

Empath-D: VR-based Empathetic App Design for Accessibility

Wonjung Kim*
wjkim@nclab.kaist.ac.kr
KAIST

Kenny Tsu Wei Choo
kenny.choo.2012@smu.edu.sg
Singapore Management University

Youngki Lee
youngkilee@smu.edu.sg
Singapore Management University

Archan Misra
archanm@smu.edu.sg
Singapore Management University

Rajesh Krishna Balan
rajesh@smu.edu.sg
Singapore Management University

ABSTRACT

With app-based interaction increasingly permeating all aspects of daily living, it is essential to ensure that apps are designed to be *inclusive* and are usable by a wider audience such as the elderly, with various impairments (e.g., visual, audio and motor). We propose *Empath-D*, a system that fosters empathetic design, by allowing app designers, *in-situ*, to rapidly evaluate the usability of their apps, from the perspective of impaired users. To provide a truly authentic experience, *Empath-D* carefully orchestrates the interaction between a smartphone and a VR device, allowing the user to experience simulated impairments in a virtual world while interacting naturally with the app, using a real smartphone. By carefully orchestrating the VR-smartphone interaction, *Empath-D* tackles challenges such as preserving low-latency app interaction, accurate visualization of hand movement and low-overhead perturbation of I/O streams. Experimental results show that user interaction with *Empath-D* is comparable (both in accuracy and user perception) to real-world app usage, and that it can simulate impairment effects as effectively as a custom hardware simulator.

CCS CONCEPTS

• **Human-centered computing** → **Systems and tools for interaction design; Ubiquitous and mobile computing systems and tools; Accessibility design and evaluation methods; Accessibility systems and tools; Ubiquitous and mobile computing design and evaluation methods;**

KEYWORDS

empathetic design; accessibility; mobile design; virtual reality; multi-device, distributed user interfaces

ACM Reference Format:

Wonjung Kim, Kenny Tsu Wei Choo, Youngki Lee, Archan Misra, and Rajesh Krishna Balan. 2018. Empath-D: VR-based Empathetic App Design for Accessibility. In *MobiSys '18: The 16th Annual International Conference on Mobile*

*This work was done while the author was on an internship at Singapore Management University

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than ACM must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from [permissions@acm.org](https://doi.org/10.1145/3210240.3210331).

MobiSys '18, June 10–15, 2018, Munich, Germany
© 2018 Association for Computing Machinery.
ACM ISBN 978-1-4503-5720-3/18/06...\$15.00
<https://doi.org/10.1145/3210240.3210331>



Figure 1: Overview of *Empath-D*

Systems, Applications, and Services, June 10–15, 2018, Munich, Germany. ACM, New York, NY, USA, 13 pages. <https://doi.org/10.1145/3210240.3210331>

1 INTRODUCTION

Digital interactions have become increasingly commonplace and immersive. We now constantly interact with our personal devices and computing-enhanced ambient objects (such as coffeemakers, home automation systems and digital directories), while engaging in everyday activities, such as commuting, shopping or exercising. Given the ubiquity of such interactions, it is important to ensure that the associated computing interfaces remain accessible to segments of the population, such as the *elderly*, who suffer from various impairments. The global elderly population is projected to reach 16.7% by 2050 [33], and such users suffer disproportionately from impairments (e.g., vision) that hinder accessibility.

To support more accessible design, our earlier work [11] introduced the vision of *Empath-D*, which uses a virtual reality (VR) device to provide mobile application/object designers with a *realistic emulation* of the interaction experience that impaired users would encounter. In this work, we present the design, implementation and validation of the *Empath-D* system inspired by this vision. *Empath-D*'s goal is to allow unimpaired application designers to *step into the shoes* of impaired users and rapidly evaluate the usability of alternative prototypes. While we shall principally focus on empathetic evaluation of mobile applications (apps), *Empath-D*'s design is generic enough to permit emulation of other real-world interactions—e.g., how an elderly user with *cataracts* and *hearing loss* would experience a traffic-light controlled pedestrian intersection.

Empath-D's¹ key idea is to present the user with an impairment-augmented view of the smartphone interface (or other digital objects) in a virtual world, while allowing the non-impaired user to perform natural interactions, using a physical smartphone, with a real-world instance of the smartphone app. At a high-level, *Empath-D* works as follows (see Figure 1): The (unimpaired) user uses a physical smartphone to perform real-world interactions (such as scrolls, taps or gestures) with the app, while wearing a VR device. The results of such interactions are projected *instantaneously* through the I/O interfaces (e.g. screen, speaker) of a 'virtual smartphone' visible in the VR display, but only after those I/O streams have been appropriately degraded by the specified impairment. For example, in Figure 1, the virtual phone's display (and the world view) has been appropriately vignetted, to mimic the experience of a user suffering from *glaucoma*.

Key Challenges: To mimic impairments with adequate fidelity and usability, *Empath-D* must support the following features:

- **Fast, Accurate Multi-device Operation:** *Empath-D* utilizes a split-interaction paradigm: a user interacts with an app using a real-world handheld smartphone, while perceiving (viewing, hearing) the app responses through the VR interface. To faithfully replicate the real-world experience, this split-mode interaction must have tight time coupling and visual fidelity (of the virtual phone's screen), comparable to direct interactions with a standalone smartphone.
- **Real-time Tracking:** To preserve a user's perception of naturalistic interactions, *Empath-D* must not only capture explicit phone events, but also mirrors physical actions taken by the user (e.g., swinging the phone around or having one's hand hover over the phone). Thus, *Empath-D* must also track and *render*, in real-time, the orientation/location of both the phone and the user's hand within the VR device's field-of-view.
- **Lightweight Impairment Execution:** To preserve the feel of natural interaction, *Empath-D* must insert the impairment-specific perturbations into the input/output streams with imperceptible latency or computational overhead (e.g., no reduction in video frame rate).

Key Contributions: We make the following major contributions:

- **3-Tier Virtualisation Model:** We design a novel 3-tier architecture where (i) the real-world smartphone serves merely as a *tracker*, forwarding user interaction events (e.g., screen touch and gestures) to a computationally powerful intermediary, after which (ii) the intermediary device perturbs those events by blending in specific *input* impairments (e.g., hand tremors) and passes them to an app instance running on a smartphone emulator, and finally (iii) the VR device receives the redirected outputs from this app instance and renders an appropriately-impaired (by blending in the *output* impairments) virtual world, including a virtual smartphone.
- **Real-time Hand and Phone Tracking:** We use an RGB-Depth camera, mounted on the head-worn VR device, to track the outline of a user's hand, and subsequently perform a light-weight but realistic 3-D rendering of the hand on the VR

display. We also use fiducial marker tracking [14] by the camera to track the position/orientation of the real-world smartphone. We demonstrate our ability to achieve both high-fidelity (pointing error ≤ 5 mm) and low-latency (end-to-end delays below 120 msec) hand tracking and display.

- **Usability of Virtualized Phone, in Use Environments:** We show that *Empath-D* is not just usable, but that user performance (absent impairments) using *Empath-D*'s virtual smartphone is equivalent to real-world interaction with a smartphone. In addition, we allow usability testing of apps in their use environments, a key enabler for design of mobile applications which may be used anywhere. Our Samsung Gear VR-based prototype has end-to-end latency low-enough (only 96.3 msec of latency, excluding the mobile app emulation) to permit faithful reproduction of direct smartphone usage.
- **Validation of Impairment Fidelity and Overall System:** We implement two distinct vision (glaucoma & cataract) and one audio (high-frequency hearing loss) impairment in our *Empath-D* prototype. We then conduct a set of studies using the vision impairments, where 12 participants perform a series of standardised activities (e.g., add an alarm), using both our *Empath-D* prototype (test) and a commercial hardware vision impairment simulator (control) and establish that the performance of users is *equivalent* across the test and control groups. Finally, we conduct a small-scale study to provide preliminary evidence that our empathetic approach allows developers to design accessible mobile UIs faster and better.

