

This is the accepted version of the following article:

Egocentric Network Exploration for Immersive Analytics,
J. Sorger, A. Arleo, P. Kán, W. Knecht, and M. Waldner,
Pacific Graphics 2021, Vol. 40 (2021), Nr. 7

which has been published in final form at
<http://onlinelibrary.wiley.com>.

This article may be used for non-commercial purposes in
accordance with the Wiley Self-Archiving Policy
[<http://olabout.wiley.com/WileyCDA/Section/id-820227.html>].

Egocentric Network Exploration for Immersive Analytics

J. Sorger¹ , A. Arleo² , P. Kán² , W. Knecht¹ , and M. Waldner² 

¹Complexity Science Hub, Vienna, Austria

²TU Wien, Vienna, Austria



Figure 1: Exploring the local neighborhood of a hub node ('UI46') in three immersive interfaces. Baseline (left): the user inspects 'UI46' from close-distance. Ego-Highlight (middle): the user takes the position of 'UI46'; direct neighbors are highlighted in yellow. Ego-Bubble (right): direct neighbors are arranged equidistantly to the user; edges that cross the user's bounding sphere are clipped.

Abstract

To exploit the potential of immersive network analytics for engaging and effective exploration, we promote the metaphor of “egocentrism”, where data depiction and interaction are adapted to the perspective of the user within a 3D network. Egocentrism has the potential to overcome some of the inherent downsides of virtual environments, e.g., visual clutter and cyber-sickness. To investigate the effect of this metaphor on immersive network exploration, we designed and evaluated interfaces of varying degrees of egocentrism. In a user study, we evaluated the effect of these interfaces on visual search tasks, efficiency of network traversal, spatial orientation, as well as cyber-sickness. Results show that a simple egocentric interface considerably improves visual search efficiency and navigation performance, yet does not decrease spatial orientation or increase cyber-sickness. An occlusion-free Ego-Bubble view of the neighborhood only marginally improves the user's performance. We tie our findings together in an open online tool for egocentric network exploration, providing actionable insights on the benefits of the egocentric network exploration metaphor.

CCS Concepts

• **Human-centered computing** → **Visualization techniques; Empirical studies in visualization;**

1. Introduction

Immersive analytics (IA) investigates “engaging and embodied analysis tools to support data understanding and decision making” [MSD*18]. Particularly in network analytics, prior research has empirically demonstrated many benefits of three-dimensional immersive displays over traditional two-dimensional graph visualization methods [WF96, BBHS03, WM08, HZBK08, GPK11, KMLM16, KKM*20]. In virtual reality (VR), users can interact with the data as if they were a physical entity that they can approach and inspect. They can navigate *within* a network instead of just looking *at* it.

Immersion in VR can reduce visual clutter [BM07] and provide insightful complementary perspectives on a network [SWKA19]. Room-scale immersive networks can lead to more user engagement compared to a table-scale network coupled with a zooming interface [YCB*20]. Indeed, there are numerous examples of immersive network analytics interfaces presented in literature that allow users to “dive into” and subsequently walk or fly through a 3D network [ZHF*16, DCW*17, EMP19, SWKA19, YCB*20]. While being surrounded by a room-sized network can be interesting, it can also be problematic. Navigation and interaction in abstract 3D data spaces

is more challenging compared to 2D [MCH*18], because abstract data spaces do not have a natural scale and orientation [YCB*20]. An immersive perspective can also require more positional changes by the user to resolve occlusions and blind spots [KWO*19] – a challenge for data analysts who prefer to seamlessly enter and leave the virtual environment while remaining at their desks [ZWB*17]. Flying through a cluttered visualization can then lead to cyber-sickness [SWKA19] caused by the sensory conflicts between the static user and the visual sensation of a moving environment [LJ00]. We believe that with careful design around the user, which we refer to as *egocentric*, these challenges can be alleviated while still exploiting the beneficial capabilities of VR.

In graph visualization, egocentric network analysis techniques support users in analyzing the part of the social network that is closely related to the node that represents themselves in the network [Fis05, Gol13]. We translate this principle to immersive network analytics: by applying the egocentric network metaphor to 3D node-link diagrams, users can take the perspective of any node – they can “be the node” – and thereby gain an optimized view of their local neighborhood, i.e., data points that are closely related to “them”. Associating the user with a dedicated node allows us to optimize the visibility of the local neighborhood, thereby resolving occlusion problems and removing visual clutter. The egocentric approach also has the potential for more efficient navigation within the network: by using data points (nodes) as anchors for exploration, users can change their egocentric perspective on the network by “jumping” from node to node. Yet, the restriction of the user’s movement and the introduction of local adaptations of the data representation may cause disorientation during network traversal. In this paper, we set out to systematically explore the trade-offs that the concept of egocentrism introduces to immersive network exploration. Understanding the relevance of these trade-offs in analysis tasks is crucial to conceive methods that provide an effective and pleasant exploration experience to users. The contributions of this paper are:

- Egocentric detail views as a new metaphor for immersive network analytics, exemplified by two novel VR interfaces with varying degrees of egocentrism;
- Results of a controlled user study systematically investigating the effect of the proposed interfaces on visual search efficiency, navigation performance, spatial orientation, cyber-sickness, and user preference;
- A publicly available web application for immersive egocentric network exploration synthesizing the insights from the study.

2. Related Work

The potential of network analytics in immersive environments has been investigated since the nineties [CLFZ95], with first experimentation on stereoscopic depth cues in the exploration of networks by Ware and Franck [WF96]. The results of this study suggest that the use of immersive technology could substantially increase the size of the networks a user could grasp. These findings were later confirmed and extended to AR [BBHS03], different VR display technologies [RBLN04, CDK*16], and different analytical tasks, such as cluster detection [GPK11, KWO*19] or memorizing previously highlighted nodes [KKM*20]. As VR devices increased in affordability

and popularity, so did the number of approaches for network visualization and exploration. For example, Huang et al. [HFL*17] and Erra et al. [EMP19] presented gesture-based interaction techniques for network exploration and manipulation. In terms of navigation and visualization, immersive network techniques generally rely on an outside overview perspective onto a 3D node-link diagram of the network [OAS00, HFL*17, YCB*20, ZHF*16, KKM*20, SWKA19], and/or they allow the user to fly through the node-link diagram using conventional 6DOF controllers [DCW*17, EMP19, SWKA19, YCB*20, ZHF*16].

Immersive displays can also be used for effective node highlighting. For example, Alper et al. [AHKMF11] introduce the principle of *stereoscopic highlighting*, where stereoscopic depth is used to draw relevant nodes “closer” to the user. Altarawneh et al. further develop this concept with the *Expand* approach [AHE14] for compound graphs, in which depth encodes the structural relationships between nodes. The “deadeye” technique [KCWK19] renders relevant nodes only on one of the stereoscopic displays, which strongly attracts the user’s attention. While stereoscopic highlighting can be more accurate and subtle than static visual cues [AHKMF11, KCWK19], these techniques are not applicable for immersive 3D network exploration.

Other research efforts target specialized graph presentation and navigation techniques for VR. In their early immersive graph visualization system, Osawa et al. [OAS00] introduced “heat models” that locally distort the representation of the force-directed graph layout. Nodes in the graph are selected to be the “foci”, that can absorb or emit heat. Halpin et al. [HGBK08] displayed graphs in 2D, but allowed users to extrude selected nodes into the third dimension. Kwon et al. [KMLM15] presented spherical graph layouts in VR for seated exploration with mouse interaction, where the entirety of the graph is laid out in the user’s field of view. Clutter is reduced by routing edges “away” from the user. The authors further demonstrated how spherical projections outperform a 2D layout on typical graph exploration tasks, such as finding common neighbors, path finding, and recall of node locations [KMLM16]. Fish-eye techniques are another way to expand the user’s field-of-view in immersive environments [OWK*14, DPHB15]. A case study by Sorger et al. [SWKA19] suggests that, for real-world tasks, data scientists appreciate complementary perspectives when analyzing their own, well-known network data: an external overview as stable reference point when starting the network exploration in VR and a detail view as a novel perspective to explore local neighborhoods.

