

Path Tracing in 2D, 3D, and Physicalized Networks

Michael J. McGuffin , Ryan Servera, and Marie Forest

Abstract—It is common to advise against using 3D to visualize abstract data such as networks, however Ware and Mitchell’s 2008 study showed that path tracing in a network is less error prone in 3D than in 2D. It is unclear, however, if 3D retains its advantage when the 2D presentation of a network is improved using edge-routing, and when simple interaction techniques for exploring the network are available. We address this with two studies of path tracing under new conditions. The first study was preregistered, involved 34 users, and compared 2D and 3D layouts that the user could rotate and move in virtual reality with a handheld controller. Error rates were lower in 3D than in 2D, despite the use of edge-routing in 2D and the use of mouse-driven interactive highlighting of edges. The second study involved 12 users and investigated data physicalization, comparing 3D layouts in virtual reality versus physical 3D printouts of networks augmented with a Microsoft HoloLens headset. No difference was found in error rate, but users performed a variety of actions with their fingers in the physical condition which can inform new interaction techniques.

Index Terms—3D printing, augmented reality, data physicalization, graph visualization, path finding, path following, tangible.

I. INTRODUCTION

ADVANCES in virtual reality (VR) and augmented reality (AR) headsets have fueled interest in 3D graphics for information visualization and ‘immersive analytics’ [1], [2], [3], [4]. For datasets with a natural 3D embedding, such as 3D medical images or 3D models of buildings, there is clear value in 3D visualization. On the other hand, for abstract data such as networks [5] or multidimensional multivariate data [6], the use of 3D is often advised against [7] due to previous studies that have found 2D to be better (e.g., [8], [9]). One counter-example is the task of path tracing in networks, which was shown in a carefully designed experiment [10] to be less error-prone when using a 3D layout with stereo and motion parallax depth cues. Practical implications remain unclear: should networks be embedded in 3D? The lack of clear implications is partly because the previous study did not allow the user to control their view

Manuscript received 25 July 2022; revised 20 December 2022; accepted 4 January 2023. Date of publication 23 January 2023; date of current version 26 June 2024. This work was supported by Microsoft Research and NSERC. Recommended for acceptance by J. Fekete. (Corresponding author: Michael J. McGuffin.)

This work involved human subjects or animals in its research. Approval of all ethical and experimental procedures and protocols was granted by The Research Ethics Committee at ÉTS under Application No. H20190102.

The authors are with the École de technologie supérieure, Montreal, QC H3C 1K3, Canada (e-mail: michael.mcguffin@etsmtl.ca; servera.ryan@gmail.com; marie.forest@etsmtl.ca).

This article has supplementary downloadable material available at <https://doi.org/10.1109/TVCG.2023.3238989>, provided by the authors.

Digital Object Identifier 10.1109/TVCG.2023.3238989

in 3D, nor leverage interaction with an input device, nor benefit from modern edge-routing [11] in 2D (edges in [10] were simply drawn as straight line segments, resulting in more occlusion). We extend this previous work by experimentally comparing path tracing under new conditions that are more relevant to modern VR/AR headsets, and find that 3D remains advantageous over 2D in terms of error rate.

Our work is also informed by the recent trend of physicalization of data [12], [13], e.g., via 3D printing. The ability to touch a tangible rendering of data can yield advantages over an equivalent virtual 3D visualization [9]. This is likely in part because the user’s fingers can mark elements in a physicalization, to “remember” a location and facilitate comparison with other elements. To date, however, there have been no empirical evaluations of physicalizations of networks with 3D layouts. Physicalizations also open the intriguing possibility of being augmented with virtual information displayed using an AR headset. Our 2nd study is the first to experimentally evaluate a physical network with a 3D layout, and also the first to use AR to augment physicalized networks, an example of what we call **augmented physicalization**. We observed users touching physical networks in a variety of ways, suggesting new interaction techniques.

A companion video demonstrating our experimental conditions is at <https://youtu.be/T3cRVuxvkOo>

II. BACKGROUND

The choice between visualizing data in 2D and 3D is most controversial when the data is abstract, having no intrinsic embedding. Previous works offer much advice [14, Section 3], [7], [15], [16], [17], [18]. Problems with 3D [8], [15] include occlusion hiding information, ambiguous depth, distortion due to perspective, complex navigation, and difficulty reading text. Advantages of 3D include having an additional visual channel for encoding a variable, ability to have multiple views in 3D space [1], [19], and depth cues sometimes making it easier to find information [20].

VR and AR [21] headsets provide immersion, 3D input, and enhanced depth cues, in particular head-coupled motion and stereo disparity [20], [22], [23]. Recent uses of these platforms for information visualization include ImAxes [24] and Data-Hop [25]. Input to such systems is often via handheld controllers or whole hands. Some systems use tangible input devices [26] designed to better match visualization tasks. Recent examples include a tangible cutting plane [27], tangible axes [28], [29], and a globe of the earth [30], which is both a tangible input device and a physicalization [12] of geographic data. Other systems

present virtual information on top of a tangible physicalization *without* using a headset [31], [32], [33].

The next two sections focus on previous empirical evaluations of 2D vs 3D embeddings of networks, and visualization vs physicalization of data.

A. Comparing Networks in 2D and 3D

Network visualization constitutes a large literature [5], [34], [35], [36], [37]. Some previous works have evaluated networks embedded in 3D [38], [39], [40] but without focusing on the question of comparing a flat 2D layout (on a plane) vs fully 3D layout (with nodes distributed throughout a volume). In Kwon et al. [38], the nodes of the network were laid out on a curved surface, whereas the other works [39], [40] did not employ a flat 2D layout. Irani and Ware [41] compared network-like structures rendered with 2D and 3D depth cues, but always with a flat layout of the nodes.

Other works have compared flat 2D and fully 3D layouts for tasks related to highlighted subsets of nodes [42] and counting clusters of nodes [43], [44].

Our work extends previous studies of *path tracing*, also called path finding or path following, where users identify a sequence of nodes. This is a standard task [45] with networks, used in multiple previous studies [38], [39], [40], [46], [47] that were not focused on comparing flat 2D vs fully 3D layouts. Path tracing has also been used to compare monoscopic and stereoscopic viewing of structures resembling angiograms [48], [49]. Studies comparing 2D and 3D layouts for path tracing within networks are reported in [10], [50], [51]. Two of these [50], [51] used networks with random layouts, making them less relevant to real visualizations. The most recent [10], summarized below, is also the most carefully designed.

Unlike our current work, none of the previous studies involving path tracing employed edge-routing in their 2D network visualizations.

1) *W+M: Ware and Mitchell (2008)*: Ware and Mitchell [10] report two studies, and we focus on the first of these, which we abbreviate as W+M. For each trial, two nodes were highlighted. Users had to indicate if the shortest path between the two nodes was 2 or 3 edges, a forced choice response. There were 5 viewing conditions: 2D layout, or a 3D layout with {monoscopic, stereoscopic} projection \times {no motion, motion in the form of automatic rotation at 10° per second}. Users could not actively change their view, either by moving their head nor through any input device. Viewing time was limited to 5 seconds per trial (i.e., a rotation of 50° in the conditions with ‘motion’). Results showed that the highest error rate occurred in the 2D and 3D monoscopic conditions; and the lowest error rate was with 3D stereoscopic + motion, demonstrating an advantage of the fully 3D condition over 2D.

Because W+M was more focused on low level perception, there is room for studies of path tracing in more realistic settings, where the user controls their view and their exposure time, and benefits from interaction.

B. Evaluating Data Physicalizations

Jansen et al. [9] evaluated physical barcharts. Their first study compared 4 conditions: 2D virtual barcharts, 3D virtual barcharts displayed monoscopically and stereoscopically (rotation performed with a mouse in both 3D virtual conditions), and 3D physical barcharts that users could touch. In terms of time, 2D was the best, but more interestingly, 3D physical was the 2nd best. Their second study investigated why 3D physical might be better than 3D virtual, comparing 4 conditions: (1) virtual 3D monoscopic with mouse for rotation, (2) virtual 3D monoscopic with a prop for more direct rotation, (3) physical 3D without being allowed to touch, and (4) physical 3D with touch allowed. In terms of time, the 4th condition was best, and the 3rd condition was 2nd best.

Danyluk et al. [52] compared physical barcharts that could be touched versus virtual barcharts of different sizes, and found that physical models allowed tasks to be completed faster and were preferred by users. They also demonstrated that adding interactive techniques for annotation and filtering boosted the performance of virtual models, and allude to future hybrid approaches that would combine physical models with virtual visuals.

In a path tracing task with sets of “spaghetti” noodles, Bérard and Louis [53] found no advantage of physical models over virtual representations. Their physical models were encased in plastic cubes and could not be directly touched.

Drogemuller et al. [47] evaluated networks with a flat, 2D layout, ranging from 16 to 24 nodes in size, with 3 tasks, comparing 4 conditions: virtual on-screen (“graphical-only”), and physical printouts that could be seen (“visual-only”) or touched (“haptic-only”) or both (“visual-haptic”). Users preferred the physical printouts that could be seen and touched, but within the path tracing task, no differences are reported in error rates between graphical-only, visual-only, or visual-haptic.

Our work is the first to empirically evaluate physicalized networks with 3D layouts. Also, unlike previous work, our physicalized networks were augmented with an AR headset to indicate the task.

III. OVERVIEW OF BOTH STUDIES

The following questions motivate our work: is path tracing easier in networks presented in 3D than in 2D when edge-routing is used in 2D, and when the user can interact with the network using a pointing device? Also, is path tracing easier with a physical 3D representation?