2 THE EMPATH-D VISION

We use an example to illustrate the use of *Empath-D*:

Designing for Visual Impairment. *Alice is designing a mobile app that automatically magnifies text from real environments seen through its rear camera to aid people who suffer from cataracts (a condition that dims and blurs vision). Alice starts Empath-D and is presented with a web interface that allows her to customise impairments (e.g., specify the intensity of visual blur). After customising the environment, Alice clicks in the Empath-D web interface to (1) compile the environment to her phone used for VR display (VR-phone)² and (2) connect an input/output service to a separate phone (IO-phone). She then plugs the VR-phone into the VR headset.*

*Alice then compiles her Android app, and runs it in the Android emulator. She puts on the VR headset and holds the IO-phone in her hands. A virtual smartphone (Virt-phone) shows up in VR, tracking the real-world motion of the IO-phone. Alice now navigates through the virtual world, experiencing it as an "impaired user, with cataracts". She holds up IO-phone on a street corner (in the real world), and notices that the magnified text (as seen in the virtual phone in the virtual world) is not clear enough to be legible to a cataract-impaired user. She can now iteratively and rapidly modify her app, recompile it, and execute it in the Android emulator, until she is satisfied with the output. This scenario demonstrates the ease-of-use for *Empath-D*, with no need for special instrumentation of the app.*

¹Video of *Empath-D* in action at https://is.gd/empath_d

²The VR-phone is needed only for VR devices that require a smartphone—e.g., Samsung Gear VR

3 SYSTEM OVERVIEW

3.1 Design Goals and Implications

Empath-D has the following key goals, which directly influence the salient implementation choices.

- Holistic emulation of impairments:** For a truly empathetic experience, the app designer must perceive the effects of impairments not just while using the mobile app, but throughout her immersion in the virtual world. Consider a user, suffering from cataract, who is interacting with her smartphone while attending a dimly lit dinner gathering. Simply blurring the phone display, while leaving the background illumination and focus unchanged, might not replicate challenges in visual contrast that an impaired user would face in reality. This requirement precludes the straightforward use of I/O redirection techniques such as Rio [8], which can potentially perturb the I/O streams of only the mobile device. Instead, the impairment must be applied holistically, to the *entire virtual world*.
- Realistic emulation of smartphone and mobile apps in the virtual world:** *Empath-D* aims at realistically emulating mobile apps within the virtual world rendered by a commodity VR headset. Realistic emulation of mobile apps imposes two requirements. (a) First, the virtual smartphone should have sufficient visual resolution, corresponding to typical usage where the smartphone is held $\approx 30\text{cm}$ away from the eye. We shall see (in Section 6.3) that this requirement, coupled with differences in display resolutions between smartphones and VR devices, requires careful magnification of the virtual smartphone to provide *legibility* without hampering *usage fidelity*. (b) Second, the user should not perceive any *lag* between her user input and the rendered view of the app, seen through the VR device. Quantitatively, we thus require that the task completion time, experienced by a user interacting with the emulated application in the virtual world, should be comparable to real-world app usage on a real smartphone.
- Use of unmodified app** For easy and low-overhead adoption by app designers, *Empath-D* should support the emulation of mobile applications using the original, unmodified binaries (e.g., .apks for Android). *Empath-D*'s requirement to support empathetic emulation without app modifications implies that app designers would be able to adopt *Empath-D* with minimal impact to existing development practices.
- Low-latency, accurate finger tracking:** This goal is an extension of the holistic emulation objective. In the real-world, users utilise instantaneous visual feedback and proprioception to move their fingers around the smartphone display, even when they are hovering but not actually touching the display. To ensure consistency between the user's tactile, visual and proprioceptive perceptions of her hand movement, *Empath-D* should also realistically render, in the virtual world, the user's hand movements and any changes in the position/orientation of the real-world smartphone, *without any perceptible lag*. In Section 6, we shall see how the *Empath-D* implementation meets these stringent performance bounds.




Figure 2: *Empath-D* architecture

- Light-weight, effective emulation of impairments:** *Empath-D* will need to emulate impairments, at different levels of severity. For high-fidelity empathetic emulation, the insertion of such impairments in the I/O streams of the smartphone should not add generation any additional artefacts (e.g., increased latency, reduction in display refresh rate, etc.).

3.2 System Overview

We now present the overview of the *Empath-D* system (illustrated in Figure 2).

Using *Empath-D* in VR. To immersively evaluate the application, the developer (or the tester) starts by installing her developed application binaries (i.e., Android .apks) to run on the emulated smartphone. The developer then adjusts the profile settings for the impairment using *Empath-D*'s web dashboard and selects a use case scenario (e.g., in office, in the street, etc.). She holds her physical smartphone and puts on the VR headset, earphones (when hearing impairments are involved) and experiences the immersive reality (where she can use the app - now mapped onto the physical smartphone - with the configured impairment under the designated use case scenario) that *Empath-D* generates. She then tests out various interfaces and functionalities of the app in the immersive VR environments.

Components of *Empath-D*. *Empath-D* runs across three different physical devices: a physical smartphone, a computer, and a VR device (see Figure 2).

Smartphone: In *Empath-D*, the user interacts with the app using a real smartphone held in her hand. Interestingly, this smartphone does not run the app itself, but functions as a tracking device, helping to preserve the user's realistic sense of smartphone interaction. The smartphone simply *redirects* the user interaction events (e.g., touch events such as clicks and swipes on the display and motion events captured by inertial sensors) to the computer, which is in

charge of the app emulation. This smartphone also displays a fiducial marker array [14] on its display, to help in efficient, real-time tracking of the phone's location.

Computer: The computer is at the heart of *Empath-D*'s ability to fuse the real and virtual world. It consists of two major components: *Phone and Hand Tracker* and *Mobile Emulator*, as well as a *Web Dashboard* (see Figure 6), which allows the user to select the impairment profile to be applied. In addition, as we shall discuss shortly, this computer may run an *Impairment Generator cum Virtual World Renderer*. Key functions include:

- The *Phone and Hand Tracker*, uses image captured by the VR headset-mounted camera, to track the position and pose of the smartphone (relative to the VR device), and create the virtual phone image at the correct position in the virtual world. It also uses the same camera to track the user's hand, as it interacts with the smartphone, and then renders it in the virtual world.
- The *Mobile Emulator* executes the app being tested, using the redirected stream of user interaction events transmitted by the smartphone. The resulting visual output of the app is then transmitted as a sequence of images to the VR device, where these images will be integrated into the virtual phone object; likewise, audio output (if any) is directly streamed to the VR device.

The overall *Empath-D* framework includes an *Impairment Generator* that is typically applied as one or more filters over the Virtual World Renderer (an engine such as Unity [44]) which is responsible for combining various virtual objects and rendering the virtual world). The *Impairment Generator* effectively perturbs/modifies the audio/video feeds of the virtual world, before they are displayed on the VR device. For example, to emulate cataracts, it applies an appropriate 'blurring/dimming' filter on the video feed; similarly to emulate high-frequency hearing loss (an audio impairment), this generator will apply a low-pass filter on the output audio stream. These two components are placed inside a dotted-line rectangle in Figure 2, to reflect the reality that these components run on either the Computer or the VR device, depending on whether the VR device is *tethered* or not. In untethered VR devices (such as the Samsung Gear VR), the *Impairment Generator* and the *Virtual World Renderer* run on the VR device itself. In contrast, tethered devices such as the HTC Vive will run on the computer, and typically offer higher graphics quality, frame rates, faster execution.

VR Device: Finally, the VR device is used to display the synthesised virtual world to the user. This synthesis involves the fusion of the virtual smartphone, the user's hand and the ambient virtual world, all subject to the impairment filter.

4 VR-BASED EMULATION OF MOBILE INTERACTION

Empath-D follows a split-interaction paradigm: for realistic immersion, *Empath-D* renders the visual and audio output of the target app in the virtual world (i.e., via VR headset's display and speakers), while allowing the user to interact naturalistically with a real-world physical phone. A major challenge in this paradigm is to enable natural, low-latency tracking and display of the real-world motion of both the phone and the user's hands, so as to ensure consistency

across the user's visual, tactile and proprioceptive experience. We achieve this by performing three distinct steps: (a) smartphone tracking, (b) hand tracking, and (c) hand rendering in VR, using an RGB-Depth (RGB-D) camera mounted on the VR headset. *Empath-D* first tracks the position and orientation of the physical smartphone and synchronises the position of the virtual phone to the physical smartphone (See Section 4.1). Separately, *Empath-D* also captures fingers in the real world and displays them at the correct position (relative to the virtual smartphone) in the virtual world (See Section 4.2 and 4.3).




Figure 3: Tracking physical phone with fiducial markers

Empath-D uses the headset-mounted RGB-D camera to capture the colour image along with the depth values, relative to the camera. The camera's position is always fixed, relative to the user's head. Its three axes are thus aligned to a user's head: z -axis to the user's forward (gaze) direction, and x and y axes capturing the vertical and horizontal displacement.

4.1 Tracking the physical smartphone

Empath-D uses fiducial markers, displayed on the physical smartphone's screen, to localise the smartphone efficiently. It takes a colour image as an input, and returns the transformation relative to the camera's coordinate system: translation and rotation, i.e., x , y , z , roll, pitch, yaw from the RGB-D camera's coordinate system. We employ a technique proposed and detailed in [14].

