

GAN'SDA Wrap: Geographic And Network Structured DAta on surfaces that Wrap around

Kun-Ting Chen
 Monash University
 Melbourne, VIC, Australia
 kun-ting.chen@monash.edu

Tim Dwyer
 Monash University
 Melbourne, VIC, Australia
 tim.Dwyer@monash.edu

Yalong Yang
 Virginia Tech
 Blacksburg, VA, USA
 yalongyang@vt.edu

Benjamin Bach
 University of Edinburgh
 Edinburgh, UK
 bach@ed.ac.uk

Kim Marriott
 Monash University
 Melbourne, VIC, Australia
 kim.marriott@monash.edu

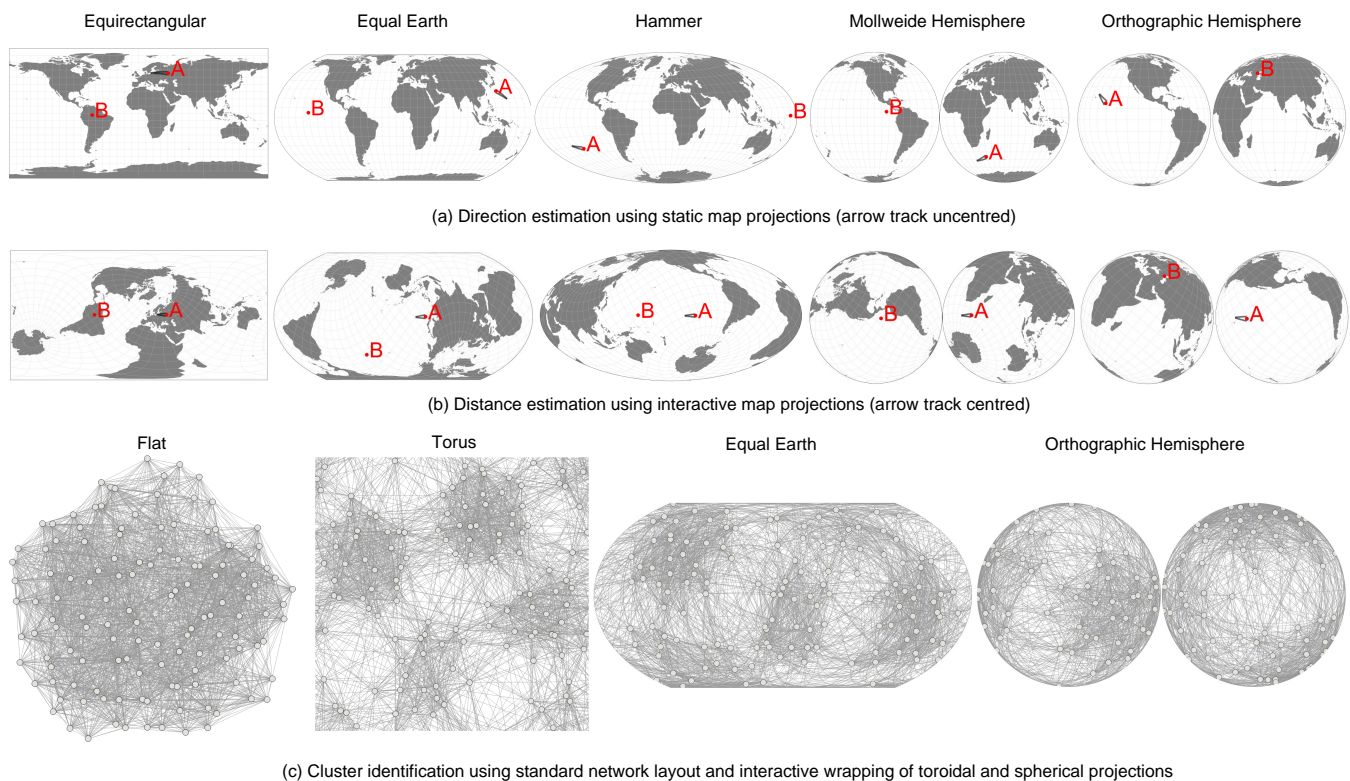


Figure 1: Top row: Direction estimation task of geographic data using the five different spherical projections evaluated in Study 1, with (a) and without (b) interaction. Bottom row: Study 2 compares the two interactive spherical projections found most effective in Study 1, and compares them to toroidal and standard ‘flat’ layout for network data.

ABSTRACT

There are many methods for projecting spherical maps onto the plane. Interactive versions of these projections allow the user to centre the region of interest. However, the effects of such interaction have not previously been evaluated. In a study with 120 participants we find interaction provides significantly more accurate area, direction and distance estimation in such projections. The surface of 3D sphere and torus topologies provides a continuous surface for uninterrupted network layout. But how best to project spherical

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than ACM must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from permissions@acm.org.
 CHI '22, April 29-May 5, 2022, New Orleans, LA, USA
 © 2022 Association for Computing Machinery.
 ACM ISBN 978-1-4503-9157-3/22/04...\$15.00
<https://doi.org/10.1145/3491102.3501928>

network layouts to 2D screens has not been studied, nor have such spherical network projections been compared to torus projections. Using the most successful interactive sphere projections from our first study, we compare spherical, standard and toroidal layouts of networks for cluster and path following tasks with 96 participants, finding benefits for both spherical and toroidal layouts over standard network layouts in terms of accuracy for cluster understanding tasks.

CCS CONCEPTS

• **Human-centered computing** → **Graph drawings; Empirical studies in visualization; Geographic visualization.**

KEYWORDS

Geographic visualization; Map projection; Graph visualization; Crowdsourced study.

ACM Reference Format:

Kun-Ting Chen, Tim Dwyer, Yalong Yang, Benjamin Bach, and Kim Marriott. 2022. GAN'SDA Wrap: Geographic And Network Structured DAta on surfaces that Wrap around. In *CHI Conference on Human Factors in Computing Systems (CHI '22)*, April 29-May 5, 2022, New Orleans, LA, USA. ACM, New York, NY, USA, 16 pages. <https://doi.org/10.1145/3491102.3501928>

1 INTRODUCTION

While most people agree that the world is spherical, most computer displays remain stubbornly flat. To show maps of the entire earth's surface on a screen we have to somehow cut, stretch and squash it. While cartographers and mathematicians have developed many methods to *project* the surface of a sphere to a plane, distortion or discontinuity at the edges of the projection are inevitable; with different methods achieving different trade-offs between (e.g.) area, distance, or direction preservation and discontinuities [56]. As a consequence distances and areas closer to the north and south poles may appear larger and the shapes more distorted than those at the centre [33]. Moreover, distances between geographical locations or areas of continents or seas that are wrapped across the projection boundary are also easily misinterpreted [26].

Most of these projections were originally developed for *static* display of the earth in printed maps and atlases. However, with modern computer graphics we can create interactive versions of these projections, which can be panned, reoriented and recentred with simple mouse or touch drags [6, 17], thus allowing the viewer to centre a region of interest so as to minimise distortion and discontinuities.

But spherical projection has application beyond geographic data. There can be advantages to laying out abstract data (data which has no inherent geometry) on the surface of a sphere such that there is no arbitrary edge to the display or privileged centre [49]. In particular, node-link representations of highly connected network data (which is ubiquitous in the world around us, from protein-protein interaction networks, to social, communication, trade or electrical networks) have no real “inside” or “outside”. Arguably, they can be better distributed across the surface of a topology which “wraps around” continuously.

While the relationship between readability of geographic and network data on spherical projections has not been studied directly,




there are commonalities in the analysis tasks that might be applicable for each. For example, understanding network clusters may involve comparing the relative size of their boundaries, similar to map area and shape comparison. Network path following tasks require the user to trace links (AKA ‘edges’) between nodes in the network while maps also require understanding how regions connect, and in both case splits or distortion due to spherical projection may present a challenge.

As we discuss in detail in Section 2, while there has been much work to develop the algorithms for spherical projection of maps and algorithms for layout of networks on a 3D sphere, we find that interactive 2D displays of such spherical layouts of data (whether geographic or network) are not well studied. Also projections of 3D spherical layouts of network data have not been compared to projections of networks arranged on the surface of other 3D geometries, in particular toroidal layouts which have recently been studied by Chen et al. [12, 13].

Therefore, **this paper has two aims that are interlinked:**

- (1) To evaluate the effect of interaction on readability of different spherical projection techniques and identify the projection techniques which best support geographic comprehension tasks, such as distance, area and direction estimation (Study 1, Section 4).
- (2) To evaluate readability of networks laid out spherically and then projected to a flat surface using the interactive techniques found most effective in Study 1, compared to conventional flat layout, and interactive projections of toroidal layout (Study 2, Section 6).

To achieve these aims many gaps in past work had to be addressed leading to seven distinct contributions:

- (1) To our knowledge, we are the first to systematically evaluate the effect of introducing interactive panning on different geographic spherical projection techniques. We find that interaction overwhelmingly improves accuracy and subjective user preference compared to static projections across all projection methods and tasks considered, at the cost of increased time due to interaction (Section 4).
- (2) For interactive projections, we find that the best of those tested depends on task, however, ORTHOGRAPHIC HEMISPHERE  and EQUAL EARTH  had advantages in terms of speed, accuracy and qualitative feedback while EQUIRECTANGULAR  may be a poor choice even with the ability to pan (Section 4).
- (3) We are also the first to compare spherical network projections against toroidal and flat layouts (Section 6).
- (4) We adapt a pairwise gradient descent algorithm for flat and toroidal layout [12] to consider spherical distance between two nodes when laying them out on a 3D spherical surface. Cartographic projection to a 2D diagram results in 2D drawings whose links wrap around the edge of the display (Section 5).
- (5) We present algorithms for computing how best to automatically rotate the spherical layout to minimise the number of links wrapped when projected (Section 5.2).
- (6) For cluster identification tasks we find that all toroidal and spherical projections tested outperform traditional flat network layout for accuracy, while toroidal and spherical equal earth outperform orthographic projection for accuracy, speed and subjective user rank.
- (7) For path following tasks we find that toroidal and traditional flat network layout outperform both spherical equal earth and

orthographic hemisphere projections for accuracy and subjective user rank (Section 6).

Our results suggest that **interactive panning should be routinely provided in online maps to alleviate misconceptions arising from distortions of map projection**. Our results also confirm the benefits of topologically closed surfaces, such as the surface of a torus or sphere, when using node-link diagrams to investigate network structure. This finding suggests that **interactive projections of networks arranged on 3D surfaces should be more commonly used for cluster analysis tasks** and further, that **toroidal layouts may be a good general solution**, being not only more accurate than standard flat layout for cluster tasks but also at least as accurate for path following tasks.






Collectively, we refer to the family of techniques for visualising geographic and network structured data on surfaces that wrap around by the acronym “GAN'SDA Wrap”, in a rhythmic head nod to Tupac Shakur. The full set of map and network study including training materials, instructions, study trials are available in the supplementary materials and associated OSF repository: <https://osf.io/73p8w/>.




2 BACKGROUND




Researchers have recently investigated the utility of visualisations based on projecting the surface of a torus [12, 13] and a cylinder [11] onto a flat 2D plane. These have shown benefits for understanding the structure of network diagrams [12] and understanding cyclic time series data [11]. Crucially, these visualisations are interactive and the viewer can pan the projection so that it “wraps” around the plane. Providing interactive panning has previously been found to be of benefit in understanding network layouts based on torus projection [13]. Here we investigate the utility of projections of the surface of a different geometric object, the sphere, on to a 2D plane. Again we allow the user to interactively pan the projection around the plane.

