

Wormholes in virtual space: From cognitive maps to cognitive graphs

William H. Warren^a

Daniel B. Rothman^a

Benjamin H. Schnapp^a

Jonathan D. Ericson^a

^aDepartment of Cognitive, Linguistic and Psychological Sciences

Brown University, Box 1821

190 Thayer St.

Providence, RI, 02912, USA

Corresponding author:

Dr. William H. Warren

Department of Cognitive, Linguistic and Psychological Sciences

Brown University, Box 1821

190 Thayer St.

Providence, RI 02912

+1 401 863 3980 (office)

+1 401 524 9294 (mobile)

Bill_Warren@brown.edu

RUNNING HEAD: Wormholes in virtual space

Abstract

Humans and other animals build up spatial knowledge of the environment on the basis of visual information and path integration. We compare three hypotheses about the geometry of this navigational knowledge: (a) 'cognitive map' with metric Euclidean structure and a consistent coordinate system, (b) 'topological graph' or network of paths between places, and (c) 'labelled graph' incorporating local metric information about path lengths and junction angles. In two experiments, participants walked in a non-Euclidean environment, a virtual hedge maze containing two 'wormholes' that visually rotated and teleported them between locations. During training, they learned the metric locations of eight target objects from a 'home' location, visible individually. During testing, shorter wormhole routes to a target were preferred, and novel shortcuts were directional, contrary to the topological hypothesis. Shortcuts were strongly biased by the wormholes, with mean constant errors of 37° and 41° (45° expected), revealing violations of the metric postulates in spatial knowledge. In addition, shortcuts to targets near wormholes shifted relative to flanking targets, revealing 'rips' (86% of cases), 'folds' (91%), and ordinal reversals (66%) in spatial knowledge. Moreover, participants were completely unaware of these geometric inconsistencies, reflecting a surprising insensitivity to Euclidean structure. The probability of the shortcut data under the Euclidean map model and labelled graph model indicated decisive support for the latter ($\text{BF}_{\text{GM}} > 100$). We conclude that knowledge of navigation space is best characterized by a labelled graph, in which local metric information is approximate, geometrically inconsistent, and not embedded in a common coordinate system. This class of 'cognitive graph' models supports route finding, novel detours, and rough shortcuts, and has the potential to unify a range of data on spatial navigation.

KEY WORDS: Human navigation, cognitive map, cognitive graph, spatial cognition

1.0 Introduction

As humans and other animals explore their environments, they build up spatial knowledge based on visual, idiothetic, and other sensory information. The underlying geometry of the resulting knowledge might take a number of forms (Tobler, 1976; Trullier, Wiener, Berthoz, & Meyer, 1997; Tversky, 1993). At one end of the spectrum (Figure 1) lies a Euclidean cognitive map, which preserves metric information about the locations of known places in a common coordinate system (Gallistel, 1990; O'Keefe & Nadel, 1978; Tolman, 1948). At the other end lies weak topological structure, such as a graph that only preserves a network of paths connecting known places (R. W. Byrne, 1979; Kuipers, Tecuci, & Stankiewicz, 2003; Erner, Riegler, & Herrmann, 2000). Various combinations of metric and topological knowledge have also been proposed, capitalizing on the advantages of each (Chown, Kaplan, & Kortenkamp, 1995; Kuipers, 2000; Mallot & Basten, 2009; Meilinger, 2008; Poucet, 1993). After decades of research on this issue, researchers still hold opposing views and the question remains unresolved. In this article we report two experiments on navigation in a non-Euclidean environment that challenge both extremes. We argue that the evidence is best accounted for by a *labelled graph* that incorporates local metric information but has no globally consistent coordinate system.

Figure 1 about here

1.1 Euclidean maps

Euclidean knowledge (Figure 1A) would be advantageous because it supports flexible navigation, including novel as-the-crow-flies shortcuts and the integration of separately learned routes. An influential theory holds that a metric Euclidean map is constructed on the basis of path integration (Gallistel & Cramer, 1996; McNaughton, Battaglia, Jensen, Moser, & Moser, 2005; O'Keefe & Nadel, 1978). Specifically, as an animal explores the environment, the path integrator registers idiothetic (i.e. proprioceptive, motor,

vestibular) information about the angles turned and distances travelled from a home location, and assigns salient places to coordinates in this inertial coordinate system. Grid, place, and head-direction cells in the hippocampal formation have been interpreted as a system for encoding metric maps from path integration (Derdikman & Moser, 2010; McNaughton et al., 2005). Indeed, mammals and insects have been observed to take shortcuts between known locations (Chapuis, Durup, & Thinus-Blanc, 1987; Cheeseman et al., 2014; Gould, 1986; Menzel et al., 2006), and humans are able to estimate the directions and distances between familiar places (Chrastil & Warren, 2013; Holmes & Sholl, 2005; Ishikawa & Montello, 2006; Schinazi, Nardi, Newcombe, Shipley, & Epstein, 2013; Waller & Greenauer, 2007; Weisberg, Schinazi, Newcombe, Shipley, & Epstein, 2014), consistent with a Euclidean cognitive map.

On closer examination, however, the evidence appears inconclusive. Apparently novel shortcuts may be more simply explained by knowledge of familiar routes, landmarks, or views (Benhamou, 1996; Bennett, 1996; Cheung et al., 2014; Dyer, 1991; Foo, Warren, Duchon, & Tarr, 2005). Directional estimates in humans are highly unreliable, with absolute angular errors of 20°-100° and angular standard deviations on the order of 30° (Chrastil & Warren, 2013; Foo et al., 2005; Ishikawa & Montello, 2006; Meilinger, Riecke, & Bühlhoff, 2014; Schinazi et al., 2013; Waller & Greenauer, 2007; Weisberg et al., 2014), while junctions tend to be orthogonalized to 90° (R. W. Byrne, 1979).

Distance estimates are biased by the number of intervening junctions, turns, and boundaries, and are asymmetric between more and less salient places (Burroughs & Sadalla, 1979; R. W. Byrne, 1979; Cadwallader, 1979; Kosslyn, Pick, & Fariello, 1974; McNamara & Diwadkar, 1997; Sadalla & Magel, 1980; Sadalla & Staplin, 1980; Tversky, 1992). People often fail to integrate learned routes, and cross-route estimates are generally poor (Golledge, Ruggles, Pellegrino, & Gale, 1993; Ishikawa & Montello, 2006; Moeser, 1988; Schinazi et al., 2013; Weisberg et al., 2014).

The *sine qua non* of a Euclidean map is a distance metric that satisfies the metric postulates. A metric space is defined by a distance metric that must satisfy the postulates of *positivity* ($AA=0$, $AB>0$), *symmetry* ($AB=BA$), segmental *additivity* ($AB+BC=AC$), and the *triangle inequality* ($A + C \geq AC$), where pairs of letters denote distances between pairs of points (Beals, Krantz, & Tversky, 1968). Yet the human distance and direction estimates reviewed above imply violations of the postulates of symmetry, additivity, and the triangle inequality.

Given its supposed ubiquity, there is thus a surprising lack of convincing evidence for a metric Euclidean map. However, it is difficult to reject the hypothesis because the expected level of performance is not well-specified. The view thus remains influential and has many prominent advocates (P. Byrne, Becker, & Burgess, 2007; Cheeseman et al., 2014; McNaughton et al., 2005; Nadel, 2013).

1.2 Topological graphs

At the other extreme lies a topological graph (Figure 1B), a network of connections in which nodes correspond to familiar places and edges to known paths between them.¹ Graph knowledge would be advantageous because available routes and detours are explicitly specified in one compact structure, and do not need to be derived from a coordinate map via additional operations. Nodes may have associated place information such as views, landmarks, or surface layout (local geometry), enabling self-localization and orientation (Epstein & Vass, 2014). It has been suggested that place cell activity might reflect a hippocampal graph (Dabaghian, Mémoli, Frank, & Carlsson, 2012; Muller, Stead, & Pach, 1996). Place fields are anchored to environmental features and

¹ A graph can also describe connections between other entities, such as views (Hübner & Mallot, 2007) or neighborhoods (Wiener & Mallot, 2003).

their metric locations shift with transformations of the layout (Dabaghian, Brandt, & Frank, 2014; Muller & Kubie, 1987; O'Keefe & Burgess, 1996).

Graph knowledge is richer than route knowledge, but weaker than 'survey' or map knowledge. Whereas routes are typically characterized as chains of place-action associations (Trullier et al., 1997), a graph can express multiple routes between two places and multiple paths intersecting at one place. Thus, graph knowledge provides the basis for finding novel routes and detours (Chrastil & Warren, 2014). On the other hand, a purely topological graph is insufficient to explain behavior such as taking shorter routes or novel shortcuts, implying that topological knowledge may be augmented by metric information.