The *Empath-D* Hand Tracker component tracks the physical phone using markers captured by the camera. Each marker, displayed on the phone screen, has a distinct pattern. The tracker knows the position of each marker (e.g., top-left, top-right, bottom-left and bottom-right) in the physical smartphone screen's coordinate system. The system first detects these markers in a given colour image, identifying them based on their unique patterns (see Figure 3). In particular, the system recognises the coordinates of each of the four corners of each marker. Moreover, the system knows the true size of, and separation between, each marker. It then uses an object pose estimation algorithm (provided by openCV's *solvePnP* function [6]), along with the array of fiducial marker points, to compute the 3-D position and orientation of the smartphone. Past

Algorithm 1 Hand Segmentation

```

1: Input:  $T \leftarrow$  Phone's translation (3-D vector)
2: Input:  $R \leftarrow$  Phone's orientation (3×3 rotation matrix),
3: Input:  $F \leftarrow$  RGBD Frame, 2-D array that each entry  $F_{i,j}$  holds
   a color value and 3-D position relative to the camera.
4: Input:  $V \leftarrow$  3-D region of interest (relative to the phone)
5: Output:  $fgMask$ , 2D bool array whose dimension equals to  $F$ 
6:
7:  $fgMask[i,j] \leftarrow$  false for all  $(i,j)$ 
8: for point  $(i,j)$  in  $F$  do
9:   if  $(i,j)$  in screen_border then
10:    /* Case A: Blue background segmentation */
11:     $fgMask[i,j] \leftarrow 1 - Blue(F_{i,j}) + 0.5 \cdot Red(F_{i,j}) > \tau$ 
12:   else
13:    /* Case B: Depth-based segmentation */
14:     $pos_{phone} \leftarrow R^{-1} \cdot (Position(F_{i,j}) - T)$ 
15:     $fgMask[i,j] \leftarrow (pos_{phone} \in V)$ 
16:   end if
17: end for
18: return  $fgMask$ 

```

results [14] show that this technique can compute an object's position and orientation with sub-cm level accuracy.

This fiducial marker-based algorithm would fail under two conditions: (a) when the markers are occluded by the user's hand, and (b) if the ambient illumination levels are too low or too high, reducing the contrast level of the markers. To tackle (a), the smartphone screen uses an entire array of markers displayed across the scene, thereby ensuring correct smartphone tracking as long as some part of the phone is visible. Contrast concerns are not particularly relevant in our scenario, as we assume that the user is testing the app in a regularly lit work/office environment.


4.2 Hand Segmentation

Empath-D uses the frames captured by the RGB-D camera to track and segment the user's hand. For each frame, we extract the segment (polygon of pixels) that represents the user's hand, and render that segment in the virtual world. As the goal of hand-tracking is to provide the user with a natural view of her smartphone interactions, we restrict the tracking technique to a 3-D region of interest (ROI) that is centred at the phone, with a depth of 2cm and a planar boundary of 6cm. In other words, we only track the hand while it is $\leq 2cm$ s away from the smartphone screen, and within $\leq 6cms$ of the smartphone edges.

A straightforward approach is to apply a depth-based segmentation strategy, where we first isolate only the foreground points which lie within a depth=2cm of the smartphone surface. However, we empirically observed that, due to the glossy surface of the smartphone, such depth estimation was inaccurate for points located on the smartphone's screen. Accordingly, we implemented two separate segmentation methods (detailed in Algorithm 1): (case A) a colour-based segmentation approach to identify points which are directly over the smartphone, and (case B) a depth-based approach to identify points which are near, but not over, the smartphone's



Figure 4: Mesh of hand

Figure 5: *Empath-D* hand segmentation

screen. We apply the colour-based segmentation to the points inside the screen's border (thick orange contour in Figure 3) and the depth-based approach to the points outside.

Colour-based segmentation: We adopt the colour-based technique proposed in [41]. The approach tests RGB values to segment foreground (hand) from background, coloured in blue. In our scenario, we target human skin as the foreground. Human skin has a property common in all races: its R value has about twice the value of G and B ($R \approx 2G \approx 2B$). Given the property of human skin, we obtain a formula that discriminates the foreground from the background whose B value is 1 (line 11 in Algorithm 1). τ is a user-tunable threshold which allows it adapt to different lighting conditions.

However, note that, to enable tracking of the phone, the phone's screen cannot be completely blue, but will need to contain the array of fiducial markers. We tackle both problems simultaneously by using blue ($R=0, G=0, B=1$) to colour the markers, over a cyan ($R=0, G=1, B=1$) background. Here we modified only G value, which is unused in the colour-based segmentation.

Points outside the smartphone's screen are segmented using the depth-based approach. After identifying the points corresponding to the user's hand, the system translates these points to 3-D coordinates in the camera's coordinate system, using the associated depth values.

4.3 Rendering the hand in the virtual world

After detecting the hand segment, the *Empath-D* system renders it in the virtual world. The system passes the tracked hands to the *Virtual World Renderer*, sharing the (i) 3D structure of the hands (surface mesh), (ii) colour image of the RGB-D frame (texture), and (iii) mapping between the surface mesh and the colour image (UV map). In common rendering engines (e.g. Unity), the 3D structure of the hand is represented by a *triangle mesh*—i.e., a set of vertices, constituting individual small triangles. The mesh is rendered at the same location as the user's hand in the real world. As the user's hand is localised in the coordinates of the RGB-D depth camera, the location is offset by an additional depth value (7cm in our implementation), to reflect the additional distance between the centre of the user's eyes and the depth camera. An important characteristic of our algorithm is that we render the *actual image* of the user's hands over this triangle mesh. Figure 4 illustrates the Delaunay




Figure 6: Screenshot of *Empath-D* impairment configuration dashboard

triangulation of a set of points. The mesh is combined with the hand’s image (Figure 5), and rendered in the VR display. Extracting and rendering the actual image of the user’s finger enhances the immersive feeling of real-life smartphone navigation in the virtual world.

The complexity of the mesh—i.e., the number of vertices (or triangles) in the rendered hand—is an important parameter in the rendering process. A larger number of vertices captures the contours of the hand more precisely, resulting in a more life-like image. However, this also results in added rendering latency in the rendering engine. To support the twin objectives of low-latency and life-like rendering, we utilise a sub-sampling technique to construct the mesh. Specifically, *Empath-D* preserves all the points on the edges of the segment, to preserve the precise contours of the hand. However, it performs a 32-fold downsampling of the interior points (prior to constructing the Delaunay triangulation), along both the row and column axes, to reduce the computational time significantly, without materially affecting the reconstructed hand image. We shall show, in Section 6, how our prototype *Empath-D* implementation uses this technique to achieve our twin objectives.

5 IMPAIRMENT SIMULATION

Empath-D aims to enable evaluation of the usability of app designs under visual, auditory and haptic impairment simulation. Realistic simulation of various impairments in the VR world is the essential requirement to achieve this goal.

There has been a thread of research to simulate impairments through physical simulator devices [1, 13, 29, 39, 49]. For instance, Zimmerman et al. use goggles and enclosing materials to simulate low vision impairments [49]. These hardware simulators generalise the impairment of interest and enable simulation of specific aspects of the impairment pathology rather than emulate exactly how an impairment is. However, impairments can vary greatly between individuals. For instance, glaucoma generally progresses in deterioration from the periphery towards the centre of vision, but in reality, it comes in different shapes and severity, affecting usability of applications in different ways. Existing physical impairment simulators simply approximate this as a central circle of



Figure 7: Simulated cataract (left) and simulated glaucoma (right)

clarity, with blur through to the periphery. *Empath-D* is advantageous over existing physical simulators in the following ways, it allows: 1) impairments to be customised, 2) simultaneous manifestation of multiple impairments, 3) the addition of new impairments easily. Figure 6 shows the web interface for designers to customise impairments for the target user group.

5.1 Simulating Visual Impairments

Vision is the dominant sensory system by which humans perceive the world, and is a key focus for *Empath-D*. Vision impairment is one of the most common causes of accessibility problems that comes with age. Common vision impairments include cataracts, glaucoma, and age-related macular degeneration. Such vision impairments present as reduced visual acuity, loss of central/peripheral vision, or decreased contrast sensitivity. It is widely studied that these symptoms can affect the interaction with various desktop and mobile applications; for example, humans use peripheral vision to pre-scan text ahead of his/her point of focus. As the peripheral vision narrows, the scanning becomes less effective, which slows reading [23]. In this work, we examine and simulate two commonly found visual impairments - cataracts and glaucoma.