To evaluate navigation within 3D networks, Drogemuller et al. [DCW*18] evaluated teleportation and one handed flying against two-handed flying [MBJS97] and “worlds in miniature” locomotion techniques for local detail view navigation in 3D in networks. Their results show that flying is more efficient for finding nodes or paths between nodes in the graph, and that teleportation frequently causes loss of orientation. In contrast, our proposed approach provides data-adaptive egocentric navigation, which optimizes the network representation for the user’s current perspective. In this work, we present first evidence that this metaphor indeed facilitates local network analysis in room-scale visualizations when compared to traditional interfaces for network exploration in VR.

3. Egocentrism in Immersive Network Analytics

We promote *egocentrism* in the context of immersive analytics to better exploit the capabilities of VR technology for locally optimized visual representation and interaction capabilities. In contrast to traditional 2D viewing and interaction modalities, a VR application can be made aware of the users' position and scale in the virtual scene, their field of view, and potentially even focus of attention with respect to the data they are exploring. Adaptions to an interface to support egocentrism can thus affect navigation, interaction, rendering, or the spatial layout of the data.

We categorize network exploration techniques in VR along two orthogonal axes, as illustrated in Figure 2: On the vertical axis, we differentiate whether users explore the network from an external/global "overview" perspective, from where the entire graph should be visible (see Fig. 3), or from an internal/local "detail" perspective within the network (see Fig. 1). On the horizontal axis, we distinguish whether or not the network visualization and the interaction techniques follow an egocentric approach, i.e., are optimized to the user's position and orientation in the virtual environment. The spherical layout proposed by Kwon et al. [KMLM15], for instance, can be classified as an egocentric overview interface, as it offers an overview that takes the user's position into account. The overview+detail interface by Sorger et al. [SWKA19] and the table-and-room-scale visualizations by Yang et al. [YCB*20] combine traditional overview and detail views (note that Yang et al. [YCB*20] refer to the traditional overview (Fig. 2 top left) as "exocentric" and to the traditional detail view (bottom left) as "egocentric view"). To the best of our knowledge, *egocentric detail views* (Fig. 2 bottom right) have not been investigated yet. They are the focus of this paper. We argue that an overview alone will lead to little benefit of using VR. Especially when viewing large networks from an external perspective, nodes and clusters in the center could be occluded. Having a detail view alone, on the other side, lacks the necessary overview to identify interesting regions in the network topology or to put the local environment into a global context. In line with the information seeking mantra, egocentric network exploration interfaces should therefore be coupled with an overview perspective that can serve




Figure 2: Categorization of network exploration techniques. The node used for the egocentric detail perspective (bottom right) is colored orange in the other networks.




Figure 3: Overview perspective of a large graph used in the study: the starting condition of the Find Path (FiP) task, showing the start (yellow) and end (red) node of the path. Hubs and outliers can be spotted more easily here, but local connectivity is hard to explore.

as a stable reference for analysts switching between immersive and traditional desktop analytics [SWKA19].

4. A Study on Egocentrism for Immersive Network Exploration

We expect that network exploration interfaces that support egocentrism can lead to a trade-off affecting several factors, such as visual search efficiency, navigation performance, spatial orientation, and cyber-sickness. To quantify these trade-offs, as well as to assess the overall benefits and limitations of egocentrism for network exploration in IA, we conducted a controlled user study where we compared the effectiveness and efficiency of typical low-level network analytics tasks (see Sec. 4.3) between a traditional interface (the "baseline") and two interfaces with an increasing level of egocentrism ("Ego-Highlight" and "Ego-Bubble").

Our proposed interfaces target traditional data analysts that are accustomed to network analysis in 2D and thus are also aware of the limitations of conventional two-dimensional approaches. Considering the space requirements in conventional offices, our interfaces are designed for desktop-bound VR setups rather than room-scale ones. Free room-scale movement is possible but not required for navigating the networks in VR. This also facilitates frequent switching between desktop work and immersive analytics, which is required for data analysis [ZWB*17]. The types of networks commonly analyzed in network analytics often correspond to small-world networks, e.g., social or financial networks, co-morbidity networks, or metabolic pathways (see Sec. 4.4).

4.1. Egocentric Network Exploration Interfaces

The three interfaces that we investigate in our study (the baseline interface and the two egocentric interfaces), represent a continuum of egocentrism, i.e., from no alterations in the baseline, over adaptations to the representation (Ego-Highlight), to local layout adaptations

(Ego-Bubble). This continuum also leads to tension between two conflicting aspects: 1) freedom of navigation vs. ease of navigation and 2) layout fidelity vs. legibility. Both aspects are assumed to have an impact on visual search efficiency, navigation performance, spatial orientation, and cyber-sickness for topological graph analysis tasks, as discussed in Section 4.3.

The **baseline** interface displays a traditional 3D representation of the force-directed node-link diagram and features 6 degrees-of-freedom (DoF) free-fly navigation (see Fig. 1, left). Nodes are displayed as spheres, and links between the nodes are represented as tubes. The interface serves as a baseline condition in our experiments, so no egocentric adaptations are implemented. Navigation is handled via touch-sensitive input on the two axes of a trackpad, like the one on the HTC Vive controller. The user flies along and perpendicular to his/her current viewing direction, as defined by the HMD view-vector – just like moving and strafing in first-person ego-shooters. Since the layout remains unaltered, nodes and edges can serve as stable landmarks to support spatial orientation. However, edges can easily obstruct the user's view, especially in close proximity to nodes.

Ego-Highlight: In line with the principle of egocentric network analytics [Fis05], the user can associate him/herself with a node (i.e., the user-node), and view the network from this particular perspective. The user-node, as an explicit user defined reference point, allows us to introduce egocentric adaptations to the interface: we explicitly highlight direct neighbor nodes with a yellow halo (see Fig. 1, middle). This in turn, allows us to remove the now redundant links to direct neighbors to reduce visual clutter. This metaphor can further simplify navigation by restricting movement to jumps between nodes. By selecting a node with the controller's laser pointer and initiating a jump with the trigger button, the user is automatically transported to the center of the selected node with an animated translation of the user's position. The user maintains control of the camera orientation throughout, as it is linked to the HMD. This adaption trades flexibility in navigation for ease of use, as navigation no longer requires constant user input. Due to the nature of the force-directed layout, some neighbor nodes can be far away from the user-node, and might still be occluded through nodes and links that are not part of the direct neighborhood. We aim to further address this issue with additional egocentric adaptations in the next interface.

The **Ego-Bubble** interface extends the Ego-Highlight interface by locally optimizing the network layout in respect to the current user-node. In order to further alleviate potential occlusion in the user's local view, direct neighbor nodes are shifted to be evenly distributed around the user. We use the Fibonacci sphere algorithm on a bounding sphere centered around the user's head. Edges that cross the sphere are clipped (see Fig. 1, right) to avoid view obstructions through links passing by the vicinity of the user. The overall layout is thereby still preserved and only locally distorted, aiming to facilitate the local neighborhood analysis of the current user-node. This type of node displacement is inspired by well-established fish-eye techniques [SB94], employed in 2D visualizations of large graphs [AKY04]. A trade-off for this cleaner display of local neighborhoods is a potential negative impact on the user's orientation through the displacement of node positions. We therefore animate the local node displacement when jumping between nodes, result-

ing in a morphing effect during navigation that is designed to help users keep and update their mental map, partially remediating the potential loss of orientation.