These considerations have led to hybrid theories in which topological and map knowledge are represented in parallel or hierarchical systems (P. Byrne et al., 2007; Chown et al., 1995; Kuipers, 2000; Thrun, 1998; Trullier et al., 1997). However, such models are less parsimonious and are compromised by the lack of evidence for metric maps.

1.3 Labelled graphs

We suggest that spatial knowledge is more appropriately characterized by a labelled graph, a single structure that incorporates local metric information (Figure 1C). Specifically, path lengths are denoted by edge weights and the angles between paths at junctions are denoted by node labels. This local metric information is typically noisy, biased, and geometrically inconsistent. Yet such a *cognitive graph* is sufficient to find generally shorter routes and detours, and to generate approximate shortcuts, although their accuracy would depend on the level of error in the graph.

What distinguishes a labelled graph from a metric map is that the local information is not embedded into a common coordinate system, a 'global metric embedding' in which

places are assigned coordinates in a globally consistent map. Apparently Euclidean behavior does not necessarily imply a metric map, for it could be underwritten by a labelled graph together with adaptive navigation strategies. For example, approximate shortcuts could be generated by vector addition along a path through the graph, which may be sufficient to bring the navigator within sight of local beacons, yielding successful shortcuts between familiar places (Foo et al., 2005).

A labelled graph is reminiscent of previous theories that combine metric and topological information in a single graph structure (Mallot & Basten, 2009; Meilinger, 2008; Poucet, 1993). These models are based on the concept of local reference frames (or coordinate systems) that are linked by vectors specifying the metric distance and direction between them, and are designed to solve the problem of novel shortcuts by representing separate places in a common coordinate frame. In contrast, a labelled graph shows that a common coordinate system is not required to account for shortcuts. We will compare these theories in more detail in the General Discussion (Section 4).

1.4 The present study

The geometry of spatial knowledge thus remains controversial. Although humans are able to estimate distances and directions in a Euclidean environment, they make large errors, and it is not clear whether they rely on a consistent metric map or some weaker form of knowledge. We focus on the prevailing view of learning in navigation space, where places are viewed individually and a cognitive map is constructed by path integrating from one place to another.

We decided to approach the issue from a different direction: rather than asking how accurately navigators learn a Euclidean environment, we asked what they learn from walking in a *non*-Euclidean environment. Participants explored a virtual hedge maze containing two 'wormholes', which instantly teleported and rotated them from one visual

place to another.² During training, the metric locations of target objects with respect to a 'home' location were preserved, providing information about their relative physical positions. We reasoned that if the navigation system attempts to construct a Euclidean map based on path integration, people would have more difficulty learning the inconsistent non-Euclidean world than a matched Euclidean world, but shortcuts between target objects should be similar because their metric locations were specified during training. In contrast, if spatial knowledge resembles a labelled graph, then learning would be comparable in the two environments, but shortcuts should be systematically biased by the wormholes. Finally, if knowledge resembles a topological graph, then participants should not prefer shorter routes and their shortcuts should be unsystematic due to the absence of metric information about distances and directions.

Figure 2 about here

2.0 Experiment 1: Route-finding and shortcuts

In the first experiment, we investigated graph knowledge and survey knowledge in Euclidean and non-Euclidean environments. We created matched versions of a virtual hedge maze with the same graph, one of which contained two 'wormholes' (Figure 2). Each environment had a 'home' location linked by radial corridors to eight 'places' marked by distinctive objects. The objects were only visible one at a time, so participants had to learn their locations by path integrating between them. We probed graph knowledge by asking participants to find routes between objects within the maze, and probed survey knowledge by asking them to take novel shortcuts between the same objects.

² Note that analogous 'wormholes' can occur in everyday life as when riding the subway.

The non-Euclidean maze was identical to the Euclidean control maze, except for the insertion of two bi-directional wormholes that visually translated the participant by 6m or 10m and rotated them by 90° when they walked through an invisible portal in a corridor (Figure 2B). This was achieved by rotating the virtual environment 90° in the opposite direction. The wormholes thus dissociated the *visual reference frame*, which rotated with the maze, from the *inertial reference frame*, which was fixed in physical space. In the Euclidean maze, these two reference frames were always congruent. In the non-Euclidean maze, walking through a wormhole instantly teleported the participant to a different location in the visual reference frame, while they remained in place in the inertial (physical) reference frame. Distances and angles traversed in the inertial frame were specified by idiothetic path integration, and in the visual frame by distance perception and visual odometry based on optic flow (Kearns, Warren, Duchon, & Tarr, 2002). Thus, referring to Figure 1, if one started at the cactus and walked a loop to Home, to the bookcase, through Wormhole 1, and back to the cactus, there would be consistent visual and idiothetic information available to reveal that the cactus had changed its location in the inertial reference frame.

Participants first freely explored the environment and experienced the wormholes. They were then trained from home to each object with congruent visual and idiothetic information, so the Euclidean locations of objects with respect to home were specified in both environments. Thus, if the navigation system builds a metric map, shortcuts should be taken in the ‘Euclidean direction’ of the target object. In contrast, if participants only learn a topological graph of the environment, one would expect the direction of shortcuts to be uniformly distributed (see Chrastil & Warren, 2013). However, if participants learn local metric information about the path to a given object, shortcuts would be systematically biased toward the ‘wormhole direction’ of the target, that is, toward the target location reached via the nearest wormhole.

In sum, the metric map hypothesis predicts that the non-Euclidean maze should be harder to learn because it is globally inconsistent, shorter routes should be preferred within the maze, and shortcuts in the two mazes should be similar, based on metric knowledge. Conversely, the topological graph hypothesis predicts comparable learning in the two environments, no preference for shorter routes, and unsystematic shortcuts. A preference for shorter routes, and shortcuts that are biased by wormholes, would be indicative of a labelled graph that includes local knowledge of path lengths and junction angles.

2.1 Material and Methods

Experiments were conducted in the Virtual Environment Navigation Laboratory (VENLab) at Brown University. Experiment 1 had a 2 (maze) x 2 (task) between-subject design.

2.1.1 Participants

A total of 39 participants completed the first experiment. One group of 20 learned the wormhole maze, half of which was tested on the route task (5M, 5F) and half on the shortcut task (5M 5F). A control group of 19 learned the Euclidean maze, with nine tested on the route task (5M, 4F) and ten on the shortcut task (7M 3F). Seven additional participants failed to achieve the learning criterion in the wormhole maze, and six failed to do so in the Euclidean maze; the data thus represent competent navigators from the top three quartiles of participants. Informed consent was obtained from all participants, who were paid for their participation. The protocol was approved by Brown University's Institutional Review Board.

2.1.2 Apparatus

Participants walked freely in a 12m x 12m area while wearing a head-mounted display (HMD, Kaiser Pro-View 80, 63° H x 53° V field of view, 640 x 512 pixels per eye, 100% binocular overlap). Stereoscopic displays were generated on a graphics workstation, calibrated to the participant's measured inter-ocular distance, and presented at 60 frames/s. Head position and orientation were recorded with an ultrasonic/inertial tracking system (Intersense IS-900, 1mm linear and 1° angular RMS error 60 Hz sampling rate), and used to update the display with a total latency of 50-70 ms. The HMD was connected to a control box by a 15m cable, which was handled by 'wrangler' who followed the participant. Participants wore earphones playing masking noise (evening sounds), and the peripheral view of the real environment was occluded by a cloth hood.

2.1.3 Displays

The virtual environment consisted of an 11m x 11m hedge maze (Figure 2A), which contained a home location (home plate), eight objects (well, cactus, etc.), and three landmarks to aid orientation (familiar paintings on the walls). The eight objects were connected to home by direct radial paths that did not pass through a wormhole. Wormhole 1 instantly translated the participant by 6m and rotated them by 90° in the visual reference frame, and Wormhole 2 by 10m and 90° respectively (Figure 2B). The maze walls (2.13m high) were mapped with a foliage texture and the paths (1m wide) with a gravel texture. The views at the entrance and exit of each wormhole were matched so the transition was visually seamless, unaccompanied by rotational or translational optic flow.

2.1.4 Procedure

Participants first walked in a practice environment for 3-5 min to adapt to virtual reality (Mohler, Creem-Regehr, & Thompson, 2006). The experiment proper consisted of three

phases. In the *exploration phase*, a participant freely explored the maze for 8 min, visiting each object at least once and passing through each wormhole at least twice; a recorded voice named each object as it was approached (e.g. "This is the bookcase"). If the participant did not visit all eight objects within the first 6 min, they were verbally guided from home to each remaining object on the direct path.