Our approach is to apply an image effect at the “eye” (i.e., a camera pair of view renderers) of the VR scene. From this camera pair, the image effect will apply to all other objects in the scene (e.g., smartphone, fingers, scene), just as how impaired users would experience it. We employed various image filters for different impairments, which 1) provide realism of impairments to help designers to find out usability issues and take corrective actions, and 2) have small computational overhead not to add noticeable delays to our entire emulation.

The approach is flexible and lightweight. Impairment simulator’s intensity is configurable at runtime. The image effects are applied at the last stage of the rendering pipeline. **Glaucoma** presents functionally as a *loss in peripheral vision*. To simulate glaucoma, we use a vignette with a clear inner circle, blurred inner-outer circle, and black extending outwards from the outer circle (see Figure 7). **Cataracts** presents functionally as *reduced visual acuity* and *reduced contrast sensitivity*. We use a blur filter to simulate reduced visual acuity, and a contrast reduction filter to simulate reduced contrast sensitivity (see Figure 7).

Table 1: Hardware of Empath-D

VR headset	Samsung Gear VR [5]
VR smartphone	Samsung Galaxy S7 [4]
RGB-D camera	Intel RealSense SR300 [20]
PC	CPU: 4 cores, 3.4 GHz RAM: 16 GB GPU: GeForce GTX 1080 [32]
Physical IO smartphone	Samsung Galaxy S5 [40]

The functional aspects of vision impairments are straightforward to create in VR, which give *Empath-D* high extensibility to implement other types of visual impairments. While we just described two impairments pertaining to our studies, it is easy to create other impairments such as colour filters to simulate colour blindness. However, we leave the effect of eye movements on impairments as the future work. Since eye-tracking is currently not supported in *Empath-D*, a user will need to move his head to achieve the same effect.

5.2 Simulating Other Modalities

We discuss how other modalities may be simulated in *Empath-D*.

Hand Tremors. Hand tremors are a common symptom of Parkinson’s disease or Essential tremor and make it hard for one to precisely point on a touchscreen. A hand tremor may be characterised by the frequency and amplitude of oscillatory movement. Since we present virtual representations of the user’s hand (i.e., as a 3D mesh) to enable his interaction with the virtual mobile phone, *Empath-D* similarly perturbs this 3D mesh in VR to create hand tremors. While a user may physically not experience hand movement, the visual perturbation would be sufficient to hinder accurate touch to simulate hand tremors.

Hearing Loss. High-frequency hearing loss is a common symptom for the elderly population. People diagnosed with high-frequency hearing loss are unable to hear sounds between 2,000 Hz and 8,000 Hz. These people often struggle to understand or keep up with daily conversations (missing consonants in higher registers, such as the letters F and S or female voices). *Empath-D* applies a bandpass filter over the output sound of the target application to diminish the sound signals between 2kHz and 8kHz and plays the filtered audio feed through the VR device.

6 IMPLEMENTATION

6.1 Hardware

We implemented our current *Empath-D* prototype using the hardware described in Table 1. We used the Samsung Gear VR fitted with the Samsung Galaxy S7 as the VR headset. We used the Intel RealSense SR300 RGB-D camera for finger tracking, selecting this among alternatives as: 1) its small size and low weight allowed us to easily attach it to the VR headset, and 2) its minimum sensing range is low enough to permit hand tracking at a distance of 30cm. We employed the Samsung Galaxy S5 as the physical I/O device, and a powerful laptop (4 core 3.4 GHz CPU, 16GB RAM) as the intermediary device. The choice of the VR headset itself was deliberate. We chose a Samsung Gear VR headset (an untethered




Figure 8: Rendering frame rate under varying virtual display resolution (width : height = 9 : 16, default resolution of Android emulator is 1080x1920)

smartphone-powered VR device) over more powerful PC-tethered VR devices such as the HTC Vive or Oculus Rift. This was mainly because PC-tethered devices such as HTC Vive use IR lasers to localise the headset, which interferes with the IR laser emitted by the RGB-D camera used for depth sensing in hand tracking.

6.2 Rendering an Emulated App

We used empirical studies to determine an appropriate screen resolution and frame rate to render the emulated app (and the smartphone) in the VR headset. *Empath-D* obtains screenshots of its mobile emulator using the Android virtual display [35] and transmits these screenshots over WiFi to the Gear VR device. The overhead of transmitting and rendering these emulated screenshots is proportional to their resolution. The default 1080p resolution could sustain a frame rate of only 18 *fps*, which causes visible jerkiness. To reduce this overhead, we reduced the resolution (using `setDisplayProjection()` method), and applied *differential transmissions*, sending a screenshot only when the emulated app’s display changes.

Figure 8 shows the experimental results on the tradeoff between the resolution and the rendering frame rate, obtained while playing a video to ensure continuous change of the screen content. The frame rate saturates at 57 *fps*, at a screen resolution of 485x863. Moreover, through another user study (described next) to understand the minimum resolution to read an app’s contents, we empirically verified that the participants had no issues in reading the app’s content at the resolution of 485x863. Hence, we choose this resolution as our default, although this setting can be modified (e.g., we can pick a higher resolution, and a lower frame rate, for an app with mostly static content).

If *Empath-D* displays the virtual smartphone at its original size in the virtual world (portrait position), its display becomes illegible. For example, the Samsung Galaxy S7 (in the Gear VR) has a resolution of 2560x1440 and an $\approx 101^\circ$ horizontal field of view yielding a horizontal pixel density of ≈ 25.3 pixels/degree. When a virtual phone is held at 30cm away, the horizontal pixel density drops below 25.3 pixels/degree due to downsampling of the virtual phone screen as seen through the VR display. This presents a problem for viewing the content of the virtual phone — in particular, text — as its pixel density is significantly lower than when viewing a physical




Figure 9: Readable font size of the virtual smartphone at a magnification ratio

phone. For instance, the Galaxy S5 gives ≈ 89.4 pixels/degree at 30cm distance.

We tackle this issue by scaling up the virtual phone’s size by a factor that ensures that the phone’s display text remains legible. To determine this factor, we recruited three participants and asked them to record the minimum readable font sizes, while showing them a virtual smartphone (at a distance of 30 cm) with various magnification ratios (increased by 0.1 from 1.0 to 2.7). Figure 9 shows that participants could read text with the font size= 12sp (the commonly used minimum font size for mobile apps) for magnification factors ≥ 1.5 . Accordingly, we used 1.5 as the default magnification ratio for the smartphone and its display. We also proportionately scaled up the user’s rendered hand. User studies (Section 7) show that users found this configuration highly usable.

6.3 Rendering Virtual Hand

As discussed in Section 4.3, the rendering latency of the virtual hand is proportional to the number of vertices in the Delaunay triangulation-based mesh. To reduce the latency, we apply a non-uniform sampling approach. Specifically, *Empath-D* preserves all the points on the edges of the segment, to preserve the precise contours of the hand. However, it performs a downsampling of the interior points (prior to constructing the Delaunay triangulation), along both the x and y axes, to reduce the computational time significantly, without materially affecting the reconstructed hand image. We empirically determined the sampling rate X , by varying X and measuring both (i) the processing latency and (ii) the SSIM [12, 46] (*Structural SIMilarity*; a metric of perceived image quality) of the hand images, using 200 RGB-D frames. Figure 10 shows the results. Without any subsampling ($X = 0\%$), the rendering latency is 311.1 msec, which is too high for our responsiveness goal. We empirically downsample the internal pixels by a factor of 32 ($X = 99.9\%$), i.e., choosing every 32^{nd} pixel on the grid. This results in a latency of 26.9 msec, while keeping the $SSIM = 0.976$, a level indistinguishable with the original as perceived by a human.

6.4 Environment Emulation

To enable holistic evaluation of app interactions, *Empath-D* emulates not just the virtual phone, but the entire virtual world as well. In our current implementation, we emulated a crowded *Urban Street* environment, which includes crosswalks, traffic lights, pedestrians and commonplace roadside obstacles. To further mimic real-world




Figure 10: Rendering latency vs. image quality of the virtual hand

movement, our implementation allows the user to navigate the virtual world by (i) rotating her head (this uses the head tracking ability of the VR device), and (ii) by ‘walking in place’, using the technique proposed in [45] as this does not require any additional hardware on the VR device.

6.5 VR Manager

This component currently executes on the VR smartphone, and is responsible for combining the output of the various components (*Hand Tracker*, *Phone Tracker* and *Virtual Phone*) in the virtual world. This component, implemented as a Unity application, renders these various components. This component is also responsible for applying the impairments on the output of the virtual world. The image effects simulating low vision impairments are defined as a script, *Shaders* in Unity.