4.2. Hypotheses

We tested six hypotheses concerning the potential benefits and negative side-effects of egocentrism in immersive network exploration:

H1: Local visual search is the most efficient in the Ego-Bubble condition and the least efficient in the baseline condition. Visual search is the task to spot a known target within a set of distractor items [WCF89]. If targets in the user's local neighborhood (i.e., neighboring nodes or edges) are occluded or outside the user's field of view, visual search will be negatively affected. We therefore expect to observe a considerable increase of visual search efficiency during local neighborhood exploration using egocentric interfaces that optimize the visibility of the user's neighborhood.

H2: Global visual search is the most efficient using the baseline. If the user's task is to search for more global structures, such as a path along multiple nodes through the network, it might be required to freely adjust the viewpoint to improve visibility. In the baseline the graph representation remains stable and navigation is not restricted to node positions. Visual search for global structures therefore can be expected to be more efficient using the baseline interface.

H3: Navigation performance is higher with egocentric network traversal than with flying in the baseline. Navigation performance is an important factor for the overall task efficiency when traversing a network in VR, e.g., when following a path. Traditional flying like in the baseline can lead to cyber-sickness [LLS18], with symptoms worsening as the movement pace increases [SLH01]. In an informal pilot study, we varied the animation speed between consecutive jumps and flying, and asked users to report their subjective well-being. As a result, the maximal flying speed is kept slow to maintain acceptable cyber-sickness rates. Contrarily, the jumping speed always measures three seconds (including ease-in/out de/acceleration). Using the baseline, flying along a path of five nodes in one of our networks (see Sec. 4.4) requires around 25 seconds, while the jump animations along five consecutive path nodes using the egocentric conditions results in around 15 seconds overall travel time. It can therefore be expected that navigation performance increases by a factor of ~ 1.7 for the egocentric network interfaces. This is also suggested by prior work, which shows that jumping leads to significantly faster travel times [WKFK18]. However, besides the motoric aspect, navigation consists also of wayfinding, which is the cognitive and stationary aspect of navigation [DP14]. We believe, however, that a potentially negative impact of the egocentric interface on wayfinding will not affect the clearly superior motoric aspect.

H4: Orientation is affected negatively through egocentric adaptations, i.e., suffers least in the baseline, and the most using Ego-Bubble. A study by Drogemüller et al. [DCW*18] has shown that flying causes less loss of orientation than teleportation when traversing a 3D network in VR. Orientation loss can also be observed for jumping, but to a lesser extent [WKFK18]. In contrast to these studies, the appearance of our scene changes after a jump, which can make it harder to keep oriented. Since the Ego-Bubble interface

also changes the layout of the nodes between jumps, we assume that users experience most orientation loss in this condition.

H5: Cybersickness symptoms are less severe when using Ego-Highlight compared to the baseline and Ego-Bubble. Potential causes for cyber-sickness are, among others, sensory conflicts of a stationary user flying through a virtual environment [LJ00]. There is prior evidence that jumping is less prone to cyber-sickness symptoms than omni-directional flying [LLS18]. However, this is based on the assumption that the scene is static during the transition. In the Ego-Bubble condition, the position of the surrounding nodes also changes during the animated jump, which could have a negative impact on the subjective well-being of participants compared to Ego-Highlight.

H6: User preference will be significantly higher for the egocentric interfaces than for the baseline. We assume that the expected benefits with respect to visual search and navigation performance will outweigh the expected draw-backs so that users prefer egocentric network exploration interfaces over the baseline. We also assume that their perceived work load will be lower.

4.3. Tasks

To enable accurate performance measurements in terms of correctness and completion times, we break the high-level task of network exploration down into low-level analytical tasks. The selected tasks are based on the well-accepted task taxonomy for graph visualization by Lee et al. [LPP*06], who categorize typical low-level graph analysis tasks into four larger categories: topology-based tasks, such as finding neighbors of a given node or finding common connections; attribute-based tasks, such as finding nodes or edges with certain associated attributes; browsing tasks, such as following a given path; and overview tasks, where the goal is to estimate values, such as the size of a network, very quickly.

We only investigate tasks that might benefit from local exploration in a detail view, as the overview perspective that serves as the starting point for certain tasks is identical across all interfaces. We therefore do not investigate tasks that are presumably much easier to carry out from an overview perspective, such as estimating the size of a network. We selected typical low-level network analytics tasks by Lee et al. [LPP*06] that fall into the categories of topology-based and browsing tasks. Below, we give a detailed description of our experiment tasks in the order in which they were presented to the users. Table 1 describes how these individual tasks relate to the high-level factors addressed by our hypotheses. Figure 4 depicts the individual tasks in the context of one of the networks used in the study, using a 2D force-directed layout for illustration purposes.

Find Neighbor (FiN): The aim of the task is to probe the ability of users to quickly query the local neighborhood of a hub node in order to find a specific node. Hub nodes for this task had between 14 and 44 neighbors. Users started this task in the detail perspective, from the center of the respective hub node. The task description specified the label of the target node that the user had to find. For reference, the target label was also displayed on the virtual representation of the VR controller. The task was automatically completed when the user found and selected the specified neighbor

Table 1: Factors addressed by the hypotheses and their associated task measures.

	High-Level Factor	Task Measures
H1	Local visual search efficiency	FiN completion time FCN completion time FCN correctness rate END judgement error
H2	Global visual search efficiency	FiP completion time FiP path correctness FiP path deviation
H3	Navigation performance	FoP completion time
H4	Spatial orientation	SO O→D angle deviation SO D→D angle deviation SO D→O angle deviation
H5	Cyber-sickness	SSQ [KLBL93]
H6	User preference	Preference ranking NASA-TLX [HS88]

node with the laser pointer. We measured task performance in terms of completion time.

Find Common Neighbors (FCN): Inspired by Kwon et al. [KMLM16], the task is to find all common neighbors between a given pair of highlighted nodes. Users started this task in the detail perspective, from the center of the first node of the pair. The respective other node of the pair was highlighted and visible when the task was initiated. Each node pair had between one to five common neighbors across all experiment configurations. The user had to pick all common neighbors using the laser pointer and signaled task completion to the co-present study instructor. We measured the completion time and logged the user-selected nodes. From the user selection, we then computed the correctness rate, the miss rate, and the false positive rate.

Estimate Node Degree (END): To assess the users' ability to identify hub nodes, we asked users to estimate the node degree of a specific node, i.e., the number of neighboring nodes. Nodes selected for this task had between 21 and 53 direct neighbors across all experiment conditions. Users started in the detail perspective, from the center of the hub node, and had to verbally report their node degree estimate to the experimenter. This task required more navigation in the baseline condition, as the user had to move to a position from where the hub node and its neighborhood could be observed. We therefore did not compare the completion time, and only analyzed the absolute deviation of the user's estimate from the ground truth node degree.

Spatial Orientation Overview → Detail (SO O→D): This task measures spatial orientation after teleportation from the overview to the detail perspective. The task was initiated from the overview, where users were presented with two highlighted nodes, i.e., the start and the end node of a path (see Fig. 3). Users had to press a button on the controller to start the task, after which the highlight of the end node was removed and the user was teleported to the detail perspective at the center of the start node. There, users were asked to point in the direction in which they estimated the end node's position. We measured the absolute angular deviation between the ray cast by the user and the ground truth direction vector between the controller position and the end node.

Find Path (FiP): Finding the shortest path between two given nodes is a common task in immersive network analytics studies [KMLM16, CDK*16, HFL*17, DCW*18, BVD19]. This task started from the end position of the previous task (SO O→D). After indicating the direction estimate, the end node of the path was highlighted again. From the detail perspective at the center of the start node, the user was now required to find the shortest connected path to the end node. The shortest path consisted of either four or five nodes across all experiment conditions. The user had to verbally report the identified nodes along the shortest path between the two highlighted start and end nodes. We counted paths as correct if all reported neighboring nodes were indeed connected by an edge. For correct paths, we computed the path deviation as the number of edges along the user-reported path divided by the number of edges along the shortest path. We also measured the time from task onset until the completed path report.