In the *training phase*, participants were trained to walk from home to each object until they reached the criterion of finding the object in less than 30s. If the participant could not find an object within 1 min, they returned home and were verbally guided to it on the direct path; if they failed to reach criterion within 20 min, they were removed from the experiment.

During the *test phase*, participants began at the home location facing 'north', were instructed to walk to a start object within the maze corridors, and were then told to walk to a target object. In the *route task*, participants walked to the target within the maze corridors. A trial ended when they arrived within 0.5m of the target, and they then walked home in the corridors for the next trial. In the *shortcut task*, when participants reached the start object the maze disappeared, leaving only a textured ground plane visible; they then turned to face the remembered location of the target object and walked straight to it. A trial ended when the participant verbally reported arriving at the target location or reached the maze boundary; no feedback was given. They then rode home in a wheelchair on a circuitous path in the dark for the next trial.

Four pairs of objects were tested in both directions: two *probe* pairs near the wormhole portals (yellow and cyan dots in Figure 2) and two *standard* pairs some distance away (purple and navy dots). There were four trials to each of the four probe targets and two trials to each of the four standard targets, for a total of 24 test trials. They were presented in a randomized order in a one-hour session.

In a post-test questionnaire, participants were asked to report their impressions of the maze and anything they noticed about it. Shortcut participants were then given a list of the eight objects and asked to draw a map of the maze on paper.

2.1.5 Data Analysis

Approximately 7% of the test trials were lost due to tracker malfunction during data collection. In the route task, errors consisted of walking within sight of an incorrect object before reaching the target. In the shortcut task, the dependent variable was the initial walking direction, defined as the unit vector from the start position to the point at which they crossed a circle with 1m radius. (Walking distance was not analysed because it was limited by the laboratory walls.) The constant error on each trial was the angle between the initial walking direction and the target's 'Euclidean direction', defined by the target's location with respect to home in the inertial coordinate frame. Responses on probe trials were normalized so increasing errors were toward the target's 'wormhole direction', defined by the target's location in the inertial coordinate frame after the maze was visually rotated. Standard trial responses were normalized so increasing errors were to the right of the target's Euclidean location (because standard objects were not near wormhole portals).

Directional data were analyzed using circular statistics (Batschelet, 1981). The circular mean of constant error was computed for each participant, and the Mean Error and between-subject Angular Deviation (AD, the circular equivalent of the standard deviation) were computed for the group. To estimate the within-subject variable error, the angular deviation was computed separately for each participant over all their trials; the circular mean of these subject values is reported as the 'Mean AD'. Angular errors were compared using two-sample Watson-Williams F-tests for circular data, and Cohen's d was used as a measure of effect size. All statistical tests were two-tailed with $\alpha=.05$.

To directly compare the Euclidean map and labelled graph hypotheses, we took a Bayesian model comparison approach. We specified both the metric map (M_M) and labelled graph (M_G) models by the von Mises probability density function (the circular normal distribution), with parameters $\theta = [\mu, \kappa]$. The mean response direction μ could differ between the two models, but the circular concentration parameter κ was held constant ($1/\kappa$ is related to AD, see Batchelet, 1981). The models were compared with the Bayes Factor, which under these assumptions reduces to the likelihood ratio,

$$BF_{GM} = \frac{L(\mu_G|x)}{L(\mu_M|x)} = \frac{p(x|\mu_G)}{p(x|\mu_M)}$$

BF=1 implies that the models are equally likely, BF>3 is considered substantial support, BF>10 strong support, BF>30 very strong support, and BF>100 decisive support for the model in the numerator (Jeffreys, 1998).

2.2 Results and Discussion

2.2.1 Exploration and training

Participants experienced the wormholes frequently during exploration, passing through each a mean of 5.6 times. Yet the wormhole maze was learned as readily as the Euclidean control maze. The mean number of trials to criterion during training was 1.77 per object (SD=0.48) for the wormhole group and 1.61 per object (SD=0.43) for the control group, $t(37)=1.12$, $p=.27$, $d=0.36$. To investigate the null hypothesis for this measure, we computed the scaled JZS Bayes factor, $BF_{01}=3.11$; this indicates substantive support for the null hypothesis, which is favored 3 to 1 over the alternative (Rouder, Speckman, Sun, Morey, & Iverson, 2009). The non-Euclidean maze thus posed no particular difficulty for learning, despite its global inconsistency.

2.2.2 Route task

In the route task, participants also found routes in the wormhole maze as easily as the Euclidean maze. First, there were no errors in walking from home to the start object. Second, when walking from the start object to the target object, the wormhole group made no errors and the control group only three (visible in Figure 2A). Third, on probe trials, participants took advantage of the shorter wormhole routes in 84% of the trials, compared to only 28% for the corresponding routes in the Euclidean maze (Figure 2), $\chi^2(1)=277.67, p<.001$, Cohen's $\omega^2=188.58$. This significantly reduced the wormhole group's mean path length (6.36m, SD=2.58m) to half that of the control group's (12.27m, SD=1.72m), $t(17)=5.92, p<.001, d=2.69$, as well as reducing their mean travel time (9.47s, SD=3.01s) compared to the control (17.31s, SD=3.51), $t(17)=5.19, p<.001, d=2.40$.

Participants thus apparently acquired knowledge about path length that enabled them to select reliably shorter routes. We note that the present experiment was not designed to dissociate metric distance (path length in meters) and topological distance (number of nodes or edges traversed). In a similar experiment, however, Chrastil and Warren (2014) found that participants took the metrically shortest route or detour to a target significantly more often than topologically equivalent but longer alternatives. This supports the conclusion that participants acquired knowledge about metric path lengths.

2.2.3 *Shortcut task*

In the shortcut task, there were again no errors in walking to the start object. Rayleigh tests on shortcut direction found that mean constant errors were not uniformly distributed on the circle in the control environment, $z(10)=9.70$ for probe trials and $z(10)=9.72$ for standard trials (both $p<.001$), or in the wormhole environment, $z(10)=9.82$ for probe trials and $z(10)=8.11$ for standard trials (both $p<.001$). This result implies that shortcuts were directional and hence relied on metric information, inconsistent with the purely topological hypothesis.

On probe trials, the control group took shortcuts toward the Euclidean location of the target (Figure 3A), with a mean constant error of only 4.38° ($AD=10.02^\circ$). In contrast, the non-Euclidean group was biased toward the wormhole location of the target (Figure 3B), with a mean constant error of 37.36° ($AD=7.74^\circ$) close to the expected error of 45° . The group difference was highly significant, with a large effect size, $F(1,18)=556.86$, $p<.001$, $d=3.68$. The wormholes thus biased the remembered direction of the target by 83%. Within-subject variability was large in both groups, with a mean within-subject AD of 27.51° in the control group and 30.41° in the wormhole group, a difference of only 2.9° with a small effect size, $F(1,18)=4.46$, $p=.049$, $d=0.33$. This indicates that uncertainty in the remembered target direction was high in both environments, consistent with previous literature.

Figure 3 about here

We compared the Euclidean map and labelled graph hypotheses by computing the probability of the observed data under each model, $p(x|\mu)$, assuming concentration parameter $\kappa=35.1$ (corresponding to $AD=10^\circ$ from the control group). For the control group, both models predict shortcuts with a mean expected error of $\mu_M=\mu_G=0^\circ$ so the Bayes Factor, $BF_{GM}=p(x|\mu_G)/p(x|\mu_M)$, is necessarily 1. For the wormhole group, in contrast, the Euclidean model predicts shortcuts toward the Euclidean target location with an expected error of $\mu_M=0^\circ$, while the labelled graph model predicts shortcuts toward the wormhole target location with an expected error of $\mu_G=45^\circ$. The Bayes Factor was large, $BF_{GM} > 100$, indicating decisive support for the labelled graph model.

On standard trials, the mean constant error was also significantly larger for the wormhole group ($CE= -16.89^\circ$ to the left $AD=25.53^\circ$) than the control group ($CE= -5.06^\circ$ to the left, $AD=9.60^\circ$), $F(1,18)=14.81$, $p=0.0012$, $d=0.61$. This result implies that the wormholes biased the remembered locations of objects that were not adjacent to wormhole exits, but in the surrounding area. Similarly, the corresponding within-

subject variable errors were significantly larger for the wormhole group (mean $AD=45.22^\circ$) than the control group (mean $AD=31.11^\circ$), $F(1,18)=59.42$, $p<.001$, $d=1.21$, implying that the wormholes also increased uncertainty in the remembered locations of surrounding target objects.