7 EVALUATION

We now present a mix of system and user experiments to evaluate the performance and efficacy of our *Empath-D* implementation. Besides micro-benchmark studies, we conducted two experiments to capture user interaction with *Empath-D*. In Experiment 1, we examine the performance of *Empath-D* vs. a real-world smartphone, in the absence of any impairments. In Experiment 2, we consider an impairment-augmented version of *Empath-D*, comparing the performance of users against the use of commercial impairment simulation hardware.

7.1 Micro-benchmark Performance of *Empath-D*

We measured the overall latency of *Empath-D*, both in terms of the delay in reflecting touch interactions in the virtual world and in terms of the hand tracking delay.

7.1.1 End-to-end Latency of Touch Interaction. As a measure of the overall responsiveness of *Empath-D*, we computed the latency between a touch input, on the physical smartphone, and the resulting change in the content of the virtual smartphone, rendered in the VR display. To measure this, we utilised a high framerate camera (operating at 240 *fps*) to concurrently record both the screen of the physical smartphone and the virtual phone (displayed in the VR). The phone screen is coloured green initially, and was programmed to turn red as soon as it received a touch input. We repeated the




Figure 11: Overhead of impairment simulation

measurement 23 times, capturing (via the video frames) the time gap between (i) the physical smartphone screen turning red and (ii) the virtual smartphone turning red in the VR display. The end-to-end latency is 237.70 msec ($SD = 20.43$).

By monitoring the intermediary computer, we obtained the breakdown of this delay: (i) smartphone responsiveness (the time from the user touching the screen till the time the phone transmits the touch event to the computer) = 0.3 msec ($SD = 0.16$); (ii) computer emulation responsiveness (the time from receiving the touch event till the time the screenshot of the modified display is sent to the VR device) = 141.37 msec ($SD = 6.6$), and (iii) the VR responsiveness (the time from receiving the screenshot till it is rendered on the VR display) = 10.46 msec ($SD = 8.36$). The remaining latency ($\approx 87 \text{ msec}$) can be attributed as the WiFi network latency. These micro-measurements suggest that the default Android emulator used in our studies was the dominant component of the latency. The default Android emulator is known to be fairly slow, and multiple third party emulators (e.g., Genymotion [16]) are reported to provide significantly lower latency. Accordingly, we anticipate that this overall latency can be reduced to $\leq 150 \text{ msec}$, without any significant architectural modification of *Empath-D*.

7.1.2 End-to-end Latency of Virtual Hand. We also evaluated the latency between the physical movement of the user’s hand and the rendering of this movement in the VR display. To capture this time difference, we displayed a small circle, at a specific point on the display, on both the smartphone and the virtual phone. Users were instructed to swipe a finger on the screen to reach the circle. We measured, over 20 experiments, the time (no. of frames from the previously used high framerate camera) between the occlusion of the circle on the physical phone and the resulting occlusion in the virtual phone, computing an average latency of 117.46 msec ($SD = 20.44$). Additionally, we measured the component delays of this rendering process as: (i) reading an RGBD frame: 4.90 msec ($SD = 0.58$); (ii) phone tracking: 4.56 msec ($SD = 0.25$); (iii) hand tracking: 8.0 msec ($SD = 1.58$), and (iv) the VR responsiveness (the time from receiving the hand mesh till it is rendered on the VR display): 26.99 msec ($SD = 5.22$). The remaining latency, attributable to the WiFi network, is $\approx 73 \text{ msec}$, consistent with the measurements reported above.

7.2 Study Design for Usability Experiments

We then conducted user studies on the usability and real-world fidelity of our *Empath-D* implementation. The user study (approved

Table 2: Study Tasks and Conditions in Experiment 1

Task	Condition	Impairment	Simulator Type	Environment
T1-T4	A	none	none	Real
	B	Cataracts	Physical	Real
	C	none	none	Virtual
	D	Cataracts	Virtual	Virtual
	E	Glaucoma	Real	Physical
	F	Glaucoma	Virtual	Virtual

Table 3: Smartphone Interaction Tasks in Experiment 1

Task Type	Task Code	Task Description
Everyday Phone Use	T1	Perform a Calculation
	T2	Add an Alarm
	T3	Search, Save Image on Browser
Controlled Pointing	T4	Number Search and Point

by our institution’s IRB) consisted of 12 users (9 males) with no pre-existing uncorrected vision impairments. Users were aged 24–39, with a mean age of 30.3 years ($SD = 5$).

Study Tasks and Measures. We adopted a repeated measures design, with participants counterbalanced for condition order (see Table 2 for the conditions). Participants were asked to perform four different tasks split into two task types; *everyday phone use*, and *controlled pointing* (see Table 3). Users were asked to perform all tasks using two-handed interaction, holding the phone at a distance that they normally would during daily use. We chose two-handed interaction to eliminate for phone balancing that is typical in one-handed interaction given the typical size of today’s smartphones.

T1-T3 are everyday tasks users perform on a smartphone. They cover smartphone touch interaction of taps, swipes, and long press, on UI widgets such as keyboards, buttons and scrolling content. Users were asked to experience performing these tasks under six conditions, including under impairments (both using the physical hardware and the VR device). At the end of all three tasks (T1-T3), users completed the NASA-TLX[18] survey to indicate their perceived workload during task performance. T4, on the other hand, is a controlled pointing task experiment. Participants were given a stimulus number and then asked to click on the button with the corresponding number, as *quickly* and as *precisely* as they could. (See Figure 12 for a screenshot of the application used in this task.) Users repeated this task 80 times in succession, for each of the six conditions (A-F; see Table 2). We recorded the touch times and positions with the task app. We conducted a short semi-structured interview at the end of the study to understand users’ experiences with, and perceptions of, the physical and virtual impairment simulations.

Instruments: We compared *Empath-D* with a commercial physical impairment simulator [13]. To calibrate for *visual acuity*, we adapted a test similar to a Snellen eye test chart [42] - showing rows of letters with each lower row having a smaller font size. We first used the physical impairment simulator to obtain the minimum acceptable font size. Using the same test page in the VR, we applied




Figure 12: Screenshot of a test application for the pointing task

the impairment and gradually adjusted the severity until we hit the minimum acceptable font size. To calibrate the inner circle of clarity for glaucoma, we implemented an app that allows us to adjust the diameter of a coloured circle. We then used the physical impairment simulator for glaucoma, and adjusted the coloured circle to the point in which the circle reaches the fringe for clarity. We then calibrated the virtual glaucoma simulation in a similar manner. Three independent measurements for visual acuity and circle of clarity were taken from the research team and averaged to determine the final calibration parameters of font size = 12 *sp* and diameter = 60 *mm*.

7.3 *Empath-D* vs. Physical Smartphone

We first investigate whether the VR-based interaction is a sufficiently faithful replica of the real-world interaction that a user would have with a regular smartphone, *in the absence of any impairments*.

Touch Accuracy: In all six conditions, users were able to achieve high levels of button touch accuracy (see Table 4), with the accuracy being 98.8% ($SD = 1.67$) when the users interacted unimpaired with the VR device. Comparing the accuracies between the physical smartphone and the VR device, we noted that the VR condition had an accuracy of 99.12% ($SD = 1.32$) (across all 6 conditions), whereas the use of the physical smartphone provided 100% accuracy. In terms of the location accuracy, we noted a difference of 2.28 *mm* ($SD = 2.98$) between the use of *Empath-D* vs. a physical smartphone. This difference is well within the uncertainty associated with finger touch interactions, and thus demonstrates that user performance was equivalent across both *Empath-D* and a physical smartphone.

Perceived Workload: NASA-TLX scores indicated that the users did perceive significant differences in their workload using *Empath-D*, compared to use of the physical smartphone ($Z = 2.824$, $p = 0.005 < 0.05$). This does suggest that the navigating an app within the VR device does require greater cognitive effort than simply interacting with a regular smartphone. However, it is difficult to

Table 4: Accuracy of Button Touch Across All Users

Impairment	Environment	Accuracy (SD) %
None	Physical	100
	Virtual	98.79 (1.67)
Cataracts	Physical	100
	Virtual	99.09 (1.36)
Glaucoma	Physical	100
	Virtual	99.49 (0.82)

decipher whether this difference is due to *Empath-D*-specific issues, or a general lack of familiarity with VR devices.

We additionally investigated the subjective feedback captured by the semi-structured interview. 83% (10) of the users reported perceiving increased latency while using *Empath-D*, while 2 users indicated that they felt no noticeable latency difference. **However, all 12 users indicated that the performance of *Empath-D* was “acceptable”, and they would be able to use the *Empath-D* system for testing the usability of apps, as long as the apps do not require extremely low-latency interactions.** (3 users indicated that the system might not be usable for testing real-time games.)

