is also missing some symbols that usually come with phone keyboards (curly braces and the tilde, to name a few). Our current dictionary was filtered to leave out any umlauts, accent marks, and so on for the sake of simplicity. It also uses case-insensitive matching when calculating Levenshtein distance.

Concerning speed, the prototype would never experience slow-downs except when we would purposely cause it to slow down and print out debug output. The one exception is when we used the Linux dictionary file, which contains around 99,000 rows of words. It would take around 3 seconds at start-up to build the dictionary, freezing the program until it finished. Perhaps multi-threading this, along with having a smaller dictionary of the most common words so the user can start typing immediately, could alleviate this drop in performance. After startup, however, searches were still fast and there were no significant framerate drops.

Other than the somewhat unreliable auto-correct performance, another issue that keeps it from truly imitating Swype is the fact that it weights letters by speed. If a user does not slow down for a few milliseconds at each letter, the program might not include that letter in the final word. At that point, it starts to resemble a hover technique more than Swype. It is still more efficient, because the waiting times for most hover-based programs is still much longer, but the user cannot simply remain at the same fast speed the entire time and still obtain their desired word, so balancing the speed check with other factors would be critical.

Conclusion and Future Work

For this project, we were able to create a prototype utilizing a Swype-like algorithm in virtual reality using a hand-tracking device to simulate typing on a 3D keyboard. We implemented both a gaze-based and a finger-based cursor control, and included a single-letter mode to compare with the Swype-like mode. A useful next step for this project would be to hook up the UI from this project to the official Swype or Swiftkey backend and then test its performance. This project has several shortcomings that cannot quite measure up to a polished, widely-used back-end that has been optimized through several iterations, so having access to something more stable would be ideal.

If the current prototype back-end code was kept, the auto-correct would have to be greatly improved. As mentioned in the Issues chapter, training it on user input would be a great way to teach it to display more accurate suggestions. As for the hand-tracking, one way to enhance it would be to combine it with a controller with built-in haptic feedback, like the Hands Omni gloves from Rice University that inflate bladders [30], or the GloveOne gloves that use vibrations [9]. While our approach sidesteps the need for haptic feedback, it would be useful to test the Swype approach in VR with an experience that better resembles the sen-

sation of using Swype on ones phone.

