

# **Microsnapshot Navigation in a Non-Euclidean Environment**

## **Embedding of Feature-Graphs**

BACHELOR THESIS IN COGNITIVE SCIENCE

submitted by

**KATJA KÖRNER**

September 2020

**Supervisor: Prof. Dr. Hanspeter Mallot**

Matrikelnummer: 4111607

Writing period: 01.05.20 - 01.09.20

**COGNITIVE NEUROSCIENCE**

DEPT. OF BIOLOGY, FACULTY OF SCIENCE

EBERHARD KARLS UNIVERSITY TÜBINGEN

EBERHARD KARLS  
UNIVERSITÄT  
TÜBINGEN



## **ERKLÄRUNG**

Hiermit versichere ich, dass ich die Arbeit selbständig verfasst habe, keine anderen als die angegebenen Hilfsmittel und Quellen benutzt habe, alle wörtlich oder sinngemäß aus anderen Werken übernommenen Aussagen als solche gekennzeichnet sind und dass die Arbeit weder vollständig noch in wesentlichen Teilen Gegenstand eines anderen Prüfungsverfahrens gewesen ist.

Außerdem versichere ich, dass ich die Arbeit weder vollständig noch in wesentlichen Teilen bereits veröffentlicht habe sowie dass das in Dateiform eingereichte Exemplar mit den eingereichten gebundenen Exemplaren übereinstimmt.

Datum, Ort, Unterschrift

---

Katja Körner

## ZUSAMMENFASSUNG

Räumliches Wissen ist für uns unerlässlich, um sich fortbewegen und orientieren zu können ohne auf unmittelbar wahrnehmbare Hinweise beschränkt zu sein. Obwohl wir uns dauernd in unserem alltäglichen Leben darauf verlassen, ist es immer noch unklar, wie Menschen räumliches Wissen erwerben und organisieren. Während der Untersuchung dieser Frage kamen mehrere Hypothesen auf, welche verschiedene Modelle vorschlagen, um zu erklären wie der biologische Ansatz zur Integration von räumlichen Erfahrungen gestaltet sein könnte. Aktuelle Erkenntnisse weisen darauf hin, dass ein Graph mit Labeln die wahrscheinlichste Struktur für räumliches Wissen ist. Dies wird durch eine Studie (Warren *et al.*, 2017) untermauert, bei der Versuchspersonen mit einer nicht-euklidischen Umgebung konfrontiert wurden. Die Ergebnisse zeigten lokale Inkonsistenzen bei der Navigation durch die Umgebung sowie räumliche Verzerrungen in der internen Repräsentation, die sich am besten mit der Hypothese eines Graphen mit Labeln erklären lassen. Die vorliegende Studie simuliert dieses Experiment mit einem topologischen Navigationsansatz, der einen Graphen mit Labeln modelliert. Nicht-euklidische Umgebungen wurden mit Hilfe von Wurmlöchern in einem Virtual-Reality-Labyrinth implementiert. Im Verlauf der Simulation wurden Graphen-Repräsentationen sowohl einer nicht-euklidischer als auch einer passenden euklidischen Umgebung erstellt. Dies geschah ausschließlich durch Extrahierung und Wiedererkennen visueller Merkmale. Darüber hinaus wird ein Modell zur Einbettung solcher Graphen vorgeschlagen, das 'Multidimensional Scaling' und Prokrustes-Transformation verwendet. Die Ergebnisse der eingebetteten Graphen zeigen eine Repräsentation der Umgebung, die mit Beobachtungen bei Menschen aus dem Warren-Experiment vergleichbar ist. Insbesondere können relative Objektpositionen allein durch die Einbettung von Graphen mit 'Multidimensional Scaling' reproduziert werden, während konkrete räumliche Verzerrungen mit zusätzlicher Prokrustes-Transformation, welche Eigenbewegungs-Informationen modelliert, reproduziert werden können. Aus den Ergebnissen wird gefolgert, dass ein Pfadintegrationssystem, welches metrische Informationen aus der Eigenbewegung liefert, für die Einbettung von Graphen basierend auf visuellen Merkmalen nützlich sein könnte. Allerdings sollte diese Art von Information auch nur eine unterstützende Rolle spielen, indem sie der visuell-basierten Einbettung untergeordnet wird. Darüber hinaus liefert die Replikation der Ergebnisse aus dem Experiment an Menschen mit dem hier vorgestellten graphenbasierten Ansatz weitere Belege für die Hypothese, dass räumliches Wissen beim Menschen wahrscheinlich mittels einer Graphenstruktur mit Labeln organisiert sein könnte.

## ABSTRACT

**KEYWORDS:** spatial knowledge; labelled graph hypothesis; feature-graphs; graph embedding; Microsnapshot navigation; Multidimensional Scaling, Procrustes transformation.