7.4 *Empath-D* vs. Hardware Impairment Simulators

We now study the performance of *Empath-D* vis-a-vis impairments generated using commercially available hardware. Figure 11 shows the overhead of *Empath-D* under impairment conditions, demonstrating that *Empath-D* is able to operate without significant performance loss even in the presence of impairments.

Touch Accuracy: Table 4 enumerates the accuracy for the pointing task (T4) for two distinct impairments (*Cataract* & *Glaucoma*), for both the VR-based *Empath-D* system and the hardware impairment simulator. We see that, in the *Cataract* condition, *Empath-D* had a mean accuracy of 99.09%, which is virtually indistinguishable from that of the hardware device (100%). A similar pattern was observed for the *Glaucoma* impairment (99.49% for *Empath-D* vs. 100% for Hardware). In terms of the location accuracy, we noted a difference of 1.7 *mm* ($SD = 1.9$) (for *Cataract*) and 1.2 *mm* ($SD = 1.6$) (for *Glaucoma*) between the use of *Empath-D* vs. the impairment hardware. Once again, this difference is well within the uncertainty associated with finger touch interactions. These results provide strong evidence that *Empath-D* is able to emulate impairment conditions that are equivalent to that of dedicated, commercial hardware.

Perceived Workload: The numerical TLX scores indicated that there was no significant difference for *Cataracts*; however, the difference for *Glaucoma* was significant ($Z = 3.061$, $p = 0.002 < 0.05$), with users indicating a higher perceived workload for the VR device.

7.5 Motion sickness

At the end of the user study, we asked each participant if they felt discomfort or unwell. Only two of the twelve participants reported slight motion sickness while using *Empath-D*. Motion sickness may

arise from: (1) the use of the VR display itself, and (2) the latency from *Empath-D*. However, it is difficult to separate the two.

The effects of motion sickness are notably minor in our current prototype of *Empath-D*. The nature of our experimentation intensifies the use of the VR display, whereas practical use of *Empath-D* is likely to be more interspersed between app redesigns. We further discuss how we may improve on latency in Section 9.2 to reduce motion sickness that may result from the latency of *Empath-D*.

8 RELATED WORK

Designing for Inclusiveness. Newell et al. [31] pointed out that traditional user-centred design techniques provides little guidance for designing interfaces for elderly and disabled users due to the large variation amongst the type and degree of impairments. They also highlighted that the standard guidelines for designing disabled-friendly UIs are too general [30] and lacked empathy for users. For instance the WCAG 2.0 lists that the use of colour “is not used as the only visual means of conveying information, indicating an action, prompting a response or distinguishing a visual element”. This requires interpretation by the designer into specific designs in his application. Over the years, various accessibility design guidelines (such as WCAG 2.0 [3], IBM Accessibility Checklist [38], US Section 508 Standards [2]) and tools (aChecker [15]) have been proposed and refined. However, the problems pointed out by Newell are remained unsolved to a large extent, which hinders elaborate design for a target user group with a specific impairment.

Simulated Design. There exists prior work on helping UI designers design better interfaces for people suffering from vision impairments. Higuchi et al. [19] proposed a tool to simulate the visual capabilities of the elderly for the design of control panels, while Mankoff et al. [26] developed a tool to simulate a user with visual and motor impairments on the desktop screen. SIMVIZ [9, 47] uses the Oculus Rift VR device to simulate visual impairments to examine reading text on a smartphone. For audio modalities, Werfel et al. [47] simulated hearing ailments by using a pair of microphones with equalised headphones.

Different from prior works, *Empath-D* uses VR as the medium for immersive evaluation to 1) flexibly support wider groups of impaired users, and 2) allow naturalistic interactions with a *mobile phone in a virtual environment*. This novel approach supports ecological validity in testing applications and is key for mobile apps which go beyond the static settings of previous work.

While previous work has focused on simulation in single modality (visual or auditory), *Empath-D* is able to flexibly combine modalities to support any application type, ailment (visual, auditory, motor) and usage environment.

System Support for Accessibility. Modern mobile OSes provide accessibility support; in particular, it allows users with farsightedness to increase fonts and users with blindness to interact through vocal interfaces. Also, Zhong et al. enhanced Android accessibility for users with hand tremor by reducing fine pointing and steady tapping [48]. We believe *Empath-D* will significantly expand basic accessibility support of commodity devices and accelerates the design and deployment of various accessibility add-ons for different impaired users.

Testing of Mobile Applications. Recently there have been many systems, such as VanarSena [37], AMC [22], Puma [17], DynoDroid [25], DECAF [24], AppsPlayground [36], for automatically testing and identifying various types of UI and systems bugs in mobile applications. *Empath-D* takes a different approach in that we do not detect bugs after the application is developed and deployed. Instead, we allow the designer to test early iterations of the designs rapidly. In this way, we hope to reduce the pain of having to make significant UI changes at the end of the design cycle – or worse, end with an application that cannot be used effectively by the target impaired demographic.

9 DISCUSSION

Our current studies indicate the considerable promise of *Empath-D*, as a mechanism for rapid and empathetic evaluation of app usability. We now discuss some additional studies and issues that we intend to explore further.

9.1 User study with Designers

We conducted a short user study with two mobile app developers to qualitatively examine *Empath-D* in actual use. Both developers have previously worked to create an Android mobile application, which was used as the baseline for the study. The developers were tasked with redesigning the mobile app for the glaucoma-impaired under two conditions: 1) without *Empath-D*, but with materials describing glaucoma and showing functionally accurate examples of glaucoma, and 2) with the same materials, and *Empath-D*. Both developers agreed that *Empath-D* helped them improve their designs over the baseline condition. The developers reported that *Empath-D* allowed them to improve their designs in two ways: 1) they can focus their attention on re-designing particular problematic parts of the UI, and 2) they are able to appropriately calibrate their modifications (for instance increasing the font size may help, but text that is too large will also cause glaucoma sufferers to visually scan more, causing fatigue).

9.2 Dealing with Latency Issues

Our experimental studies indicate that users are able to utilise *Empath-D* effectively for “conventional” apps—i.e., those that typically involve sporadic interaction by users with UI elements, such as buttons and keyboards. The current end-to-end latency (of ≈ 200 msec) is not an impediment for high-fidelity evaluation of such apps. However, the participants also indicated that this latency (lag between user actions and rendering in the VR display) would pose a problem for highly latency-sensitive applications, such as games. At present, it is thus appropriate to state that *Empath-D* potentially needs additional optimisations to support such applications. The most obvious improvement would be to replace the default Android emulator with a faster, custom emulation engine—this is likely to reduce ≈ 100 msec of the delay budget.

The current implementation streams JPEG images (hand, emulator’s screen) from the intermediary computer to the VR smartphone. We plan to adopt a low-latency video streaming codec such as H.265 HEVC [43], which would help reduce networking and rendering latency. OS-level optimisations (e.g., preemptive priority

for inter-component messages) may be needed to support even lower latency.

Recently, several works have proposed techniques for achieving high-quality VR experience on mobile devices [7, 10, 21]. *Empath-D* could borrow some techniques to improve latency and video quality.

9.3 User Performance with VR Devices

Moreover, our user studies also indicated that the time for performing tasks (T1-T4) was marginally higher when using the VR environment, compared to the direct use of a real-world smartphone. More specifically, for the pointing task T4, there was an average difference of 654 *msec* in the task completion time using *Empath-D*, compared to the smartphone. In addition, anecdotal comments suggest that continued use of the VR device, for longer-lived sessions, might pose additional usability challenges. For example, a couple of users indicated some minor muscle fatigue, most likely as a result of using a 'heavy' VR device. It is an open question whether these issues will be mitigated over time, as VR devices become lighter and more ergonomic, and as users have greater familiarity with the use of VR devices.

9.4 Advanced Uses of *Empath-D*

Our current implementation of *Empath-D* supports the virtualisation of certain output modalities (specifically the display and audio) of the emulated app. The vision of *Empath-D* can be extended to create other richer interaction modes, often blending virtual and augmented reality (AR) settings. As an example, certain emulation conditions may need to generate and integrate synthetic sensor traces, to replace the real sensor traces from the smartphone—e.g., to mimic the user's movement in locations, such as forests and mountains, the phone's real GPS trace would need to be replaced by a synthetic GPS trace as in [27, 28]. Similarly, in some cases, the app itself might need to take *inputs* from the VR world—e.g., if the app was being used to magnify certain objects embedded in the VR world. While such use cases can be supported, they will require enhancements to the current *Empath-D* framework, and it is likely that the implementation may surface additional challenges, in terms of computational overhead and latency.