2.2.4 Metric postulates

Most importantly, shortcuts revealed that spatial knowledge acquired in the wormhole environment violates the metric postulates. First, shortcuts to probe objects exhibited a large violation of the positivity postulate (Figure 4A), which states that the distance between any point and itself must be zero. When asked to walk from home H to start object A, participants succeeded on every trial. But when asked to take a shortcut from B to A – in the same inertial coordinate frame – they did not walk toward A, but were biased by a mean of 37° in the direction of A'. This response violates positivity, because the distance between A and A' in the inertial frame (6m) is much greater than zero.

Figure 4 about here

Positivity implies that shortcuts from B to A should be in the opposite direction of those from A to B, in inertial coordinates. This is because the start point for one shortcut (AB) should be the endpoint for the opposite shortcut (BA), and vice versa. Thus, if 180° is subtracted from the direction of vector AB, the mean difference should be zero under the Euclidean model. In contrast, the labelled graph model predicts a mean difference of 90° because both shortcuts should have 45° errors toward the wormhole. To compare the models, we assumed a von Mises probability density function with $\kappa=8.49$ (corresponding to $AD=20^\circ$ from the control group). For the control group both models have a mean expected difference of zero, $\mu_M=\mu_G=0^\circ$ yielding $F_{GM}=1$. But for the wormhole group, the mean expected difference under the Euclidean model is $\mu_M=0^\circ$

whereas that under the labelled graph model is $\mu_G=90^\circ$. The observed difference for the control group was 8.56° ($AD=20.09^\circ$) (refer to Figure 3A) whereas for the wormhole group it was 75.27° ($AD=15.25^\circ$) (refer to Figure 3). The Bayes factor was large $BF_{GM} > 100$, providing decisive evidence against positivity and the Euclidean model, and in favor of the labelled graph model.

Note that this behaviour also implies violation of the triangle inequality, which states that the third side of a triangle must be less than the sum of the other two sides ($AB < HA+HB$). In the present case the third side is so large that triangle ΔHA is not closed (Figure 4A).

2.2.5 Self-reports and sketch maps

Participants completely failed to detect these geometric inconsistencies. When asked to describe their impressions of the maze, no participant reported noticing the wormholes or any spatial anomalies.

Both Euclidean and non-Euclidean groups willingly drew sketch maps of their environments, even though the wormhole environment cannot be represented in the plane without crossed paths. Representative sketch maps appear in Figure 5. The sketch maps of both groups captured the radial structure of the maze, without any crossed paths. We analyzed the sketch maps from the shortcut subgroups in some detail.

To evaluate the graph structure of the maps, we first analysed the number of paths (edges) between adjacent objects (nodes) that were correctly remembered and falsely inserted. The control group had a mean hit rate of 81% of the edges, with only 5% false insertions (sensitivity: $d'=2.53$). The wormhole group had a mean hit rate of 72% of the edges, with only 2.5% false insertions ($d'=2.54$). Both groups thus preserved the topological graph of their respective mazes with comparable sensitivity. However, the wormhole group drew significantly shorter wormhole paths than the control group

(mean ratio = 0.61), $t(18)=2.78$, $p=.012$, $d=1.24$, reflecting metric knowledge of relative path lengths.

We then correlated the cyclic order of objects around the circumference of each sketch map with the cyclic order of objects in the two environments. The mean cyclic correlation with the wormhole order was significantly greater for the wormhole group, $r(8)=.97$, than for the control group, $r(8)=.77$, $t(18)=2.1195$, $p<.05$. Within the control group, the mean cyclic correlation was higher with the Euclidean order, $r(8)=.83$, than with the wormhole order, $r(8)=.77$, although the difference did not reach significance. Within the wormhole group, the cyclic correlation was higher with the wormhole order, $r(8)=.97$, than the Euclidean order, $r(8)=.91$, but not significantly so. Both groups thus remembered the cyclic order of objects in their respective environments with similar accuracy.

Figure 5 about here

2.2.6 Conclusions

The results of Experiment 1 appear to be most consistent with a labelled graph. Contrary to a purely topological graph, shorter routes were preferred and shortcuts were directional, implying that some knowledge of path lengths and angles was acquired. Contrary to a Euclidean map, learning was comparable in the two environments but shortcuts were strongly biased by experience with wormholes, leading to violations of the metric postulates. Strikingly, participants were completely unaware of these glaring inconsistencies, demonstrating a marked insensitivity to Euclidean structure.

This pattern of results can be accounted for by the labelled graph hypothesis. Local knowledge about path lengths and angles is unreliable because it derives from a noisy path integrator, is biased by experience with short wormhole paths, and is not embedded in a globally consistent coordinate system. The resulting spatial knowledge supports the

selection of shorter routes and rough shortcuts, but they are highly variable, influenced by wormholes, and geometrically inconsistent.

Consider some alternative explanations. Perhaps participants acquired a *noisy* Euclidean map, rather than a labelled graph. Spatial knowledge might be low-resolution, yet nonetheless Euclidean. While this hypothesis would account for the high variability of shortcuts in both environments, it does not explain their systematic bias in the wormhole environment.

Or perhaps an hour in the virtual maze was insufficient to learn a Euclidean map, but with more experience participants would eventually do so. This explanation is belied by the findings that shortcuts in the Euclidean maze were accurate on average (overall mean CE=2.66°) and that their variable error (mean AD=28.63°) was comparable to previous studies with extended experience in real environments (Ishikawa & Montello, 2006; Moeser, 1988; Schinazi et al., 2013). There is little evidence in the literature that further experience will yield up a Euclidean map. Other alternative explanations will be considered in the General Discussion (Section 4.2).

3.0 Experiment 2: Rips and folds

The results of Experiment 1 imply that spatial knowledge may include local discontinuities. For example, in Figure 4A, the remembered location of probe object A (star at A') appears to be 'ripped' away from neighboring object C and 'folded' onto object D, in the inertial reference frame. The first experiment demonstrated shifts in the remembered location of single probe objects, between groups. In the second experiment, we tested whether rips and folds occur within the spatial knowledge of one individual. Specifically, we tested the remembered configuration of a triad of objects by having participants take shortcuts from one start object to three target objects.

Figure 6 about here

Only the wormhole environment was tested in Experiment 2. Participants learned the maze as before. During the test phase, however, they took shortcuts from the start object to three types of target objects (Figure 6A): a *probe* (green dot) near a wormhole exit as before, a neutral *flanker_n* (blue dot) on one side of the probe, and a wormhole *flanker_w* (red-orange dot) on the other side, between the probe and its wormhole location. On half the trials, the start object was *near* the wormhole's entrance enabling a specific prediction for wormhole shortcuts; on the other half the start object was *far* from the entrance. If learning the wormhole maze creates rips and folds in spatial knowledge, then shortcuts to the probe (green vector) should shift away from those to flanker_n (blue vector) and toward those to flanker_w (red vector). Specifically, rips predict that constant errors to probe targets should be larger than those to flanker_n, and folds predict that constant errors to probes should be larger than those to flanker_w. Moreover, because flanker_w was in between the probe's Euclidean and wormhole directions, we expected that shortcuts to the probe would 'fold over' those to flanker_w – an ordinal reversal that could not be accommodated by a continuous metric map.

3.1 Material and Methods

Experiment 2 had a 3 (target type) x 2 (start object) x 4 (wormhole entrance) within-subject design. The methods were the same as for the wormhole group in Experiment 1, with the following exceptions.

3.1.1 Participants

A total of 11 participants (9 M, 2 F) completed the experiment. Two additional participants failed to achieve the learning criterion, and two others withdrew due to symptoms of motion sickness.

3.1.2 Procedure

Participants wore a Rockwell Collins SR80-A HMD (63° H x 53° V field of view, 1280 x 1024 pixels per eye, 100% binocular overlap). A ninth object (gear) was added to the ‘east’ side of the maze, in the corridor parallel to Wormhole 2. All participants learned the wormhole maze and performed the shortcut task in two one-hour sessions. The first session included the *exploration phase*, *training phase*, and a partial *test phase* (24 trials). The second session began with a refresher exploration phase (4 minutes), completed the test phase (48 trials), and ended with the questionnaire.

For each of the four wormhole entrances, there was a start object *near* and *far* from the entrance, crossed with a triad of target objects: a *probe* object near the wormhole exit, a neutral *flanker_n* on one side of the probe, and a wormhole *flanker_w* on the other side, between the probe’s Euclidean location and its wormhole location. There were three replications of these 24 conditions, yielding a total of 72 test trials, presented in a randomized order. Constant error was computed as before, with all responses normalized so that increasing errors were in the direction of the probe’s wormhole location.