Spatial knowledge is essential for travelling and orientating ourselves beyond the restrictions of our immediate perception. But as much as we rely on it in our daily lives, it is still not fully clear how humans acquire and organise spatial knowledge. In the course of investigating this question, several hypotheses came up proposing different models for the biological approach to integrate experiences of space. Current research suggests a labelled graph as the most likely underlying structure for spatial knowledge. This is substantiated by a study conducted by Warren *et al.* (2017) which confronted participants with a non-Euclidean environment. Results revealed local biases in shortcut tasks as well as spatial distortions in the internal representation of humans which can be best explained by the labelled graph hypothesis assuming a graph with local metric information as the structure of spatial knowledge. The present study simulates this experiment with a topological navigation approach modelling a labelled graph. Non-Euclidean environments were implemented by means of wormholes in a virtual reality maze environment. During the course of the simulation, graph representations of non-Euclidean as well as matching Euclidean environments were solely built based on extracting and recognising visual features. In addition, a model for embedding such graphs is proposed employing Multidimensional Scaling and Procrustes transformation. Results on embedded graphs show a representation of the environment comparable to observations from the Warren experiment conducted on humans. In particular, relative object positions can be reproduced solely based on a Multidimensional Scaling embedding of graph representations while shortcut biases are replicated with additional Procrustes transformation modelling self-motion cues. It is inferred from the results that a path integration system providing metric information from self-motion could be useful for embedding feature graphs but should also only play a supporting role by being subordinated to visual-based embedding. Furthermore, in replicating the results from an experiment on humans with the presented graph-based approach, this provides further evidence for substantiating the hypothesis that spatial knowledge in humans is most likely organised using a labelled graph structure.

# TABLE OF CONTENTS

<b>CHAPTER 1: INTRODUCTION</b> . . . . .	1
1.1 Spatial Knowledge . . . . .	1
1.1.1 Navigation . . . . .	1
1.1.2 Organisation of Spatial Knowledge . . . . .	2
1.1.3 Maps versus Graphs: Hypotheses for Spatial Knowledge . . . . .	3
1.2 Non-Euclidean Wormhole Experiments . . . . .	6
1.2.1 Methods and Predictions . . . . .	7
1.2.2 Results . . . . .	9
1.2.3 Conclusion and Implications . . . . .	10
1.2.4 Related Work . . . . .	11
<b>CHAPTER 2: METHODS</b> . . . . .	12
2.1 Microsnapshot Navigation . . . . .	12
2.2 Simulating Wormhole Experiments . . . . .	13
2.3 Graph Embedding and Visualisation . . . . .	14
2.3.1 Multidimensional Scaling . . . . .	15
2.3.2 Procrustes Transformation . . . . .	17
2.4 The Present Study . . . . .	17
<b>CHAPTER 3: SIMULATION</b> . . . . .	19
3.1 Simulation Environment . . . . .	19
3.2 Graph Processing Algorithm . . . . .	20
3.3 Simulation Procedure . . . . .	22
<b>CHAPTER 4: RESULTS</b> . . . . .	26
4.1 Graph Reduction . . . . .	26
4.2 Route Knowledge . . . . .	27
4.3 Survey Knowledge . . . . .	28
4.3.1 Mapping . . . . .	28
4.3.2 Pointing . . . . .	33
4.3.3 Folds in spatial Representations . . . . .	35
<b>CHAPTER 5: DISCUSSION</b> . . . . .	37
5.1 Evaluation of Results . . . . .	37
5.1.1 Summary of Results . . . . .	37
5.1.2 Interpretation of Results . . . . .	38
5.2 Limitations and possible Extensions . . . . .	40
5.3 Conclusion . . . . .	43
<b>APPENDIX A: FIGURES RESULTS</b> . . . . .	45
<b>APPENDIX B: SUPPLEMENTARY INFORMATION</b> . . . . .	52
B.1 Abbreviations . . . . .	52
B.2 Software references . . . . .	52
B.2.1 Unity . . . . .	52
B.2.2 Visual Studio . . . . .	52
B.2.3 MATLAB . . . . .	52
B.2.4 Code Libraries . . . . .	53
B.3 Image references . . . . .	53
<b>LIST OF TABLES</b> . . . . .	55
<b>LIST OF FIGURES</b> . . . . .	55
<b>REFERENCES</b> . . . . .	60

# CHAPTER 1

## INTRODUCTION

### 1.1 Spatial Knowledge

Motion and perception constitute fundamental pillars for interacting with the world around us. While motion entails a loss of information, perception provides a gain of information (Thrun *et al.*, 2005). In order to find a balance, it is therefore essential that motion and perception constantly interact with one another. This becomes especially apparent when exploring our spatial surroundings. For example, trying to orientate oneself in a new environment requires to physically move through it while simultaneously using perception to both get an understanding of the surroundings' layout as well as localising oneself within it. In robotics, this is known as the SLAM (Simultaneous Localization And Mapping) problem (Milford and Wyeth, 2008). However, perception itself just provides us with instant local information in our current view and therefore only enables loco-motor control (Warren, 2019). In order to expand this local knowledge, i.e. not to be limited to instantaneous feedback from perception, experiences need to be stored and organised in a way so that useful global spatial knowledge for solving e.g. the SLAM problem is available. *Navigation* (as described in Warren (2019)) can then be defined as the process which draws on this knowledge and with which we can plan, travel and orientate ourselves beyond the restrictions of immediate perception.

#### 1.1.1 Navigation

Navigating an environment poses multiple tasks with the main ones being: 1. Keeping track of where you are (localising), 2. Returning to your start location (homing), 3. Re-traversing known routes and 4. Finding new routes and shortcuts (Warren, 2019). In research it is tried to extrapolate from navigational performances to underlying strategies required to accomplish these tasks. Tasks 2 (homing) and 3 (re-traversing) for example can be reduced to strategies like learning a sequence of place-action-place associations based on recognizing environmental features (Martinet *et al.*, 2011).