In both our studies, which were planned in parallel, the task was to find the length (between 1 and 5 edges) of a shortest path between two end-nodes indicated by the system. The user’s non-dominant hand (NDH) held and repositioned the network, because this matches the use of the NDH in the kinematic chain model [54], and because there is some evidence that rotation via a handheld prop is superior to using a mouse for the same purpose [9], and because it provides an easy-to-understand way to simultaneously pan and zoom within a 2D layout, by simply translating the layout sideways or holding it closer or farther




Fig. 1. In our two studies, users had to find the distance (in edges) between the two nodes indicated in red. Left and Middle: the VR2DHilite and VR3DHilite conditions, where a mouse moved a green cursor, highlighting edges incident on the node under the cursor. Right: the AR3DTouch condition, where the user could touch a 3D printout, and a Microsoft HoloLens augmented reality (AR) headset indicated the nodes.

away. In addition, in some conditions, the user’s dominant hand (DH) could move a mouse cursor over nodes (causing incident edges to highlight) or touch a physical 3D printout of the network.

In all conditions of both studies, the NDH activated a trigger button to open a radial button to provide the user’s answer from 1 to 5. A radial menu was used so that every answer would take the same amount of time to select.

A. Comparison With W+M

W+M focused on “visual searches that could be conducted rapidly” [10]. Our studies are designed to be more realistic and relevant to VR/AR headsets. Our participants can freely change their view of the network by moving their head and hand. In W+M, the user’s field-of-view (FOV) was $\approx 26 \times 16^\circ$ per eye, much smaller than the FOV of the VR headset used in our Study 1, and slightly smaller than the AR headset in our Study 2. In our studies, users also had control over their exposure time, further increasing the realism of the task.

In addition, our 2D conditions use a state-of-the-art routing algorithm [11] (Fig. 1 (Left) and Fig. 2), to make better use of space and reduce ambiguity. Our experimental task involves paths that are longer. Our Study 1 also involves conditions with interactive highlighting, double the number of participants of W+M, and was preregistered (Section IV-C).

IV. STUDY 1: 2D AND 3D VIRTUAL

Study 1 was preregistered [55], done in VR, and crossed the dimensionality {2D, 3D} of the network layout with the use of a mouse to highlight edges.

A. Main Conditions

The previous W+M study [10] found that 3D outperformed 2D, but it is plausible that this could change if the 2D condition is improved with edge-routing, or if the user can use a pointing

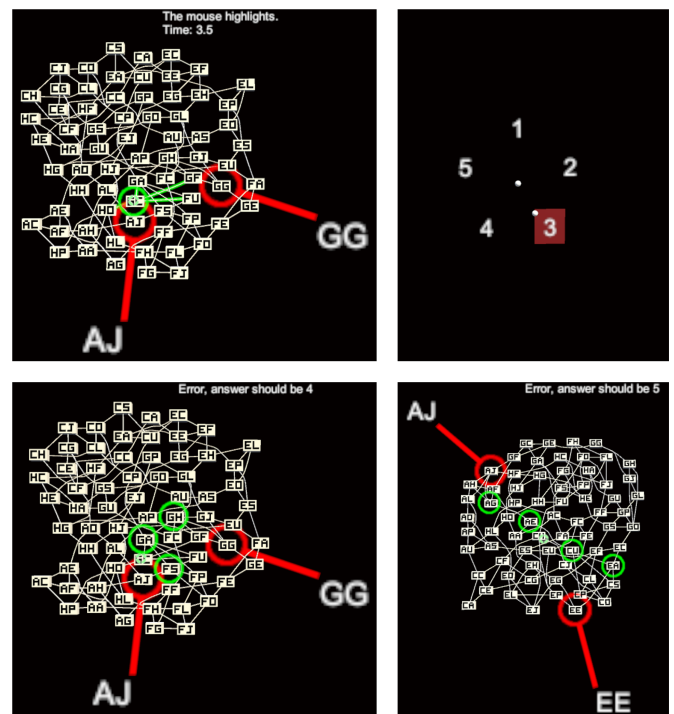


Fig. 2. In study 1, in the VR2DHilite condition, the user can move the mouse cursor over a node, causing incident edges to highlight in green (Top Left). Once the user has determined the distance in edges between the two nodes indicated in red (AJ and GG), they hold down the trigger button with their non-dominant hand (NDH), causing a radial menu to appear (Top Right). By tilting their hand down and right, they select “3” in the menu, and then release the trigger button to complete their answer. In this example, their answer is wrong, so the system displays error feedback: a correct shortest path of length 4 is highlighted (Bottom Left). Another example of error feedback (Bottom Right) is for the network shown in Fig. 1 (Left), where the shortest path has length 5.

device for simple interaction with the network. The simplest interaction we could think of that might help with path tracing is for the edges incident on a node to highlight when the user hovers over that node with a pointing device. This led to our choice of

conditions for Study 1. The independent variable *MainCondition* has four possible values:

- VR2D: Virtual network displayed in VR with 2D layout. The non-dominant hand (NDH) holds a controller to position the network with 6 degrees of freedom (DoF).
- VR2DHilite: Same as previous, but with a mouse in the dominant hand (DH) used to highlight edges.
- VR3D: Virtual network displayed in VR with 3D layout. The NDH holds a controller to position the network with 6 DoF.
- VR3DHilite: Same as previous, but with a mouse in the DH used to highlight edges.

We decided to not compare with 2D conditions on a desktop screen (without VR headset) because this would have introduced confounds in having different display hardware, and also potentially different input devices. The VR handheld controller in the user's NDH provides an easy-to-understand and quick way to reposition a 3D network layout, and also simultaneously pan and zoom when examining a 2D network layout, and there is no equivalent NDH input on standard desktop PCs.

B. Task

The experiment consisted of a sequence of trials where the task required the user to find the length, in edges, of a shortest path between two nodes (the path's *end-nodes*) in the network. The independent variable *PathLength* ranged from 1 to 5, where 1 means the end-nodes are neighbors. The shortest path was not necessarily unique.

At the start of each trial, the two end-nodes were indicated with red callout line segments, as well as with red circular rings. The user then examined the network by repositioning and rotating it with their NDH. Bringing the network closer to their eyes allowed them to effectively “zoom in” to see more detail. In some conditions (VR2DHilite and VR3DHilite), the user could also move a mouse with their DH, causing a cursor to hover over different nodes. Whichever node was under the mouse cursor was highlighted in green, as were all the edges incident on that node. Users were told that using the mouse in these conditions was not mandatory, but that it might allow them to more easily find the shortest path and confirm its length. The mouse buttons served no purpose.

Once the user thought they knew the answer, they pressed a trigger button using the index finger of their NDH to open up a radial menu containing the answers 1 through 5 (Figs. 2 (Top Right) and 3 (Top Right)). To select within this menu, the user tilted their head slightly in the direction of the desired answer and released the trigger button. Releasing the trigger button without tilting their head dismissed the radial menu and allowed the trial to continue. The trial ended only when the user made a selection within the menu. A text message appeared immediately after to inform the user if their answer was correct or not, and in the latter case, the text also indicated the correct answer, and the system highlighted a shortest path (Figs. 2 (Bottom) and 3 (Bottom)). The system then moved on to the next trial.

Users were aware that they could take their time during the warmup trials at the start of each condition, but after these




Fig. 3. Study 1, the VR3DHilite condition. The first three images show a trial ending with error feedback displaying a correct shortest path of length 3 (Bottom Left). The last image (Bottom Right) shows a correct shortest path of length 5 for the network shown in Fig. 1 (Middle).



Fig. 4. The equipment for Study 1: HTC Vive headset and controller (held in the NDH), mouse (for the DH) and keypad (to advance to the next trial after a break).

warmup trials, they were instructed to complete trials as quickly as possible with no errors. The following was also explained to each user. The correct path length is always at most 5, and, sometimes, the shortest path is quite difficult to find. After 20 seconds into a trial, the text instructions displayed by the headset turn red. Once this happens, the user is free to continue searching for a shortest path if they so desire, but a reasonable strategy after 20 seconds would also be to simply estimate an answer such as 5, or perhaps 4, even if the user has not found a path of that length. (This was explained to avoid having users spend too much time searching for difficult paths.) On the other hand, if a user answers quickly and incorrectly, the system imposes a “punishment” of a delay of up to 15 seconds before proceeding to the next trial. (If t was the time in seconds taken by the user to give an incorrect answer, the precise delay imposed after the trial was $\max(\min(20 - t, 15), 5)$, i.e., a decreasing ramp function clamped between 15 and 5 seconds.) This disincentivizes a user from answering sloppily to complete the experiment faster.

The software asked users to take breaks between conditions. Users could indicate that they were ready to proceed by pressing a key on a keypad with their DH. A Lego brick attached to the key made it easier to feel when the user was wearing the VR headset.

C. Pilot and Predictions

A pilot was performed with 6 users, after which minor tuning to the protocol was made, and a preregistration [55] was archived to declare the number of users to recruit, criteria for including participants, predictions to test, and the R script for plotting data and testing predictions.

Two predictions were preregistered: first, that the error rate (averaged over trials of *PathLength* 2, 3 and 4, and averaged over conditions with and without the mouse) would be smaller in 3D than in 2D; and second, that the error rate (averaged over trials of *PathLength* 2, 3 and 4, and averaged over 2D and 3D conditions) would be smaller with the mouse than without the mouse. The R script tests each of these predictions by computing a single error rate for each user and each subset of conditions, not including warmup trials, and then performing a paired sample *t*-test.