3.2 Results and Discussion

As expected, the remembered direction of probe objects shifted away from *flanker_n* and toward *flanker_w* (Figure 6B), reflecting rips and folds in spatial knowledge. Yet only one participant reported noticing anything unusual about the wormhole environment.

3.2.1 Shortcuts

We focus on the near condition, which allowed a specific prediction about the expected direction of wormhole shortcuts. Shortcuts to the probe had a large mean constant error of 42.1° ($AD=21.5^\circ$) in the wormhole direction (positive error), close to the expected angle of 45° . This value was significantly greater than the mean constant error of shortcuts to flanker_n, which was 15.8° ($AD=12.0^\circ$), $F(1,20)=11.58$, $p=.003$, $d=1.51$. Thus, the remembered direction of the probe significantly shifted away from that of flanker_n by a mean of 26° evidence of a ‘rip’ in spatial knowledge. On the other side, constant error to probes was also significantly greater than the mean constant error to flanker_w, which was -21.5° ($AD=17.6^\circ$) in the opposite direction (negative error), $F(1,20)=51.98$, $p<.001$, $d=3.00$. This result indicates that the remembered direction of the probe significantly shifted toward that of flanker_w by 63.6° evidence of a ‘fold’. Of the total of 44 estimates (four triads x 11 participants), 86.4% exhibited rips and 90.9% exhibited folds (sign test, both $p<.001$).

Most importantly, shortcuts to the probe actually ‘folded over’ shortcuts to flanker_w by an average of 24.58° ($AD=25.74^\circ$), evidence of an ordinal reversal in remembered target locations. For example, in Figure 6 the mean probe vector (green) shifted away from the flanker_n mean vector (blue) and crossed over the flanker_w mean vector (red). Of the 44 estimates (4 triads x 11 participants), 65.9% of them exhibited ordinal reversals (binomial test, $p<.001$).

We directly compared the Euclidean map and labelled graph hypotheses by computing the probability of the observed triad data under each model, $p(x|\mu)$, assuming $\kappa=11.66$ (corresponding to $AD=17^\circ$, the triad mean). For each triad, the map model predicts shortcuts toward the three Euclidean target locations, yielding a mean expected error of $\mu_M=0^\circ$ whereas the graph model predicts shortcuts toward the Euclidean locations of the two flankers (0°) but the wormhole location of the probe (45°), yielding a mean

expected error of $\mu_G=15^\circ$. The Bayes Factor was large, $BF_{GM} > 100$, offering decisive support for the labelled graph model.

In the far condition the mean constant error to probe objects was 29.1° ($AD=21.5^\circ$) to flanker₀ was 25.1° ($AD=20.0^\circ$) and to flanker_w was -12.3° ($AD=16.0^\circ$). These values did not differ significantly from those in the near condition, indicating that the wormholes also distorted the remembered location of targets relative to start objects remote from a wormhole entrance.

The within-subject variability was similar for all three target types (mean $AD=43.1^\circ$ for flanker_n 46.2° for flanker_w and 42.6° for probes *ns*), implying comparable uncertainty in their remembered directions.

3.2.2 Partial replication of Experiment 1

The 'near' condition included a replication of the wormhole probe trials in Experiment 1. As noted above, shortcuts to the probe objects again exhibited a mean constant error of 42.1° ($AD=21.5^\circ$), significantly greater than zero, $CI_{99\%} \pm 19.3^\circ$ $p < .01$, $d=1.96$, and close to the expected value of 45° , corresponding to a wormhole bias of 94%. We compared the Euclidean map and labelled graph hypotheses for the probe data alone, assuming $\kappa=11.66$ (corresponding to $AD=21^\circ$, from the probe data). The mean expected error under the map model is $\mu_M=0^\circ$ whereas that under the graph model is $\mu_G=45^\circ$. The Bayes Factor was large, $BF_{GM} > 100$, decisive confirmation of the labelled graph model.

This pattern of responses also violated the positivity postulate. There were no errors in walking from home to the start objects, but shortcuts to these same objects were biased by an average of 42.1° . Under the Euclidean model the expected difference between opposite shortcuts is $\mu_M=0^\circ$, while under the labelled graph model it is $\mu_G=90^\circ$; the observed difference was 85.75° ($AD=41.57^\circ$). We compared the two models, assuming

$\kappa=2.224$ (corresponding to $AD=42^\circ$ from the present difference data, to be conservative). The Bayes factor was large, $BF_{GM} > 100$, decisive evidence against positivity and in favor of the labelled graph model.

3.2.3 *Self-reports*

Only one participant reported being aware of "wrinkles" in the virtual environment. The remaining ten participants failed to notice any discontinuities or other spatial anomalies.

3.2.4 *Conclusions*

Experiment 2 offers evidence of rips, folds, and ordinal reversals in spatial knowledge. The remembered locations of probe objects were 'ripped away' from flanker_n and 'folded over' flanker_w. Such local discontinuities are contrary to a Euclidean mental map, but are easily accommodated by a labelled graph. In addition, the results replicate the finding from Experiment 1 that shortcuts are strongly biased by wormhole experience, implying that spatial knowledge violates the metric postulates.

The possibility of "tears, folds, and holes" in knowledge of large-scale urban environments has been considered previously (Golledge & Hubert, 1982). However, these anomalies were attributed to inadequate sampling of the environment or missing information in what was assumed to be a continuous mental map. The present experiment finds evidence of rips, folds, and reversals in a densely sampled medium-scale environment, induced by exposure to wormholes. They appear to reflect local discontinuities rather than missing information, and imply that the fabric of spatial knowledge is not essentially continuous.

4.0 General Discussion

Based on the present findings, we question the view that humans construct a Euclidean cognitive map of navigation space. People do learn a network of paths between familiar places, which can be described as a topological graph. In addition, they learn approximate path lengths and angles, enabling them to select shorter routes and take rough shortcuts. But this local metric information does not appear to be integrated into a globally consistent metric map. Novel shortcuts are highly variable, biased by wormholes, manifest violations of the metric postulates, and reveal rips, folds, and ordinal reversals in spatial knowledge. Perhaps most surprising, people completely fail to detect radical violations of metric geometry in their environment, revealing a marked insensitivity to Euclidean structure.

4.1 A cognitive graph

We argue that this pattern of results is best accounted for by a labelled graph. As a navigator explores a new environment, they learn a network of paths that link salient places, which can be modeled by edges and nodes in a graph. At the same time, the path integrator registers piecewise measurements of path lengths and junction angles, which are expressed by edge weights and node labels. Due to the low resolution and discontinuity of human path integration (Loomis et al., 1993; Zhao & Warren, 2015a, 2015b), this metric information is noisy, biased by experience, and globally inconsistent. Edges that are learned in opposite directions may have different weights (ie. a *directed* graph), yielding violations of the symmetry postulate. Violations of the triangle inequality result from geometrically inconsistent edge weights. Both Euclidean and non-Euclidean environments are learned in the same manner, and spatial knowledge thus incorporates discontinuities and violations of the metric postulates, without the detection of inconsistencies.

The resulting cognitive graph supports route finding and novel detours with generally shorter paths, and enables approximate shortcuts. Shortcuts may be generated on the fly

by vector addition along a route through the graph, despite the absence of a common coordinate system or reference frame.³ However, their reliability and accuracy is limited by the level of noise and bias in local metric knowledge, which explains the influence of wormholes on participants' shortcuts.

In particular, this view can account for the violations of metric postulates observed in Experiment 1. Let's return to our initial example of the cactus and the bookcase (refer to Figure 4A). Suppose that, during exploration, a participant path-integrates from Home (H) to the cactus (A) and registers its path length (HA). Subsequently, they path-integrate from H to the bookcase (B) and register its path length (HB), as well as the approximate angle between the two paths ($\angle AHB \approx 90^\circ$). If the participant at B then turns 90° and walks through wormhole 1 the maze rotates, and they register a short path length (BA') before arriving at the cactus (visual place A') – which now has different inertial coordinates from A (Figure 4B). In principle, this provides enough information to detect the violation of positivity. But when the participant recognizes A' ("Aha – the cactus again!") the path integrator is reset (Etienne, Maurer, Boulens, Levy, & Rowe, 2004; Zhao & Warren, 2015b), so the inconsistency is not detected and the wormhole goes unnoticed. Subsequent shortcuts from B to A would be based on the resulting edge weights and node labels, which are biased by experience with the wormhole.