Conversely, task 1 (localising oneself) and 4 (finding shortcuts) cannot be explained by these strategies but require more advanced navigational mechanisms. Theories about the nature of these navigational strategies propose a key role of path integration and vector addition. *Path integration* is the process of maintaining an estimate for one's location and orientation relative to a starting position. This entails obtaining any changes in position and rotation and updating the estimate accordingly (Warren, 2019; Strickrodt, 2019). Changes in position and rotation can either be derived directly from e.g. propriocep-

tion or by measuring self-motion relative to a starting point. The path integration system thus maintains estimates of travelled distances and angles which can be described as vectors and enable self-localisation. It is assumed that for every visited place a vector pointing towards it is stored, in particular this could include a 'home vector' pointing towards the starting location (Strickrodt, 2019). One theory is that during exploration or following a path, the set of stored vectors is constantly updated and synchronized e.g. upon rediscovery of places (Strickrodt, 2019). Vector addition of stored vectors then enables the computation of a direct return path (Warren, 2019) and to find shortcuts to a specific goal location. Even though the path integration system provides a solid base for explaining navigation behaviour, it only utilises piece wise local measurements of distances and angles as they are experienced (Warren, 2019). Besides this, vector estimates based on internal measurements are far from being accurate and prone to errors which accumulate with every change of position and orientation. To compensate for an erroneous path integration system, there needs to be some form of global embedding of local measurements (Strickrodt, 2019) to generate a meaningful representation of the experienced environment. This raises the question how spatial knowledge is organised and what underlying structure is used to integrate our experiences of the environment.

### 1.1.2 Organisation of Spatial Knowledge

As a current subject of research, the question how biological spatial knowledge is organised brought forth multiple theories about the internal structure supposed to enable representing and navigating our environment. In order to grasp the different hypotheses, it can be helpful to first have a look at externalised representations of space, i.e. maps.

Humans have put a lot of thought into how to most usefully represent spatial knowledge for the purpose of navigation. The first maps are assumed to date back to prehistoric times and since then, cartography has developed considerably. While prehistoric rock drawings focus on salient landmarks, historic maps from the 19th century try to reproduce the world as detailed as possible with the aspiration for scientific accuracy (Fig. 1.1a). Although exact coordinates can be read from them, historic scientific maps also slightly distort reality reflecting the human point of view. This is shown in the fact that even though the earth has a spherical shape, humans employ rules of 'flat' Euclidean geometry in order to reduce spatial representation to a two-dimensional map (Borg and Groenen, 2005).

Nowadays, there exist also more abstract kinds of maps e.g. for public transportation systems (Fig. 1.1b). These maps lack exact spatial metrics like distance or direction but only depict a topological network of connected places with instructional information like which actions to take when travelling between shown places (Baumann, 2019). A network like this can also be called a graph with nodes that represent places and that are connected by edges labelled with action instructions. While involving huge topographic distortions, modern abstract maps are still well applicable for navigation.

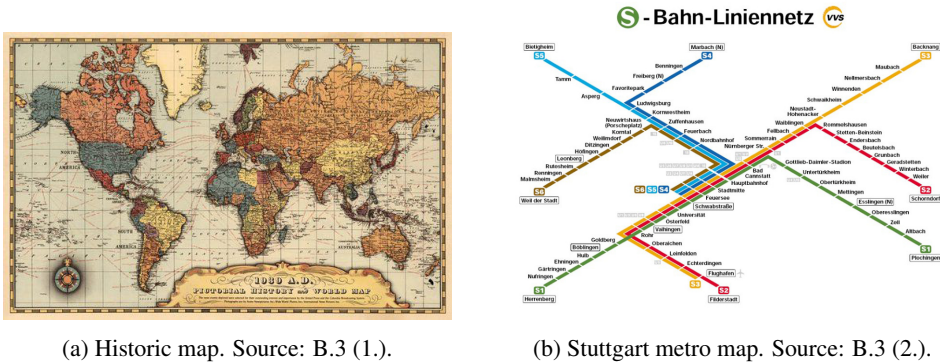


Fig. 1.1: Maps: two types of spatial representations

Just like maps created by humans can be divided into two categories, Euclidean and (labelled) topological, so can the main hypotheses about the internal organisation of spatial knowledge. More precisely, one position is proposing an Euclidean cognitive map as the underlying structure for spatial knowledge while others argue for a graph-like structure with similarities to public transportation maps. Therefore, the ongoing investigation is to determine whether one of these approaches and, if so, which one is able to explain navigational performances and is put into practice in our daily lives.

### 1.1.3 Maps versus Graphs: Hypotheses for Spatial Knowledge

The following sections give an overview of the different hypotheses describing possible structures for spatial knowledge. They differ in the nature of the structure for embedding spatial knowledge and in the degree of metric information included in the representation. Figure 1.2 presents the three models of spatial knowledge considered in the following.




Fig. 1.2: Models of spatial knowledge. Figure and description directly adopted from Warren *et al.* (2017). (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.



### Euclidean Cognitive Map

The *Euclidean cognitive map* can be thought of as being closely related to the concept of a historic map. It describes an allocentric Euclidean coordinate system serving as a fixed grid in which places are embedded by assigning coordinates to them. It is assumed that a path integration system computes local displacement measurements which are then combined through vector addition to yield the respective coordinates (Warren, 2019). Given arbitrary points A, B, C in space, the definition of an Euclidean coordinate system involves certain metric postulates that must be satisfied (Warren *et al.*, 2017):

- Positivity ( $\overline{AA} = 0, \overline{AB} > 0$ )
- Symmetry ( $\overline{AB} = \overline{BA}$ )
- Triangle inequality ( $\overline{AB} + \overline{BC} \geq \overline{AC}$ )

The insertion of places according to their estimated coordinates has to follow these metric postulates. That way, spatial representation is built upon an a priori general structure, the grid, while only the position of places within this structure is based on experience. A powerful property of this concept is that all estimated geometric relationships between represented places are preserved and exactly defined (Gallistel, 1990). This enables flexible navigation behaviour by deriving movement instructions from metric relations in the coordinate system based on vector algebra (Warren, 2019).