The W+M study [10] found that 3D yielded a smaller error rate, possibly because stereo and motion disambiguate edges. A second mechanism that could play a role is that shortest paths in 3D tend to follow a more straight line (hence, are easier to perceive) than in 2D. This 2nd mechanism may not have been at play in W+M because they “selected paths in such a way that the mean euclidean distance between start nodes and end nodes was the same” [10], regardless of whether the path was 2 or 3 edges long. However, both mechanisms could benefit 3D in our study, since we do not hold this euclidean distance constant. The edge-routing in our 2D layouts can also make 2D layouts easier to read, but the paths are still “less straight” than in 3D. The reason our predictions about error rates exclude *PathLength* 1 and 5 is that those cases tend to be easier for users: in the case of 1, the nodes are often clearly adjacent, and in the case of 5, the user knows that the length cannot be greater than 5 and can therefore guess a length of 5 when the path is difficult to find.

D. Mouse Cursor in 2D and 3D

To allow the user to hover over a node for interactive highlighting, we wanted to use the same pointing device in the 2D and 3D conditions, for simplicity and consistency across conditions. Although a 6 DoF handheld controller could have been used for pointing, the mouse is very often used for raycast pointing in 3D, whereas controllers are rarely used for pointing in 2D. Furthermore, a 2D mouse may be easier and less tiring to control than a handheld controller, because the user’s arm can partially rest on the desk, and the depth dimension is automatically handled by the software. We therefore implemented a variant of raycast pointing with a 2D mouse.

In our variant of raycast pointing, we wanted the user to be able to position the mouse cursor over a node with their DH, and then move or rotate the network with their NDH while the mouse cursor remains ‘stuck’ on the same node. Therefore, the mouse cursor’s position is stored in the network’s local 3D

space, and the DH applies relative displacements. Whenever the DH moves the mouse, a proportional translation is applied to the mouse cursor parallel to the camera plane (i.e., the plane perpendicular to the camera’s forward direction). After applying this translation, a ray is cast from the camera position through the cursor, and if this ray encounters one or more nodes, our software finds the intersection between the ray and the node closest to the camera, and moves the cursor to that intersection. The user therefore sees the cursor automatically jump forward or backward to stick to the nearest node, but these jumps only happen if the mouse is being moved. If the user is only using their NDH to reposition or rotate the network, no automatic jumps happen, and the cursor retains its position in the network’s local space.

Like EZCursorVR [56], the size of our mouse cursor is scaled to be bigger when the cursor is farther from the camera, so that the projected size of the cursor on the camera plane appears constant.

E. Virtual Networks

We generated 10 networks. For each network, we computed its layout in 3D, and projected the 3D node positions down to a plane to obtain a layout in 2D.

Each network has 70 nodes and 140 edges (hence an average degree of 4), and was generated with a Watts-Strogatz [57] small-world synthesis algorithm. The algorithm begins by constructing a regular ring lattice of 70 nodes each with degree 4. Each edge is then randomly rewired with 20% probability. The average degree distribution that resulted over the 10 networks was: 1, 15.7, 38.2, 12.9, 1.8, and 0.4 nodes of degree 2 through 7, respectively.

Each of the 70 nodes was assigned a unique 2-character string label such as “AA,” “FE,” or “HL”.

Layout of nodes was performed in two passes. The first pass uses stress majorization (Eq. 12 in [58]) to position the nodes in 3D. Projecting node positions down to a 2D plane results in overlap between labels, hence a second pass applies repulsive forces between nodes whose labels overlap in the 2D plane, pushing nodes away from each other in the horizontal plane. The new positions are saved in both 2D and 3D. Thus, the 2D layout uses the projected coordinates of the 3D layout. This was done to ensure that the user could read all node labels of the 3D networks from the “front” of the network with only tiny rotations, and to ensure that the 2D and 3D layouts were as similar as possible from the front.

Next we compute the layout of edges. In the 3D case, each node is modeled as an elongated box (7×7×15 mm) with a text label on one side, and each edge is modeled as a single segment (3 mm thick), with each edge’s endpoint connected to the center or extremity of the node’s box in such a way as to avoid extending through the labeled face of the box. In 2D, nodes are 7×11 mm rectangles, and each edge is a multi-segment polygonal line (0.5 mm thick), whose layout is computed using the MSAGL (Microsoft Automatic Graph Layout) library [59] based on [11]. The use of thinner edges in 2D was necessary to maintain the same size of nodes and same distances between nodes.

The 2D and 3D virtual networks are displayed with the same color and font and have the same width and height.

F. Hardware

We used an HTC Vive headset, which has a 2160×1200 resolution (1080×1200 per eye) and $\approx 110^\circ$ FOV. We measured a framerate of 90 fps. The handheld controller was held in the user's NDH. The headset was connected to a PC with an Intel i7-8700K 6-core CPU at 4.7 GHz, water cooling, and Nvidia GeForce GTX 1080 Ti GPU. A Logitech M100 mouse (with default acceleration settings in Microsoft Windows 10) on a gaming mouse pad was held in the user's DH.

G. Measurement of Rotation

In the 3D conditions, we measured how much the user looked at the network from different points of view. Let H_t be the position of the user's head at time t , and let N_t be the pose (position and orientation) of the network at time t . We define a direction vector $d = d(H, N)$ as a function of H and N , where d points from the network to the head, and where the components of d are computed in the local space of the network (i.e., a change in the position of the head, or in the position of the network, or in the orientation of the network, will each change the components of d). We compute $d_t = d(H_t, N_t)$ for each frame during a trial, and at the end of the trial, we compute the mean direction \bar{d} , and compute and record the standard deviation of the angles between all the d_t and \bar{d} . This standard deviation is an overall measure of how much the user looked at the network from different points of view.

We also computed how much the user rotated the network with their hand, versus how much they moved their head to look at the network from different points of view, expressed as two percentages that sum to 100%. To compute these percentages, we compute the directions $d_{t,H} = d(H_t, N_{t-1})$ and $d_{t,N} = d(H_{t-1}, N_t)$ that would have resulted if only the head, or only the network, respectively, had moved. We find the angle α_H between d_{t-1} and $d_{t,H}$, and the angle α_N between d_{t-1} and $d_{t,N}$, and then define the contribution of the head motion to the rotation as $\alpha_H / (\alpha_H + \alpha_N)$, and the contribution of the network's motion as $\alpha_N / (\alpha_H + \alpha_N)$. These fractions are computed for each frame of the trial (except the first frame), averaged over the entire trial, and recorded as percentages.

H. Protocol

Equipment was disinfected prior to each user session. At the start of each session, after signing a consent form, users had their interpupillary distance (IPD) measured, as well as their stereo acuity, which was assessed using the 'circle test' of the FLY stereo acuity test by Vision Assessment Corporation. (This resulted in a score on a scale of 10, corresponding to a disparity of 400, 200, 160, 100, 63, 50, 40, 32, 25, or 20 seconds of arc.) Users then filled out a pre-questionnaire, and were also shown several printouts of example trials, with conditions in random order, to explain the task and test their understanding of path length. The equipment for the experiment was then explained to

the user, the headset was adjusted for comfort, and the IPD of the headset was set to the value measured earlier. After the trials were completed, a post-questionnaire was filled out.

I. Users

Throughout our text, we refer to the participants as "users" for brevity. A sample size of 34 users was chosen in the preregistration, not including the pilot participants. This number was chosen to achieve a power of 0.8 at $\alpha = 0.05$ for a medium effect size of 0.5, as calculated using the G*Power software [60] and also with an online calculator [61].

Of the 34 users, 24 were male, 10 female; 28 right handed, 4 left handed, 2 ambidextrous, but all with a habit of using the mouse with their right hand; age 19 to 45 years (average 24.7); IPD 54 to 70 mm (average 63.7); stereoacuity test scores 3/10 to 10/10 (average 8.5).

J. Design

Each user experienced the 4 levels of *MainCondition* {VR2D, VR2DHilite, VR3D, VR3DHilite} in random order. For each *MainCondition*, the user performed 10 warmup trials in random order (5 levels of *PathLength* \times 2 repetitions) with the warmup network, followed by 20 trials in random order (5 levels of *PathLength* \times 4 repetitions) with another network chosen at random, followed by another 20 trials with another network. There were a total of 34 users \times 4 levels of *MainCondition* \times 2 networks \times 5 levels of *PathLength* \times 4 repetitions per trial = 5440 trials, not counting warmup trials and not counting pilot data. Each session with a user lasted ≈ 1.5 hours.

K. Results

For these results, advice was adapted from Dragicevic [62]. Although we report some p values, we do not emphasize null hypothesis significance testing (NHST) as we wish to avoid misleading, dichotomous thinking ("Tip 25" in [62]). We present effect sizes visually and with CIs (Tips 15, 16), where the CIs are computed using one (averaged) value for each (user, condition) pair (Tip 9); and we clearly distinguish between pre-experiment predictions and post-hoc exploratory data analysis, to avoid HARKING (Hypothesizing After the Results are Known) and p -hacking.

All CIs are 95%. The CIs for error rates were calculated using bootstrapping, to prevent them from falling outside the [0, 100%] range. For other variables such as time, CIs were calculated using the t -distribution.

Figs. 5 and 6 show the results of testing the preregistered predictions. The fact that the zero line falls far outside the CI of the difference in the first figure, and barely intersects the CI of the difference in the second figure, is reflected by the very small p value in the first case and a p value somewhat larger than 0.05 in the second case. We thus have strong evidence that 3D results in a lower error rate, and limited evidence that the error rate is reduced by interactive highlighting of edges with the mouse.




Fig. 5. Error rates in study 1, excluding *PathLength* 1 and 5. “2D” is the union of VR2D and VR2DHilite; “3D” is the union of VR3D and VR3DHilite. Each dot is the average for one user, and the bars show 95% confidence intervals (CIs), computed from the 34 users. A paired t -test yields $p < 0.0000005$. The 3D conditions resulted in a lower error rate than 2D, confirming the first of our preregistered predictions.