2.1 Geographical map projections

Projecting the surface of a sphere onto a 2D plane is not novel: cartographers have been doing this for centuries. Cartographers and mathematicians have devised hundreds of projection techniques [57]. The reason for this diversity is that none of these techniques can be considered optimal in depicting geographic information [55]. Rather, each projection is a trade-off between preserving shape, area, distance or direction [57] as it is not possible for a 2D projection to do all of these simultaneously.

As a consequence, cartographers have invented what are called *equal area projections* that preserve the relative area of regions on the globe. These include EQUAL EARTH , HAMMER  and MOLLWEIDE HEMISPHERE  (see Figure 1) [33, 53]. They have also invented *compromise projections* that do not preserve any of these criteria exactly but instead trade them off, creating a map that does not distort area, shape, distance or direction “too much.” These include ORTHOGRAPHIC HEMISPHERE  and EQUIRECTANGULAR  (see Figure 1).

Projections also differ on the shape of the projection. Some, such as EQUIRECTANGULAR  are rectangular, others, such as EQUAL EARTH  or HAMMER , reduce distortion at the poles

by projecting to a more ovoid shape. Some, such as the MOLLWEIDE HEMISPHERE  and the ORTHOGRAPHIC HEMISPHERE , resemble the front and back view of the 3D globe. Maps such as these in which the projection region is split are said to be interrupted. The ORTHOGRAPHIC HEMISPHERE , in particular, has a naturalistic appearance as it shows the Earth viewed from infinity [33].

There are several user studies on the readability and user preference of map projection visualisations, e.g. [10, 26, 28, 52]. In particular, these have found that viewers find it difficult to understand the distance or direction between two points if this requires reasoning about the discontinuity in the projection and mentally wrapping the projection around a globe to understand their relative position.

This suggests that allowing the viewer to interactively pan the map projection so as to centre a region of interest may improve their understanding of the inherent distortion introduced by the projection and of the Earth's underlying geography. For instance, this allows them to reposition two points so that they are no longer separated by a discontinuity. Such interactive panning, also called spherical rotation [56], has been provided in many online maps for several years, e.g. [6, 17, 32]. One recent study of pannable terrain maps found that they perform more accurately than static maps but at the cost of additional time in task completion and concluded therefore that results of existing static map reading studies are likely not transferrable to interactive maps [27]. However, surprisingly, as far as we are aware there has been no evaluation of whether interactive panning of globe projections improves performance on standard geographical tasks such as estimating the distance or direction between two points or the relative area of two regions.

The only direct research of pannable globes that we know of are two studies in virtual reality (VR) investigating different visualisations for understanding origin-destination flow between locations on the Earth's surface [64]. While it was not their main focus, the studies revealed that interactive panning improved task accuracy at the cost of time when viewing flow shown using straight lines on a flat map. However, it is likely that this was not because panning was used to reduce geographical distortion but rather that it was used to separate the flow lines which were the focus of the tasks. Furthermore, the static and interactive conditions were across different studies so comparison was between groups. Here we present a more systematic and direct study of interactive panning for geographic tasks.

2.2 Spherical representation of non-geographic data

Researchers have also explored non-planar geometries for visualising networks, with evidence that the third dimension can be used improve readability by removing link crossings [24, 59] at the cost of known issues of three-dimensional representation (e.g. occlusion, readability, perspective distortion). Visualisation researchers have also explored the benefits of laying out node-link diagrams on to the surface of a sphere. The potential benefit is that, just as for the torus, they are topologically closed surfaces: there is no centre or border to the surface and so it may allow the layout to better unravel the network and show its structure [8, 47, 49]. Such spherical network layouts are most commonly viewed in immersive VR

(e.g. [36]) or as perspective projection of a globe on a standard 2D monitor (e.g. [8, 34]). It is much less common for these to be projected onto a 2D plane using a map projection (though some static map projections of graphs were demonstrated by [49]). However, map projections have the great advantage over simple perspective rendering of a 3D globe, that the whole network can be seen at once (with Orthographic Hemisphere projection being the closest to a direct rendering of the 3D globe, but with most sides shown simultaneously).



To the best of our knowledge, while it is common to allow the viewer to interactively rotate a globe when shown in 3D we do not know of previous use of interactive panning of spherical network diagrams when they have been projected onto a 2D plane using a map projection. However, we would expect similar benefits to providing interactive panning for projections based on the torus [13].

Similar to node-link embeddings, spheres have been used to embed self-organising maps (SOMs) [62], shown on both 3D representations, and projected onto a 2D-plane. Beyond node-link diagrams and SOMs, a range of other information visualisations can potentially be projected onto a sphere to minimise artefacts that occur when trying to find an optimal embedding for a 2D plane: multi-dimensional scaling (MDS) and their extension, time-curves [1].




Rodighiero [49] preliminarily visualised networks on a variety of spherical projections, but he did not consider spherical projection with interaction, which is essential for the users to inspect different parts of the networks. Manning implemented an interactive force-directed spherical layout algorithm for a variety of projections [39]. However, this approach used Euclidean distance between the points rather than the great circle distance on the surface of the sphere, and therefore does not take full advantage of spherical layout [47]. Kobourov and Wampler give a generalisation of force-directed (based on spring-embedding) network layout to non-Euclidean topologies, including the surface of a sphere [35] while Perry et al. [47] give an algorithm based on MDS using the smacofSphere R package [18]. Most importantly, there is no user study that evaluates these proposed spherical graph projection techniques, and therefore there is no empirical evidence for their effectiveness.




3 MAP PROJECTION TECHNIQUES


We chose five representative map projections for the study. We aim to cover a wide range of distortion properties such as preservation of area, distance, shape, direction as discussed in Section 2, and user preference [52]. We demonstrate the key characteristics of the map projections in Table 1 and discuss the details as follows:



EQUIRECTANGULAR  projects the earth onto a space-filling rectangle with the north and south poles extending along the top and bottom edges, respectively. Rectangular projection is the most widely used projection, with common variations including Mercator or Plate Carrée [33, 52]. **EQUIRECTANGULAR**  projections preserve distances along all meridians and are useful when differences in latitude are measured [33]. However, it does not preserve the relative size of areas. Furthermore, it has been found confusing for tasks requiring understanding how the edges connect to each

other, such as predicting the path of air plane routes crossing the top and bottom edge of the map [26].

EQUAL EARTH  is similar to **EQUIRECTANGULAR**  but relaxes the rectangle with rounded corners, diminishing the strength of horizontal distortion near the poles. Furthermore, the similarly shaped Robinson projection has received good subjective ratings from viewers [52]. Unlike Robinson projection, **EQUAL EARTH**  preserves the relative size of the areas well.


HAMMER  further diminishes horizontal stretching by projecting the earth onto an ellipse such that the poles are points at the top and bottom. It preserves the relative size of areas and generally has similar properties to **EQUAL EARTH**  [33, 65]. The similarly shaped Mollweide projection has been shown less confusing than **EQUIRECTANGULAR**  when judging the continuity of air plane routes as described above [26]. Furthermore, map readers prefer to see the poles as points rather than lines [52]. It has also been found pleasing to many cartographers than other projections due to its elliptical shape [33].

MOLLWEIDE HEMISPHERE  projects the earth onto two circles (hemispheres). This again helps diminish horizontal stretching but introduces the cost of new tears (interruption) that a viewer needs to mentally close.

ORTHOGRAPHIC HEMISPHERE  is also hemispheric. It has a naturalistic appearance as it shows the globe (from both sides) seen from an infinite distance. However, compared with **MOLLWEIDE HEMISPHERE** , it is not area-preserving.

For each projection, we created both **STATIC** and **INTERACTIVE** versions. With **INTERACTIVE**, a user can freely move regions of interest to the centre of the projection, thus reducing their distortion. Examples of the effect of interaction are shown in the top row in Figure 1.

We used D3 libraries for creating all the map projection techniques. For **INTERACTIVE**, we follow Yang et al. [65] and use versor dragging which controls three Euler angles. This allows the geographic start point of the gesture to follow the mouse cursor [17].






We implemented **ORTHOGRAPHIC HEMISPHERE**  with two Orthographic map projections with one showing the western and the other showing the eastern hemisphere. They are placed close together as shown in the top row of Figure 1. When one hemisphere is dragged, the Orthographic projection of the other hemisphere automatically adjusts three-axis rotation angles such that it shows the correct opposite hemisphere.

Spherical rotation of all projections is demonstrated in the supplementary material video and the OSF repository¹

4 STUDY 1: MAP PROJECTION READABILITY

The goal of our first study was to understand the effectiveness of (i) different projections for geographic data as well as the (ii) benefit of interactively changing the centre point of these projections.

4.1 Techniques

The techniques are **EQUIRECTANGULAR** , **EQUAL EARTH** , **HAMMER** , **MOLLWEIDE HEMISPHERE** , and **ORTHOGRAPHIC HEMISPHERE** , described in Section 3. Each technique is given both static images without the ability to rotate, and with interactive

¹Interactive examples can be found in <https://observablehq.com/@kun-ting>.







Projection	Type	Area-preserving	Left/right edges	Top/bottom edges	Naturalistic
EQUIRECTANGULAR 	continuous	no	straight	straight	very low
EQUAL EARTH 	continuous	yes	curved	straight	low
HAMMER 	continuous	yes	curved	curved	low
MOLLWEIDE HEMISPHERE 	interrupted	yes	circular	circular	high
ORTHOGRAPHIC HEMISPHERE 	interrupted	no	circular	circular	very high


Table 1: Key characteristics of the five map projections tested in Study 1.


spherical rotation, using the mouse. A user can rotate the visualisation such that when the view is panned off one side of the display, it either reappears on the opposite side (left-right) or is horizontally mirrored (top-top, bottom-bottom), as shown in the top-row of Figure 1. The area of the rectangular bounding box of each map projection condition was fixed at 700×350 pixels.

4.2 Tasks & Datasets

We selected three representative geographic data visualisation tasks. They were also used in existing map projection studies [10, 26, 28, 65].

DISTANCE COMPARISON : Which pair of points (pair A or pair B) represents the greater geographical distance on the surface of a globe? Participants had to compare the true geographical distance (as the crow flies) between two pairs of points on the projection. Participants were provided with radio buttons to answer A, B, or not sure. We created data sets for two difficulties through extensive piloting, based on the difference of point distance across point pairs: 10% difference (easy) and 5% difference (difficult). The geographic point pairs were randomly chosen, constrained by their individual angular distance (measured in differences in geographical coordinates) between 40° (approx. 4444km) and 60° (approx. 6666km) [65], and a minimum 60° distance across point pairs. We did not set any upper-bound to not to bias any projection. We created an additional quality control trial with 40% difference of node distance to test a participant's attention. An example is shown in Figure 4(a,c).

AREA COMPARISON : Which polygon (A or B) covers the greater geographical area on the surface of a globe? Participants had to compare the size (area) of the polygons. Participants were provided with radio buttons to answer A, B, or not sure. Again, we created data sets for two difficulties based on the difference in area they cover: 10% difference (easy) and 7% difference (difficult). Eight geographic points of convex polygons were randomly chosen using the same method as Yang et al. [65], constrained by the individual geographic area between 40 and 60, and a minimum 60° angular distance between centroids of pairwise polygons. There is no upper-bound for the same reason as above. We create an additional quality control trial with 40% difference in area to test a participant's attention. An example is shown in Figure 4(b,d).