Follow Path (FoP): In this task, users were asked to follow a highlighted path through the network, as quickly as possible. The task was initiated from the overview perspective, where users were presented with a single highlighted start node. The user had to click on the highlighted node, so that the next node in the task sequence was highlighted. Depending on the interface, the click also initiated a jump to the (clicked) node in the sequence (Ego-Highlight and Ego-Bubble), or users had to manually navigate to the proximity of the node (baseline). We measured the time between clicking the start node and reaching the end node to assess the effect of jumping vs. flying through the graph.

Spatial Orientation Detail → Detail (SO D→D): This task measures spatial orientation after network traversal. The task started from the end position of the previous one (FoP). After reaching the end node, users were asked to point the laser pointer on the controller back into the direction of the start node (which was no longer highlighted). We measured the angle deviation as described for SO O→D. This task is similar to Kwon et al.'s "recall node locations" [KMLM16].

Spatial Orientation Detail → Overview (SO D→O): This task measures spatial orientation after teleportation from the detail view to the overview. The task started from the position of the previous task (SO D→D): After estimating the direction to the start node, users were teleported back to the initial overview position of the SO O→D task. From this position, users had to indicate the direction of the last node of the path they had followed. We measured the angular deviation of their estimation towards the ground truth direction.

4.4. Data

In order to achieve meaningful results, we performed our study on networks that resemble real world examples of networks of scientific interest, such as the world-wide-web and authors citation networks [AB02]. Such "scale-free" networks, have a node degree distribution that follows a power law: the probability $P(k)$ of having a node with degree k decreases exponentially with the law $P(k) \sim k^{-\gamma}$. This results in a small number of high degree nodes (hubs), while the majority of nodes have a low degree. When $2 < \gamma < 3$, the network is an "ultra-small network", with a relatively small diameter (i.e., length of the longest shortest path) $d \sim \ln \ln(N)$, and N the size

of the vertex set [CH03]. We generated our graphs with the Barabási-Albert random graph generator algorithm [AB02] (providing $\gamma = 3$) using the *NetworkX* Python library [HSSC08]. We used six graphs in our study experiments: three with 165 nodes and 326 edges ("small" in the following), and three with 415 nodes and 826 edges (referred to as "large", see Fig. 4), where the linear density (number of edges relative to number of nodes [YAD*18]) is 2.

4.5. Apparatus

The presented interfaces and the study framework are implemented as a client-only web application. The code is written in JavaScript using *three.js* [thr] and *A-Frame* [afr]. The interface between the VR application (i.e., the web-browser) and the VR hardware is handled by *SteamVR* [ste] that natively supports the HTC Vive. Our system builds on an open source library for viewing graphs in VR [for]. At the time of writing, the core library handles loading, layout, and rendering of the graph, while offering mouse and keyboard navigation. We extended this core with the presented egocentric interfaces, the study framework (i.e., scripting, scheduling task sequence), as well as to improve the core system's rendering performance. The most notable extensions concern VR controller support for the HTC Vive, perspective-dependent graph navigation and interaction modes, local layout and rendering adaptations, and support for switching between overview and detail perspectives.

All graph elements had the same color and size, i.e., node diameter and edge thickness. A single HTC Vive controller was used to handle user inputs via a trigger button, a trackpad and a virtual laser pointer. The laser pointer could be used in all conditions to investigate the direct neighborhood of the closest intersecting node by lowlighting the rest of the graph. Nodes could be discriminated from each other through unique alphanumeric labels that were always oriented towards the user (see Fig. 1). In our study, we excluded any sort of artificial, stationary landmarks from the scene to be able to better judge how well users can orient themselves based on the network topology itself (see H4). To create a controlled setting, we disabled teleportation between overview and detail view, except for the start (or end) of tasks SO O→D and D→O.

4.6. Study Design

We used a within-subjects design where *interface* is the independent variable and the individual task measures (see Table 1) are the dependent variables. To counter-balance the order of appearance of the three interface conditions between the participants, as well as the association of graph data sets to the interface conditions, we applied a Graeco-Latin square on these two factors. The task order was always presented in the sequence as listed in Section 4.3.

4.7. Procedure

Before starting the study[†], users had to sign a consent form and fill out a demographic questionnaire. A general description of the study procedure and of the VR controls were supplied as printouts

[†] The study was conducted before the global COVID-19 pandemic.




Figure 4: A 2D representation of one of the three networks used for the study, showing color-coded nodes for the different tasks: FiN (yellow = selected, green = target), FCN (purple = selected, green = target), END (blue = selected), FiP (cyan = start, salmon = target), FoP (red nodes).

to participants. Each user had to complete the interface conditions in the order specified by the Graeco-Latin square.

Before starting a condition, users were given a short textual description of the condition's interface. After this, they had to complete a short sequence of tutorial tasks that aimed at familiarizing them with the controls, i.e., navigation and interaction within the current interface. Users could remain in the tutorial until they felt comfortable with the controls and the visual encoding. Until the completion of the tutorial, users were permitted to pose questions about the interface and the study procedure. For each interface condition, the users then performed the set of tasks described in Section 4.3 twice: once on a small graph, and once on a large graph. Trials using small graphs were treated as training rounds to allow users to familiarize themselves with the tasks and the interface. Only the results of the trials using large graphs were finally analyzed, as discussed in Section 4.9, and presented in Section 5. After each interface condition, users were asked to fill out a Simulator Sickness Questionnaire (SSQ) [KLBL93], and a NASA Task Load Index (TLX) [HS88] questionnaire. At the end of the study, users were asked to rate their preference for the three conditions for solving network analytics tasks. The study was concluded by a semi-structured interview, where users were asked to describe what they liked or disliked about the individual conditions, tasks, and VR network exploration in general and report suggestions and improvements.

4.8. Participants

Based on the results of a power analysis after a pilot study, we recruited 25 subjects (19 males, six females) from a local university and an inter-university research facility. Participants were aged 23 to 64, with a median age of 31. All participants had normal or

corrected to normal vision. 17 participants had a background in computer science, others in physics or finance and economics. Four participants reported never to have played any computer games and not to have any experience with VR. Most users had tried out VR on one or multiple occasions. While 15 users reported to play computer games regularly, six users also play VR games regularly. All 25 participants completed the study. They were compensated with 10 Euros for their time.

4.9. Analysis

All obtained responses (completion times, correctness measures, and questionnaire responses) were analyzed individually per task. Completion time responses (tasks FiN, FCN, FiP, FoP) were log-transformed to reduce skewness and the effect of outliers [Rat93], while completion time charts show the untransformed data. All obtained measures were tested for normal distribution. Normally distributed responses were analyzed using a repeated measures ANCOVA with the three conditions as within-subjects factor and the order of the conditions, given by the Graeco-Latin Square, as covariate. If the normal distribution was violated, we removed outliers (samples higher or lower than $1.5 \cdot IQR$). If the filtered responses still did not follow a normal distribution, we performed a Friedman Test. All pairwise post-hoc comparisons were Bonferroni-adjusted.

5. Results

We report the results of the study with respect to the factors listed in Table 1 as F - and χ^2 -scores, as well as partial η^2 effect sizes for statistically significant results. The supplemental material contains the results of all performed tests. To put the quantitative results into a qualitative context, we also report user preferences and feedback from a semi-structured interview (see Sec. 5.6).

5.1. Local Visual Search Efficiency

For finding a specified neighbor (FiN), we found that users were significantly slower using the baseline than Ego-Highlight or Ego-Bubble ($F_{2,42} = 25.722; p < .001; \eta^2 = .551$). On average, users needed around 9 seconds to find the given neighbor using Ego-Bubble, 20 seconds using Ego-Highlight, and more than 100 seconds using the baseline (see Fig. 5).