Fig. 6. Error rates in study 1, excluding *PathLength* 1 and 5. “WithHilite” is the union of VR2DHilite and VR3DHilite; “WithoutHilite” is the union of VR2D and VR3D. A paired t -test yields $p < 0.09$, providing limited evidence of our second preregistered prediction, that the mouse improves error rates.




Fig. 7. Error rates in study 1 by *MainCondition*, including all *PathLength* values 1-5.

1) *Subjective Feedback and Exploratory Data Analysis:* In subjective comments, 10 out of the 34 users described 3D positively as compared to 2D, saying that 3D makes the task easier, simpler, less confusing, fun, making pathways clearer, only requiring rotation. One of these users mentioned longer paths as being clearer in 3D than in 2D.

Fig. 7 shows error rates by *MainCondition*, suggesting that interactive highlighting of edges with the mouse helped in 2D but not in 3D. This is further supported by comments made by users and by the subjective results in Fig. 8. (For ease of comparison, the results of Study 2, discussed later, are presented alongside several figures.) There, we see that VR2D was the least favorite condition, and that VR2DHilite required less mental effort, produced less frustration, and better enabled the task. However, VR3D was the most favorite condition, not VR3DHilite. 14 out of the 34 users described the highlighting of edges with the mouse in 3D (VR3DHilite) in negative terms, such as being

	MainCondition	Number of users		Likert score (1 to 7)			
		Most favorite	Least favorite	Mental effort	Physical effort	Enabled task	Frustration
Study 1 (34 users)	VR2D	2	19	5.1	2.8	3.9	4.7
	VR2DHilite	9	2	4.1	3.1	5.2	3.4
	VR3D	17	6	4.4	3.6	5.5	3.6
Study 2 (12 users)	VR3DHilite	6	7	4.3	3.8	5.5	3.8
	VR3D	12	0	3.6	2.2	6.3	2.3
	AR3DTouch	0	12	5.3	3.6	4.3	4.4

Fig. 8. Subjective results of each study.




Fig. 9. Error rates in more detail. Each dot is the average for one user, and the bars show 95% CIs, computed from the 34 or 12 users, in studies 1 and 2, respectively. In the 2D conditions (green and orange) of study 1, notice that edge highlighting with the mouse appears to have helped for *PathLength* 2 and 3 but not 4.

not intuitive, requiring extra motion and time, or difficult to position the cursor in the depth dimension. Note that users were given no explanation of how the mouse worked in 3D, and it is possible that several users were confused by it because they were moving their NDH and DH at the same time. Despite this, Fig. 8 also shows that $6/34 = 18\%$ of users chose VR3DHilite as their favorite condition, and 8 users described it as helpful, useful, faster, requiring less rotation of the network, and making pathways more visible.

Figs. 9 and 10 present the error rates and times in more detail. Notice that both the times and error rates appear smaller in 3D than in 2D. In Fig. 9, if we examine the 2D conditions in Study 1, it seems that the mouse helped with *PathLength* 2 and 3 but hindered performance with *PathLength* 4. This may be because the interactive highlighting sometimes misled users into following suboptimal paths. This is partially supported by comments made by 3 users, who talked about the edge highlighting causing them to focus more on those edges, inducing a different “mental exercise” than without edge highlighting, and encouraging the user to explore the network by “testing” different edges.

As mentioned earlier (Section IV-C), in the W+M study, “the mean euclidean distance between start nodes and end nodes was the same” [10]. Because our paths could be up to 5 edges long, this kind of control was not feasible in our experiment, nor would it have yielded realistic tasks. However, to check how much the euclidean distance may have influenced user responses in our




Fig. 10. Duration of trials.

Study 1, we computed two additional variables for each trial: first, $\Delta = \text{response} - \text{PathLength}$, so that Δ is negative, zero, or positive when the user's response is under, equal to, or over the correct *PathLength*, respectively. For example, if the user's response is 5 when the *PathLength* is actually 3, then $\Delta = 2$. Second, we found all shortest paths in the network and, for each *PathLength*, we found the average and standard deviation of the euclidean distances (from end-node to end-node) of those paths. This allowed us to compute a *z*-score for any pair of nodes, as a way to compare their euclidean distance to that of other pairs of nodes with the same topological distance. So in a given trial, if the shortest path between the end-nodes is 3 edges, but the *z*-score for those end-nodes is greater than 1, this means that shortest paths of 3 edges in that network tend to have end-nodes that are closer (in the euclidean sense), and we might expect such a *z*-score to bias the user toward over-estimating the *PathLength* in that trial. To test for such a bias, we checked for a correlation between Δ and the *z*-score. The correlation test yielded $R < 0.05$ and $p > 0.05$, hence no evidence of the euclidean distance biasing the user toward erroneous responses.

For each of the shortest paths in the trials of Study 1, we computed an "indirectness" ratio as the sum of the lengths of the edges in the path divided by the euclidean distance between the two end-nodes of the path. This ratio is high if the path is circuitous, but close to 1 if the path is straight. The average ratio for our 2D trials was 1.509, but for 3D it was 1.400, thus more straight.

During the conditions without mouse, where the DH was free, 3 out of the 34 users (plus 1 other user from the pilot) were either observed lifting their DH toward the controller during trials, as if trying to touch the virtual network, and/or described imagining touching or wanting to touch the virtual network during the post-questionnaire discussion. One of these users said this may have been because of their previous experience with an Oculus Quest 2 which displays the user's hands, and another user suggested using AR to allow the user to see their own hands.




Fig. 11. How much users rotated their view of the network (Section IV-G).




Fig. 12. How much of the rotation was due to hand motion, as opposed to head motion (Section IV-G).

L. Discussion

Compared to W+M, our path tracing task was more realistic (Section III-A), allowing users to control their view and their exposure time, and allowing users to benefit from interactive highlighting in 2D, but 3D still yielded a better error rate.

As discussed in Section IV-C, the lower error rate in 3D may be due in part to the shortest paths being more straight in 3D, because the layout algorithm has more freedom to position nodes. The higher indirectness ratio for 2D confirms they were less straight, and their difficulty may be related to studies finding that participants take longer to trace curves connecting two targets when the curve is longer, even when the euclidean distance between the targets is the same [63].

Fig. 12 suggests that most of the benefit of rotation (i.e., of motion parallax) comes from motion of the hand rather than of the head. Figs. 6 and 7 in [10] suggest that motion provides at

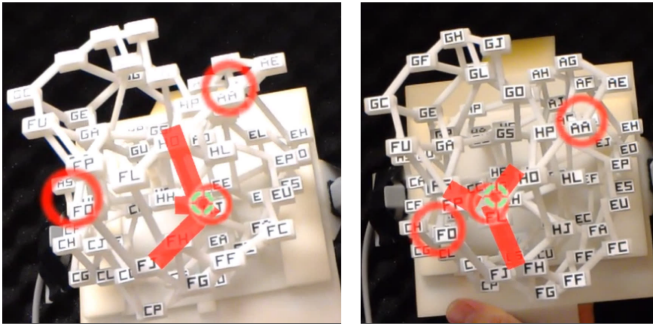


Fig. 13. Prototype implementation of highlighting of physical edges. Nodes FO and AA are indicated with red circular rings. The mouse cursor (green) hovers over HJ (Left) and FL (Right), causing three incident edges to highlight in each case. This was not used in our studies due to insufficient tracking accuracy.

least as much benefit as stereo. Taken together, these suggest that a user would benefit simply from the ability to rotate a 3D visualization with their hand, without any headset, stereo, or head-coupled perspective.

We observed that some users wished they could touch the networks. Previous work [47] found that users preferred physical networks that could be touched. This further motivates our next study which was planned in parallel with the first.

V. STUDY 2: 3D VIRTUAL AND PHYSICAL

Study 2 compared virtual and physical representations of networks, using a VR and an AR headset, respectively, and the same task as Study 1. Study 2 was more exploratory, where the number of participants was determined by convenience, and not preregistered.

A. Pilot and Choice of Main Conditions

To display virtual information on top of a 3D printed physical network, which we call **augmented physicalization**, we need some way for the AR headset to know where the network is located. We tested 3 different ways of tracking the physical network (detailed in the supplemental material, available online), including using the headset’s built-in camera and a combination of 3 external cameras. Section V-C describes our ultimate tracking method. Fig. 13 shows a prototype AR3DHilite condition, where the user’s DH moves a mouse, and the AR headset displays a virtual cursor and virtual highlighting on parts of the network. Unfortunately, we were unable to achieve the accuracy necessary to clearly highlight individual nodes or edges on a physical network. We suspect that part of the problem is due to small errors in the IPD used to generate the stereo rendering on the headset, making virtual imagery slightly misaligned with physical objects. We thus dropped the AR3DHilite condition and do not display the circular rings around nodes in any conditions of Study 2. Nevertheless, the end-nodes in Study 2 are indicated with callout line segments. In VR, these callouts are precisely located, but in AR, there are errors of $\approx 2\text{-}3\text{cm}$ in their apparent locations. Despite having only approximately correct positions in AR, the callouts do help the user find the correct physical




Fig. 14. The first three images show a trial in the AR3DTouch condition of study 2. The user touches the indicated nodes FG and FE and answers “2” (Top Right) only to be shown the error feedback “Error, answer should be 1” (Bottom Left). The last image (Bottom Right) shows the variant of the VR3D condition that was used in study 2. In both conditions, nodes are indicated with red callout line segments but not circular rings, due to insufficient tracking accuracy in AR.

end-nodes faster than if the user had no visual aid, and constitute an example of augmented physicalization. We expect that the approximate locations in AR will slow down users compared to VR, but it is plausible that this will have minimal impact on the user’s error rate in AR.