DIRECTION ESTIMATION : Does the trajectory of dot A hit or miss dot B on the surface of a globe? Participants had to assess whether the trajectory, indicated by an arrow track, from point A hits or misses point B. Participants were provided with radio buttons to answer Hit, Miss, or not sure. We randomly created pairs of geographic points (A, B, and arrow head) with a minimum angular distance of 60° between A and B. There is no upper-bound. For trials where the

trajectory of A misses B, the angular distance between trajectory and B was constrained to 40° . Examples of this task is shown in the first two rows of Figure 1 for each projection techniques, where A misses B for equal earth and orthographic hemisphere. For this task, there was only one level of difficulty. We create an additional quality control trial with dot A at the centre and arrow track hitting to a dot B aligned horizontally at the centre to test a participant's attention.





4.3 Hypotheses


Our hypotheses were pre-registered with the Open Science Foundation: <https://osf.io/vctfu>.

H1-1: INTERACTIVE has a lower error rate than all static map projections across all tasks. Our intuition is that interaction allows regions of interest to be centred and, thus, their distortion reduced.












H1-2: INTERACTIVE projections have longer task-completion time than static across all projections and tasks. Users will spend time interacting to find an optimal centre point for each projection to solve the task.

H1-3: For interactivity, users prefer INTERACTIVE to STATIC across all tasks. Intuition as above.

H1-4: INTERACTIVE ORTHOGRAPHIC HEMISPHERE  has a lower error rate than other interactive non-hemisphere (EQUIRECTANGULAR , EQUAL EARTH , HAMMER ) projections across all tasks. This is inspired by a virtual globe study by Yang et al. [65]. Our intuition is that an interactive view of the 3D globe will have similar benefits to the VR representation.

H1-5: Users prefer INTERACTIVE ORTHOGRAPHIC HEMISPHERE  over all other INTERACTIVE projections across tasks. Intuition as above.

4.4 Experimental Setup

Design is within-subject per task, where each participant performed one task (DISTANCE COMPARISON , AREA COMPARISON , DIRECTION ESTIMATION ) on all projections (EQUIRECTANGULAR , EQUAL EARTH , HAMMER , MOLLWEIDE HEMISPHERE , ORTHOGRAPHIC HEMISPHERE ) in both STATIC and INTERACTIVE and in all levels of difficulty (easy, hard). Each of these 10 conditions for AREA COMPARISON  and DISTANCE COMPARISON  was tested in 12 trials with two difficulty levels (6 easy, 6 hard). For DIRECTION ESTIMATION , we reduced the number of trials to 8 as pilot participants reported the direction estimation was too difficult for long distance. Similar to existing visualisation crowd-sourced studies [9, 11], we randomly inserted a quality control trial with low difficulty in addition to normal trials to each condition to test participants' attention. The study was blocked by interactivity. Within each block, the order of the map projection techniques was

counterbalanced using William et al.'s Latin-square design [61]. This technique resulted in 10 possible orderings for the 5 projections while the order of projections in each block was the same. Each recorded trial had a timeout of 20sec, inspired from pilot studies.

4.5 Participants and Procedures

We crowd-sourced the study via the *Prolific Academic system* [44]. Participants on Prolific Academic have been reported to produce data quality comparable to Amazon Mechanical Turk [46]. Many visualisation studies have used this platform in the past [11, 51]. We hosted the online study on Red Hat Enterprise Linux (RHEL7) system. We set a pre-screening criterion on performance that required a minimum approval rate of 95% and a minimum number of previous successful submissions of 10. We also limited our study to desktop users with larger screens. We paid £5 (i.e. £7.5/h), considered to be a good payment according to the Prolific Academic guidelines.

We recorded 120 participants who passed the attention check trials, completed the training and recorded trials. This comprised 4 full counterbalanced blocks of participants ($10 \times 4 \times 3$ tasks). 57 of our participants were females, 63 were males. The age of participants was between 18 and 55. 7 participants rank themselves as *regularly* using GIS or other tools to analyse geographical data. 105 occasionally read maps, e.g., using Google Map or GPS systems. 8 had very little experience with any sort of maps.

Before starting the study, each participant had to complete a tutorial explaining projection techniques and tasks. The tutorial material contained *Tissot's indicatrix* [56], a set of circular areas placed on both the poles the equator, indicating the type and magnitude of area, shape, and angular distortion in a given projection. The setting has been inspired by Yang et al. [65]. Specific instructions were given for each task, available in our supplementary material².

4.6 Dependent Variables

We measured **task-completion-time** (TIME) for each trial in milliseconds, counted between the first rendering of the visualisation and the mouse click of the *answer trial* button, which hid the visualisation and showed an interface for the participants to input their answers. We measured the **error rate** (ERROR) as the ratio of incorrect over all answers. We asked participants to **rank** (RANK) each map projection individually within the *static* and the *interactive* block according to their perceived effectiveness. We also ask participants to provide their justifications for the rankings as **qualitative feedback**. After they completed both blocks, we recorded participants' preference of the interactivity between STATIC and INTERACTIVE individually for each projection, and their overall preference between STATIC and INTERACTIVE.

4.7 Statistical Analysis Methods

We used *sqr*t-transformation for TIME to meet the normality assumption. We then used linear mixed modelling to evaluate the effect of independent variables on the dependent variables [3]. We modelled all independent variables (five map projections, two interaction levels and two difficulty levels) and their interactions as

fixed effects. We evaluated the significance of the inclusion of an independent variable or interaction terms using log-likelihood ratio. We then performed Tukey's HSD post-hoc tests for pairwise comparisons using the least square means [38]. We used predicted vs. residual and Q-Q plots to graphically evaluate the homoscedasticity and normality of the Pearson residuals respectively. For ERROR and RANK, as they did not meet the normality assumption, we used the *Friedman* test to evaluate the effect of the independent variable, as well as a Wilcoxon-Nemenyi-McDonald-Thompson test for pairwise comparisons. We also used Wilcoxon signed-rank test for comparing the accuracy of static and interactive map projections. The confidence intervals are 95% for all the statistical testing. We demonstrate the *error rate* and *time* in Figure 2. We show interactivity and map projection rankings as stacked bar charts in Figure 3, Section A.1 (Appendix): Figure 9, and Figure 10.

Following existing work which reports statistical results with standardised effect sizes [31, 43, 66] or simple effect sizes with confidence intervals [4, 5], we interpret the standardised effect size for a parametric test using Cohen's *d* classification, which is 0.2, 0.5, and 0.8 or greater for small, moderate, and large effects, respectively [15]. For non-parametric test, we interpret the standardised effect size for a Wilcoxon's signed-rank test using Cohen's *r* classification, which is 0.1, 0.3, and 0.5 or greater for small, moderate, and large effects, respectively [15, 45].

4.8 Key Findings and Discussion

We report on the most significant findings for DISTANCE COMPARISON ●●●, AREA COMPARISON ◆◆, and DIRECTION ESTIMATION ➔ visually in Figure 2 and Figure 3. Detailed statistic results and pairwise effect sizes are provided in the appendix: Section A.1 in the supplementary materials.

Interaction improved error rate and is preferred over static map projections while taking participants longer time to complete (across all tasks). This main effect was found statistically significant with moderate and large effect sizes in all three tested tasks (Figure 2, Section A.1: Figure 10, Figure 11-13). For DISTANCE COMPARISON ●●● and DIRECTION ESTIMATION ➔, INTERACTIVE is always better (significantly less error and more preferred) than STATIC. For AREA COMPARISON ◆◆, INTERACTIVE has less error than the corresponding non-interactive projection but not necessarily all other non-interactive projections. INTERACTIVE was significantly preferred (moderate effects) for HAMMER ○ and EQUAL EARTH ○ but not necessarily other projections for AREA COMPARISON ◆◆. However, INTERACTIVE was significantly preferred in the overall user rank over STATIC (top horizontal bar in Section A.1: Figure 10) for DISTANCE COMPARISON ●●● (large effects), DIRECTION ESTIMATION ➔ (large effects) and AREA COMPARISON ◆◆ (moderate effects). INTERACTIVE is always slower (large effects) than STATIC (Figure 2-Time). Therefore, we accept H1-1, H1-2, H1-3.

These results provide strong evidence that, with interactions, participants were able to find a better projection centre than the default one in a static map. Some participants explicitly mentioned the benefits of having interaction and their preference, e.g., "when moved [the interactive conditions makes] it easier to judge when [the areas] were both placed in the middle in the least distorted part of the map." (P20, Area-INTERACTIVE). and "The fact I couldn't move the

²An online demonstration of the study is available: <https://observablehq.com/@kunting/gansdawrap>