Other recent findings are consistent with this account. Purely visual exposure to a new environment enables successful route-finding and detours (Chrastil & Warren, 2015; Heft, 1983) but not novel shortcuts, which are barely above the chance level (Chrastil & Warren, 2013). This result demonstrates that humans can acquire graph knowledge, but not map knowledge, from a sequence of views. The addition of idiothetic information

³ Formally, vector addition can be performed in a coordinate-free space by iterative application of the parallelogram law and cosine and sine rules.

during learning enables the selection of shorter routes and detours (Chrastil & Warren, 2014) and improves shortcut accuracy, although shortcuts remain highly variable (Chrastil & Warren, 2013). These findings imply that people learn path lengths (edge weights) and junction angles (node labels) based on idiothetic path integration. After walking in an ‘impossible’ virtual world, such as a triangular corridor with two 90° corners, people can reproduce traversed distances and turn angles and find short routes⁴ (Kluss, Marsh, Zetsche, & Schill, 2015; Zetsche, Wolter, Galbraith, & Schill, 2009). They thus acquire local metric knowledge without detecting global inconsistency. These results lend support to a labelled graph with local metric information that is not integrated into a globally consistent map.

4.2 *Alternative explanations*

A number of alternative explanations might be offered for the present results. First and foremost, some theories propose that the navigation system has mechanisms that acquire both Euclidean and topological structure (P. Byrne et al., 2007; Chown et al., 1995; Trullier et al., 1997). Perhaps participants in the Euclidean maze learned a metric map, whereas those in the wormhole maze learned a labelled graph. Experiment 1 offers no support for two different mechanisms, however, as both the Euclidean and non-Euclidean groups had similar trials-to-criterion, route-finding performance, and shortcut variability. Moreover, even though wormhole participants were trained on the metric locations of targets, sufficient to build both a graph and a map, their shortcuts only reflected non-Euclidean spatial knowledge. We believe that the data are more parsimoniously explained by a single labelled graph structure.

Second, perhaps participants constructed a *distorted* cognitive map of the wormhole maze, but one with an intrinsic metric geometry nonetheless (Tobler, 1976). For

⁴ These authors did not test shortcuts.

instance, multi-dimensional scaling of distance estimates in urban environments yields distorted Euclidean maps⁵ (Golledge & Spector, 1978). In the wormhole maze, such distortions might result from attempting to assign visual places to inertial coordinates, as in the prevailing view (Gallistel & Cramer, 1996; McNaughton et al., 2005).

Specifically, Wormhole 1 shortcuts (Figure 3B) might be explained by a cognitive map that pinches and warps the 'northwest' corner of the maze so that A is congruent with A' and B is congruent with B'; the resulting map may be inaccurate, but it would be metric.

A distorted Euclidean map must still satisfy the metric postulates, however, such that each visual place occupies a unique location in the inertial coordinate system (positivity). To the contrary, shortcuts show that the same visual place is reported in two widely separated physical locations (A and A') in inertial coordinates. Similarly, paths in opposite directions through a wormhole deviate from the positivity postulate by 75° (Exp. 1) to 86° (Exp. 2) in the inertial coordinate frame. Averaging the two inertial locations would yield shortcuts mid-way between A and A', rather than the observed wormhole biases of 83% (Exp. 1) to 94% (Exp. 2). The present data thus cannot be accounted for by a distorted metric map in inertial coordinates.

Conversely, perhaps vision dominates spatial learning, and participants constructed a Euclidean map by assigning visual places to coordinates defined by visual odometry. This *purely visual map* must also satisfy the metric postulates, such that each place occupies a unique location in the visual coordinate system. Yet wormhole shortcuts violate positivity on this view as well, for the same object is reported at two inconsistent locations in coordinates given by visual odometry, as illustrated in Figure 3B. It is theoretically possible to produce a distorted visual map of the wormhole maze in some other coordinate system (e.g. by multi-dimensional scaling of visually perceived

⁵ Classical MDS assumes a Euclidean distance metric and produces a Euclidean solution.

distances). However, as Chrastil and Warren (2013) showed, purely visual exposure during learning does not yield metric map knowledge, for shortcuts are nearly random. Thus, the present data cannot be explained by a purely visual Euclidean map.

Finally, consider a more exotic hypothesis: wormhole participants may have constructed a metric map with a non-zero genus.⁶ Specifically, each wormhole might correspond to a short 'bridge' across a cusp in a two-dimensional surface, linking the entrance and exit locations. Wormhole 1, for example, would bridge over the path from home to the clock, while preserving other metric relations on the surface. However, the present sketch maps gave no indication of such bridges or path crossings in the plane. This remains a logical, but we believe unlikely, possibility.

4.3 Comparison with previous theories

Several previous theories have also combined topological and metric information in a single graph structure. Although the present experiments were not explicitly designed to distinguish these proposals, we can make some comparisons. Poucet (1993) first proposed a model in which places are described by local reference frames (metric coordinate systems), with vectors indicating the metric distance and direction to neighboring places (i.e. between local coordinate systems). In parallel, a topological network represents connections between places. With learning, the local vectorial and topological network information is integrated into a global representation with a common reference direction, enabling novel shortcuts. This process operates at two levels: linking places within a vista space into a local 'chart' and linking charts into an overall representation with a reference direction.

Meilinger (2008) proposed a related 'network of reference frames' theory in which a vista space is described by a local reference frame (i.e. a metric map of visible places,

⁶ Thanks to Matthew Collett for this suggestion.

analogous to a ‘chart’). Local frames are linked by edges that specify the ‘perspective shift’ (metric rotation and translation) required to move between them. These perspective shifts may become more precise with learning, but in contrast to Poucet (1993), local reference frames are not integrated into a consistent overall framework. Instead, shortcuts are planned by imagining a sequence of perspective shifts (coordinate transformations) from the current location to the goal location, so the goal is represented in the current reference frame.

Most recently, Mallot and Basten (2009) described a ‘space graph’ in which place nodes are assigned metric coordinates and are linked by action rules (“turn left 120° walk 2 km”) which implicitly express local metric information. The place nodes are, at least partially, embedded in a consistent global framework. Like Poucet (1993), the authors resist calling this global framework a metric ‘map’ because it does not explicitly represent all spatial positions.

While there are meaningful differences between these three theories, they share the premise that places are associated with local reference frames or coordinate systems, and try to solve the problem of novel shortcuts by representing separate places in a common coordinate frame. In contrast, the present cognitive graph theory argues, first, that spatial navigation does not require a common reference frame or coordinate system, for novel shortcuts can be generated from a labelled graph by vector addition (see footnote 3). A cognitive graph is a more parsimonious structure that calls upon less cognitive apparatus.

Second, whereas Poucet (1993) and Mallot and Basten (2009) propose a consistent overall representation or a global metric embedding, both the present theory and Meilinger’s (2008) network of reference frames maintain that spatial knowledge is not in general, geometrically consistent. Indeed, the observed violations of metric

postulates, and rips, folds, and reversals, imply that local metric information is not generally integrated into a consistent global coordinate system.

The present results bear on spatial learning in navigation space, when places are viewed individually and their relations are acquired via path integration. In contrast, there is evidence that an array of objects viewed together in vista space are remembered in one reference frame (Easton & Sholl, 1995; Rieser, 1989; Shelton & McNamara, 2001), and that objects in overlapping arrays or views are related in a common reference frame (Greenauer & Waller, 2010; Kelly & McNamara, 2010; Mou, McNamara, & Zhang, 2013). It has been proposed that this principle generalizes across environmental scales (Greenauer & Waller, 2010). However, recent research indicates that vista space and navigation space differ markedly in this respect. Whereas an array of objects in vista space is remembered in a single reference frame, the same array in navigation space is not (Meilinger, Strickrodt, & Bühlhoff, 2016). Moreover, evidence for a common reference frame does not entail a metric Euclidean map. Even within vista space, the visual perception of layout does not preserve Euclidean structure (Koenderink, van Doorn, & Lappin, 2000; Loomis, da Silva, Philbeck, & Fukushima, 1996; Norman, Todd, Perotti, & Tittle, 1996; Wagner, 2006); the geometric structure of memory for vista space is unknown.

In summary, the present findings cast doubt on the view that the navigation system builds a Euclidean cognitive map, which guides much research on human and animal navigation. They are also inconsistent with alternative explanations that seek to preserve the construct of a metric map, as well as with purely topological knowledge. We conclude that knowledge of navigation space is better described by a labelled graph, in which local metric information is approximate, geometrically inconsistent, and not embedded in a common coordinate system. This class of ‘cognitive graph’ models

supports route finding, novel detours, and rough shortcuts, and has the potential to unify a range of behavioral and neurophysiological data on spatial navigation.