As in Study 1, the NDH activated a trigger button to open the radial menu. Had we instead used the DH to open the radial menu, there would have been a “homing time” for the user to move their DH between the physical network and a button to open the radial menu.

We ran a pilot with 3 users and 4 conditions: VR3D, VR3DHilite, AR3D, and AR3DTouch. In the AR3D condition, the user repositions the network with their NDH but may not touch the network with their DH. In AR3DTouch (Fig. 14), the user is encouraged to touch the network with their DH. For each of the 4 conditions, the user performed 10 warmup trials, followed by 20 trials with each of two networks. The ordering of headsets was random, as was the ordering of the pair of main conditions within each headset. In contrast with Study 1, each user was given more explanation of how the mouse worked in 3D for the VR3DHilite condition. Sessions lasted 2 hours per user. All users chose AR3D as their least favorite condition; one user found it uncomfortable to wear the HoloLens for so long; and two users reported that the tracking accuracy got worse over time (this was probably due to the users holding the network in different positions during calibration and during trials). Therefore, some changes were made for the final Study 2 experiment: we eliminated the AR3D condition, which was the least favorite condition of the users and less realistic than AR3DTouch; we also eliminated the VR3DHilite condition, because most users in Study 1 did not find the mouse in VR3DHilite useful, and we wanted Study 2 to involve the same number of trials with each

headset. We also increased the number of opportunities for the user to take breaks in both conditions, and opportunities to redo the calibration during the AR3DTouch condition. With only 2 main conditions, we could also slightly increase the number of trials per condition while also decreasing the total duration of each user session (Section V-F).

Thus, Study 2 had two values for *MainCondition*: VR3D, and AR3DTouch.

As detailed in the next section, the AR condition suffered from a smaller FOV, latency, and tracking error, compared with the VR condition. This creates confounds, however whichever condition outperforms the other, the results could be informative: if AR3DTouch yields a lower error rate, despite the shortcomings of the AR system, this demonstrates the importance of physical realism and/or the ability to touch the network; and if VR3D yields a lower error rate, despite affording no way to interact with the network, this shows how important it is for AR systems to be improved to reach their full potential.

B. Physicalized Networks

The 3D form of each network used in Study 1 was 3D printed using stereolithography (SLA) with a white plastic. The virtual 3D and physical networks have the same geometry, same node dimensions, same edge thickness, same white color, and same font used for labels.

When planning both studies, the number of nodes for the networks was chosen based on physical limits of common 3D printers. First, we wanted networks that would fit within a $6 \times 6 \times 6$ inch volume, to accommodate lower-end printers. In addition, a 6-inch width fits within the HoloLens's FOV at a distance of 30 cm, well within arm's reach. Second, we wanted edges to be 3 mm thick to avoid fragility. Third, we wanted nodes big enough to accommodate embossed text labels that would be clear even on lower-end FDM (Fused Deposition Modeling) printers. We implemented a custom font (each character a 5×4 bitmap) allowing us to print text labels on nodes 7 mm high with 1 mm stroke thickness. We found that 70 nodes resulted in a network of reasonable complexity that fit the size constraints. Testing revealed that embossed text on mono-color 3D printouts is both difficult to read and to paint, and because multi-color 3D printing is much more expensive, we finally used 2D printed stickers with the same custom font to label each node.

C. Hardware

For the VR3D condition, the same headset and controller were used as in the previous study. For AR3DTouch, we used a Microsoft HoloLens headset, which has 1268×720 pixels per eye. The FOV for displaying virtual information is limited to $\approx 31 \times 17^\circ$, however the physical world is visible through a much wider FOV. The virtual images are rendered at a fixed focal distance of ≈ 2 meters [64] however stereo and vergence depth cues create the illusion of virtual imagery at any distance. We measured a framerate of 30 fps.

3D printed networks were held in a “network holder” by the user's NDH. A Polhemus Patriot reported the positions and orientations (in the Patriot's local coordinate system) of two




Fig. 15. The equipment for the AR3DTouch condition of Study 2. Top: the network holder, held in the NDH, with a trigger button. Bottom left: registration rig with 2 fake hands, Microsoft HoloLens headset, network holder, and keypad. Bottom right: during a trial.

sensors attached to the network holder. Because the HoloLens has built-in functionality to detect hands and report their 3D position in the headset's coordinate system, we constructed a rig with two fake hands (Fig. 15 (Bottom Left)) whose positions were fixed with respect to the Patriot, allowing us to determine the position of the network with respect to the headset with an accuracy of ≈ 5 cm. The user also performed a simple calibration procedure to further improve the accuracy of the tracking to ≈ 2 -3 cm.

The Patriot was connected to the same PC mentioned in Section IV-F. This PC processed the position and orientation information from the Patriot and transmitted it via UDP packets over wifi to the HoloLens. Although the Patriot can read information at 60 Hz, we only transmitted 10 packets per second to the HoloLens, because a higher rate led to dropped packets. By studying footage recorded through the HoloLens, and also by using a method similar to Cattani et al.'s [65], we estimate that the latency between moving the network holder and the HoloLens updating its rendered virtual imagery was ≈ 150 -250 ms, despite using a dedicated high bandwidth wifi router (ASUS RT-AC5300).

Having two different headsets in Study 2, with differences in FOV and other characteristics, necessarily introduces confounds. An alternative approach would have been to use the HoloLens in both conditions of Study 2, however this would have meant imposing a limited FOV that is not representative of

the state-of-the-art in VR. Although the differences in headsets will certainly create differences in the time taken for trials, we are more interested in differences in error rate between the main conditions, and it is plausible that error rate will be less affected by the differences in the headsets.

D. Use of Blur in Virtual Feedback

Like most headsets, our VR and AR headsets each reproduce correct stereo disparity and vergence depth cues, but not correct accommodation depth cues, due to the virtual imagery being rendered at a fixed focal distance. In VR, this results in the well known “vergence-accommodation conflict,” which is often barely noticeable. However, with see-through AR headsets like the HoloLens, there is a further challenge: if virtual imagery (such as a virtual highlight) is rendered at the same location as a physical object (such as part of a physical network), it is impossible for the user to accommodate (i.e., “focus on”) both simultaneously. Although both may appear to be 30 cm from the user’s eyes in terms of stereo disparity and vergence depth cues, the focal distance of all virtual imagery rendered by the HoloLens is ≈ 2 meters [64].

To avoid having users focus on such virtual feedback, making the physical network appear blurry, we render callouts with a blur effect (Fig. 14), i.e., without sharp edges. For consistency, this was done in both VR3D and AR3DTouch.

E. Users

12 new users were recruited: 9 male, 3 female; all right handed; age 20 to 42 years (average 27.3); IPD 57 to 70 mm (average 62.7); stereoacuity test scores 3/10 to 10/10 (average 8.9). As in the previous study (Section IV-H), users were shown several printouts of example trials, with conditions in random order, to explain the task, and each headset was adjusted to the user’s IPD before beginning warmup trials.

F. Design

Each user experienced the 2 levels of *MainCondition* {VR3D, AR3DTouch} in random order. For each *MainCondition*, the user performed 10 warmup trials in random order (5 levels of *PathLength* \times 2 repetitions) with the warmup network, followed by 15 trials in random order (5 levels of *PathLength* \times 3 repetitions) with another network chosen at random, followed by another 5 warmup trials with the warmup network, 15 trials with another network, 5 warmup trials with the warmup network, and another 15 trials with another network. Each subsequence of warmup trials gave the user an opportunity to take a break, adjust the headset, and redo the calibration if they wished. There were a total of 12 users \times 2 levels of *MainCondition* \times 3 networks \times 5 levels of *PathLength* \times 3 repetitions per trial = 1080 trials, not counting warmup trials and not counting pilot data. Each session with a user lasted ≈ 1.5 hours.

G. Results

We find no evidence of a difference in error rate between the AR3DTouch and VR3D conditions (Fig. 9). In AR3DTouch,




Fig. 16. How far users held the network from their head.

users took more time (Fig. 10), which is explained by the difficulty that users had in finding the end-nodes at the start of each trial. Users also rotated their view less (Fig. 11) and rotated less with their hand and more with head motion (Fig. 12), which is explained by user feedback indicating that the network holder in AR3DTouch was not as easy to rotate as the handheld controller in the VR3D condition. Users also held the network further away from their head (Fig. 16), which is explained by the more limited FOV in AR.

All 12 users preferred the VR3D condition to AR3DTouch (Fig. 8). The reasons given by users for this preference were: the inaccurate positioning of callouts in AR (mentioned by 9 out of the 12 users); the limited FOV of the AR headset (mentioned by 4 users); the AR headset being less comfortable (4); the network holder in the AR3DTouch condition being more difficult to rotate than the VR controller (4); the VR3D condition providing better visual contrast between the network and the black background (4); and the latency with the AR system (1).

Based on user feedback, users often performed each trial of AR3DTouch in two stages: first, identifying the end-nodes, and second, finding the shortest path between them. The first was difficult because of the limited FOV and the positional error in the callouts. Once the end-nodes were identified, users would often “mark” them by touching their fingertips to one or both end-nodes and maintaining contact while looking for the shortest path. Several other behaviors were observed with users’ fingers: touching intermediate nodes along a path (observed in 9 out of the 12 users), touching multiple nodes simultaneously (8 users), pivoting the network around the hand while maintaining contact with one or more DH fingers (7), pointing at nodes without touching them (6), and grabbing a node with 2 fingers (3).