9.5 Developing Impairment Filters and Profiles

To demonstrate the viability of *Empath-D*, we focused on demonstrating the ability to simulate visual impairments and in particular cataracts and glaucoma. As we explored, these impairments have functional aspects that are commonly employed to characterise them, such as visual acuity or contrast sensitivity, and are often accompanied by standard tests such as the Snellen eye test chart [42] and Pelli-Robson contrast sensitivity chart [34] respectively. From examining the commercial physical impairment simulator and our experimentation, we believe that *Empath-D* has the ability to functionally simulate other impairments.

We recognise two important directions that *Empath-D* needs address to improve impairment simulation and use. First, impairment filters have to be developed in concert with medical professionals who are subject matter experts in the areas of the specific pathologies. This helps to develop a library of impairment filters. Second,

with verified impairment filters, we may create *impairment profiles*, which characterise groups of users with possibly overlapping requirements. For instance, a hypothetical impairment profile may calibrate for a demographic of a range of ages, sex, and percentage of the population who may have myopia and cataracts—both which affect visual acuity. With impairment profiles, app developers may easily select and understand the demographic to which they are designing for.

10 CONCLUSION

We presented the design and evaluation of *Empath-D*, a framework that allows app developers to 'step into the shoes' of impaired users, and perform an empathetic evaluation of their app interfaces. Our key idea is to utilise a virtual world (using a commodity VR device) to present an *impaired* view of the app's interface, while allowing the user to interact naturally with a real commodity smartphone in the physical world. Overcoming the current computational limitations (of the VR device and the Android emulator) required us to make careful system choices, such as (i) appropriate tradeoffs between the resolution and frame rate for rendering the virtual smartphone, (ii) subsampling of the mesh representing the user's hand and (iii) scaling up the size of the virtual smartphone to overcome the lower resolution of the VR device. User studies show that *Empath-D* is effective in (a) providing usability that is equivalent to using a real app (on a real smartphone), for applications that do not require ultra-low latency and (b) emulating impairments in a similar fashion to custom hardware devices. We believe that *Empath-D* can be a powerful new paradigm for effective bidirectional integration between real-world user actions and virtual worlds, and that this can enable additional immersive applications beyond just 'impairment emulation'.

11 ACKNOWLEDGEMENT

We are thankful to our shepherd Prof. Xia Zhou and all anonymous reviewers for their valuable reviews. This research is supported partially by Singapore Ministry of Education Academic Research Fund Tier 2 under research grant MOE2014-T2-1063, and by the National Research Foundation, Prime Minister's Office, Singapore under its IDM Futures Funding Initiative. All findings and recommendations are those of the authors and do not necessarily reflect the views of the granting agency, or SMU.

REFERENCES

- [1] [n. d.]. AGNES (Age Gain Now Empathy Systems). ([n. d.]). Retrieved 2018-04-13 from <http://agelab.mit.edu/agnes-age-gain-now-empathy-system>
- [2] [n. d.]. US Section 508 Standards. ([n. d.]). Retrieved 2018-04-13 from <https://www.section508.gov/>
- [3] 2008. Web Content Accessibility Guidelines (WCAG) 2.0. (11 December 2008). Retrieved 2018-04-13 from <https://www.w3.org/TR/WCAG20/>
- [4] 2016. Samsung Galaxy S7 Specifications. (2016). Retrieved 2018-04-13 from <http://www.samsung.com/global/galaxy/galaxy-s7/#/spec>
- [5] 2017. Samsung Gear VR Specifications. (2017). Retrieved 2018-04-13 from <http://www.samsung.com/global/galaxy/gear-vr/specs/>
- [6] 2018. SolvePnP, Camera Calibration and 3D Reconstruction, OpenCV. (2018). Retrieved 2018-04-13 from https://docs.opencv.org/2.4/modules/calib3d/doc/camera_calibration_and_3d_reconstruction.html
- [7] Omid Abari, Dinesh Bharadia, Austin Duffield, and Dina Katabi. 2017. Enabling High-Quality Untethered Virtual Reality. In *14th USENIX Symposium on Networked Systems Design and Implementation (NSDI 17)*. USENIX Association, Boston, MA, 531–544. <https://www.usenix.org/conference/nsdi17/technical-sessions/presentation/abari>