Figure 5: The completion times in seconds for the FiN task.

Similarly, we found a large and significant difference in terms of completion time for finding common neighbors (FCN): $F_{1,455,29,101} = 20.853; p < .001; \eta^2 = .510$. On average, users required 55 seconds to complete the task using the baseline, which is significantly slower than using Ego-Highlight (32 seconds) or Ego-Bubble (29 seconds), which are not significantly different, as shown in Figure 6. In terms of correctness, we did not find any significant differences for the FCN task. On average, though, users could reach

the highest correctness using Ego-Highlight (92%) and the lowest using the baseline (74%). The high error rate of the baseline was mainly caused by missed common neighbors (28% miss rate), while Ego-Bubble had the highest false positive rate (18%), as shown in Figure 6.

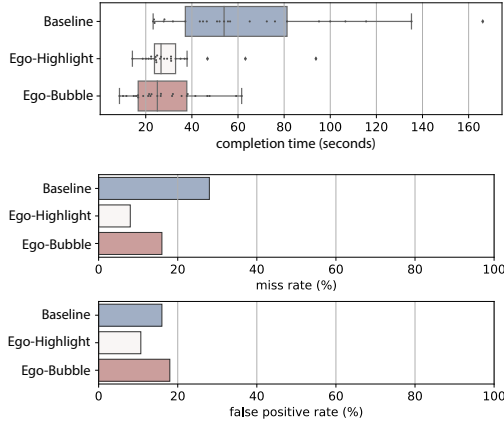


Figure 6: Completion times in seconds, miss rates, and false positive rates for the FCN task.

We also found a significant difference between node degree estimation (END) errors ($\chi^2(2) = 6.720; p = .035$), where user reports were significantly more deviating from the ground truth node degrees using the baseline (26%) than when using Ego-Bubble (15%). For all three conditions, users tended to underestimate the number of neighbors rather than overestimate them (see Fig. 7).




Figure 7: Judgement error (deviation from ground truth, END task).

Overall, our hypothesis concerning the effect of egocentrism on local visual search efficiency is partially supported (**H1**): Indeed, *users are less efficient when analyzing the local neighborhood using the baseline*. The benefit of the Ego-Bubble interface in comparison to the Ego-Highlight interface regarding efficiency, however, is rather small and does not reach significance.

5.2. Global Visual Search Efficiency

For finding the shortest path, we first removed those users who did not report a correct path in either condition (i.e., they reported non-adjacent nodes or they gave up before reaching the target node). In total, nine responses by six users were incorrect paths, evenly distributed across the three conditions, and nine of the correct paths were longer than the shortest path. The average deviation from the ground truth path length was highest for Ego-Bubble (7%) and lowest for the baseline (3%), but this difference is not statistically significant. There is a large, yet statistically insignificant difference

in terms of completion time for those users who reported correct paths: $F_{2,32} = 2.956; p = .066; \eta^2 = .156$ (see Fig. 8).




Figure 8: Completion times in seconds for the FiP task.

Therefore, our results do not support **H2**: *Users were only insignificantly more accurate but also slower when trying to find the shortest path between two nodes using the baseline compared to the egocentric interfaces*.

5.3. Navigation Performance

As expected, for the path following task (FoP), we found a large and significant effect of the interface on completion time ($F_{2,44} = 49.243; p < .001; p \eta^2 = .691$). However, the observed speed differences between the interfaces were larger than expected: Using the baseline, users required significantly more time to follow the highlighted path (40 seconds, on average) than with both, Ego-Highlight (23 seconds) and Ego-Bubble (21 seconds); see Figure 9.




Figure 9: Completion times in seconds for the FoP task.

This supports our hypothesis (**H3**) that *jumping through the network along a highlighted path is faster (almost twice as fast) than flying through it*. This indicates that egocentrism does not add extra effort for wayfinding.

5.4. Orientation

In general, the measured angle deviations for node-direction estimations were quite large, ranging from 19° average deviation after jumping from the overview to the detail view (SO O→D) to an average error of 35° when estimating the direction of the start node after following the path in task FoP (SO D→D). However, for none of the three spatial orientation tasks, we could find any statistically significant difference between the conditions.

Therefore our hypothesis (**H4**) that local adaptations introduced through egocentrism lead to decreased spatial orientation in the network is not supported. However, the high angle deviation clearly shows that *explicit spatial orientation cues are required for immersive network exploration, irrespective of the detail interface*.

5.5. Cyber-sickness

We derived three metrics from the SSQ, as described by Kennedy et al. [KLBL93]: nausea (e.g., sweating and stomach awareness), oculomotor (e.g., fatigue, headache, and eye strain), and disorientation (e.g., vertigo and dizziness). Note that, due to the within-subjects study design, we compared the scores directly to each other, not to a pre-exposure baseline. While we did not find any differences between the interface conditions in terms of nausea, we found a significant effect on the oculomotor-related responses ($\chi^2(2) = 8.195; p = .017$) and disorientation symptoms ($\chi^2(2) = 8.617; p = .013$). Contrary to our hypothesis (**H5**), post-hoc comparisons for both metrics revealed a significant difference between the baseline and Ego-Bubble. On average, *Ego-Bubble received slightly lower oculomotor and disorientation symptom scores than the baseline, but the difference between Ego-Highlight and the other two conditions is small and insignificant* (see Fig. 10).

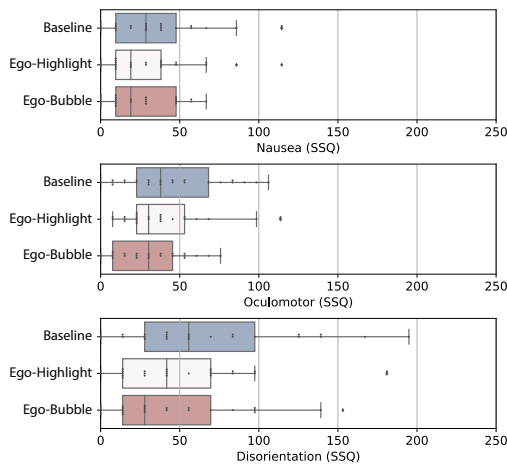


Figure 10: SSQ scores related to nausea (top), oculomotor (middle), and disorientation (bottom).

5.6. User Preferences and Feedback

We computed an aggregated task load index, as described by Rubio et al. [RDMP04], to assess the perceived work load of users per condition. The difference between the interface conditions with respect to the perceived task load is significant ($F_{2,48} = 6.546; p = .003; \eta^2 = .214$). The baseline received significantly higher task load scores (3.8 on average) than Ego-Bubble (3.3) and Ego-Highlight (3.4).

Analogously to the perceived task load, users also ranked their preference for the baseline considerably lower than for Ego-Highlight and Ego-Bubble. Only four users ranked the baseline first; these users appreciated the fact that they were able to move and look around freely and that the interface was very simple. Ego-Highlight and Ego-Bubble were ranked first by eleven participants each (with one vote tied between them). Users appreciated Ego-Highlight mainly because of its implicit highlighting of direct neighbors. Three users mentioned that they felt less nausea or vertigo in the Ego-Bubble condition, and five described it as “tidier” or less cluttered compared to the other conditions. Two users preferring

Ego-Highlight stated that they did not find any advantage of the Ego-Bubble condition compared to Ego-Highlight. One user preferring the baseline found it hard to orient in the Ego-Bubble view, presumably due to the “nodes moving towards oneself”.