We also noted which fingers of the DH were employed during the AR3DTouch condition. Letters t, i, m, r, p denote thumb, index, middle, ring, and pinky fingers, respectively, and multiple letters indicate combinations, such as ti for thumb + index. Fingers i, m, p were often employed individually. We also

observed simultaneous uses of fingers: im (used by 7 out of the 12 users), ti (6 users), tm (5), ip (2), ir (1), mp (1), tim (1), imr (1).

Other notable behaviors observed were: touching edges in addition to nodes; tapping a sequence of nodes along a path; sliding a fingertip along the edges of a path; walking two fingers (index and middle) along the nodes of a path like legs; touching one node and pivoting the hand around the network, while maintaining contact with the finger; touching three nodes at once (either with tim or with imr).

One user explained that they used 2 fingers simultaneously to trace 2 different paths between end-nodes. Another said they would have preferred the AR3DTouch condition over VR3D if the problems with FOV and tracking accuracy were fixed. Another user said that touch had value for tracing the path with the finger. Another stated that they really wanted to touch the network in the VR3D condition, and wanted to see their fingers in VR, and that they were running their fingers through space where the path would have been during the VR3D condition. Another user (during the pilot) actually preferred the AR3DTouch condition to VR3D despite the limitations of the AR system, stating that the physical network allowed the task to be done faster and more simply.

H. Discussion

In the AR3DTouch condition, users employed their fingers in a variety of manners. The different uses of hands has been studied before in visualizations [66] and physicalizations [33], [47]. Other work [67] has proposed a general taxonomy of uses of hands. We are unaware of a taxonomy of uses of hands specifically relevant to data physicalizations.

Users preferred VR3D over AR3DTouch, but the reasons given for this are almost entirely due to technological limitations of the AR platform. Physicalizations that can be touched have been shown to be preferred [47] or yielded better performance [9] than 3D visualizations. Hence, next steps could be to either improve the AR platform, or improve the VR platform to better support the way users leverage their fingers. We now consider each of these.

Our AR platform suffers from inaccurate alignment of virtual and physical objects. Before using the Polhemus Patriot for tracking, we tried using the RGB camera on the HoloLens as well as multiple external cameras for object tracking, but none of these methods achieved acceptable accuracy. An alternative approach would be to use video pass-through AR (either with a headset or not, as in [31]) to avoid the need for highly accurate tracking, and enable virtual highlighting of physical nodes and edges (Fig. 13) with pixel-precise alignment, correct occlusion cues, and no conflicting accommodation (focal) distances.

To improve the VR platform to better support finger interaction, we notice that users in our AR3DTouch condition primarily used fingers to mark parts of the network. For example, users would often touch one or more fingertips against parts of the network, and then pivot a hand (to examine the network from a different view) while *maintaining contact* with fingertips, which was *made easier by friction*. This suggests an interaction

technique in VR where the user can use their DH to define one or more “sticky marks” on a virtual network, that remain even while the hands continue to move. We also observed users sliding their fingers along edges, which suggests support for “slippery marks” whose positions are updated to slide along edges, maybe as if being pulled by strings attached to the DH. Although physical fingers are limited to maintaining one mark per finger, this needn’t be the case with virtual marks: different fingers might be pinched against the thumb to invoke different commands to create, modify, or delete each of many marks. For example, when the user’s DH approaches a virtual network, a rope cursor [68] from the DH’s thumb could extend to the nearest node, and a DH pinch against the thumb by the index, middle, or ring finger could create a sticky mark, create a slippery mark, or delete the mark, respectively, at that node. Such finger-based interaction might be useful for tasks beyond finding shortest paths, and could be adapted to 2D layouts, perhaps in combination with other multitouch techniques [69], [70].

VI. CONCLUSION

Our Study 1 provides strong evidence that path tracing is less error prone in 3D layouts than in 2D layouts (Fig. 5), despite the use of edge-routing in 2D. The use of mouse-driven interactive highlighting in 2D reduces the error rate in 2D, but not as much as using a 3D layout (Fig. 7). 3D was also the most preferred layout (Fig. 8)

Our Study 2 found no evidence of different error rates between the virtual (VR) and physical (AR) conditions (Fig. 9). VR was preferred by users, but this was largely due to technological limitations of the AR platform, and users touched the physicalized networks in a variety of ways.

VII. FUTURE DIRECTIONS

Section V-H described ideas for future work. In addition, to benefit from the advantages of 2D and 3D, techniques could be designed allowing a user to rapidly switch between 2D and 3D layouts, possibly mixing styles in a hybrid that is 2D for most of the network but 3D in certain parts where the user is more interested in perceiving shortest paths. New fabrication methods might also enable physicalized networks that contain embedded buttons, touch sensors, and/or lights, for richer interaction.

ACKNOWLEDGMENT

Thanks to Dylan R. McGuffin for mounting the trigger button and building the registration rig; to Alicia Servera for help building an earlier rig; to Steve Haroz and Pierre Dragicevic for detailed advice on experiment design and statistical analysis; to Nathalie Henry Riche for the idea of comparing with physical networks; to Eyal Ofek, Ken Hinckley, and Jean-Daniel Fekete for general discussion; to Lev Nachmanson and Tim Dwyer for advice on edgerouting; to François Bérard for ideas on measuring latency; to Mar Gonzalez Franco and Bill Buxton for pointers to related work; to Peter Vale and Erwan Normand for programming tips; and to Minda Miloff for coaching. This research was supported by Microsoft Research and NSERC.