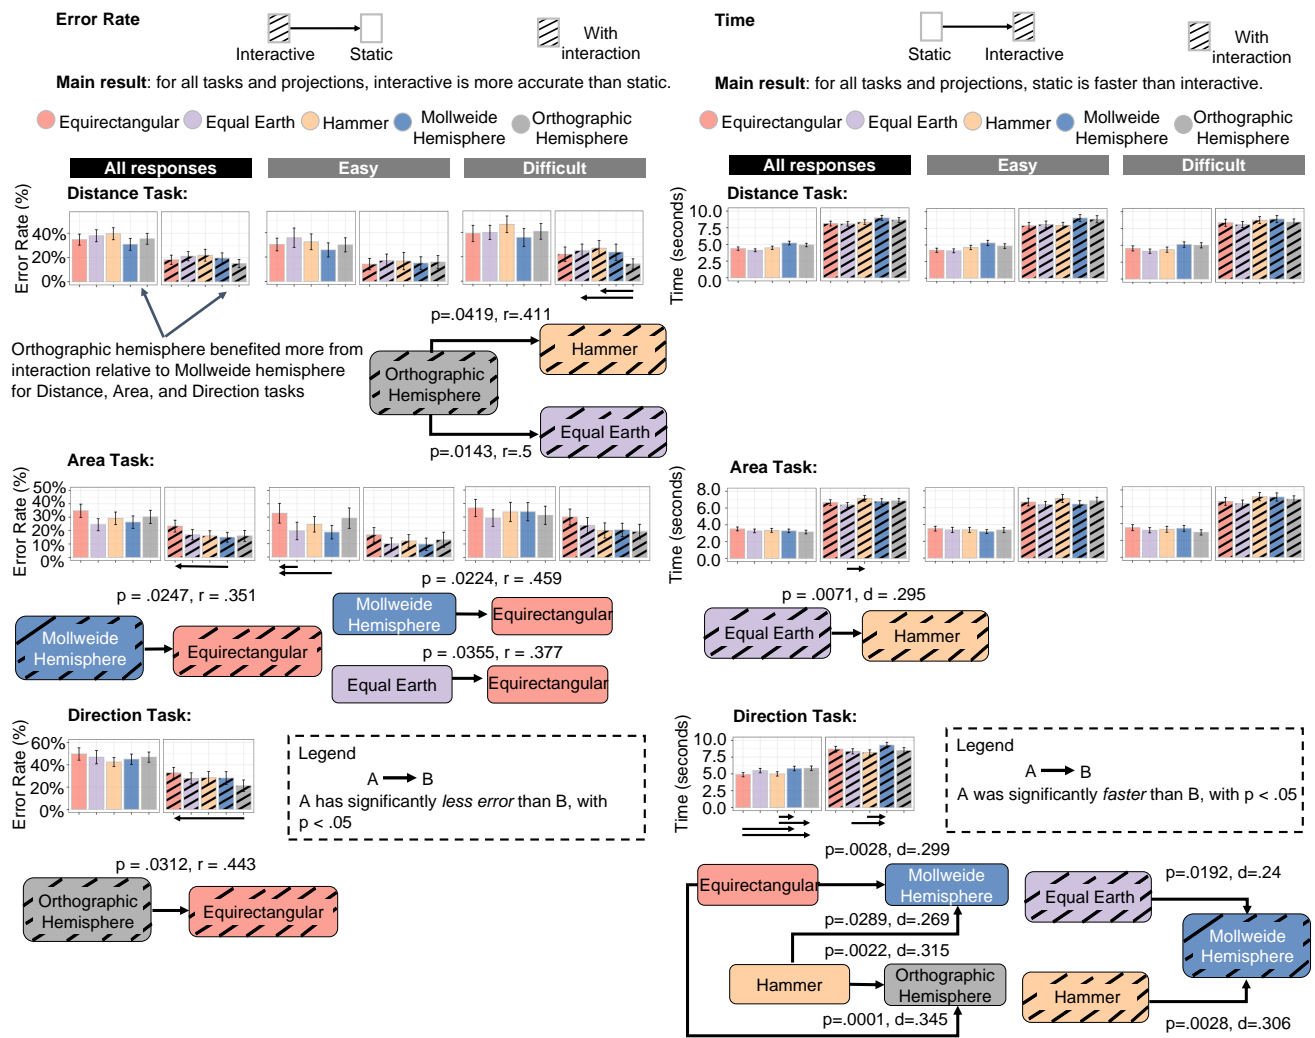


Figure 2: Error rate (left) and Time (right) results. Significant differences between projections are shown as arrows. Significant differences between interactive and non-interactive conditions are omitted to improve readability. Error bars indicate 95% confidence intervals. Effect size results for Cohen’s r and Cohen’s d [15] are presented for ERROR and TIME, respectively. Statistically significant results are highlighted in flow diagrams below the bar charts. Bars and boxes with stripe pattern refer to interactive conditions. Overall, equal earth performed well for TIME and ERROR for some tasks. Orthographic hemisphere benefited more from interaction than mollweide hemisphere for ERROR.

pictures was frustrating [for the static conditions] and I think I didn’t get many of the guesses right.” (P4, Direction-STATIC).

The choice of projections makes less difference and depends on tasks. However, overall, we found that equal earth and orthographi hemisphere performed well, while equirectangular may be a poor choice, organised in the following key findings.

In continuous projections, EQUAL EARTH performed well in terms of ERROR for AREA COMPARISON and TIME for AREA COMPARISON and DIRECTION ESTIMATION, and RANK-STATIC for both DISTANCE COMPARISON and AREA COMPARISON. EQUAL EARTH was not significantly

worse than any other continuous projections. We found that for static projections, EQUAL EARTH tended to be more accurate (moderate effects) than EQUIRECTANGULAR for Area-Easy. With interaction, EQUAL EARTH tended to be faster (small effects) than MOLLWEIDE HEMISPHERE for DIRECTION ESTIMATION, and HAMMER for Area-All (Figure 2-Time). Though the time results are statistically significant, the small effect sizes seem to indicate that the choice of projection makes a slight difference [14, 25, 31, 54]. For RANK-STATIC, EQUAL EARTH was significantly preferred over HAMMER (moderate effect), MOLLWEIDE HEMISPHERE (moderate effect), and ORTHOGRAPHIC HEMISPHERE (large effect) for DISTANCE COMPARISON (Figure 3-Distance). Furthermore,

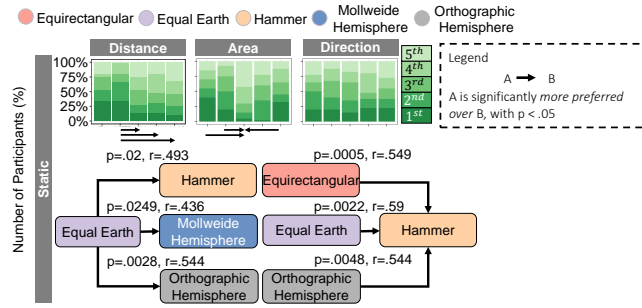


Figure 3: Subjective user rank of map projections within the static group for three tested tasks. Lower rank indicates stronger preference. Arrows indicate statistical significance with $p < 0.05$. Overall, equal earth was found to be preferred over hemispheric projections for the distance comparison task, while hammer was preferred over by other projections for area comparison task.

there is a strong evidence with important effects for RANK-STATIC that EQUIRECTANGULAR \square , EQUAL EARTH \circ , and ORTHOGRAPHIC HEMISPHERE \circ were significantly preferred over HAMMER \circ (large effects) for AREA COMPARISON \blacklozenge (Figure 3-Area). With interaction, there is no significant differences in user preference. This result is omitted from the paper and is available in Section A.1: Figure 9.

User preference of STATIC EQUAL EARTH \circ partially confirms Šavrič et al. [52], who found Robinson projection (which is similar to EQUAL EARTH \circ) was preferred over interrupted projections such as MOLLWEIDE HEMISPHERE \circ and Goode Homolosine. This is also supported by our participants' feedback where continuous maps are preferred over interrupted ones for STATIC (Section A.1.1).

Surprisingly, although HAMMER \circ is an equal-area projection, participants did not like it for area comparison in static maps (Figure 3-Area). This result partially differs from existing studies [52] where poles represented as points were preferred over poles represented as lines. We conjecture this effect is because when the target area is at the edges, the shapes are severely distorted, which makes it difficult to accurately accumulate its represented area, as participants mentioned "Equirectangular only was distorted from top to bottom, while [hammer was] also distorted on the sides." (P31, Area-STATIC). More quotes can be found in Section A.1.1.

In hemispheric projections, interaction reduced ERROR of ORTHOGRAPHIC HEMISPHERE \circ to a point that it tended to have a lower error rate than some interactive continuous projections for Distance-Hard and Direction, and not significantly slower than any interactive projections across all tasks. We found ORTHOGRAPHIC HEMISPHERE \circ performed well. ORTHOGRAPHIC HEMISPHERE \circ benefited more from interaction (large effects) for ERROR than MOLLWEIDE HEMISPHERE \circ (moderate effects) for DISTANCE COMPARISON $\bullet\bullet\bullet$, while ORTHOGRAPHIC HEMISPHERE \circ benefited slightly more from interaction with similar effect sizes for ERROR than MOLLWEIDE HEMISPHERE \circ for AREA COMPARISON \blacklozenge and DIRECTION ESTIMATION \blackrightarrow (Figure 2-ERROR, Section A.1: Figure 11-Figure 13). This is also

supported by a strong evidence that INTERACTIVE ORTHOGRAPHIC HEMISPHERE \circ has a lower error rate (large effects) than EQUAL EARTH \circ and a lower error rate (moderate effects) than HAMMER \circ for Distance-Hard (Figure 2-Error). By contrast, even with interaction, MOLLWEIDE HEMISPHERE \circ was still slower (small effects) than non-hemisphere for DIRECTION ESTIMATION \blackrightarrow (Figure 2-Time). INTERACTIVE ORTHOGRAPHIC HEMISPHERE \circ did not have a significantly lower error rate for AREA COMPARISON \blacklozenge than any interactive continuous projections, nor was there any significant difference in RANK-INTERACTIVE (Figure 9). Therefore, we reject H1-4, H1-5.



It was surprising that ORTHOGRAPHIC HEMISPHERE \circ was comparable to other interactive projections for AREA COMPARISON \blacklozenge since it is not an equal-area map projection. Meanwhile, the other non-equal-area map projection, i.e., EQUIRECTANGULAR \square tended to produce more errors (moderate effects). We believe ORTHOGRAPHIC HEMISPHERE \circ were perceived as less distorted than the other projections due to the "natural" orthographic distortion, which is similar to viewing the sphere at infinite distance (e.g. as if through a telescope), echoed by our participants (Section A.1.1).










Despite being hemispheric, MOLLWEIDE HEMISPHERE \circ was found to be less intuitive for some tasks by participants (Section A.1.1). We conjecture that there is a slight distortion near the edge of two circles which may make it confusing when centring the region of interest. However, there were also participants who did not like the hemispheric projections due to the need to inspect two separated spheres and instead they preferred the continuous maps in the non-hemispheric group (Section A.1.1).







Even with interaction, EQUIRECTANGULAR \square still tended to perform poorly in terms of ERROR for AREA COMPARISON \blacklozenge and DIRECTION ESTIMATION \blackrightarrow . Overall, for static projections, EQUIRECTANGULAR \square tended to have a higher error rate (moderate effects) than EQUAL EARTH \circ and MOLLWEIDE HEMISPHERE \circ (Figure 2-Error). To our surprise, with the ability to rotate to centre the region of interest, EQUIRECTANGULAR \square still tended to be outperformed for ERROR by MOLLWEIDE HEMISPHERE \circ (moderate effect) for Area-All and by ORTHOGRAPHIC HEMISPHERE \circ (moderate effect) for DIRECTION ESTIMATION \blackrightarrow (Figure 2-Error). This partially confirms Hennerdal et al.'s static map study where STATIC EQUIRECTANGULAR \square was found confusing when estimating the airplane route that wraps around [26]. We conjecture that EQUIRECTANGULAR \square features the most distortion of all tested projections due to the high level of stretching near the poles, supported by participants' feedback (Section A.1.1).

4.9 Limitations

The statistically significant results with large effect sizes provide a strong evidence that adding the spherical rotation interaction to static maps improves the accuracy, is strongly preferred, but at the cost of longer completion time across all tasks. However, despite being statistically significant, the differences between projections within static or interactive groups are of small sizes for TIME, medium-sized for ERROR and large-sized for RANK-Static-Area (Figure 2, Figure 3). Although medium-sized differences in ERROR are likely to be noticeable in practical applications, these

results only allow us to say EQUAL EARTH  and ORTHOGRAPHIC HEMISPHERE  performed well for some tasks [14, 25, 31, 54].

Although the results show statistically significant differences between the selected hemispheric and non-hemispheric projections across all tasks for ERROR, the results do not allow us to say that hemispheric projections are always more accurate than non-hemispheric projections. Arguably, the poor performance of non-hemispheric projections for AREA COMPARISON  and DIRECTION ESTIMATION  is entirely based on EQUIRECTANGULAR  compared with either MOLLWEIDE HEMISPHERE  or ORTHOGRAPHIC HEMISPHERE . If the results from the EQUIRECTANGULAR  were not considered, the hemispheric and non-hemispheric projections seem to have very similar error rates for AREA COMPARISON  and DIRECTION ESTIMATION . Similarly, the results do not allow us to say that non-hemispheric projections are always faster, as this seems only based on MOLLWEIDE HEMISPHERE  being significantly slower (small effects) than some non-hemispheric projections for STATIC and INTERACTIVE for the direction tasks.

Surprisingly, unlike the study by Yang et al. [65], we did not identify superior performance of ORTHOGRAPHIC HEMISPHERE  compared to other map projections. We conjectured that rendering ORTHOGRAPHIC HEMISPHERE  on a 2D flat display produced different effectiveness in perception and interaction compared to Yang et al's 3D globes in VR. Meanwhile, although not in all tasks, ORTHOGRAPHIC HEMISPHERE  demonstrated some advantages for DISTANCE COMPARISON , DIRECTION ESTIMATION  and benefited more from interaction than MOLLWEIDE HEMISPHERE  for ERROR (Section 4.8).