Biological support for the Euclidean cognitive map hypothesis is argued to be found in the discovery of place, grid, and head-direction cells in the hippocampus (Fyhn *et al.*, 2004) encoding a metric grid (Derdikman and Moser, 2010). Grid cells can be considered as representing a coordinate system while place cell populations encoding specific locations constitute entries in the grid and head-direction cells provide global directional information Knierim *et al.* (1995).

However, the fact that maps make sense of the world by imposing an Euclidean geometry onto it does not imply that internal mapping mechanisms are also based on Euclidean principles. In fact, empirical studies on humans rather contradict the assumption of a strictly Euclidean representation by showing large errors in direction estimates (Chrastil and Warren, 2013; Foo *et al.*, 2005) and distance estimates biased by visual landscape properties (McNamara, 1986; Sadalla and Staplin, 1980). Based on that evidence, it could still be argued for an imprecise, distorted Euclidean cognitive map. However, researchers also found severe violations of metric postulates which are fundamental for the Euclidean structure of such a map. For one thing, symmetry seems to be violated as participants show asymmetric distance estimates between a pair of points dependent on the target point (Moar and Bower, 1983). Furthermore, distance estimates appear to be biased by the number of junctions and other geometric features along a path which contradicts the postulate of additivity (Sadalla and Staplin, 1980). Finally, an observed tendency to assume intersections to be orthogonal (Moar and Bower, 1983) is in conflict with the triangle inequality. Because empirical findings raise many doubts about the use of Euclidean geometry and a coordinate system in spatial knowledge, transferring the idea of an Euclidean spatial representation like in historical maps to the internal representation used in biology lacks valid justification.

### Topological Graph

The *topological graph* hypothesis constitutes an opposite position to the Euclidean cognitive map with regards to stored metric relations between places and also the underlying structure for spatial representation. This hypothesis proposes a network of paths - a graph - as the structure according to which places are saved (Warren *et al.*, 2017) and is therefore similar to the structure of public transportation maps. The nodes of the graph represent visited places (or views, etc.) and can also hold further information about these places. Edges are built when travelling between two places, connecting the respective nodes when a traversable path was found between them. In comparison to the Euclidean cognitive map representation where there first exists a grid of predefined angles and distances and places are later inserted according to their relations, this approach first represents places individually and then defines neighbourhood relations between them. The topological graph representation explicitly holds discovered paths and intersections. Navigation is then realised by choosing suitable edges representing a route as well as recombining known paths.

Even though this type of navigation might be sufficient for deciding what metro line to take, certainly, human navigation cannot rely on simple place-to-place actions when solving more complicated tasks. When looking at navigational performances like finding shortest paths and novel shortcuts, it is apparent that a purely topological representation without any metric information cannot explain this behaviour. For determining the shortest path, some kind of distance information is necessary and novel shortcuts cannot be obtained by solely recombining known paths (Warren, 2019).

### Labelled Graph

The *labelled graph* hypothesis addresses the issues entailed in the topological graph hypothesis and includes metric knowledge in its representation. The overall structure for spatial knowledge is still assumed to be a graph with nodes representing places connected by edges based on experienced transitions between them. Distinguishing it from the topological graph, in this approach local metric information additionally puts places in relation to each other. Metric information is thereby stored as labels where edge labels, also known as weights, hold distance estimates between places and node labels store angles between paths (Warren *et al.*, 2017). As with the topological graph, no rigid Euclidean structure is used, instead the graph and included local metric information are acquired together during exploration (Warren, 2019).

A commonality with the Euclidean cognitive map on the other hand is that metric information is obtained by estimates from a path integration system. Due to the fact that these estimates are local measurements and errors accumulate with increasing displacement from assured locations, the derived information is usually imprecise and often biased (Warren, 2019). A significant difference to the Euclidean cognitive map, however, is that the present approach does not utilise a common global coordinate system for local

spatial information. Together with the imprecise path integration measurements this generally yields a geometrically inconsistent representation (Warren *et al.*, 2017). Nevertheless, it is argued that seemingly Euclidean behaviour can be explained by the labelled graph approach since navigation strategies like vector addition and homing by landmarks are still possible without a global coordinate system (Warren, 2019) and can be adaptively combined.

Empirical findings in humans seem to go well with the labelled graph hypothesis as it explains recalling local metrics, navigating along known routes, and novel shortcuts while allowing for a globally inconsistent metric representation prone to biases (Warren, 2019; Strickrodt, 2019). Thus, the assumption so far on human cognitive mapping is that places are viewed separately in a sequence where vision alone provides topological knowledge while metric relations like path lengths and junction angles are learned by path integration between those places (Warren *et al.*, 2017) building a graph-like representation of the environment.

Nevertheless, the question remains how humans can cope with metric discontinuities in their internal representation caused by noisy and erroneous measurements and still organise spatial knowledge in a way that enables navigation and makes sense of the world. In research, many approaches for obtaining a coherent representation (Duckett *et al.*, 2002) propose optimising for geometric consistency. Thereby, local metric information is adjusted to optimally fit on a global scale. Implementations of such approaches can be found for example in Hübner and Mallot (2007), Lu and Milios (1997) and Duckett *et al.* (2002). Some authors propose global embedding of a labelled graph by updating upon rediscovery of places (Strickrodt, 2019; Rothkegel, 2000). Biological support for this approach can be found in the fact that firing patterns of hippocampal cells are tied to and updated by visual landmarks in the environment (Milford and Wyeth, 2008). When considering this updating mechanism, a current matter of research is what role internal cues (like self-motion) and external visual cues (like landmarks) play relative to each other, i.e. what weighting is given to them respectively regarding their influence on updating.