Users suggested several improvements for the immersive interfaces. Almost half of the users (10) suggested to enable permanent highlighting of user-selected nodes. Three users wanted to have color-coding based on the geodesic distance towards their current location. Three other users suggested to have additional information about the graph nodes and edges (such as their node degree), for instance as a details-on-demand label on the controller. Six users explicitly mentioned that they do not want to be (exclusively) inside the network, but would like to gain some overview from an outside perspective as well. Eight users requested persistent landmarks in the scene to support spatial orientation, and three users would have preferred to have a combination of jumping and flying navigation. Overall, our hypothesis **H6** is supported: *Egocentric immersive network exploration is clearly preferred by the users and also leads to a lower subjective task load.*

6. Discussion

The results of our study support our basic assumption: immersive egocentric network exploration interfaces are more efficient and more effective than their non-egocentric counter-part. Below, we discuss our findings in regard to the aspects investigated in our study.

Visual search in the local neighborhood is clearly facilitated by a decluttered egocentric interface. For instance, finding a neighbor node with a given label takes around five times as long using the baseline interface. Clearly, the implicit highlighting of neighbors in the egocentric conditions facilitates the identification of connections in the local neighborhood. However, such highlighting is only possible if the user is associated with a chosen node. While the implicit highlighting of neighbors in Ego-Highlight has a large effect, the local layout optimization of the Ego-Bubble interface has only a small effect, indicated by slightly less underestimation in the degree estimation task. Interestingly, only a few participants found Ego-Bubble “tidier”, while many users did not notice the difference between the two egocentric interfaces, or they found it irrelevant.

Global visual search for paths between two distant nodes was considered a hard task. The baseline interface was insignificantly more accurate for finding paths, presumably due to edges being always visible and due to unconstrained navigation. From informal user feedback, we conclude that path finding in a detail perspective is generally not efficient without dedicated interaction support.

Navigation performance was significantly lower in the baseline than in the egocentric interface conditions, where users could jump between nodes. The navigation performance was effectively doubled when jumping through the network, comparable to results of a prior study using natural scenes [WKFK18]. Due to the lower navigation speed and manual steering in the baseline, this was expected, yet not to the observed extent. It has to be noted that different flying methods for immersive network exploration have been investigated [DCW*18, ZHF*16], which could slightly improve navigation performance compared to our classic ego-shooter-like navigation in the baseline interface.

In our study, we observed **orientation** loss after teleportation between the overview and the detail view for all interfaces. Orientation loss after teleportation can be expected [MBR18, DCW*18]. However, we also observed considerable loss of orientation after traversing the network in the detail view – irrespective of whether the user was continuously flying or jumping between nodes. Contrary to our expectations, orientation loss is not higher in the egocentric views than in the baseline condition. This is a strong indication that in abstract detail views additional orientation aids, such as static landmarks, are required to support orientation.

As expected, the baseline received the highest average **cyber-sickness** ratings. The low cyber-sickness ratings for the Ego-Bubble interface were surprising, despite having additional morphing animations during navigation. A speculative explanation could be the reduced scene complexity [SHL01] and the less pronounced optical flow effect, caused by a constant, farther distance of the neighboring nodes to the camera. In the future, it might be interesting to investigate the effect of visualization decluttering on the users' well-being during immersive analytics.

7. Actionable Insights From Our Study

Our study allowed us to investigate the effects of egocentrism on immersive network exploration, but it also helped us to identify open challenges that could not be sufficiently resolved by this metaphor alone. Many aspects of network exploration remain challenging in an immersive detail perspective, despite egocentric optimizations: node degrees tend to get underestimated, common neighbors can be missed, non-existing edges are sometimes inferred, and maintaining spatial orientation is difficult. In an effort to alleviate these issues, we developed an online tool for immersive network exploration, leveraging the following insights from our study [Ego]. In our web application, overview and detail perspectives can be switched at any point during the exploration. The overview is represented by a camera object that can be selected with the laser pointer. The switch is executed through a smooth transition, to avoid potential disorientation after instant teleportation (see Sec. 6).

Permanent node selection: Almost half of the users suggested such a feature as a means to “bookmark” potentially interesting nodes, which could, for instance, help to identify paths.

Geodesic distance encoding: Path identification was challenging in all interfaces. However, by leveraging information on the node, the user is currently associated with, we can visually encode the geodesic distance to all other nodes to better support such tasks. We encode the distance in a color range from yellow to red: the closer a node's geodesic distance, the closer is its color to red – clearly showing connected nodes and their path length in respect to the user, which does not necessarily correspond to the Euclidean distance in a force-directed graph layout.

Node Info Display: As many users requested more quantitative information during their exploration, we added a head-up display attached to the controller displaying information about the current user node's node degree.

Flexible Navigation: In our study, we enabled flying *or* jumping for network navigation. The feedback of some study participants

indicates that, in practice, a combination of both modalities could be beneficial for more experienced users. As users are not uniquely associated with a dedicated node during free-flying, our implementation currently applies egocentric adaptations automatically to the closest node in respect to their position within the network. However, informal feedback revealed that fish-eye distortions in combination with free flying induces considerable cyber-sickness.

Landmarks: Since orientation in the 3D network proved challenging during navigation and after teleportation across all interfaces, we provide a skybox displaying a ground- and sky-plane as a static landmark. As additional optional aid we provide visual cues by highlighting all visited nodes when viewing the graph from the overview perspective. It will be necessary to investigate how such measures will fare in supporting spatial orientation in the future.

8. Study Limitations and Future Work

In our study, we only investigated a small fraction of the vast design space of egocentric network exploration interfaces. Here, we discuss some of the limitations of our study and suggest specific future work. We deliberately limited our tasks to those that would potentially benefit from an immersive detail perspective, as the **overview** perspective was identical in all three tasks. It is clear that many tasks not investigated in this study could be solved more efficiently from an overview perspective, such as estimating the number of communities (similarly as performed by Greffard et al. [GPK11]) or finding the most central node in the graph. Investigating the benefits of egocentric optimizations for network overview- and hybrid-visualizations would be an interesting avenue for future work. We currently only use the user-selected node (or the user's position in the network) as information on which egocentric optimizations are based on. Further research is needed to evaluate the potential benefits of including other user-centered information, such as the field-of-view or even the user's gaze-focus.

Our goal was *not* to formally evaluate the benefit of 3D network representations in VR in respect to classic 2D networks. We relied on ample prior evidence suggesting benefits of stereo and immersive 3D networks – at least as complementary analysis environment – for path finding, especially with a larger number of nodes [WF96, BBHS03], community detection in complex graphs [GPK11, KWO*19], for distance assessments between 3D points [WFRFN18], and for a generally improved mental model of the graph [KKM*20].

In our study, we assumed that many users do not have the required space to use a room-scale VR system. Such a setup would allow users to navigate the graph simply by **walking** through it. Room-scale navigation potentially reduces the amount of cyber-sickness [UAW*99, LLS18]. Yang et al. [YCB*20] recently conducted a study where users could walk in room-sized immersive networks. A formal comparative study between walking, flying, and jumping through a node-link diagram has not yet been conducted.

According to a survey by Yoghoudjian et al. [YAD*18], the networks we used in our evaluation were very large in terms of the number of nodes but sparse in terms of their **density**. We deliberately limited the number of edges after some informal pilot tests to be able solve local visual search tasks in the baseline condition, which would

have been increasingly difficult in denser graphs (cf. Fig. 1, left). In the future, it will be important to conduct a similar experiment comparing egocentric clutter reducing interfaces in denser networks. Our expectation is that clutter reduction techniques, such as the fish-eye distortion or edge bundling [Hol06, HVW09], will become more important with increasing graph density.

9. Conclusions

In this work, we promoted *egocentrism* as a metaphor for user-centered optimizations in immersive network exploration interfaces. We introduced two simple instances of egocentrism and systematically investigated their effects in a user study in comparison to a traditional network exploration interface. We obtained initial insights about the metaphor's performance on typical network exploration tasks and potential caveats. The study showed how interfaces implementing egocentrism considerably increase local visual search efficiency by locally de-cluttering and highlighting the direct node neighborhood. While additional local layout optimizations yield comparably fewer advantages, users reported slightly lower symptoms of cyber-sickness compared to the traditional interface. Users prefer the egocentric de-cluttered perspective of the network. Sacrificing free navigation is reported to be only a minor trade-off. We speculate that the benefit of a local distortions will be evident in denser networks, which needs to be studied in the future.