Supplementary Material

The data supporting this article may be accessed from the Brown University Digital Repository (<https://repository.library.brown.edu/studio/item/bdr:581513/>), doi:10.7301/Z0JS9NC5.

Acknowledgments

This work was supported by the National Science Foundation, United States (BCS-0214383 and BCS-0843940). Thanks to Matthew Van Wormer, Aaron Mandle, Katherine Costa, and Joost de Nijs for their assistance.

Figure Captions

Figure 1. Models of spatial knowledge. (A) Euclidean map: places are assigned locations in a common coordinate system. (B) Topological graph: nodes correspond to places and edges to paths between them. (C) Labelled graph: edge weights denote approximate path lengths and node labels denote approximate junction angles.

Figure 2. The virtual hedge maze. (A) Euclidean environment. (B) Non-Euclidean environment with two wormholes (yellow arrows). Object pairs are represented by dots of same color. Red traces represent all probe trials for the route task in Experiment 1, plotted in the visual reference frame.

Figure 3. Shortcuts for Wormhole 1. (A) In the Euclidean environment, participants took shortcuts from the start object in the approximate direction of the target object (yellow dots, A and B), in both directions. (B) In the non-Euclidean environment, shortcuts are shifted toward the 'wormhole direction' of the target (yellow stars A' and B'). Red traces represent all probe trials for the cactus-bookcase pair in Experiment 1, plotted in the inertial reference frame.

Figure 4. (A) Violation of metric postulates: Participant walks successfully from H to A, and from H to B, yet takes shortcut from B toward A' (axes represent inertial coordinates). Positivity is violated ($AA' > 0$) as is the triangle inequality (ΔHA is not closed). The remembered location of A (e.g. the star at A') is 'ripped' away from neighbor C and 'folded' onto neighbor D. (B) If participant walks from B through Wormhole 1, the maze rotates, and they arrive at visual place A' with different inertial

coordinates from A.

Figure 5. Sample sketch maps from (a) the Euclidean control group, and (b) the wormhole group. Note ordinal reversals in the cyclic order of objects: the bookcase and clock switch positions, as do the sink and moon.

Figure 6. Rips, folds, and ordinal reversals for a triad of targets in Wormhole 2. (A) Predicted Euclidean shortcuts from a near start object (well) to the probe (sink, green), flanker_n (moon, blue), and flanker_w (gear, red). (B) Observed shortcuts to probe (green traces and mean vector) are 'ripped' away from flanker_n (blue) and 'folded' over flanker_w (red). Yellow star represents the wormhole location of the probe. Traces represent all test trials from the well in Experiment 2.

References

- Batschelet, E. (1981). *Circular Statistics in Biology*. London: Academic Press, Inc.
- Beals, R., Krantz, D. H., & Tversky, A. (1968). Foundations of multidimensional scaling. *Psychological Review*, 75, 127-142.
- Benhamou, S. (1996). No evidence for cognitive mapping in rats. *Animal Behaviour*, 52, 201-212.
- Bennett, A. T. D. (1996). Do animals have cognitive maps? *Journal of Experimental Biology*, 199(1), 219-224.
- Burroughs, W. J., & Sadalla, E. K. (1979). Asymmetries in distance cognition. *Geographical Analysis*, 11, 414-421.
- Byrne, P., Becker, S., & Burgess, N. (2007). Remembering the past and imagining the future: a neural model of spatial memory and imagery. *Psychological Review*, 114, 340-375.
- Byrne, R. W. (1979). Memory for urban geography. *Quarterly Journal of Experimental Psychology*, 31, 147-154.
- Cadwallader, M. T. (1979). Problems in cognitive distance: Implications for cognitive mapping. *Environment and Behavior*, 11, 559-576.
- Chapuis, N., Durup, M., & Thinus-Blanc, C. (1987). The role of exploratory experience in a shortcut task by golden hamsters (*Mesocricetus auratus*). *Animal Learning and Behavior*, 15, 174-178.
- Cheeseman, J. F., Millar, C. D., Greggers, U., Lehmann, K., Pawley, M. D. M., Gallistel, C. R., . . . Menzel, R. (2014). Way-finding in displaced clock-shifted bees proves bees use a cognitive map. *Proceedings of the National Academy of Sciences*, 111(24), 8949-8954.

- Cheung, A., Collett, M., Collett, T. S., Dewar, A., Dyer, F., Graham, P., . . . Stürzl, W. (2014). Still no convincing evidence for cognitive map use by honeybees. *Proceedings of the National Academy of Sciences*, *111*(42), E4396-E4397.
- Chown, E., Kaplan, S., & Kortenkamp, D. (1995). Prototypes location, and associative networks (PLAN): Towards a unified theory of cognitive mapping. *Cognitive Science*, *19*, 1-51.
- Chrastil, E. R., & Warren, W. H. (2013). Active and passive spatial learning in human navigation: Acquisition of survey knowledge. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, *39*(5), 1520.
- Chrastil, E. R., & Warren, W. H. (2014). From cognitive maps to cognitive graphs. *PLoS One*, *9*(11), e112544.
- Chrastil, E. R., & Warren, W. H. (2015). Active and passive spatial learning in human navigation: Acquisition of graph knowledge. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, *41*, 1162-1178.
- Dabaghian, Y., Brandt, V. L., & Frank, L. M. (2014). Reconceiving the hippocampal map as a topological template. *eLife*, *3*, e03476.
- Dabaghian, Y., Mémoli, F., Frank, L. M., & Carlsson, G. (2012). A topological paradigm for hippocampal spatial map formation using persistent homology. *PLoS Computational Biology*, *8*(8), e1002581.
- Derdikman, D., & Moser, E. I. (2010). A manifold of spatial maps in the brain. *Trends in Cognitive Sciences*, *14*(12), 561-569.
- Dyer, F. (1991). Bees acquire route-based memories but not cognitive maps in a familiar landscape. *Animal Behaviour*, *41*, 239-246.

- Easton, R. D., & Sholl, M. J. (1995). Object-array structure, frames of reference, and retrieval of spatial knowledge. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, 21(2), 483.
- Epstein, R. A., & Vass, L. K. (2014). Neural systems for landmark-based wayfinding in humans. *Philosophical Transactions of the Royal Society B*, 369, 20120533.
- Etienne, A. S., Maurer, R., Boulens, V., Levy, A., & Rowe, T. (2004). Resetting the path integrator: A basic condition for route-based navigation. *Journal of Experimental Biology*, 207(1491-1508).
- Foo, P., Warren, W. H., Duchon, A., & Tarr, M. (2005). Do humans integrate routes into a cognitive map? Map- vs. landmark-based navigation of novel shortcuts. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, 31, 195-215.
- Gallistel, C. R. (1990). *The organization of learning*. Cambridge, MA: The MIT Press.
- Gallistel, C. R., & Cramer, A. E. (1996). Computations on metric maps in mammals: getting oriented and choosing a multi-destination route. *Journal of Experimental Biology*, 199, 211-217.
- Golledge, R. G., & Hubert, L. J. (1982). Some comments on non-euclidean cognitive maps. *Environment and Planning A*, 14, 107-118.
- Golledge, R. G., Ruggles, A. J., Pellegrino, J. W., & Gale, N. D. (1993). Integrating route knowledge in an unfamiliar neighborhood: Along and across route experiments. *Journal of Environmental Psychology*, 13(4), 293-307.
- Golledge, R. G., & Spector, A. N. (1978). Comprehending the urban environment: Theory and practice. *Geographical Analysis*, 10(4), 403-426.

- Gould, J. L. (1986). The locale map of honey bees: Do insects have cognitive maps? *Science*, 232, 861-863.
- Greenauer, N., & Waller, D. (2010). Micro-and macreference frames: Specifying the relations between spatial categories in memory. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, 36(4), 938.
- Heft, H. (1983). Way-finding as the perception of information over time. *Population and Environment*, 6(3), 133-150.
- Holmes, M. C., & Sholl, M. J. (2005). Allocentric coding of object-to-object relations in overlearned and novel environments. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, 31, 1069-1087.
- Hübner, W., & Mallot, H. A. (2007). Metric embedding of view graphs. A vision and odometry-based approach to cognitive mapping. *Autonomous Robots*, 23, 183-196.
- Ishikawa, T., & Montello, D. R. (2006). Spatial knowledge acquisition from direct experience in the environment: Individual differences in the development of metric knowledge and the integration of separately learned places. *Cognitive Psychology*, 52, 93-129.
- Jeffreys, H. (1998). *Theory of probability* (3 ed.): Oxford: Oxford University Press.
- Kearns, M. J., Warren, W. H., Duchon, A. P., & Tarr, M. (2002). Path integration from optic flow and body senses in a homing task. *Perception*, 31, 349-374.
- Kelly, J. W., & McNamara, T. P. (2010). Reference frames during the acquisition and development of spatial memories. *Cognition*, 116(3), 409-420.
- Kluss, T., Marsh, W. E., Zetsche, C., & Schill, K. (2015). Representation of impossible worlds in the cognitive map. *Cognitive Processing*, 1-6.