## 1.2 Non-Euclidean Wormhole Experiments

As described in the previous section, it is an ongoing matter of research to consolidate the inquired theories about spatial knowledge and its embedding. Distinguishing the Euclidean map hypothesis from the (labelled) graph hypothesis, the main characteristic of the Euclidean map representation is its global embedding in a common coordinate system satisfying Euclidean metrics. Therefore, investigating the role of Euclidean metrics constitutes a promising approach to gain insight into the internal organisation of spatial knowledge.

A study described in Warren *et al.* (2017) confronted participants with an environment where Euclidean metrics are violated. In this non-Euclidean environment, a correct global embedding in an Euclidean

map is impossible without conflict given the encountered physical circumstances. When assuming an Euclidean map as the underlying structure of spatial knowledge, this is expected to cause difficulties and inconsistencies in both mapping and navigating the environment.

On that basis, it is of interest how the spatial navigation system deals with a non-Euclidean environment and in what way metric violations influence the constructed representation of the environment. This approach addressing the main characteristic distinguishing the two hypotheses could lead to decisive insights about the organisational structure of spatial knowledge and provide further evidence in evaluating the labelled graph hypothesis.

### 1.2.1 Methods and Predictions

Putting this idea into action, Warren *et al.* (2017) conducted a virtual reality experiment aiming to find decisive support for either the labelled graph hypothesis or an Euclidean map based on path integration. They let one group of participants explore an Euclidean maze environment while the other group was presented with a non-Euclidean version of the same maze (see Fig. 1.3 for the maze layout). Non-Euclidean characteristics were realised by adding two wormholes to the maze which at entry relocated and rotated the participant relative to the original maze layout (arrows in figure 1.3 B). From the perspective of the participants, the transition was visually concealed. The experiment began with an exploration and training phase to get to know the environment in depth before starting the experimental tasks.

The first kind of task (*route task*) was designed to investigate the participants' acquired route knowledge by testing travelled paths between objects within the maze corridors. A second task implemented the investigation of the acquired survey knowledge (*shortcut task*). For this purpose, the authors let participants take novel shortcuts after removing the corridor structure from the maze and evaluated the initial walking direction. In both route and shortcut tasks standard object pairs as well as object pairs near wormhole endpoints were tested.

In a second experiment (*rips and folds*), the authors set up specific object sets for the shortcut task to investigate eventual local discontinuities like rips and folds in the internal representation. These sets consisted of object triplets including the target object near a wormhole exit, a wormhole flanker (located between target and wormhole location), and a neutral flanker (located on the other side of the probe).

In Warren *et al.* (2017), specific predictions for the participants' behaviour based on either the Euclidean map or the labelled graph hypothesis were made. The predictions for each task address several aspects:

1. Route task:
  - a. the knowledge acquisition process
  - b. the length of chosen routes
2. Shortcut task:
  - a. the existence of a shortcut biases
  - b. the specific direction of a potential shortcut bias
3. Rips and folds: the type of discontinuity

Table 1.1 summarizes the predictions for the topological, the labelled and the Euclidean map hypothesis.

Aspects of Investigation	Topological Graph Hypothesis	Labelled Graph Hypothesis	Euclidean Map Hypothesis
1a. <i>Knowledge acquisition</i>	Learning is comparable between (non-)Euclidean conditions		Learning is more difficult due to global inconsistency
1b. <i>Route lengths</i>	Shorter routes not preferred	Shorter routes preferred	
2a. <i>Shortcut bias</i>	Bias		No bias due to same object locations
2b. <i>Shortcut direction</i>	Unsystematic, uniformly distributed	Systematically biased by wormholes (45°)	-
3. <i>Rips and Folds</i>	Shortcuts 'fold over' to wormhole flanker		-

Table 1.1: Extracted from Warren *et al.* (2017): authors' predictions for the topological graph, the labelled graph and the Euclidean map hypothesis.




Fig. 1.3: Experimental environment, directly adopted from Warren *et al.* (2017). Red lines correspond to paths taken in the route task.

### 1.2.2 Results

The authors report the following results from their experiments for the respective tasks (for task labels, see table 1.1). As to aspect 1a, no increased learning difficulty was observed in the non-Euclidean maze compared to the Euclidean version. Due to the fact that an impossible embedding did not affect the acquisition of spatial knowledge, the authors argue that this indicates an insensitivity to Euclidean metrics during the construction of an internal representation. The route task yielded an accurate mapping of the environment by the participants, enabling them to walk towards the target object without any errors in both the Euclidean and non-Euclidean condition. Regarding aspect 1b of investigation, the route task demonstrated that 84% of the participants in the non-Euclidean maze took metrically shorter routes containing a wormhole (see Fig. 1.3B).