REFERENCES

- [1] K. Marriott et al., "Immersive analytics: Time to reconsider the value of 3D for information visualisation," in *Immersive Analytics*. Berlin, Germany: Springer, 2018, pp. 25–55.
- [2] A. Fonnet and Y. Prie, "Survey of immersive analytics," *IEEE Trans. Vis. Comput. Graph.*, vol. 27, no. 3, pp. 2101–2122, Mar. 2021.
- [3] B. Ens et al., "Grand challenges in immersive analytics," in *Proc. CHI Conf. Hum. Factors Comput. Syst.*, 2021, pp. 1–17.
- [4] M. Kraus et al., "Immersive analytics with abstract 3D visualizations: A survey," *Comput. Graph. Forum*, vol. 41, no. 1, pp. 201–229, 2022.
- [5] T. von Landesberger et al., "Visual analysis of large graphs: State-of-the-art and future research challenges," *Comput. Graph. Forum*, vol. 30, no. 6, pp. 1719–1749, 2011. [Online]. Available: <https://doi.org/10.1111/j.1467-8659.2011.01898.x>
- [6] P. C. Wong and R. D. Bergeron, "30 years of multidimensional multivariate visualization," in *Scientific Visualization: Overviews, Methodologies, and Techniques*, G. M. Nielson, H. Hagen, and H. Müller eds., Los Alamitos, CA, USA: IEEE Comput. Soc. Press, 1997.
- [7] T. Munzner, *Visualization Analysis and Design*. Boca Raton, FL, USA: CRC Press, 2014, pp. 117–130.
- [8] M. Sedlmair, T. Munzner, and M. Tory, "Empirical guidance on scatterplot and dimension reduction technique choices," *IEEE Trans. Vis. Comput. Graph.*, vol. 19, no. 12, pp. 2634–2643, Dec. 2013.
- [9] Y. Jansen, P. Dragicevic, and J.-D. Fekete, "Evaluating the efficiency of physical visualizations," in *Proc. CHI Conf. Hum. Factors Comput. Syst.*, 2013, pp. 2593–2602.
- [10] C. Ware and P. Mitchell, "Visualizing graphs in three dimensions," *ACM Trans. Appl. Percep.*, vol. 5, no. 1, pp. 1–15, 2008.
- [11] T. Dwyer and L. Nachmanson, "Fast edge-routing for large graphs," in *Proc. Symp. Graph Drawing*, 2009, pp. 147–158.
- [12] Y. Jansen et al., "Opportunities and challenges for data physicalization," in *Proc. CHI Conf. Hum. Factors Comput. Syst.*, 2015, pp. 3227–3236.
- [13] "Data physicalization wiki," 2021. [Online]. Available: <http://dataphys.org>
- [14] T. Munzner, "Process and pitfalls in writing information visualization research papers," in *Information Visualization: Human-Centered Issues and Perspectives*. Berlin, Germany: Springer, 2008, pp. 134–153.
- [15] R. Brath, "3D InfoVis is here to stay: Deal with it," in *Proc. IEEE VIS Int. Workshop 3DVis*, 2014, pp. 25–31.
- [16] C. Ware, "Designing with a 2.5D attitude," *Inf. Des. J.*, vol. 10, no. 3, pp. 258–265, 2000.
- [17] B. Lee, M. Cordeil, A. Prouzeau, B. Jenny, and T. Dwyer, "A design space for data visualisation transformations between 2D and 3D in mixed-reality environments," in *Proc. CHI Conf. Hum. Factors Comput. Syst.*, 2022, pp. 1–14.
- [18] B. Shneiderman, "Why not make interfaces better than 3D reality?," *IEEE Comput. Graph. Appl.*, vol. 23, no. 6, pp. 12–15, Nov./Dec. 2003.
- [19] C. Collins and S. Carpendale, "VisLink: Revealing relationships amongst visualizations," *IEEE Trans. Vis. Comput. Graph.*, vol. 13, no. 6, pp. 1192–1199, Nov./Dec. 2007.
- [20] B. Zou, Y. Liu, and J. M. Wolfe, "Top-down control of attention by stereoscopic depth," *Vis. Res.*, vol. 198, 2022, Art. no. 108061.
- [21] K. Kim, M. Billinghurst, G. Bruder, H. B.-L. Duh, and G. F. Welch, "Revisiting trends in augmented reality research: A review of the 2nd decade of ISMAR (2008–2017)," *IEEE Trans. Vis. Comput. Graph.*, vol. 24, no. 11, pp. 2947–2962, Nov. 2018.
- [22] J. P. McIntire and K. K. Liggett, "The (possible) utility of stereoscopic 3D displays for information visualization: The good, the bad, and the ugly," in *Proc. IEEE VIS Workshop 3DVis*, 2014, pp. 1–9.
- [23] J. P. McIntire, P. R. Havig, and E. E. Geiselman, "Stereoscopic 3D displays and human performance: A comprehensive review," *Displays*, vol. 35, no. 1, pp. 18–26, 2014.
- [24] M. Cordeil, A. Cunningham, T. Dwyer, B. H. Thomas, and K. Marriott, "ImAxes: Immersive axes as embodied affordances for interactive multivariate data visualisation," in *Proc. Annu. ACM Symp. User Interface Softw. Technol.*, 2017, pp. 71–83.
- [25] D. Hayatpur, H. Xia, and D. Wigdor, "DataHop: Spatial data exploration in virtual reality," in *Proc. Annu. ACM Symp. User Interface Softw. Technol.*, 2020, pp. 818–828.
- [26] M. Cordeil, B. Bach, Y. Li, E. Wilson, and T. Dwyer, "Design space for spatio-data coordination: Tangible interaction devices for immersive information visualisation," in *Proc. IEEE Pacific Vis. Symp.*, 2017, pp. 46–50.
- [27] B. Bach, R. Sicut, J. Beyer, M. Cordeil, and H. Pfister, "The hologram in my hand: How effective is interactive exploration of 3D visualizations in immersive tangible augmented reality?," *IEEE Trans. Vis. Comput. Graph.*, vol. 24, no. 1, pp. 457–467, Jan. 2018.
- [28] M. Cordeil et al., "Embodied axes: Tangible, actuated interaction for 3D augmented reality data spaces," in *Proc. CHI Conf. Hum. Factors Comput. Syst.*, 2020, pp. 1–12.
- [29] J. Smiley et al., "The MADE-Axis: A modular actuated device to embody the axis of a data dimension," *Proc. ACM Hum.-Comput. Interact.*, vol. 5, no. ISS, pp. 1–23, 2021.
- [30] K. A. Satriadi et al., "Tangible globes for data visualisation in augmented reality," in *Proc. CHI Conf. Hum. Factors Comput. Syst.*, 2022, pp. 1–16.
- [31] A. Gillet, M. Sanner, D. Stoffler, and A. Olson, "Tangible interfaces for structural molecular biology," *Structure*, vol. 13, no. 3, pp. 483–491, 2005.
- [32] F. Taher et al., "Exploring interactions with physically dynamic bar charts," in *Proc. CHI Conf. Hum. Factors Comput. Syst.*, 2015, pp. 3237–3246.
- [33] F. Taher, Y. Jansen, J. Woodruff, J. Hardy, K. Hornbæk, and J. Alexander, "Investigating the use of a dynamic physical bar chart for data exploration and presentation," *IEEE Trans. Vis. Comput. Graph.*, vol. 23, no. 1, pp. 451–460, Jan. 2017.
- [34] G. Di Battista, P. Eades, R. Tamassia, and I. G. Tollis, *Graph Drawing: Algorithms for the Visualization of Graphs*, Englewood Cliffs, NJ, USA: Prentice Hall, 1999.
- [35] M. Kaufmann and D. Wagner, Eds., *Drawing Graphs: Methods and Models*. Berlin, Germany: Springer, 2001.
- [36] V. Yoghoudjian et al., "Exploring the limits of complexity: A survey of empirical studies on graph visualisation," *Vis. Inform.*, vol. 2, no. 4, pp. 264–282, 2018.
- [37] M. Burch, W. Huang, M. Wakefield, H. C. Purchase, D. Weiskopf, and J. Hua, "The state of the art in empirical user evaluation of graph visualizations," *IEEE Access*, vol. 9, pp. 4173–4198, 2020.
- [38] O.-H. Kwon, C. Muedler, K. Lee, and K.-L. Ma, "A study of layout, rendering, and interaction methods for immersive graph visualization," *IEEE Trans. Vis. Comput. Graph.*, vol. 22, no. 7, pp. 1802–1815, Jul. 2016.
- [39] M. Cordeil, T. Dwyer, K. Klein, B. Laha, K. Marriott, and B. H. Thomas, "Immersive collaborative analysis of network connectivity: CAVE-style or head-mounted display?," *IEEE Trans. Vis. Comput. Graph.*, 2017.
- [40] A. Drogemuller, A. Cunningham, J. Walsh, B. H. Thomas, M. Cordeil, and W. Ross, "Examining virtual reality navigation techniques for 3D network visualisations," *J. Comput. Lang.*, vol. 56, 2020, Art. no. 100937.
- [41] P. Irani and C. Ware, "Diagramming information structures using 3D perceptual primitives," *ACM Trans. Comput.-Hum. Interact.*, vol. 10, no. 1, pp. 1–19, 2003.
- [42] B. Alper, T. Höllerer, J. Kuchera-Morin, and A. Forbes, "Stereoscopic highlighting: 2D graph visualization on stereo displays," *IEEE Trans. Vis. Comput. Graph.*, vol. 17, no. 12, pp. 2325–2333, Dec. 2011.
- [43] N. Greffard, F. Picarougne, and P. Kuntz, "Visual community detection: An evaluation of 2D, 3D perspective and 3D stereoscopic displays," in *Proc. Symp. Graph Drawing*, 2011, pp. 215–225.
- [44] N. Greffard, F. Picarougne, and P. Kuntz, "Beyond the classical monoscopic 3D in graph analytics: An experimental study of the impact of stereoscopy," in *IEEE VIS Int. Workshop 3DVis*, 2014, pp. 19–24.
- [45] B. Lee, C. Plaisant, C. S. Parr, J.-D. Fekete, and N. Henry, "Task taxonomy for graph visualization," in *Proc. AVI Workshop BEYOND Time Errors: Novel Eval. Methods Inf. Vis.*, 2006, pp. 1–5.
- [46] R. James, A. Bezerianos, O. Chapuis, M. Cordeil, T. Dwyer, and A. Prouzeau, "Personal+context navigation: Combining AR and shared displays in network path-following," in *Proc. Graph. Interface*, 2020, pp. 1–12.
- [47] A. Drogemuller, A. Cunningham, J. A. Walsh, J. Baumeister, R. T. Smith, and B. H. Thomas, "Haptic and visual comprehension of a 2D graph layout through physicalisation," in *Proc. CHI Conf. Hum. Factors Comput. Syst.*, 2021, pp. 1–16.
- [48] R. L. Sollenberger and P. Milgram, "Effects of stereoscopic and rotational displays in a three-dimensional path-tracing task," *Hum. Factors*, vol. 35, no. 3, pp. 483–499, 1993.
- [49] M. H. van Beurden, A. Kuijsters, and W. A. IJsselstein, "Performance of a path tracing task using stereoscopic and motion based depth cues," in *Proc. 2nd Int. Workshop Qual. Multimedia Experience*, 2010, pp. 176–181.
- [50] C. Ware and G. Franck, "Evaluating stereo and motion cues for visualizing information nets in three dimensions," *ACM Trans. Graph.*, vol. 15, no. 2, pp. 121–140, 1996.
- [51] D. Belcher, M. Billinghurst, S. Hayes, and R. Stiles, "Using augmented reality for visualizing complex graphs in three dimensions," in *Proc. 2nd IEEE ACM Int. Symp. Mixed Augmented Reality*, 2003, pp. 84–93.
- [52] K. Danyluk, T. Ulusoy, W. Wei, and W. Willett, "Touch and beyond: Comparing physical and virtual reality visualizations," *IEEE Trans. Vis. Comput. Graph.*, vol. 28, no. 4, pp. 1930–1940, Apr. 2022.