- Koenderink, J. J., van Doorn, A. J., & Lappin, J. S. (2000). Direct measurement of the curvature of visual space. *Perception*, 29, 69-79.
- Kosslyn, S. M., Pick, H. L., & Fariello, G. R. (1974). Cognitive maps in children and men. *Child Development*, 707-716.
- Kuipers, B. (2000). The spatial semantic hierarchy. *Artificial Intelligence*, 119, 191-233.
- Kuipers, B., Tecuci, D. G., & Stankiewicz, B. J. (2003). The skeleton in the cognitive map: A computational and empirical exploration. *Environment and Behavior*, 35, 81-106.
- Loomis, J. M., da Silva, J. A., Philbeck, J. W., & Fukushima, S. S. (1996). Visual perception of location and distance. *Current Directions in Psychological Science*, 3, 72-77.
- Loomis, J. M., Klatzky, R. L., Golledge, R. G., Cicinelli, J. G., Pellegrino, J. W., & Fry, P. A. (1993). Nonvisual navigation by blind and sighted: assessment of path integration ability. *Journal of Experimental Psychology: General*, 122(1), 73-91.
- Mallot, H. A., & Basten, K. (2009). Embodied spatial cognition: Biological and artificial systems. *Image and Vision Computing*, 27(11), 1658-1670.
- McNamara, T. P., & Diwadkar, V. A. (1997). Symmetry and asymmetry of human spatial memory. *Cognitive Psychology*, 34, 160-190.
- McNaughton, B. L., Battaglia, F. P., Jensen, O., Moser, E. I., & Moser, M.-B. (2005). Path integration and the neural basis of the 'cognitive map'. *Nature Reviews: Neuroscience*, 7, 663-678.
- Meilinger, T. (2008). The network of reference frames theory: a synthesis of graphs and cognitive maps. In C. Freksa, N. S. Newcombe, P. Gärdénfors, & S. Wölfl (Eds.), *Spatial Cognition VI* (pp. 344-360). Berlin: Springer.


- Meilinger, T., Riecke, B. E., & Bühlhoff, H. H. (2014). Local and global reference frames for environmental spaces. *The Quarterly Journal of Experimental Psychology*, 67, 542-569.
- Meilinger, T., Strickrodt, M., & Bühlhoff, H. H. (2016). Qualitative differences in memory for vista and environmental spaces are caused by opaque borders, not movement or successive presentation. *Cognition*, 155, 77-95.
- Menzel, R., Greggers, U., Smith, A., Berger, S., Brandt, R., Brunke, S., . . . Watzl, S. (2006). Honey bees navigate according to a map-like spatial memory. *Proceedings of the National Academy of Sciences*, 102, 3040-3045.
- Moeser, S. D. (1988). Cognitive mapping in a complex building. *Environment and Behavior*, 20, 21-49.
- Mohler, B. J., Creem-Regehr, S. H., & Thompson, W. B. (2006). The influence of feedback on egocentric distance judgments in real and virtual environments. *Proceedings of the 3rd symposium on applied perception in graphics and visualization, ACM*, 9-14.
- Mou, W., McNamara, T. P., & Zhang, L. (2013). Global frames of reference organize configural knowledge of paths. *Cognition*, 129(1), 180-193.
- Muller, R. U., & Kubie, J. L. (1987). The effects of changes in the environment on the spatial firing of hippocampal complex-spike cells. *Journal of Neuroscience*, 7, 1951-1968.
- Muller, R. U., Stead, M., & Pach, J. (1996). The hippocampus as a cognitive graph. *The Journal of General Physiology*, 107(6), 663-694.
- Nadel, L. E. (2013). Cognitive maps. In D. E. Waller & L. E. Nadel (Eds.), *Handbook of spatial cognition*. Washington, D.C.: American Psychological Association.

- Norman, J. F., Todd, J. T., Perotti, V. J., & Tittle, J. S. (1996). The visual perception of three-dimensional length. *Journal of Experimental Psychology: Human Perception and Performance*, 22, 173-186.
- O'Keefe, J., & Burgess, N. (1996). Geometric determinants of the place fields of hippocampal neurons. *Nature*, 381, 425-428.
- O'Keefe, J., & Nadel, L. E. (1978). *The hippocampus as a cognitive map*. Oxford: Clarendon Press.
- Poucet, B. (1993). Spatial cognitive maps in animals: New hypotheses on their structure and neural mechanisms. *Psychological Review*, 100(2), 163-183.
- Rieser, J. J. (1989). Access to knowledge of spatial structure at novel points of observation. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, 15(6), 1157-1165.
- Rouder, J. N., Speckman, P. L., Sun, D., Morey, R. D., & Iverson, G. (2009). Bayesian tests for accepting and rejecting the null hypothesis. *Psychonomic Bulletin & Review*, 16(2), 225-237.
- Sadalla, E. K., & Magel, S. G. (1980). The perception of traversed distance. *Environment and Behavior*, 12, 65-79.
- Sadalla, E. K., & Staplin, L. J. (1980). The perception of traversed distance: Intersections. *Environment and Behavior*, 12, 167-182.
- Schinazi, V. R., Nardi, D., Newcombe, N. S., Shipley, T. F., & Epstein, R. A. (2013). Hippocampal size predicts rapid learning of a cognitive map in humans. *Hippocampus*, 23(6), 515-528.
- Shelton, A. L., & McNamara, T. P. (2001). Systems of spatial reference in human memory. *Cognitive Psychology*, 43(4), 274-310.


- Thrun, S. (1998). Learning metric-topological maps for indoor mobile robot navigation. *Artificial Intelligence*, 99(1), 21-71.
- Tobler, W. R. (1976). The geometry of mental maps. In R. G. Golledge & G. Rushton (Eds.), *Spatial choice and spatial behavior* (pp. 69-82). Columbus, OH: Ohio State University Press.
- Tolman, E. C. (1948). Cognitive maps in rats and men. *Psychological Review*, 55, 189-208.
- Trullier, O., Wiener, S. I., Berthoz, A., & Meyer, J.-A. (1997). Biologically based artificial navigation systems: Review and prospects. *Progress in Neurobiology*, 51, 483-544.
- Tversky, B. (1992). Distortions in cognitive maps. *Geoforum*, 23(2), 131-138.
- Tversky, B. (1993). Cognitive maps, cognitive collages, and spatial mental models *Spatial Information Theory A Theoretical Basis for GIS* (pp. 14-24): Springer.
- Wagner, M. (2006). *The geometries of visual space*: Psychology Press.
- Waller, D., & Greenauer, N. (2007). The role of body-based sensory information in the acquisition of enduring spatial representations. *Psychological Research*, 71(3), 322-332.
- Weisberg, S. M., Schinazi, V. R., Newcombe, N. S., Shipley, T. F., & Epstein, R. A. (2014). Variations in cognitive maps: Understanding individual differences in navigation. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, 40(3), 669.
- erner S. r ieg- r ckner, B., & Herrmann, T. (2000). Modelling navigational knowledge by route graphs. In C. Habel, W. Brauer, C. Freksa, & K. F. Wender

- (Eds.), *Lecture Notes in Computer Science: Spatial Cognition 2000* (Vol. 1849, pp. 295-316). Heidelberg: Springer.
- Wiener, J. M., & Mallot, H. A. (2003). 'Fine-to-coarse' route planning and navigation in regionalized environments. *Spatial Cognition and Computation*, 3, 331-358.
- Zetsche, C., Wolter, J., Galbraith, C., & Schill, K. (2009). Representation of space: image-like or sensorimotor? *Spatial Vision*, 22(5), 409-424.
- Zhao, M., & Warren, W. H. (2015a). Environmental stability modulates the role of path integration in human navigation. *Cognition*, 142, 96-109.
- Zhao, M., & Warren, W. H. (2015b). How you get there from here: Interaction of visual landmarks and path integration in human navigation. *Psychological science*, 26, 915-924.


Figure(s)



B




C



A





B





Figure(s)
[Click here to download high resolution image](#)



A





B




Figure(s)
[Click here to download high resolution image](#)





A



B