The evaluation of the shortcut task in the Euclidean condition showed a mean error of  $4.38^\circ$  ( $a$  to  $b$ ) and  $8.56^\circ$  ( $b$  to  $a$ ) with a rather high variation (see Fig. 1.4a). Analysing the results from the non-Euclidean condition according to aspect 2a yielded a biases initial walking direction by  $37.36^\circ$  and  $42.1^\circ$ , i.e. close to the expected  $45^\circ$ , towards the target's wormhole location (see Fig. 1.4b). Therefore, some underlying metric knowledge is shown regarding aspect 2b instead of an undirected distortion that would take shape as an uniformly distributed bias direction. Also, combining the different trials of the shortcut task leads to violations of metric postulates, as shortcuts from  $a$  to  $b$  and vice versa are not opposite to each other but show an angle of  $75.27^\circ$  between them which disagrees with the notion of positivity. A bias was also found for target objects not in close proximity to a wormhole, indicating a distorted embedding deviating from the Euclidean coordinates on a global level.

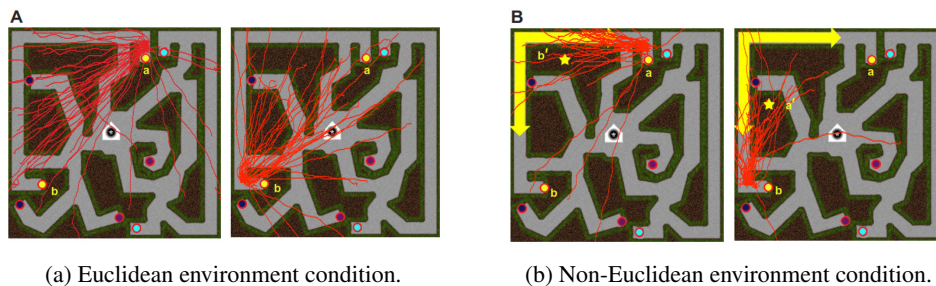


Fig. 1.4: Results of shortcut tasks from Warren *et al.* (2017).

The existence of such a bias violating metric postulates contradicts the assumption of an Euclidean map as the underlying organisation of spatial knowledge. Because specific shortcut biases as well as finding the shortest path in the route task requires establishing some kind of metric representation, the authors further take this as an indication for favouring the labelled graph hypothesis over the topological graph hypothesis. However, they recorded a high uncertainty in both conditions of the shortcut task, questioning the accuracy of the represented metric information in humans in general.

Investigating the influence of wormhole biases in detail (aspect 3), experiment two yielded some insight about the structure of the constructed internal representation. The authors observed a shift for the probe object towards the wormhole, altering its relation to the flanker objects' locations. Specifically, shortcuts to the probe object were 'folded over' shortcuts to the flanker object near a wormhole by  $24.58^\circ$ , while ripped away from the neutral flanker by  $26^\circ$  (Fig. 1.5a).

This modification in the represented maze layout could visually be observed in sketches by participants when asked to reproduce the experienced environment. The drawn maps showed ordinal reversals of target object locations aligning with folds from experiment two (see Fig. 1.5b). Apart from that, participants were able to reproduce the overall cyclic structure of the maze and to sketch relative path lengths reflecting the true distances in the Euclidean maze as well as in the non-Euclidean version with its shorter wormhole paths. The vast majority of participants also declared that no inconsistencies concerning the maze layout were detected, which implies an insensitivity to Euclidean structure at least on a conscious level.

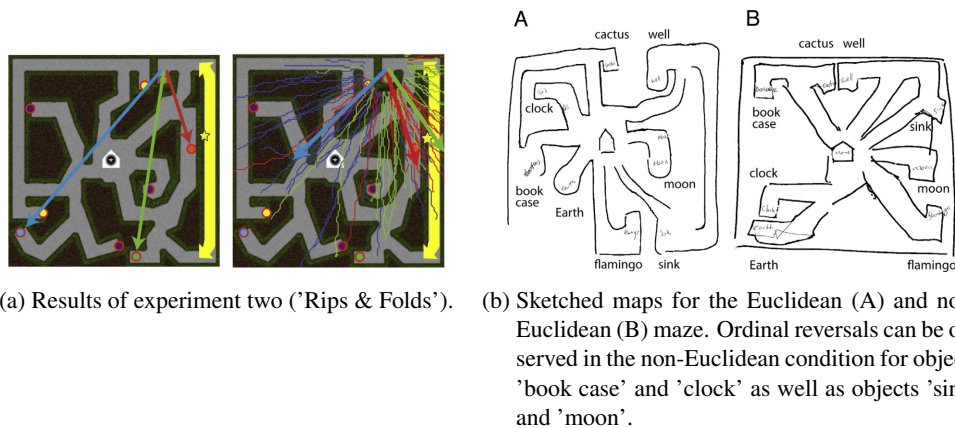


Fig. 1.5: Results for experiment two and maze sketches both adopted from Warren *et al.* (2017).

### 1.2.3 Conclusion and Implications

Based on the outlined results, Warren *et al.* (2017) argue that their study provides evidence strengthening the labelled graph hypothesis. They also point out the difficulty to determine what degree and type of error in navigation tasks is sufficient to dismiss the Euclidean map hypothesis. However, the authors come to the conclusion that the observed systematic biases induced by a wormhole cannot be explained with an Euclidean map representation, even with distortions thereof. Considering that participants demonstrated local metric knowledge about relative path lengths and shortcut directions both through their navigation behaviour as well as when sketching the maze, this is regarded to be convincing evidence against a topological graph missing metric information. The authors further conclude that a labelled graph best explains their results and is regarded the most likely structure of spatial knowledge in humans.

For one thing, assuming a labelled graph as the internal structure of spatial knowledge can account for various navigation behaviours like pointing errors being uncorrelated with the metric distance but rather being connected to the number of intervening nodes on the route (Meilinger *et al.*, 2016). From another perspective, humans do derive local metric knowledge from a noisy path integration system, which is prone to biases and therefore not fully reliable for building an accurate Euclidean map (Warren *et al.*, 2017). On that basis, it seems that organising spatial knowledge in a flexible graph-like representation entails many advantages like easily adapting to new experiences by updating the affected local metric information without major global conflicts.