5 SPHERICAL NETWORK LAYOUT

In Study 1, we found that interaction (panning by spherical rotation) makes spherical geographic projections overwhelmingly more accurate for distance, area and direction tasks. A question, however, is whether such 2D interactive spherical projections are also useful for abstract data. As discussed in Section 2, there has been a number of systems using immersive environments to visualise network data on 3D spherical surfaces or straightforward perspective projections of spheres. Various advantages have been claimed to the opportunities for embedding a network layout in the surface of a sphere—without boundary—such as centring any node of interest in the layout [8, 47, 49], improving readability by reducing link crossings using the third dimension [59], and stereoscopy outperforms standard 2D layouts for highly overlapping clustered graphs [24]. However, the visualisation design literature cautions against such use of 3D if there are layout approaches better suited to the plane [41, Ch. 6].

Further, there are obvious disadvantages to projection, since all projections introduce some degree of distortion and discontinuity. There are therefore three questions:

RQ1: Which of the most promising projections from our first study are the best for visualising the layout of a node-link diagram?

RQ2: Does a spherical projection have advantages in supporting network understanding tasks compared to conventional 2D layouts?

RQ3: Does a spherical projection provide perceptual benefits compared with arrangements on other 3D topologies, such as a torus?

Before we can answer these questions, we need techniques to create effective layouts of complex network data on a spherical surface and to orient the projections optimally in 2D.

5.1 Plane, Spherical and Toroidal Stress Minimisation

The tasks we investigate are cluster understanding tasks and path following. Network clusters are loosely defined as subsets of nodes within the graph that are more highly connected to each other than would be expected from chance connection of edges within a graph of that density. More formally, a clustered graph has disjoint sets of nodes with positive *modularity*, a metric due to Newman which directly measures the connectivity of given clusters compared to overall connectivity [42]. To support cluster understanding tasks we need a layout method which provides good separation between these clusters.

To support path following tasks, we need a layout method which spreads the network out relatively uniformly according to connectivity. This will help minimise crossings between edges.

We follow recent work [12, 13, 67] which adopted a stress minimising approach. Stress-minimisation is a commonly used variant of a general-purpose force-directed layout and does a reasonable job of satisfying both of these readability criteria [30, 48]. The *stress* metric (σ) for a given layout of a graph with n nodes in a 2D plane is defined (following Gansner et al. [23]) as:

$$\sigma_{\text{plane}} = \sum_{i=1}^{n-1} \sum_{j=i+1}^n w_{ij} (\delta_{ij} - d_{ij})^2, \quad d_{ij} = |x_i - x_j|$$

where: δ_{ij} is the ideal separation between the 2D positions (x_i and x_j) of a pair of nodes (i, j) taken as the all-pairs shortest path length between them; d_{ij} is the actual distance between nodes i and j (in the plane this is Euclidean distance); and w_{ij} is a weighting which is applied to trade-off between the importance of short and long ideal distances, we follow the standard choice of $w_{ij} = 1/d_{ij}^2$.

We follow previous recent work by Perry et al. [47], in adapting stress-based graph layout to a spherical surface by redefining d_{ij} to arc-length on the sphere surface, or (assuming a unit sphere):

$$\sigma_{\text{sphere}} = \sum_{i=1}^{n-1} \sum_{j=i+1}^n w_{ij} (\delta_{ij} - d_{ij})^2, \quad d_{ij} = \arccos(x_i \cdot x_j)$$

where x_i and x_j are the 3D vector offsets of nodes i and j respectively from the sphere centroid and (\cdot) is the inner product. For the layout to be reasonable, the ideal lengths δ must be chosen such that the largest corresponds to the largest separation possible on the unit sphere surface, which is π . Thus, we set the ideal length of all edges to $\pi/\text{graphdiameter}$.

The other layout against which we compare is a projection of a 3D torus, which, as discussed in Section 2 has recently been shown to provide better separation between clusters than a flat (conventional) 2D layout by Chen et al. [12]. We use the same layout method which is also based on stress in the 2D plane but which, for each node pair, requires selecting the stress term from the set A of 9 possible torus adjacencies for that pair which contributes the least to the overall stress:

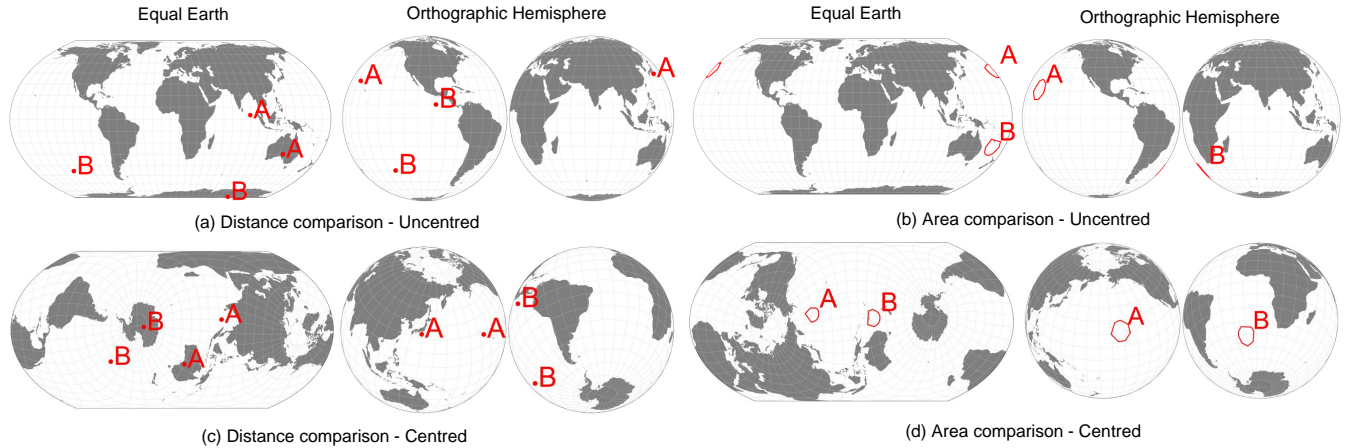




Figure 4: Sample trials of Equal Earth and Orthographic Hemisphere for distance and area comparison tasks. STATIC (a,b) may result in uncentred view. INTERACTIVE (c,d) allows a user to drag the map to find the best angle (centred) to answer the task. Direction tasks are shown in Figure 1.



$$\sigma_{\text{torus}} = \sum_{i=1}^{n-1} \sum_{j=i+1}^n w_{ij} \arg \min_{\alpha \in A} (\delta_{ij} - d_{ij\alpha})^2, d_{ij\alpha} = |x_i - x_{j\alpha}|$$



While Perry et al. follow the multi-dimensional scaling literature in using a *majorization* method to minimise σ_{sphere} , we follow Chen et al. [12] in using the stochastic pairwise gradient descent approach developed by [67] which can be adapted straightforwardly and effectively to obtain layout for all of σ_{plane} , σ_{sphere} , σ_{torus} .

Note that the layouts which result from minimising σ_{plane} and σ_{torus} are already in 2D. For the spherical layout obtained by minimising σ_{sphere} , we use either ORTHOGRAPHIC HEMISPHERE  or EQUAL EARTH  projections to generate the stimuli for our study. The detailed pseudocode of our algorithms are available in <https://github.com/Kun-Ting/gansdawrap>.





5.2 Auto-pan Algorithms

For toroidal network layout, Chen et al. introduced an algorithm to automatically pan the toroidal layout horizontally and vertically to minimise the number of edges which wrap around at the boundaries [12]. Spherical projections can also suffer when too many edges are split across the boundaries. Furthermore, edges are more distorted in spherical projections when they are near the edges. Therefore, for fair comparison with toroidal layouts, it was necessary to find a method to auto-rotate the sphere to reduce the numbers of such edges. However, while the toroidal auto-panning algorithm can be done with horizontal and vertical scans (linear time in the number of edges), the spherical layout does not permit such a trivial search algorithm. We therefore develop heuristics to perform auto-rotate for the spherical projections. For both, we choose a simple stochastic method of randomly selecting a large number (e.g., 1000 iterations) of three-axis spherical rotation angle triples (λ, ϕ, γ) and choosing the triple for which edges crossing (or near) boundaries is minimised.

For ORTHOGRAPHIC HEMISPHERE  projection this crossing number is trivial to count precisely. Simply, for all pairs of nodes if the nodes are not on the same face, then they must cross a boundary. A suboptimal ORTHOGRAPHIC HEMISPHERE  projection rotation, and the result of autorotation to minimise this crossing count is shown in Figure 5(a) and Figure 5(b), respectively.

For EQUAL EARTH  projection the edges are curved and so determining those that cross the boundary for a given geo-rotation is more complex. Further, in this projection edges near the periphery are significantly more distorted than those near the centre, so even visually determining if an edge that comes close to the boundary continues on the same side of the projection or wraps around to the other side is not easy. Thus, instead of counting boundary crossings, we penalise all edges which come close to the boundaries. To compute the penalty we analyse the periphery of a monochrome bitmap of the projected edge paths. The penalty is then simply the number of black pixels. Example masked bitmaps are shown for sub-optimal and more-optimal rotations of an EQUAL EARTH  projection in Figure 5(c) and Figure 5(d), respectively.

5.3 Automatic Panning Results

We conducted a small empirical analysis of the Auto-Pan algorithm to assess the numbers of links wrapped for ORTHOGRAPHIC HEMISPHERE  and the number of pixels (higher values indicate less wrappings) at the boundary of EQUAL EARTH  projections compared to 10 random rotations. Across 10 study graphs (SMALL+EASY, SMALL+HARD) used in our cluster understanding tasks, we found for ORTHOGRAPHIC HEMISPHERE , the mean crossing count for random rotations was 262.16. With automatic panning, this number was improved by 25.6% and was reduced to 208.7. For EQUAL EARTH , our automatic panning increased number of pixels at the boundary region by 12.1% (Section A.2: Table 2). We found these auto-pan algorithms resulted in a noticeable improvement in keeping clusters from being separated, as evidenced in Figure 5.

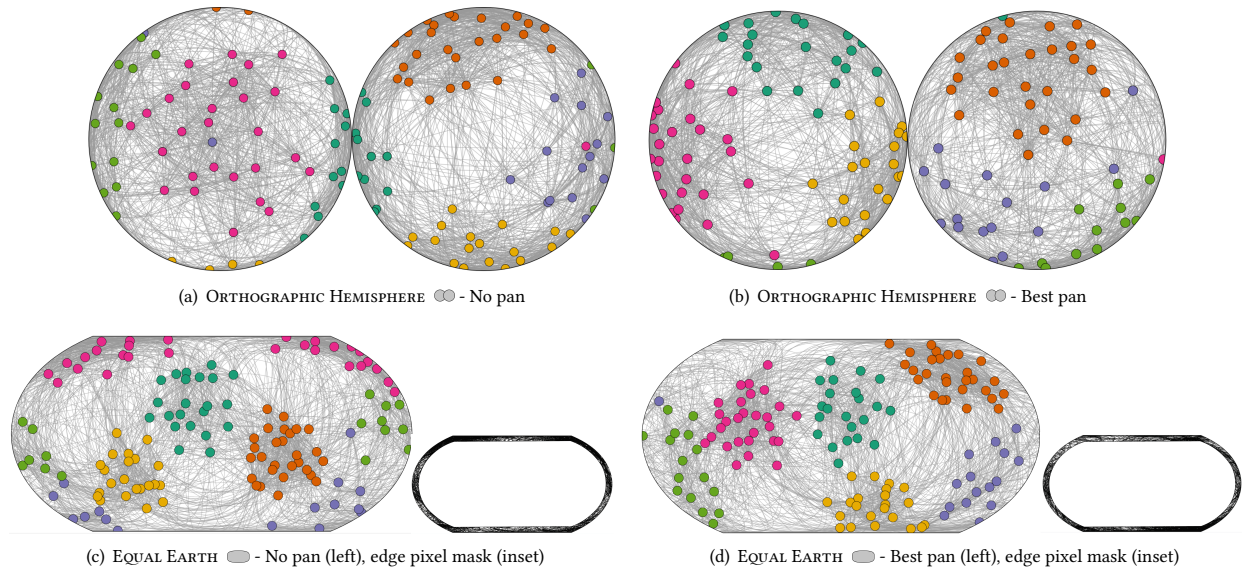


Figure 5: Before and after demonstration of our auto-pan algorithms of a graph with six clusters differentiated by colour (note: study graphs did not have colour). Without auto-pan, clusters can be split across the hemispheres in ORTHOGRAPHIC HEMISPHERE (a) or at the boundaries in EQUAL EARTH (c). Auto-pan reduces the number of wrapped edges and thereby brings the clusters together (b) and (d).