- [53] F. Bérard and T. Louis, “The object inside: Assessing 3D examination with a spherical handheld perspective-corrected display,” in *Proc. CHI Conf. Hum. Factors Comput. Syst.*, 2017, pp. 4396–4404.
- [54] Y. Guiard, “Asymmetric division of labor in human skilled bimanual action: The kinematic chain as a model,” *J. Motor Behav.*, vol. 19, no. 4, pp. 486–517, 1987.
- [55] M. J. McGuffin, “Preregistration,” 2022. [Online]. Available: <https://osf.io/tdcng>
- [56] A. Ramcharitar and R. Teather, “EZCursorVR: 2D selection with virtual reality head-mounted displays,” in *Proc. Graph. Interface*, 2018, pp. 123–130.
- [57] D. J. Watts and S. H. Strogatz, “Collective dynamics of ‘small-world’ networks,” *Nature*, vol. 393, no. 6684, 1998, Art. no. 440.
- [58] E. R. Gansner, Y. Koren, and S. North, “Graph drawing by stress majorization,” in *Proc. Symp. Graph Drawing*, 2004, pp. 239–250.
- [59] Accessed: 2019. [Online]. Available: <https://github.com/Microsoft/automatic-graph-layout>
- [60] F. Faul, E. Erdfelder, A. Buchner, and A.-G. Lang, “Statistical power analyses using G*Power 3.1: Tests for correlation and regression analyses,” *Behav. Res. Methods*, vol. 41, pp. 1149–1160, 2009.
- [61] N. K. Dhand and M. S. Khatkar, “Statulator: An online statistical calculator. Sample size calculator for comparing two paired means,” 2014. [Online]. Available: <http://statulator.com/SampleSize/ss2PM.html>
- [62] P. Dragicevic, “Fair statistical communication in HCI,” in *Modern Statistical Methods for HCI*. Berlin, Germany: Springer, 2016, pp. 291–330.
- [63] D. Crundall, G. G. Cole, and G. Underwood, “Attentional and automatic processes in line tracing: Is tracing obligatory?,” *Percep. Psychophys.*, vol. 70, no. 3, pp. 422–430, 2008.
- [64] Accessed: 2017. [Online]. Available: <https://docs.microsoft.com/en-us/windows/mixed-reality/design/comfort>
- [65] E. Cattani, A. Rochet-Capellan, and F. Bérard, “A predictive approach for an end-to-end touch-latency measurement,” in *Proc. Int. Conf. Interactive Tabletops Surf.*, 2015, pp. 215–218.
- [66] J. Walny, S. Huron, C. Perin, T. Wun, R. Pusch, and S. Carpendale, “Active reading of visualizations,” *IEEE Trans. Vis. Comput. Graph.*, vol. 24, no. 1, pp. 770–780, Jan. 2018.
- [67] N. Kamakura, *Postures and Movement Patterns of the Human Hand: A Framework for Understanding Hand Activity for Clinicians and Engineers*. Boca Raton, FL, USA: Brown Walker Press, 2022.
- [68] M. Guillon, F. Leitner, and L. Nigay, “Investigating visual feedforward for target expansion techniques,” in *Proc. CHI Conf. Hum. Factors Comput. Syst.*, 2015, pp. 2777–2786.
- [69] S. Schmidt, M. A. Nacenta, R. Dachselt, and S. Carpendale, “A set of multi-touch graph interaction techniques,” in *Proc. ACM Int. Conf. Interactive Tabletops Surf.*, 2010, pp. 113–116.
- [70] J. Thompson, A. Srinivasan, and J. Stasko, “Tangraphe: Interactive exploration of network visualizations using single hand, multi-touch gestures,” in *Proc. Int. Conf. Adv. Vis. Interfaces*, 2018, pp. 1–5.



Michael J. McGuffin is a full professor with ETS, a French-language engineering school in Montreal, Canada, where his students do research in HCI and visualization. His recent interests include visual programming and artificial life. In 2009, his paper at the IEEE Information Visualization Conference (InfoVis 2009) received an Honorable Mention.



Ryan Servera received the bachelor’s of engineering degree in electrical engineering with a minor in software engineering from McGill University, in 2020. He is a software developer. In 2019, he interned as a research assistant with ETS where he was able to explore VR and AR. He is particularly interested in human computer interaction in video games.



Marie Forest received the DPhil degree in statistics from the University of Oxford. She is a lecturer in Statistics with ETS. She is interested in the development of web application teaching tools to help students understand probability and statistics through simulations. In the past, her research areas were in the analysis of genomic data and stochastic modelling applied to genetics.

Path Tracing in 2D, 3D, and Physicalized Networks: Supplemental

Michael J. McGuffin, Ryan Servera, Marie Forest

Hardware

We measured the masses of the devices used:

HTC Vive headset: $\approx 615\text{g}$

One HTC Vive handheld controller: $\approx 200\text{g}$

HoloLens headset: $\approx 577\text{g}$

3D printed network + network holder: $\approx 370\text{-}378\text{g}$, depending on the network

Additional Results from Study 1

We measured the hand span (distance between tips of thumb and pinky with fingers spread out) of each user. We also computed the overall error rate of each user in 3D (computed from trials in 3D, not including warmups, and only including path lengths of 2, 3, 4, because path lengths of 1 and 5 yielded lower error rates that are less interesting), and the overall amount of rotation performed by users in 3D (not including warmups).

We found little to no evidence of a correlation between IPD and error rate, IPD and stereo acuity, age and error rate, hand span and rotation, hand span and error rate, stereo acuity and error rate. We also found no evidence of a difference in the error rates between male and female participants.

Rendering on the HoloLens

“Focus point” hinting was used to stabilize the image rendered by the HoloLens (<https://docs.microsoft.com/en-us/windows/mixed-reality/develop/unity/focus-point-in-unity>), and rendering of polygons was done with a flat shader to optimize the framerate.

Tracking of Networks in Study 2

Before deciding to use the Polhemus Patriot for tracking, we previously implemented and tested two alternative tracking methods, first using the RGB camera on the HoloLens with Vuforia, and second using three simultaneous external cameras with OpenCV and ArUco, to track fiducials on the network holder. The accuracy achieved with these methods was never better than 1cm, and suffered from other problems such as loss of tracking during motion blur and occlusion.

First, we tested using the RGB camera in the HoloLens, for which there is functionality provided to use Vuforia to track fiducials. We found this method sometimes exhibited latency of more than 1 second and sometimes reported 180° errors in the orientation of the fiducials.

Next, we implemented a tracking system using 3 external Logitech HD (1920×1080) resolution USB cameras all connected to one PC (the same PC described in the main paper, with a 6-core, 4.7 GHz CPU). We wrote custom code using OpenCV and ArUco to track fiducials and also to compute 3D-to-3D homographies to transform from the coordinate system of each camera to the coordinate system of the HoloLens. We placed 4 fiducials on the network holder (Figure 1), and also built a physical rig marked with 8 fiducials (Figure 2) for registration. We found that tracking of fiducials with this method could fail if there were shadows, which we alleviated using softbox lighting, and also whenever the fiducials moved too quickly causing a blurry image. An additional problem was the tradeoff between spatial and temporal resolution. Despite using a fast PC with a Vantac “quad chip” USB card (with a dedicated chip for each USB port), we could only read the 3 cameras at $\approx 5\text{-}10$ fps at full HD resolution. Downgrading the camera resolution to 640×480

allowed reading at 30 fps, but because each camera has a FOV of 66° horizontally, the spatial resolution was $(\text{arc length covered by one pixel}) = (\text{FOV in radians}) \times (\text{distance to camera}) / (\text{pixels}) = ((66^\circ/180^\circ)\pi) \times (100\text{cm}) / (640 \text{ pixels}) \approx 0.2 \text{ centimeters}$, assuming the camera is 100cm away from the fiducial. ArUco's estimation of a fiducial's center seemed very sensitive to image noise, moving 2-3 pixels from frame-to-frame even when the fiducial is not moving. So this could result in an error of more than 0.5cm, even assuming no error from the estimated position of the HoloLens.




Figure 1: Network holder with 4 fiducials.




Figure 2: Registration rig with 8 fiducials.

Our 3rd attempt to track the network holder was done with the Polhemus Patriot, as reported in the main paper.




Figure 3: In the ARTouch condition of study 2, the user performs calibration by moving the network holder onto a white virtual square.

Calibration of HoloLens

We implemented calibration with a homography matrix that transforms from the Patriot’s 3D space to the 3D coordinate system of the HoloLens. Ideally, the calibration would have been done only once by a researcher, and the resulting homography matrix saved, so that users wouldn’t need to perform any calibration. However, there are nonnegligible errors in the positions reported by the Patriot and the HoloLens, and the best homography to use changes from one session to the next. A further difficulty arises from the coordinate system of the HoloLens being different for each session: its origin and axes are defined by however the user’s head happens to be positioned and oriented when our client code begins executing. We now describe how we kept the calibration as simple as possible for users, while estimating a good quality homography.

We extended a 2D-to-2D homography calculation given by Kriegman¹ and implemented a subroutine that takes as input a list of points whose 3D coordinates are known in two different coordinate systems, and outputs a 4×4 homography matrix to transform between the two coordinate systems. Because the HoloLens reports the positions of the fake hands in its own coordinate system, and because we know (roughly) where the hands are located in the Patriot’s coordinate system, we can use our subroutine to compute an initial homography h_0 . It is possible, however, to obtain a better homography H_0 by having the user go through a 9-point calibration procedure that we implemented: in this procedure, the HoloLens displays a sequence of 9 virtual points (located at the corners and center of a cube occupying the user’s working volume), and the user holds a Patriot sensor at each of the displayed locations, allowing our client code to obtain the coordinates of each of the 9 points in both coordinate systems, and then use the same subroutine to compute the improved H_0 . However, we did not want each participant in our experiment to have to perform this 9-point calibration, as it is time-consuming to explain and perform carefully.

We speculated that each homography $h = vf$ can be factored into two components, a variable component v that changes from one session to the next depending on noise and on how the headset initializes its coordinate system, and a fixed component f that depends on the placement of the fake hands with respect to the Patriot. We won’t try to extract these components, but there is an indirect

¹David Kriegman, Homography Estimation, Course notes for Computer Vision I, CSE 252A, Winter 2007, University of California, San Diego, 2007

way of improving the estimate of f and therefore improving h . Consider the h_0 and H_0 defined in the previous paragraph, where the former is a lower quality homography, and the latter is a higher quality one obtained with 9-point calibration. Each has factors $h_0 = v_0 f$ and $H_0 = v_0 F$, where F is a better estimate of the fixed parameters that don't change from one session to the next. Compute and save $Q = h_0^{-1} H_0$. In any future session i , compute h_i (based on the positions of the fake hands) and then multiply by the previously saved Q to obtain $H_i = h_i Q$ which is a better homography than h_i . To see why, notice that $H_i = h_i Q = h_i h_0^{-1} H_0 = h_i (v_0 f)^{-1} v_0 F = v_i f (f^{-1} v_0^{-1}) v_0 F = v_i F$. In other words, the multiplication by Q effectively replaces the lower quality f with the higher quality F . The 9-point procedure is only performed once by a researcher to obtain Q and needn't be repeated by each participant. We implemented this and found that it noticeably improved the initial tracking accuracy. However, to improve accuracy even more, each participant also performed the simple calibration procedure shown in Figure 3.