#### **1.2.4 Related Work**

Another study utilizing the concept of wormholes was conducted by Strickrodt (2019). Participants of this study had to perform a pointing task in either an Euclidean (heptagon) or non-Euclidean (decagon with teleportation across a gap between two corridors) version of a circular maze. Thereby, both versions had the same amount of corridor turns in contrast to the Warren maze where each wormhole skipped a 90° turn. The results of the study by Strickrodt (2019) also agree with the notion of a labelled graph, showing precise local metric knowledge as well as systematic pointing biases towards the wormhole location of objects in the non-Euclidean condition yielding an inconsistent global representation.



## CHAPTER 2

### METHODS

#### 2.1 Microsnapshot Navigation

The results from Warren *et al.* (2017) and several other works suggest a labelled graph, i.e. a topological structure extended with metric knowledge as the most likely model for biological spatial representation. Employing refined navigation strategies, this approach is able to reproduce many of the tested navigation behaviours.

In order to further explore this idea, Baumann (2019) developed an algorithm modelling a labelled graph as described e.g. in Warren *et al.* (2017), called *Microsnapshot Navigation* (hereafter abbreviated as MSN algorithm). The objective for this model was to provide a simplistic structure for spatial knowledge while being able to emulate parts of navigation behaviour found in biological systems. In replicating navigation experiments on humans and animals, the MSN algorithm provides a source for further understanding of the involved mechanisms. On the other hand, the MSN algorithm poses a possible navigation policy in mobile robotics.

The MSN algorithm is characterized by its topological nature without metric embedding, generating a graph supplemented with local metric estimates. Inspiration came for example from visually guided navigation in insects (Cartwright and Collett, 1983), where it is assumed that the memorization of visual-based place representations corresponding to view-images of the environment is combined with various strategies to produce the observable navigation behaviour. The following section summarizes key characteristics of the MSN algorithm extending this approach:

*View-based.* Acquiring spatial knowledge in the algorithm is solely based on recognition of visual features and of the area from where features can be detected, so called place fields. Visual features (microsnapshots) are extracted from a panoramic image (snapshot) by the 'upright speeded-up robust features' (U-SURF) algorithm. Prioritising the most salient ones, a fixed number of features for each snapshot is stored as nodes in a graph. A respective descriptor vector for each feature based on its intensity gradient is used to compare and match extracted features with existing ones to be able to recognize rediscovered features. The feature matching process also assesses neighbourhood similarity and only considers two features to be identical if they have enough common neighbours. Because comparison of features is solely based on their visual similarity ignoring metric relations, two distinct features could get confused for being identical which is known as aliasing (see edges crossing the grey space in figure 2.1). Places in the environment are therefore solely distinguished by the subset of features visible from each place (Baumann, 2019).

*Edge formation.* Edges describe place field transitions, i.e. connect features across the border of a place field. When a feature was visible in the last frame but disappears in the current frame, it is connected to one of the currently visible features with an edge of weight 1. Therefore, connectivity in the graph encodes neighbouring place-fields with overlapping common features instead of physical transitions between features.

*Local metric estimates.* The MSN algorithm resigns from relying on explicit metric information in favour of utilizing local measurements. Metric information is stored as labels in both nodes and edges of the graph. Nodes are labelled with a local compass direction estimate constantly updated upon sight of known features. This compass is closely related to the concept of head-direction cells in the hippocampus (Baumann, 2019). Edges between two nodes are labelled with directional movement instructions relative to the stored compass direction describing how to move from one feature's place field to the other. As an exception from a purely local representation but solely for visualisation purposes, x-, y-coordinates corresponding to the ground truth location from where a feature was first detected are stored for each node.

*Dual population coding.* Even though the graph structure of the MSN algorithm's representation entails discrete nodes, an approximation of continuity is provided by a dual population coding approach. This concept is inspired by the discovery of a place cell system in the hippocampus where populations of cells with overlapping place-fields encode locations in an environment by their firing patterns (O'Keefe and Dostrovsky, 1971; Wilson and McNaughton, 1993). Accordingly, the present algorithm represents place-fields by a population code of all nodes associated with features visible from that place. This means, every location is characterized by the activity of a unique set of multiple nodes in contrast to a representation where one location is associated with only one node. The second part of dual population coding concerns the navigation procedure and is described in the following.

*Navigation.* Path planning is implemented by *Dijkstra shortest-path-search* (Dijkstra, 1959) which - in a graph with uniform edge weights like the present one - finds the path with the least number of edges between two given points. The MSN algorithm calculates a bundle of paths between the given start and goal nodes (see Fig. 2.1) and applies a voting scheme to obtain a consensual trajectory. Thereby, the planned trajectory generally does not follow the edges of the graph and neither does it need to on the grounds that these edges do not represent traversable paths but rather neighbourhood relations between place fields.

## 2.2 Simulating Wormhole Experiments

Several research demonstrated the applicability of graph-based topological navigation approaches related to the MSN algorithm in the field of robotics (for example Franz *et al.* (1998) and Hübner and Mallot (2007)). Beyond that, a primary objective of this work is to assess the MSN algorithm's potential for




Fig. 2.1: Example graph generated by the MSN algorithm, image from Baumann (2019). Blue lines are edges, long blue lines across non-traversable space of the environment are alias edges. Orange lines show path bundles found by the navigation mechanism.

modelling human navigational performances. Therefore, it is of interest to validate the MSN algorithm with regards to different behaviours to obtain a more comprehensive understanding of its explanatory power for the underlying structure of spatial knowledge in biology.