- [8] Ardalan Amiri Sani, Kevin Boos, Min Hong Yun, and Lin Zhong. 2014. Rio: A System Solution for Sharing I/O Between Mobile Systems. In *Proceedings of the 12th Annual International Conference on Mobile Systems, Applications, and Services (MobiSys '14)*. ACM, New York, NY, USA, 259–272. <https://doi.org/10.1145/2594368.2594370>
- [9] Halim Cagri Ates, Alexander Fiannaca, and Elke Folmer. 2015. Immersive Simulation of Visual Impairments Using a Wearable See-through Display. In *Proceedings of the Ninth International Conference on Tangible, Embedded, and Embodied Interaction (TEI '15)*. ACM, New York, NY, USA, 225–228. <https://doi.org/10.1145/2677199.2680551>
- [10] Kevin Boos, David Chu, and Eduardo Cuervo. 2016. FlashBack: Immersive Virtual Reality on Mobile Devices via Rendering Memoization. In *Proceedings of the 14th Annual International Conference on Mobile Systems, Applications, and Services (MobiSys '16)*. ACM, New York, NY, USA, 291–304. <https://doi.org/10.1145/2906388.2906418>
- [11] Kenny Tsu Wei Choo, Rajesh Krishna Balan, Tan Kiat Wee, Jagmohan Chauhan, Archan Misra, and Youngki Lee. 2017. Empath-D: Empathetic Design for Accessibility. In *Proceedings of the 18th International Workshop on Mobile Computing Systems and Applications (HotMobile '17)*. ACM, New York, NY, USA, 55–60. <https://doi.org/10.1145/3032970.3032981>
- [12] Eduardo Cuervo, Alec Wolman, Landon P. Cox, Kiron Lebeck, Ali Razeen, Stefan Saroiu, and Madanlal Musuvathi. 2015. Kahawai: High-Quality Mobile Gaming Using GPU Offload. In *Proceedings of the 13th Annual International Conference on Mobile Systems, Applications, and Services (MobiSys '15)*. ACM, New York, NY, USA, 121–135. <https://doi.org/10.1145/2742647.2742657>
- [13] Marshall Flax. 2018. Low Vision Simulators. (2018). Retrieved 2018-04-13 from <https://www.lowvisionsimulators.com/>
- [14] S. Garrido-Jurado, R. Mu noz Salinas, F.J. Madrid-Cuevas, and M.J. Marín-Jiménez. 2014. Automatic generation and detection of highly reliable fiducial markers under occlusion. *Pattern Recognition* 47, 6 (2014), 2280–2292. <https://doi.org/10.1016/j.patcog.2014.01.005>
- [15] Greg Gay and Cindy Qi Li. 2010. AChecker: Open, Interactive, Customizable, Web Accessibility Checking. In *Proceedings of the 2010 International Cross Disciplinary Conference on Web Accessibility (W4A) (W4A '10)*. ACM, New York, NY, USA, Article 23, 2 pages. <https://doi.org/10.1145/1805986.1806019>
- [16] Genymotion. [n. d.]. Genymotion Android Emulator. ([n. d.]). Retrieved 2018-04-13 from <https://www.genymotion.com/>
- [17] Shuai Hao, Bin Liu, Suman Nath, William G.J. Halfond, and Ramesh Govindan. 2014. PUMA: Programmable UI-automation for Large-scale Dynamic Analysis of Mobile Apps. In *Proceedings of the 12th Annual International Conference on Mobile Systems, Applications, and Services (MobiSys '14)*. ACM, New York, NY, USA, 204–217. <https://doi.org/10.1145/2594368.2594390>
- [18] Sandra G Hart and Lowell E Staveland. 1988. Development of NASA-TLX (Task Load Index): Results of empirical and theoretical research. In *Advances in psychology*. Vol. 52. Elsevier, 139–183.
- [19] Kazunori Higuchi, Yasuo Sakaguchi, Kazuhiko Sugiyama, and Tomoaki Nakano. 1999. Simulating the human vision of elderly for designing control panels. In *Systems, Man, and Cybernetics, 1999. IEEE SMC '99 Conference Proceedings. 1999 IEEE International Conference on*, Vol. 5. IEEE, 703–708.
- [20] Intel. 2016. Intel®RealSense™ Camera SR300 Product Specifications. (2016). Retrieved 2018-04-13 from <https://ark.intel.com/products/92329/Intel-RealSense-Camera-SR300>
- [21] Zeqi Lai, Y. Charlie Hu, Yong Cui, Linhui Sun, and Ningwei Dai. 2017. Furion: Engineering High-Quality Immersive Virtual Reality on Today's Mobile Devices. In *Proceedings of the 23rd Annual International Conference on Mobile Computing and Networking (MobiCom '17)*. ACM, New York, NY, USA, 409–421. <https://doi.org/10.1145/3117811.3117815>
- [22] Kyungmin Lee, Jason Flinn, T.J. Giuli, Brian Noble, and Christopher Peplin. 2013. AMC: Verifying User Interface Properties for Vehicular Applications. In *Proceedings of the 11th Annual International Conference on Mobile Systems, Applications, and Services (MobiSys '13)*. ACM, New York, NY, USA, 1–12. <https://doi.org/10.1145/2462456.2464459>
- [23] Gordon E Legge, Sing-Hang Cheung, Deyue Yu, Susana TL Chung, Hye-Won Lee, and Daniel P Owens. 2007. The case for the visual span as a sensory bottleneck in reading. *Journal of Vision* 7, 2 (2007), 9–9.
- [24] Bin Liu, Suman Nath, Ramesh Govindan, and Jie Liu. 2014. DECAF: Detecting and Characterizing Ad Fraud in Mobile Apps. In *11th USENIX Symposium on Networked Systems Design and Implementation (NSDI 14)*. USENIX Association, Seattle, WA, 57–70. https://www.usenix.org/conference/nsdi14/technical-sessions/presentation/liu_bin
- [25] Aravind Machiry, Rohan Tahiliani, and Mayur Naik. 2013. Dynodroid: An Input Generation System for Android Apps. In *Proceedings of the 2013 9th Joint Meeting on Foundations of Software Engineering (ESEC/FSE 2013)*. ACM, New York, NY, USA, 224–234. <https://doi.org/10.1145/2491411.2491450>
- [26] Jennifer Mankoff, Holly Fait, and Ray Juang. 2005. Evaluating accessibility by simulating the experiences of users with vision or motor impairments. *IBM Systems Journal* 44, 3 (2005), 505–517.
- [27] Chulhong Min, Seungchul Lee, Changhun Lee, Youngki Lee, Seungwoo Kang, Seungpyo Choi, Wonjung Kim, and Junehwa Song. 2016. PADA: Power-aware Development Assistant for Mobile Sensing Applications. In *Proceedings of the 2016 ACM International Joint Conference on Pervasive and Ubiquitous Computing (UbiComp '16)*. ACM, New York, NY, USA, 946–957. <https://doi.org/10.1145/2971648.2971676>
- [28] Chulhong Min, Youngki Lee, Chungkuk Yoo, Seungwoo Kang, Sangwon Choi, Pillsong Park, Inseok Hwang, Younghyun Ju, Seungpyo Choi, and Junehwa Song. 2015. PowerForecaster: Predicting Smartphone Power Impact of Continuous Sensing Applications at Pre-installation Time. In *Proceedings of the 13th ACM Conference on Embedded Networked Sensor Systems (SenSys '15)*. ACM, New York, NY, USA, 31–44. <https://doi.org/10.1145/2809695.2809728>
- [29] Produkt + Projekt Wolfgang Moll. [n. d.]. Age simulation suit GERT - the GERontologic Test suit. ([n. d.]). Retrieved 2018-04-13 from <http://www.age-simulation-suit.com/>
- [30] Alan Newell and Peter Gregor. 1988. Human computer interaction for people with disabilities. (1988).
- [31] Alan F Newell, Peter Gregor, Maggie Morgan, Graham Pullin, and Catriona Macaulay. 2011. User-sensitive inclusive design. *Universal Access in the Information Society* 10, 3 (2011), 235–243.
- [32] Nvidia. 2016. GeForce GTX 1080 Specifications. (2016). Retrieved 2018-04-13 from <https://www.geforce.com/hardware/desktop-gpus/geforce-gtx-1080/specifications>
- [33] National Institute on Aging. 2016. World's older population grows dramatically. (28 March 2016). Retrieved 2018-04-13 from <https://www.nih.gov/news-events/news-releases/worlds-older-population-grows-dramatically>
- [34] DG Pelli, JG Robson, et al. 1988. The design of a new letter chart for measuring contrast sensitivity. In *Clinical Vision Sciences*. Citeseer.
- [35] Android Open Source Project. 2017. SurfaceFlinger and Hardware Composer. (March 2017). Retrieved 2018-04-13 from <https://source.android.com/devices/graphics/arch-sf-hwc>
- [36] Vaibhav Rastogi, Yan Chen, and William Enck. 2013. AppPlayground: Automatic Security Analysis of Smartphone Applications. In *Proceedings of the Third ACM Conference on Data and Application Security and Privacy (CODASPY '13)*. ACM, New York, NY, USA, 209–220. <https://doi.org/10.1145/2435349.2435379>
- [37] Lenin Ravindranath, Suman Nath, Jitendra Padhye, and Hari Balakrishnan. 2014. Automatic and Scalable Fault Detection for Mobile Applications. In *Proceedings of the 12th Annual International Conference on Mobile Systems, Applications, and Services (MobiSys '14)*. ACM, New York, NY, USA, 190–203. <https://doi.org/10.1145/2594368.2594377>
- [38] IBM Accessibility Research. 2017. IBM Accessibility Checklist 7.0. (18 July 2017). Retrieved 2018-04-13 from http://www-03.ibm.com/able/guidelines/ci162/accessibility_checklist.html
- [39] Justin B. Rousek, Sonja Konecny, and M. Susan Hallbeck. 2009. Simulating Visual Impairment to Detect Hospital Wayfinding Difficulties. *Proceedings of the Human Factors and Ergonomics Society Annual Meeting* 53, 8 (Oct. 2009), 531–535.
- [40] Samsung. 2014. Samsung Galaxy S5 Specifications. (2014). Retrieved 2018-04-13 from <http://www.samsung.com/uk/smartphones/galaxy-s5-g900f/SM-G900FZKABTU/>
- [41] Alvy Ray Smith and James F. Blinn. 1996. Blue Screen Matting. In *Proceedings of the 23rd Annual Conference on Computer Graphics and Interactive Techniques (SIGGRAPH '96)*. ACM, New York, NY, USA, 259–268. <https://doi.org/10.1145/237170.237263>
- [42] Herman Snellen. 1873. *Probuchstaben zur bestimmung der sehschärfe*. Vol. 1. H. Peters.
- [43] G. J. Sullivan, J. R. Ohm, W. J. Han, and T. Wiegand. 2012. Overview of the High Efficiency Video Coding (HEVC) Standard. *IEEE Transactions on Circuits and Systems for Video Technology* 22, 12 (Dec 2012), 1649–1668. <https://doi.org/10.1109/TCSVT.2012.2221191>
- [44] Unity Technologies. [n. d.]. Unity. ([n. d.]). Retrieved 2018-04-13 from <https://unity3d.com/>
- [45] Sam Tregillus and Elke Folmer. 2016. VR-STEP: Walking-in-Place Using Inertial Sensing for Hands Free Navigation in Mobile VR Environments. In *Proceedings of the 2016 CHI Conference on Human Factors in Computing Systems (CHI '16)*. ACM, New York, NY, USA, 1250–1255. <https://doi.org/10.1145/2858036.2858084>
- [46] Zhou Wang, Alan C Bovik, Hamid R Sheikh, and Eero P Simoncelli. 2004. Image quality assessment: from error visibility to structural similarity. *IEEE transactions on image processing* 13, 4 (2004), 600–612.
- [47] Fabian Werfel, Roman Wiche, Jochen Feitsch, and Christian Geiger. 2016. Empathizing Audiovisual Sense Impairments: Interactive Real-Time Illustration of Diminished Sense Perception. In *Proc. of AH*.
- [48] Yu Zhong, Astrid Weber, Casey Burkhardt, Phil Weaver, and Jeffrey P. Bigham. 2015. Enhancing Android Accessibility for Users with Hand Tremor by Reducing Fine Pointing and Steady Tapping. In *Proceedings of the 12th Web for All Conference (W4A '15)*. ACM, New York, NY, USA, Article 29, 10 pages. <https://doi.org/10.1145/2745555.2747277>
- [49] George J. Zimmerman. 1979. Zimmerman Low Vision Simulation Kit. (1979). Retrieved 2018-04-13 from <http://www.lowvisionsimulationkit.com/>