We conclude from our study that egocentric visualization and interaction design that is carefully tailored towards the user's relation to the data can yield significant benefits for immersive detail views in network exploration. We therefore see the design and investigation of further effective egocentric metaphors as a great open research challenge in the field of immersive analytics.

Acknowledgments

This work is supported by the Austrian Research Agency (FFG) project #882184. We would like to thank all study participants for their time and effort.

References

- [AB02] ALBERT R., BARABÁSI A.-L.: Statistical mechanics of complex networks. *Reviews of modern physics* 74, 1 (2002), 47. 6
- [afr] A-frame: Hello webvr. <https://aframe.io>. Online; accessed 30-August-2021. 6
- [AHE14] ALTARAWNEH R., HUMAYOUN S. R., EBERT A.: Expand: A stereoscopic expanding technique for compound graphs. In *Proc. of the 20th ACM Symposium on Virtual Reality Software and Technology* (2014), pp. 223–224. 2
- [AHKM11] ALPER B., HOLLERER T., KUCHERA-MORIN J., FORBES A.: Stereoscopic highlighting: 2d graph visualization on stereo displays. *IEEE trans. on Visualization and Computer Graphics* 17, 12 (2011), 2325–2333. 2
- [AKY04] ABELLO J., KOBOUROV S. G., YUSUFOV R.: Visualizing large graphs with compound-fisheye views and treemaps. In *International Symposium on Graph Drawing* (2004), Springer, pp. 431–441. 4
- [BBHS03] BELCHER D., BILLINGHURST M., HAYES S., STILES R.: Using augmented reality for visualizing complex graphs in three dimensions. In *The Second IEEE and ACM International Symposium on Mixed and Augmented Reality, 2003. Proc.* (2003), IEEE, pp. 84–93. 1, 2, 10
- [BM07] BOWMAN D. A., MCMAHAN R. P.: Virtual reality: how much immersion is enough? *Computer* 40, 7 (2007), 36–43. 1
- [BVD19] BÜSCHEL W., VOGT S., DACHSELT R.: Augmented reality graph visualizations. *IEEE computer graphics and applications* 39, 3 (2019), 29–40. 6
- [CDK*16] CORDEIL M., DWYER T., KLEIN K., LAHA B., MARRIOTT K., THOMAS B. H.: Immersive collaborative analysis of network connectivity: Cave-style or head-mounted display? *IEEE trans. on visualization and computer graphics* 23, 1 (2016), 441–450. 2, 6
- [CH03] COHEN R., HAVLIN S.: Scale-free networks are ultrasmall. *Physical review letters* 90, 5 (2003), 058701. 6
- [CLFZ95] CRUTCHER L. A., LAZAR A. A., FEINER S. K., ZHOU M. X.: Managing networks through a virtual world. *IEEE Parallel & Distributed Technology: Systems & Applications* 3, 2 (1995), 4–13. 2
- [DCW*17] DROGEMULLER A., CUNNINGHAM A., WALSH J., ROSS W., THOMAS B. H.: Vrige: exploring social network interactions in immersive virtual environments. In *Proc. of the international symposium on big data visual analytics (BDVA). IEEE NJ, USA* (2017). 1, 2
- [DCW*18] DROGEMULLER A., CUNNINGHAM A., WALSH J., CORDEIL M., ROSS W., THOMAS B.: Evaluating navigation techniques for 3d graph visualizations in virtual reality. In *2018 International Symposium on Big Data Visual and Immersive Analytics (BDVA)* (2018), IEEE, pp. 1–10. 2, 4, 6, 9, 10
- [DP14] DARKEN R. P., PETERSON B.: Spatial orientation, wayfinding, and representation., 2014. 4
- [DPHB15] DEBARBA H. G., PERRIN S., HERBELIN B., BOULIC R.: Embodied interaction using non-planar projections in immersive virtual reality. In *Proc. of the 21st ACM Symposium on Virtual Reality Software and Technology* (2015), pp. 125–128. 2
- [Ego] Egocentric vr: online immersive network exploration. <https://vis.csh.ac.at/egocentricvr/>. Online; accessed 30-August-2021. 10
- [EMP19] ERRA U., MALANDRINO D., PEPE L.: Virtual reality interfaces for interacting with three-dimensional graphs. *International Journal of Human-Computer Interaction* 35, 1 (2019), 75–88. 1, 2
- [Fis05] FISHER D.: Using egocentric networks to understand communication. *IEEE Internet Computing* 9, 5 (2005), 20–28. 2, 4
- [for] Aframe forcegraph component. <https://github.com/vasturiano/aframe-forcegraph-component>. Online; accessed 30-August-2021. 6
- [Gol13] GOLBECK J.: Chapter 2 - nodes, edges, and network measures. In *Analyzing the Social Web*, Golbeck J., (Ed.). Morgan Kaufmann, Boston, 2013, pp. 9 – 23. 2
- [GPK11] GREFFARD N., PICAROUGNE F., KUNTZ P.: Visual community detection: An evaluation of 2d, 3d perspective and 3d stereoscopic displays. In *International Symposium on Graph Drawing* (2011), Springer, pp. 215–225. 1, 2, 10
- [HFL*17] HUANG Y.-J., FUJIWARA T., LIN Y.-X., LIN W.-C., MA K.-L.: A gesture system for graph visualization in virtual reality environments. In *2017 IEEE Pacific Visualization Symposium (PacificVis)* (2017), IEEE, pp. 41–45. 2, 6
- [Hol06] HOLTEN D.: Hierarchical edge bundles: Visualization of adjacency relations in hierarchical data. *IEEE trans. on visualization and computer graphics* 12, 5 (2006), 741–748. 11
- [HS88] HART S. G., STAVELAND L. E.: Development of nasa-tlx (task load index): Results of empirical and theoretical research. In *Advances in psychology*, vol. 52. Elsevier, 1988, pp. 139–183. 5, 7
- [HSSC08] HAGBERG A., SWART P., S CHULT D.: *Exploring network structure, dynamics, and function using NetworkX*. Tech. rep., Los Alamos National Lab.(LANL), Los Alamos, NM (United States), 2008. 6
- [HVW09] HOLTEN D., VAN WIJK J. J.: Force-directed edge bundling for graph visualization. In *Computer graphics forum* (2009), vol. 28, Wiley Online Library, pp. 983–990. 11