6 STUDY 2: NETWORK PROJECTION READABILITY

In Study 1, we found EQUAL EARTH (c), had advantages in terms of error rate, time, and subjective user feedback. ORTHOGRAPHIC HEMISPHERE (a), despite being an interrupted projection, benefited more from interaction than MOLLWEIDE HEMISPHERE (b) and performed well in terms of error rate for distance and direction tasks (Section 4.8). Based on these findings, we chose interactive EQUAL EARTH (c) and ORTHOGRAPHIC HEMISPHERE (a), to understand their performance in visualising networks. We compare these spherical projections to standard flat graph layout (FLAT (d)) and a projection of a TORUS (e) geometry [12]. Our study was run in a similar way to Study 1 and included a new set of 96 participants through the Prolific platform.

6.1 Techniques & Setup

The techniques in our study are FLAT (d), TORUS (e), EQUAL EARTH (c), and ORTHOGRAPHIC HEMISPHERE (a). Layouts are computed as described in Section 5. All the techniques support interactive panning except for FLAT (d) which does not wrap around. The area of the rectangular bounding box of each technique condition is the same. For FLAT (d) and TORUS (e), the resolution is fixed at 650×650 pixels, and for EQUAL EARTH (c) and ORTHOGRAPHIC HEMISPHERE (a), the resolution is fixed at 900×317 pixels. For CLUSTER NUMBER (f), we did not color clusters to not reveal any graph structure.

6.2 Tasks

We selected two representative network visualisation tasks, inspired by existing work comparing 2D layout of 3D surface topology [12, 13] and task taxonomies [37, 50].

CLUSTER NUMBER (f): Please count the number of clusters in this graph. Participants answered through radio buttons: choices ranged from 1 to 10. We created a quality control trial for each condition with two clusters with clearly marked boundaries to assess participants' attention. Participants who did not answer these trials correctly were excluded. An example of 5 clusters from LARGE+EASY graphs (Section 6.3) using 4 different layouts is shown in Figure 1-Bottom.

SHORTEST PATH NUMBER (g): What is the shortest path length between the red nodes?: Participants were required to count the smallest number of links between two red nodes. Radio buttons allowed them to answer between 1 to 6. Again, we created a quality control trial for each layout condition, with shortest path length two and links on the path highlighted in red. An example of EASY graphs (Section 6.3) using 4 different layouts is shown in Figure 6.

6.3 Data Sets

We prepared a separate graph corpus for each task, full details and stimuli are presented in our supplementary material. For CLUSTER NUMBER (f), we use graphs from Chen et al. [12], generated using algorithms designed to simulate real-world community structures in graphs [7, 21]. Graphs are grouped by two variables: difficulty in terms of graph modularity [42] (EASY: modularity=0.4, HARD: modularity=0.3) and size (2 levels: SMALL: 68-80 nodes, 710-925 links, and LARGE: 126-134 nodes, 2310-2590 links). The number of clusters is between 4 and 7. For SHORTEST PATH NUMBER (g) the clustered graphs were too dense, so we generated sparser graphs using scale-free models [2, 60]. We chose graphs with two levels of density (EASY: 0.075, and HARD: 0.11) with 50 to 57 nodes. The shortest path length varied between 1 and 4.




Figure 6: Example of one study graph laid out using 4 techniques described in Section 5.1 for SHORTEST PATH NUMBER \curvearrowright . The shortest path length is 3 in this example. Participants were provided with interactive panning for spherical and toroidal layouts to explore the network.

6.4 Hypotheses

Hypotheses were pre-registered with the Open Science Foundation: <https://osf.io/equhp>.

H2-1: EQUAL EARTH \ominus and TORUS $\omin�$ have better task effectiveness for CLUSTER NUMBER \clubsuit (in terms of time and error) than ORTHOGRAPHIC HEMISPHERE $\omin�$ (RQ1, RQ3) or FLAT \curvearrowright (RQ2). While ORTHOGRAPHIC HEMISPHERE $\omin�$ performed well in Study 1, our pilot studies for CLUSTER NUMBER \clubsuit revealed that cuts and distortion of clusters at the borders made them hard to count. For (RQ2), the inspiration was based on prior cluster readability studies [12].

H2-2: Participants will prefer EQUAL EARTH \ominus and TORUS $\omin�$ to ORTHOGRAPHIC HEMISPHERE $\omin�$ (RQ1, RQ3) or FLAT \curvearrowright (RQ2) for CLUSTER NUMBER \clubsuit . Our early pilots indicating this preference — perhaps for the same reasons as above and inspiration from existing studies [12].

H2-3: TORUS $\omin�$ has better task effectiveness (in terms of time and error) for SHORTEST PATH NUMBER \curvearrowright than FLAT \curvearrowright , EQUAL EARTH \ominus or ORTHOGRAPHIC HEMISPHERE $\omin�$ (RQ3). This assumption is based on pilot studies and prior studies indicating curved links might hamper path tracing tasks [19, 63].

H2-4: Participants will prefer TORUS $\omin�$ to FLAT \curvearrowright , EQUAL EARTH \ominus or ORTHOGRAPHIC HEMISPHERE $\omin�$ (RQ3) for SHORTEST PATH NUMBER \curvearrowright . Again, it was assumed due to distortion of links.

6.5 Experimental Design

We use a within-subjects design for each task with 4 techniques. Each participant was randomly assigned one of the tasks by the experimental software. For CLUSTER NUMBER \clubsuit , we used 2 levels of difficulty (Easy, Hard) \times 2 sizes (Small, Large) \times 5 repetitions. We randomly inserted one additional quality control trial to each layout condition. Each recorded trial had a timeout of 20 seconds to prevent participants from trying to perform precise link counting. This leaves us with a total of 80 recorded trials per participant. We counterbalanced the order of the techniques using a full-factorial design. The order of each level of difficulty and size in each technique was the same: SMALL+EASY \rightarrow LARGE+EASY \rightarrow SMALL+HARD \rightarrow LARGE+HARD. The order of trials for each technique within each level was randomised.

For SHORTEST PATH NUMBER \curvearrowright , we used 2 levels of difficulty (Easy, Hard) \times 8 repetitions. There were 2 repetitions of each shortest path length per level. One additional quality control trial was

added for each layout condition. This leaves us with a total of 64 trials per participant. Each recorded trial had a timeout of 30 seconds, informed by pilot studies. The order of each level of difficulty in each technique was the same: EASY \rightarrow HARD. The order of trials for each technique within each level was randomised.

6.6 Participants and Procedures

Setup and inclusion criteria for participants were the same as for Study 1. We recorded 96 participants (37 female, 59 male, age range [18,50]) who passed the attention check trials, completed the training and recorded trials. This comprised 2 fully counterbalanced blocks of participants (24 \times 2 \times 2 tasks). They ranked their familiarity with network diagrams as: 9 often seeing network diagrams; 62 occasionally; and 25 never.

For CLUSTER NUMBER \clubsuit , each participant was presented with a tutorial explaining the concept of network clusters. For TORUS $\omin�$, EQUAL EARTH \ominus , and ORTHOGRAPHIC HEMISPHERE $\omin�$, animated videos were given to demonstrate the interactive panning or rotation. Each participant then completed training trials, similar to the recorded trials. For SHORTEST PATH NUMBER \curvearrowright , each participant was presented with a tutorial explaining the concept. For TORUS $\omin�$, EQUAL EARTH \ominus , ORTHOGRAPHIC HEMISPHERE $\omin�$, animated videos were given to demonstrate the interactive panning or rotation. Training trials were similar to the recorded trials.

Specific instructions given for each task is available in the supplementary material³.

6.7 Dependent Variable and Statistical Analysis Methods

We recorded task-completion time (TIME), task-error (ERROR), and subjective *confidence* as RANK (the smaller the more confident). We calculated ERROR as the normalised absolute difference between the correct answer and response. We used the same statistical analysis methods and standardised effect sizes from the first study (Section 4.7).

³An online demonstration of the study is available: <https://observablehq.com/@kunting/gansdawrap>

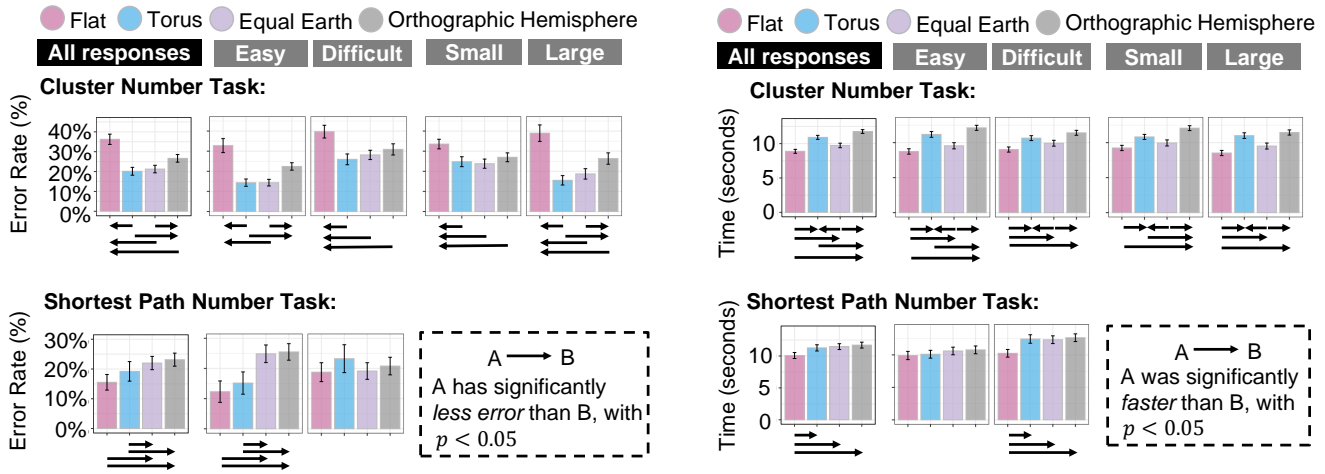


Figure 7: Error rate (left) and Time (right) results of Study 2. Error bars are 95% confidence intervals. Significant differences between projections are shown as arrows. Detailed statistical results and effect sizes are available in Section A.2.