References

- [1] Edgar Alvarez. Samsung made a web browser for the Gear VR. <https://www.engadget.com/2015/12/01/samsung-internet-for-gear-vr/>, 2015. Engadget, Accessed: 2017-02-20.
- [2] Jon Axworthy. The origins of virtual reality. <https://www.wearable.com/wearable-tech/origins-of-virtual-reality-2535>, 2016. Wearable, Accessed: 2017-02-22.
- [3] Atman Binstock. Powering the Rift. <https://www3.oculus.com/en-us/blog/powering-the-rift/>, 2015. Oculus Blog, Accessed: 2017-03-15.
- [4] Andy Borrell. Unitys UI system in VR. <https://developer3.oculus.com/blog/unitys-ui-system-in-vr/>, 2015. Oculus VR, Accessed: 2016-08-15.
- [5] Joshua Corvinus. VR hex keyboard using Leap Motion. <https://www.youtube.com/watch?v=ZERwYJVZOgk>, 2016. Youtube, Accessed: 2017-02-20.
- [6] CrashCourse. The future of virtual reality: Crash course games. <https://www.youtube.com/watch?v=BfcBjJ3c9lg>, 2016. PBS Digital Studios, Accessed: 2017-02-15.
- [7] Benj Edwards. Unraveling the enigma of Nintendos Virtual Boy, 20 years later. <https://www.fastcompany.com/3050016/unraveling-the-enigma-of-nintendos-virtual-boy-20-years-later>, 2015. Fast Company, Accessed: 2017-03-1.
- [8] Word Frequency. Corpus of contemporary American English. <http://www.wordfrequency.info>, 2017. Word Frequency data, Accessed: 2017-03-07.
- [9] GloveOne. Gloveone: Feel virtual reality. <https://www.kickstarter.com/projects/gloveone/gloveone-feel-virtual-reality>, 2015. Kickstarter, Accessed: 2017-03-15.
- [10] David Jagneaux. Why Epics Tim Sweeney is fine with more devs using Unity. <https://uploadvr.com/tim-sweeney-on-unreal-vs-unity-priority-on-shipping/-first-and-foremost-with-ease-of-use-accessibility-being-second/>, 2016. Accessed: 2017-03-15.
- [11] Nick Johnson. Damn cool algorithms, part 1: BK-trees. <http://blog.notdot.net/2007/4/Damn-Cool-Algorithms-Part-1-BK-Trees>, 2007. Nicks Blog, Accessed: 2016-08-15.
- [12] Zach Kinstner. Hoverboard VR interface. <https://www.youtube.com/watch?v=hFpdHjA9uR8>, 2015. Youtube, Accessed: 2017-02-20.
- [13] Patrick Kolan. Ign retro: Virtual Boy revisited. <http://www.ign.com/articles/2008/01/14/ign-retro-virtual-boy-revisited>, 2008. IGN, Accessed: 2017-02-22.
- [14] Henry E. Lowood. Virtual Reality (VR). <https://www.britannica.com/technology/virtual-reality>, 2015. Encyclopaedia Britannica, Accessed: 2017-02-26.
- [15] Palmer Luckey. Oculus Rift: Step into the game. <https://www.kickstarter.com/projects/1523379957/oculus-rift-step-into-the-game>, 2012. Kickstarter, Accessed: 2017-02-22.
- [16] Leap Motion. Leap Motion developer. <https://developer.leapmotion.com/#101>, 2017. Leap Motion, Accessed: 2017-03-15.
- [17] Leap Motion. Oculus Rift DK2 setup. <https://developer.leapmotion.com/vr-setup/dk2>, 2017. Leap Motion, Accessed: 2017-03-15.
- [18] Myo. Connect with Myo. <https://www.myo.com/connect>, 2016. Myo, Accessed: 2017-03-15.
- [19] Oculus. Oculus. <https://www.oculus.com/>, 2017. Oculus, Accessed: 2017-03-15.
- [20] Vangos Pterneas. Finger tracking using Kinect V2. <http://pterneas.com/2016/01/24/kinect-finger-tracking/>, 2016. Vangos Pterneas Blog, Accessed: 2017-03-15.
- [21] Adi Robertson. HTC's hands-free VR phone tool is as clever and frustrating as the Vive itself. <http://www.theverge.com/2016/4/6/11377740/htc-vive-vr-bluetooth-phone-notifications-hands-on>, 2016. The Verge, Accessed: 2017-02-20.
- [22] Sascha Segan. Swype working on typing by waving hands in the air. <http://www.pcmag.com/article2/0,2817,2380227,00.asp#fbid=bidQg4EBxv7>, 2011. PC Mag, Accessed: 2017-01-25.
- [23] Sid Shuman. Playstation VR: The ultimate faq. <http://blog.us.playstation.com/2016/10/03/playstation-vr-the-ultimate-faq/>, 2016. Playstation, Accessed: 2017-02-20.
- [24] Lachlan Sleight. Leap keyboard demo. <https://www.youtube.com/watch?v=ckAGpmf21a8>, 2014. Youtube, Accessed: 2017-02-20.
- [25] Swype. Swype home. <http://www.swype.com/>, 2017. Swype, Accessed: 2017-02-20.
- [26] Unity. VR devices. <https://docs.unity3d.com/Manual/VRDevices.html>, 2017. Unity Documentation, Accessed: 2017-03-15.
- [27] Unknown. The history of stereo photography. http://www.arts.rpi.edu/rui/z/stereo_history/text/historystereog.html, 1996. Rensselaer Polytechnic Institute, Accessed: 2017-02-28.
- [28] Wikipedia. Damerau Levenshtein distance. https://en.wikipedia.org/wiki/Damerau_Levenshtein_distance, 2017. Wikipedia, Accessed: 2016-08-15.
- [29] Wikipedia. Levenshtein distance. https://en.wikipedia.org/wiki/Levenshtein_distance, 2017. Wikipedia, Accessed: 2016-08-15.
- [30] Mike Williams. Gamers feel the glove from Rice engineers. <http://news.rice.edu/2015/04/22/gamers-feel-the-glove-from-rice-engineers-2/>, 2015. Rice University News and Media, Accessed: 2017-03-15.
- [31] Chris Wren. www.wrenar.com VR keyboard w/Oculus Rift and Leap Motion. <https://www.youtube.com/watch?v=67Hyb2w1xFs>, 2014. Youtube, Accessed: 2017-02-20.
- [32] Xenopax. The BK-tree - a data structure for spell checking. <https://nullwords.wordpress.com/2013/03/13/the-bk-tree-a-data-structure-for-spell-checking/>, 2013. Xenopax, Accessed: 2016-08-15.
- [33] Adrian Zugaj. Test: Using Leap Motion with Swype keyboard. <https://www.youtube.com/watch?v=-4k0PnU46o>, 2013. Youtube, Accessed: 2017-02-20.

Author Biography

Janis Jimenez received her Master's degree in computer science in March 2017 from the University of California San Diego. Her research focus was on novel user interfaces for head-mounted display based virtual reality systems.

Dr. Jurgen Schulze is an Adjunct Professor in the computer science department and a Research Scientist at UC San Diego's Qualcomm Institute. He is the Director of Immersive Visualization Lab at the Qualcomm Institute and he teaches computer graphics and 3D user interfaces. His research is in scientific visualization, immersive virtual environments, human-computer interaction. He holds an M.S. degree from the University of Massachusetts and a Ph.D. from the University of Stuttgart, Germany.