- [HZBK08] HALPIN H., ZIELINSKI D. J., BRADY R., KELLY G.: Exploring semantic social networks using virtual reality. In *International Semantic Web Conf.* (2008), Springer, pp. 599–614. 1, 2
- [KCKW19] KREKHOV A., CMENTOWSKI S., WASCHK A., KRÜGER J.: Deadeye visualization revisited: Investigation of preattentiveness and applicability in virtual environments. *IEEE transactions on visualization and computer graphics* 26, 1 (2019), 547–557. 2
- [KKM*20] KOTLAREK J., KWON O.-H., MA K.-L., EADES P., KERREN A., KLEIN K., SCHREIBER F.: A study of mental maps in immersive network visualization. *arXiv preprint arXiv:2001.06462* (2020). 1, 2, 10
- [KLBL93] KENNEDY R. S., LANE N. E., BERBAUM K. S., LILIEN-THAL M. G.: Simulator sickness questionnaire: An enhanced method for quantifying simulator sickness. *The international journal of aviation psychology* 3, 3 (1993), 203–220. 5, 7, 9
- [KMLM15] KWON O.-H., MUELDER C., LEE K., MA K.-L.: Spherical layout and rendering methods for immersive graph visualization. In *2015 IEEE Pacific Visualization Symposium (PacificVis)* (2015), IEEE, pp. 63–67. 2, 3
- [KMLM16] KWON O.-H., MUELDER C., LEE K., MA K.-L.: A study of layout, rendering, and interaction methods for immersive graph visualization. *IEEE trans. on visualization and computer graphics* 22, 7 (2016), 1802–1815. 1, 2, 5, 6
- [KWO*19] KRAUS M., WEILER N., OELKE D., KEHRER J., KEIM D. A., FUCHS J.: The impact of immersion on cluster identification tasks. *IEEE trans. on visualization and computer graphics* 26, 1 (2019), 525–535. 2, 10
- [LJ00] LAVIOLA JR J. J.: A discussion of cybersickness in virtual environments. *ACM Sigchi Bulletin* 32, 1 (2000), 47–56. 2, 5
- [LLS18] LANGBEHN E., LUBOS P., STEINICKE F.: Evaluation of locomotion techniques for room-scale vr: Joystick, teleportation, and redirected walking. In *Proc. of the Virtual Reality International Conference-Laval Virtual* (2018), ACM, p. 4. 4, 5, 10
- [LPP*06] LEE B., PLAISANT C., PARR C. S., FEKETE J.-D., HENRY N.: Task taxonomy for graph visualization. In *Proc. of the 2006 AVI workshop on BEyond time and errors: novel evaluation methods for information visualization* (2006), ACM, pp. 1–5. 5
- [MBSJ97] MINE M. R., BROOKS JR F. P., SEQUIN C. H.: Moving objects in space: exploiting proprioception in virtual-environment interaction. In *SIGGRAPH* (1997), vol. 97, pp. 19–26. 2
- [MBR18] MOGHADAM K. R., BANIGAN C., RAGAN E. D.: Scene transitions and teleportation in virtual reality and the implications for spatial awareness and sickness. *IEEE trans. on visualization and computer graphics* (2018). 10
- [MCH*18] MARRIOTT K., CHEN J., HLAWATSCH M., ITOH T., NACENTA M. A., REINA G., STUERZLINGER W.: Immersive analytics: Time to reconsider the value of 3d for information visualisation. In *Immersive Analytics*. Springer, 2018, pp. 25–55. 2
- [MSD*18] MARRIOTT K., SCHREIBER F., DWYER T., KLEIN K., RICHEL N. H., ITOH T., STUERZLINGER W., THOMAS B. H.: *Immersive Analytics*, vol. 11190. Springer, 2018. 1
- [OAS00] OSAWA N., ASAI K., SUGIMOTO Y. Y.: Immersive graph navigation using direct manipulation and gestures. In *Proc. of the ACM symposium on Virtual reality software and technology* (2000), ACM, pp. 147–152. 2
- [OWK*14] ORLOSKY J., WU Q., KIYOKAWA K., TAKEMURA H., NITSCHKE C.: Fisheye vision: peripheral spatial compression for improved field of view in head mounted displays. In *Proc. of the 2nd ACM symposium on Spatial user interaction* (2014), pp. 54–61. 2
- [Rat93] RATCLIFF R.: Methods for dealing with reaction time outliers. *Psychological bulletin* 114, 3 (1993), 510. 7
- [RBLN04] RAJA D., BOWMAN D., LUCAS J., NORTH C.: Exploring the benefits of immersion in abstract information visualization. In *Proc. Immersive Projection Technology Workshop* (2004), pp. 61–69. 2
- [RDMP04] RUBIO S., DÍAZ E., MARTÍN J., PUENTE J. M.: Evaluation of subjective mental workload: A comparison of swat, nasa-tlx, and workload profile methods. *Applied psychology* 53, 1 (2004), 61–86. 9
- [SB94] SARKAR M., BROWN M. H.: Graphical fisheye views. *Communications of the ACM* 37, 12 (1994), 73–85. 4
- [SHL01] SO R. H., HO A., LO W.: A metric to quantify virtual scene movement for the study of cybersickness: Definition, implementation, and verification. *Presence: Teleoperators & Virtual Environments* 10, 2 (2001), 193–215. 10
- [SLH01] SO R. H., LO W., HO A. T.: Effects of navigation speed on motion sickness caused by an immersive virtual environment. *Human factors* 43, 3 (2001), 452–461. 4
- [ste] Steamvr. <https://store.steampowered.com/steamvr>. Online; accessed 30-August-2021. 6
- [SWKA19] SORGER J., WALDNER M., KNECHT W., ARLEO A.: Immersive analytics of large dynamic networks via overview and detail navigation. In *2019 IEEE International Conf. on Artificial Intelligence and Virtual Reality (AIVR)* (2019), IEEE, pp. 144–1447. 1, 2, 3
- [thr] three.js – javascript 3d library. <https://threejs.org/>. Online; accessed 30-August-2021. 6
- [UAW*99] USOH M., ARTHUR K., WHITTON M. C., BASTOS R., STEED A., SLATER M., BROOKS JR F. P.: Walking> walking-in-place> flying, in virtual environments. In *Proc. of the 26th annual conf. on Computer graphics and interactive techniques* (1999), ACM Press/Addison-Wesley Publishing Co., pp. 359–364. 10
- [WCF89] WOLFE J. M., CAVE K. R., FRANZEL S. L.: Guided search: an alternative to the feature integration model for visual search. *Journal of Experimental Psychology: Human perception and performance* 15, 3 (1989), 419. 4
- [WF96] WARE C., FRANCK G.: Evaluating stereo and motion cues for visualizing information nets in three dimensions. *ACM trans. on Graphics (TOG)* 15, 2 (1996), 121–140. 1, 2, 10
- [WFRFN18] WAGNER FILHO J. A., REY M. F., FREITAS C. M., NEDEL L.: Immersive visualization of abstract information: An evaluation on dimensionally-reduced data scatterplots. In *2018 IEEE Conf. on Virtual Reality and 3D User Interfaces (VR)* (2018), IEEE, pp. 483–490. 10
- [WKFK18] WEISSKER T., KUNERT A., FRÖHLICH B., KULIK A.: Spatial updating and simulator sickness during steering and jumping in immersive virtual environments. In *2018 IEEE Conf. on Virtual Reality and 3D User Interfaces (VR)* (2018), IEEE, pp. 97–104. 4, 9
- [WM08] WARE C., MITCHELL P.: Visualizing graphs in three dimensions. *ACM trans. on Applied Perception (TAP)* 5, 1 (2008), 2. 1
- [YAD*18] YOGHOORDJIAN V., ARCHAMBAULT D., DIEHL S., DWYER T., KLEIN K., PURCHASE H. C., WU H.-Y.: Exploring the limits of complexity: A survey of empirical studies on graph visualisation. *Visual Informatics* 2, 4 (2018), 264–282. 6, 10
- [YCB*20] YANG Y., CORDEIL M., BEYER J., DWYER T., MARRIOTT K., PFISTER H.: Embodied navigation in immersive abstract data visualization: Is overview+ detail or zooming better for 3d scatterplots? *IEEE Transactions on Visualization and Computer Graphics* 27, 2 (2020), 1214–1224. 1, 2, 3, 10
- [ZHF*16] ZIELASKO D., HORN S., FREITAG S., WEYERS B., KUHNEN T. W.: Evaluation of hands-free hmd-based navigation techniques for immersive data analysis. In *2016 IEEE Symposium on 3D User Interfaces (3DUI)* (2016), IEEE, pp. 113–119. 1, 2, 9
- [ZWB*17] ZIELASKO D., WEYERS B., BELLGARDT M., PICK S., MEIBNER A., VIERJAHN T., KUHNEN T. W.: Remain seated: towards fully-immersive desktop vr. In *2017 IEEE 3rd Workshop on Everyday Virtual Reality (WEVR)* (2017), IEEE, pp. 1–6. 2, 3