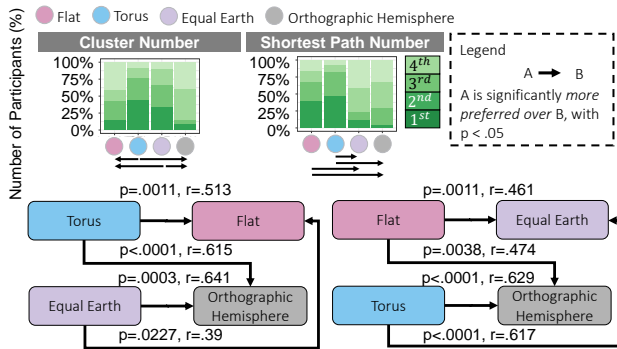


Figure 8: User confidence ranking of each condition for two tested tasks. Arrows indicate statistical significance with $p < 0.05$. Effect size results of Cohen's r [15] are presented along with the arrows.

6.8 Key Findings and Discussion

We report on the most significant findings for CLUSTER NUMBER ♣, SHORTEST PATH NUMBER ♠ visually in Figure 7 and Figure 8. Detailed statistical results and pairwise effect sizes are provided in Section A.2.


For CLUSTER NUMBER ♣, EQUAL EARTH ○ and TORUS ☉ performed equally well and EQUAL EARTH ○ was slightly faster than TORUS ☉, while they both significantly outperformed ORTHOGRAPHIC HEMISPHERE ☉ in terms of ERROR, TIME, and RANK. The poor performance of ORTHOGRAPHIC HEMISPHERE ☉ compared with EQUAL EARTH ○ and TORUS ☉ was found statistically significant with moderate and large effect sizes for ERROR (All, Easy, Large), TIME (Easy, Large), and large effects for RANK (Figure 7-top, Figure 8, Section A.2: Figure 15, Figure 16, and Figure 17), confirming H2-1, H2-2 for (RQ1, RQ3). Surprisingly, while results from Study 1 (Section 4.8) showed that interactive ORTHOGRAPHIC HEMISPHERE ☉ tended to be more accurate than


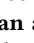



EQUAL EARTH ○ for distance comparisons and not worse than any other projections for area comparisons, it turned out to be significantly worse than EQUAL EARTH ○ and TORUS ☉ for reading CLUSTER NUMBER ♣.







While it has the advantage of being a straightforward mapping from the sphere, ORTHOGRAPHIC HEMISPHERE ☉ is an interrupted and non-area-preserving projection. Therefore, we conjecture that this discontinuity caused participants to struggle to make out cluster boundaries as compared with continuous representations in TORUS ☉ and EQUAL EARTH ○. Furthermore, there is no link distortion in TORUS ☉. This was supported by participants feedback, e.g., “The ‘Equal Earth’ method felt much easier to distinguish every [individual] cluster.” (P20), and “Orthographic Hemisphere uses 2 maps so it is much more difficult to interpret than the others [...] making it harder to isolate clusters.” (P12). More quotes can be found in Section A.2.1.



We also note that although being statistically significant, the small effect sizes indicate that EQUAL EARTH ○ is slightly faster than TORUS ☉ [14, 25].







EQUAL EARTH ○, ORTHOGRAPHIC HEMISPHERE ☉ and TORUS ☉ significantly resulted in less error than FLAT ☆. EQUAL EARTH ○ and TORUS ☉ were significantly preferred over FLAT ☆ but TORUS ☉ and ORTHOGRAPHIC HEMISPHERE ☉ took longer time than FLAT ☆. This was found statistically significant with moderate and large effect sizes (Figure 7-top, Figure 8, Section A.2: Figure 15, Figure 16, and Figure 17), leading us to reject H2-1 for (RQ2) and confirm H2-2 for (RQ2). These results provide strong evidence that with automatic panning and interaction, participants were able to better identify the high-level network structures using spherical projections. They also confirm Chen et al.'s results where toroidal layouts with automatic panning significantly outperformed FLAT ☆ in terms of error for cluster understanding tasks [12]. Some participants explicitly mentioned that good separation of clusters in continuous surfaces such as EQUAL EARTH ○ and TORUS ☉ helped understanding (Section A.2.1).

Participants also mentioned the automatic panning helped them “see” without the need to interact e.g., “[with EQUAL EARTH 


For SHORTEST PATH NUMBER , FLAT  and TORUS  performed equally well, both significantly outperforming EQUAL EARTH  and ORTHOGRAPHIC HEMISPHERE 

We conjecture that although TORUS  involves broken links across the boundaries, it appears similar to FLAT  using straight links while the distortion of paths in EQUAL EARTH  might hamper path tracing tasks, as participants mentioned “Flat surface is easier to read, it helps sometimes when you can also move it. Earth-like is just hard to use, especially when it’s an equal globe.” (P9), and “the warping in [EQUAL EARTH ] made my eyes hurt a little therefore its in 4th place, and the torus being the most straight forward gets 1st with [FLAT ] in second as they’re very similar.” (P13). On the other hand, ORTHOGRAPHIC HEMISPHERE  has less link distortion but the interruption between two hemispheres and the strong perspective distortion near the boundary of the maps may make participants confused, supported by participants’ feedback (Section A.2.1).




Although automatic panning (Section 5.2) provides some benefits for reducing split of clusters across the boundaries for CLUSTER NUMBER , it seems it has less benefits for SHORTEST PATH NUMBER , as participants mentioned they still need to position the image to see the full path to identify the shortest path connecting the nodes within the time limit (Section A.2.1).






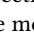
Overall, these findings suggest that TORUS  presents a general solution being not only less error prone than FLAT  or ORTHOGRAPHIC HEMISPHERE  for CLUSTER NUMBER , but also comparable to FLAT  for SHORTEST PATH NUMBER .

6.9 Limitations

A limitation of our study is that we only tested with one layout algorithm. While many different algorithms exist for FLAT  and it is possible that other layout algorithms could be adapted to spherical and torus embeddings, doing so is not necessarily trivial. Also, it should be noted that there are algorithms which can optimise layout for a known set of clusters (i.e. where the cluster labelling is known in advance) [16, 22, 29, 40]. However, for the tasks tested in this paper we do not preidentify the clusters but rather leave cluster identification as the user task. Similar layouts [20] have been used before for cluster understanding tasks on torus layouts [12, 13].

7 CONCLUSION AND FUTURE WORK

The overwhelming finding in Study 1 that interactive panning led to significantly lower error rate and more preferred over standard map projections across all tasks tested suggests that such panning should be routinely provided in online world maps, for example in education or in reporting of world events, climate patterns, and so on. While less clear, our results indicated that EQUAL EARTH  was the best performing continuous projection, while the straight-forward ORTHOGRAPHIC HEMISPHERE  was an effective hemispheric projection. Our results also indicate that even with interaction, EQUIRECTANGULAR  may be a poor choice for comparing areas and estimating directions.

Our Study 2 results also confirm the benefits of topologically closed surfaces, such as the surface of a torus or sphere, when using node-link diagrams to investigate network structure. All of EQUAL EARTH , ORTHOGRAPHIC HEMISPHERE , and TORUS  outperformed FLAT  for cluster understanding tasks. While the spherical projections impeded path following tasks, it seems TORUS  may be a good general solution, being as accurate as FLAT  for path following. Although they have been explored in research, 2D projections of toroidal and spherical network layouts are rarely seen in practice. This may be because until recently effective layout and projection methods for such geometries were not easily available. We intend to make all our algorithms and extensions of existing open-source tools available for easy consumption in web applications via GitHub and npm packages.

A limitation of our work is that there are many more spherical map projections than those evaluated here, although we tried to select the most representative techniques. Further work may extend such evaluation to other projection types. Another interesting topic for future investigation would be investigating spherical and torus projections of other types of abstract data representation, such as multi-dimensional scaling techniques of high-dimensional data. Another family of interactive techniques for exploring graphs involve applying spatial distortion around regions of interest to achieve a kind of “structure aware zooming”, e.g. [58]. It would be interesting to compare the efficacy of these techniques against interactive sphere and torus projections.

ACKNOWLEDGMENTS

We thank Bernhard Jenny, Lonni Besançon, Sarah Schöttler, Hong Gui, Mengxing Li, Ishwari Bhade, Umair Afzal, Sunny Singh for helpful discussion about the study design, anonymous participants for user study feedback, and our reviewers for helpful suggestions to improve this manuscript. We thank Stay Healthy and Monash eSolution crew for their technical support. This research is supported in part by Monash FIT Graduate Research Candidature Funding Scheme.

REFERENCES

- [1] Benjamin Bach, Conglei Shi, Nicolas Heulot, Tara Madhyastha, Tom Grabowski, and Pierre Dragicevic. 2015. Time curves: Folding time to visualize patterns of temporal evolution in data. *IEEE transactions on visualization and computer graphics* 22, 1 (2015), 559–568.
- [2] Albert-László Barabási and Réka Albert. 1999. Emergence of scaling in random networks. *science* 286, 5439 (1999), 509–512.
- [3] Douglas Bates, Martin Mächler, Ben Bolker, and Steve Walker. 2015. Fitting Linear Mixed-Effects Models Using **lme4**. *Journal of Statistical Software* 67, 1 (2015), 47 pages. <https://doi.org/10.18637/jss.v067.i01>

- [4] Lonni Besançon, Mehdi Ammi, and Tobias Isenberg. 2017. Pressure-based gain factor control for mobile 3D interaction using locally-coupled devices. In *Proceedings of the 2017 CHI Conference on Human Factors in Computing Systems*. 1831–1842.
- [5] Lonni Besançon, Paul Issartel, Mehdi Ammi, and Tobias Isenberg. 2017. Mouse, Tactile, and Tangible Input for 3D Manipulation. In *Proceedings of the ACM Conference on Human Factors in Computing Systems (CHI)*. Denver, United States, 4727–4740. <https://doi.org/10.1145/3025453.3025863>
- [6] Michael Bostock and Jason Davies. 2013. Code as Cartography. *The Cartographic Journal* 50, 2 (May 2013), 129–135. <https://doi.org/10.1179/0008704113Z.00000000078>
- [7] Ulrik Brandes, Marco Gaertler, and Dorothea Wagner. 2003. Experiments on graph clustering algorithms. In *European Symposium on Algorithms*. Springer, 568–579.
- [8] Richard Brath and Peter MacMurchy. 2012. Sphere-based information visualization: Challenges and benefits. In *2012 16th International Conference on Information Visualisation*. IEEE, 1–6.
- [9] Matthew Brehmer, Bongshin Lee, Petra Isenberg, and Eun Kyoung Choe. 2018. Visualizing ranges over time on mobile phones: a task-based crowdsourced evaluation. *IEEE transactions on visualization and computer graphics* 25, 1 (2018), 619–629.
- [10] Claus-Christian Carbon. 2010. The Earth is flat when personally significant experiences with the sphericity of the Earth are absent. *Cognition* 116, 1 (2010), 130–135.
- [11] Kun-Ting Chen, Tim Dwyer, Benjamin Bach, and Kim Marriott. [n.d.]. Rotate or Wrap? Interactive Visualisations of Cyclical Data on Cylindrical or Toroidal Topologies. ([n. d.]).
- [12] Kun-Ting Chen, Tim Dwyer, Benjamin Bach, and Kim Marriott. 2021. It's a Wrap: Toroidal Wrapping of Network Visualisations Supports Cluster Understanding Tasks. In *Proceedings of the 2021 CHI Conference on Human Factors in Computing Systems*. 1–12.
- [13] Kun-Ting Chen, Tim Dwyer, Kim Marriott, and Benjamin Bach. 2020. DoughNets: Visualising Networks Using Torus Wrapping. In *Proceedings of the 2020 CHI Conference on Human Factors in Computing Systems*. 1–11.
- [14] Andy Cockburn, Pierre Dragicevic, Lonni Besançon, and Carl Gutwin. 2020. Threats of a replication crisis in empirical computer science. *Commun. ACM* 63, 8 (2020), 70–79.
- [15] Jacob Cohen. 2013. *Statistical power analysis for the behavioral sciences*. Academic press.
- [16] Christopher Collins, Gerald Penn, and Sheelagh Cpendale. 2009. Bubble sets: Revealing set relations with isocontours over existing visualizations. *IEEE Transactions on Visualization and Computer Graphics* 15, 6 (2009), 1009–1016.
- [17] Jason Davies. 2013. *Rotate the World*. <https://www.jasondavies.com/maps/rotate/>
- [18] Jan De Leeuw and Patrick Mair. 2009. Multidimensional scaling using majorization: SMACOF in R. *Journal of statistical software* 31, 1 (2009), 1–30.
- [19] Fan Du, Nan Cao, Yu-Ru Lin, Panpan Xu, and Hanghang Tong. 2017. isphere: Focus+ context sphere visualization for interactive large graph exploration. In *Proceedings of the 2017 CHI Conference on Human Factors in Computing Systems*. 2916–2927.
- [20] Tim Dwyer, Kim Marriott, and Michael Wybrow. 2008. Topology preserving constrained graph layout. In *International Symposium on Graph Drawing*. Springer, 230–241.
- [21] Santo Fortunato. 2010. Community detection in graphs. *Physics reports* 486, 3-5 (2010), 75–174.
- [22] Emden R Gansner, Yifan Hu, and Stephen G Kobourov. 2009. Gmap: Drawing graphs as maps. In *International Symposium on Graph Drawing*. Springer, 405–407.
- [23] Emden R Gansner, Yehuda Koren, and Stephen North. 2004. Graph drawing by stress majorization. In *International Symposium on Graph Drawing*. Springer, 239–250.
- [24] Nicolas Greffard, Fabien Picarougne, and Pascale Kuntz. 2011. Visual community detection: An evaluation of 2d, 3d perspective and 3d stereoscopic displays. In *International Symposium on Graph Drawing*. Springer, 215–225.
- [25] Jouni Helske, Satu Helske, Matthew Cooper, Anders Ynnerman, and Lonni Besançon. 2021. Can visualization alleviate dichotomous thinking Effects of visual representations on the cliff effect. *IEEE Transactions on Visualization and Computer Graphics* (2021).
- [26] Pontus Hennerdal. 2015. Beyond the periphery: Child and adult understanding of world map continuity. *Annals of the Association of American Geographers* 105, 4 (2015), 773–790.
- [27] Lukáš Herman, Vojtěch Juřík, Zdeněk Stachoň, Daniel Vrbík, Jan Rusnák, and Tomáš Rezník. 2018. Evaluation of user performance in interactive and static 3D maps. *ISPRS International Journal of Geo-Information* 7, 11 (2018), 415.
- [28] Florian Hruby, Mónica Chico Avelino, and Raymundo Montoya Ayala. 2016. Journey to the end of the world Map—How edges of world maps shape the spatial mind. *Journal for Geographic Information Science* 1 (2016), 314–323.
- [29] Yifan Hu, Emden R Gansner, and Stephen Kobourov. 2010. Visualizing graphs and clusters as maps. *IEEE Computer Graphics and Applications* 30, 6 (2010), 54–66.
- [30] Weidong Huang, Peter Eades, and Seek-Hee Hong. 2009. Measuring effectiveness of graph visualizations: A cognitive load perspective. *Information Visualization* 8, 3 (2009), 139–152.
- [31] Transparent Statistics in Human-Computer Interaction Working Group. 2019. Transparent Statistics Guidelines. <https://doi.org/10.5281/zenodo.1186169> (Available at <https://transparentstats.github.io/guidelines>).
- [32] Bernhard Jenny, Johannes Liem, Bojan Šavrič, and William M. Putman. 2016. Interactive video maps: A year in the life of Earth's CO2. *Journal of Maps* 12, sup1 (Nov. 2016), 36–42. <https://doi.org/10.1080/17445647.2016.1157323>
- [33] Bernhard Jenny, Bojan Šavrič, Nicholas D Arnold, Brooke E Marston, and Charles A Preppernau. 2017. A guide to selecting map projections for world and hemisphere maps. In *Choosing a map projection*. Springer, 213–228.
- [34] Stephen G Kobourov and Matthew Landis. 2008. Morphing planar graphs in spherical space. (2008).
- [35] Stephen G Kobourov and Kevin Wampler. 2005. Non-Euclidean spring embedders. *IEEE Transactions on Visualization and Computer Graphics* 11, 6 (2005), 757–767.
- [36] Oh-Hyun Kwon, Chris Muelder, Kyungwon Lee, and Kwan-Liu Ma. 2016. A study of layout, rendering, and interaction methods for immersive graph visualization. *IEEE transactions on visualization and computer graphics* 22, 7 (2016), 1802–1815.
- [37] Bongshin Lee, Catherine Plaisant, Cynthia Sims Parr, Jean-Daniel Fekete, and Nathalie Henry. 2006. Task Taxonomy for Graph Visualization. In *Proceedings of the 2006 AVI Workshop on BEyond Time and Errors: Novel Evaluation Methods for Information Visualization* (Venice, Italy) (BELIV '06). ACM, New York, NY, USA, 1–5. <https://doi.org/10.1145/1168149.1168168>
- [38] Russell V. Lenth. 2016. Least-Squares Means: The R Package **lsmeans**. *Journal of Statistical Software* 69, 1 (2016), 33 pages. <https://doi.org/10.18637/jss.v069.i01>
- [39] Christopher Manning. 2017. *Spherical Force-Directed Layout*. <http://bllocks.org/christophermanning/4208494>
- [40] Wouter Meulemans, Nathalie Henry Riche, Bettina Speckmann, Basak Alper, and Tim Dwyer. 2013. Kelfusion: A hybrid set visualization technique. *IEEE transactions on visualization and computer graphics* 19, 11 (2013), 1846–1858.
- [41] Tamara Munzner. 2014. *Visualization analysis and design*. CRC press.
- [42] Mark EJ Newman. 2006. Modularity and community structure in networks. *Proceedings of the national academy of sciences* 103, 23 (2006), 8577–8582.
- [43] Mershack Okoe, Radu Jianu, and Stephen Kobourov. 2018. Node-link or adjacency matrices: Old question, new insights. *IEEE transactions on visualization and computer graphics* 25, 10 (2018), 2940–2952.
- [44] Stefan Palan and Christian Schitter. 2018. Prolific. ac—A subject pool for online experiments. *Journal of Behavioral and Experimental Finance* 17 (2018), 22–27.
- [45] Julie Pallant. 2013. *SPSS survival manual*. McGraw-hill education (UK).
- [46] Eyal Peer, Laura Brandimarte, Sonam Samat, and Alessandro Acquisti. 2017. Beyond the Turk: Alternative platforms for crowdsourcing behavioral research. *Journal of Experimental Social Psychology* 70 (2017), 153–163.
- [47] Scott Perry, Mason Sun Yin, Kathryn Gray, and Stephen Kobourov. 2020. Drawing Graphs on the Sphere. In *Proceedings of the International Conference on Advanced Visual Interfaces*. 1–9.
- [48] Helen C Purchase. 2002. Metrics for graph drawing aesthetics. *Journal of Visual Languages & Computing* 13, 5 (2002), 501–516.
- [49] Dario Rodighiero. 2020. Drawing Network Visualizations on a Continuous, Spherical Surface. In *2020 24th International Conference Information Visualisation (IV)*. IEEE, 573–580.
- [50] Bahador Saket, Paolo Simonetto, and Stephen Kobourov. 2014. Group-level graph visualization taxonomy. *arXiv preprint arXiv:1403.7421* (2014).
- [51] Kadek Ananta Satriadi, Barrett Ens, Tobias Czauderna, Maxime Cordeil, and Bernhard Jenny. 2021. Quantitative Data Visualisation on Virtual Globes. In *Proceedings of the 2021 CHI Conference on Human Factors in Computing Systems*. 1–14.
- [52] Bojan Šavrič, Bernhard Jenny, Denis White, and Daniel R Strebe. 2015. User preferences for world map projections. *Cartography and Geographic Information Science* 42, 5 (2015), 398–409.
- [53] Bojan Šavrič, Tom Patterson, and Bernhard Jenny. 2019. The equal earth map projection. *International Journal of Geographical Information Science* 33, 3 (2019), 454–465.
- [54] Thomas Schäfer and Marcus A Schwarz. 2019. The meaningfulness of effect sizes in psychological research: Differences between sub-disciplines and the impact of potential biases. *Frontiers in Psychology* 10 (2019), 813.
- [55] Sarah Schöttler, Yalong Yang, Hanspeter Pfister, and Benjamin Bach. 2021. Visualizing and interacting with geospatial networks: A survey and design space. In *Computer Graphics Forum*. Wiley Online Library.
- [56] John Parr Snyder. 1987. *Map projections—A working manual*. Vol. 1395. US Government Printing Office.
- [57] John P Snyder. 1997. *Flattening the earth: two thousand years of map projections*. University of Chicago Press.
- [58] Yunhai Wang, Yanyan Wang, Haifeng Zhang, Yinqi Sun, Chi-Wing Fu, Michael Sedlmair, Baoquan Chen, and Oliver Deussen. 2018. Structure-aware fisheye

- views for efficient large graph exploration. *IEEE transactions on visualization and computer graphics* 25, 1 (2018), 566–575.
- [59] Colin Ware and Peter Mitchell. 2008. Visualizing graphs in three dimensions. *ACM Transactions on Applied Perception (TAP)* 5, 1 (2008), 1–15.
- [60] Duncan J Watts and Steven H Strogatz. 1998. Collective dynamics of 'small-world' networks. *nature* 393, 6684 (1998), 440–442.
- [61] Evan James Williams. 1949. Experimental designs balanced for the estimation of residual effects of treatments. *Australian Journal of Chemistry* 2, 2 (1949), 149–168.
- [62] Yingxin Wu and Masahiro Takatsuka. 2006. Spherical self-organizing map using efficient indexed geodesic data structure. *Neural Networks* 19, 6-7 (2006), 900–910.
- [63] Kai Xu, Chris Rooney, Peter Passmore, Dong-Han Ham, and Phong H Nguyen. 2012. A user study on curved edges in graph visualization. *IEEE transactions on visualization and computer graphics* 18, 12 (2012), 2449–2456.
- [64] Yalong Yang, Tim Dwyer, Bernhard Jenny, Kim Marriott, Maxime Cordeil, and Haohui Chen. 2019. Origin-Destination Flow Maps in Immersive Environments. *IEEE Transactions on Visualization and Computer Graphics* 25, 1 (2019), 693–703. <https://doi.org/10.1109/tvcg.2018.2865192>
- [65] Yalong Yang, Bernhard Jenny, Tim Dwyer, Kim Marriott, Haohui Chen, and Maxime Cordeil. 2018. Maps and globes in virtual reality. In *Computer Graphics Forum*, Vol. 37. Wiley Online Library, 427–438.
- [66] Vahan Yeghouchian, Yalong Yang, Tim Dwyer, Lee Lawrence, Michael Wybrow, and Kim Marriott. 2020. Scalability of network visualisation from a cognitive load perspective. *IEEE Transactions on Visualization and Computer Graphics* 27, 2 (2020), 1677–1687.
- [67] Jonathan X Zheng, Samraat Pawar, and Dan FM Goodman. 2018. Graph drawing by stochastic gradient descent. *IEEE transactions on visualization and computer graphics* 25, 9 (2018), 2738–2748.