Investigating the influence of non-Euclidean structures on navigation behaviour plays an important role in arguing for a labelled graph as the possible underlying structure for human spatial knowledge. That is why a simulation of the experiment from Warren utilising a non-Euclidean environment seems to pose a promising approach for assessing the MSN algorithm's biological plausibility. Not to neglect that reproducing the results from experiments on humans with the MSN algorithm bears potential for providing further evidence strengthening the labelled graph hypothesis in the first place.

Hence, the following questions guiding this work can be deduced:

- How does the MSN algorithm handle an impossible non-Euclidean environment with respect to the construction of the graph representation?
- Can the MSN algorithm predict human performances observed in Warren *et al.* (2017) and, in particular, is the graph representation affected by wormholes in the same way as the mental representation in humans?
- Based on that, what can be stated about the inherent local metric information included in the graph?

### 2.3 Graph Embedding and Visualisation

In the course of comparing results from experiments with humans in Warren *et al.* (2017) to the graph generated by the MSN algorithm, some considerations have to be made about how to visualise the graph. It is vital to utilise a reference frame according to which similarity and potential distortions can be characterized in a meaningful way. Figures visualising the graph representation so far, like shown in figure 2.1, were drawn according to the stored ground truth position for each node. Applied to the Euclidean

maze, displaying the graph according to the ground truth coordinates yields an acceptable representation of the environment. However, a comprehensible visualization based on the ground truth position reaches its limits when it comes to a non-Euclidean environment due to overlapping of distinct paths and different rotation of same paths. Apart from the visualisation, it was aspired to analyse the information captured by the graph's inherent structure and without imposing artificial structural properties on it, i.e. independently from ground truth positions. Hence, an embedding of the graph has to be found which both enables meaningful assessment of the represented information while also providing a basis for comparing it to the experimental results from Warren. Thereby, the embedding should also be visualised in a way that is easily accessible to the human observer.

### 2.3.1 Multidimensional Scaling

In search for a metric embedding model satisfying these requirements, the approach of *Multidimensional Scaling* (abbreviated MDS) seems to be suitable. MDS is specialised in analysing and visualising dissimilarities of a given set of data (Pich, 2009).

For each individual application of MDS, it has to be determined what dissimilarity means with respect to the data. Considering the notion of a graph, dissimilarity of two data points, i.e. nodes, can be described by the distance between those nodes, likewise similarity corresponds to their proximity to each other. Distances in a graph are generally defined as the length of the shortest path between two points (Brandes and Pich, 2011), also called topological distance.

MDS operates on a complete set of pairwise distances  $\delta_{ij}$ , i.e. the matrix of shortest-path lengths between all pairs of nodes in the graph. Mathematically, this distance matrix is defined by

$$\Delta = (\delta_{ij}) = \begin{bmatrix} \delta_{11} & \dots & \delta_{1n} \\ \dots & & \dots \\ \delta_{n1} & \dots & \delta_{nn} \end{bmatrix}; i, j \in 1, \dots, n$$

where  $n$  denotes the number of nodes. The problem to be solved by MDS is the following: given the distance matrix  $\Delta$ , find a configuration of  $p$ -dimensional coordinates of points  $\vec{x}_i$

$$C = (x_{ik}) = \begin{bmatrix} x_{11} & \dots & x_{1p} \\ \dots & & \dots \\ x_{n1} & \dots & x_{np} \end{bmatrix}; i \in 1, \dots, n, k \in 1, \dots, p$$

satisfying  $\|\vec{x}_i - \vec{x}_j\| \approx \delta_{ij}$ .

In the course of finding a suitable configuration  $C$ , MDS minimises a cost function called *stress*:

$$S(\Delta, C) = \sum_{i,j} (\|\vec{x}_i - \vec{x}_j\| - \delta_{ij})^2$$

Definitions are adopted from Mallot (2018) and Pich (2009). For a detailed step-by-step computation of MDS, see Brandes and Pich (2011), Mallot (2018) or Pich (2009).

Using metric - or classical - MDS, the input distances are considered to be Euclidean geometric information (Borg and Groenen, 2005). For a requested  $p$ -dimensional embedding, the input data is reduced

to the first  $p$  most informative dimensions, i.e. the  $p$  Eigenvectors with the highest Eigenvalues, trying to reflect as much of the variability of the distances as possible (Brandes and Pich, 2011; Mallot, 2018). This process is closely related to the concept of principal component analysis.

In solving the embedding problem, the found configuration  $C$  of node positions represents the underlying structure of e.g. the graph up to rotation, reflection, and translation, implying that MDS solutions are not unique with respect to these transformations (Mallot, 2018). Also, when embedding input data of dimension  $n$  into  $p < n$  dimensions, there will be a residual error (stress) (Mallot, 2018). This indicates that MDS cannot fully preserve all given distance relations but rather results in a slightly distorted version of the original data. The remaining stress is defined by:

$$E(p) = \left( \sum_{i=1}^n \sum_{j \neq i} (\delta_{ij} - \|\vec{x}_i - \vec{x}_j\|_p)^2 \right)^{\frac{1}{2}} \text{ with } \|\vec{x}_i\|_p = \left( \sum_{k=1}^p x_k^2 \right)^{\frac{1}{2}} \text{ (Mallot, 2018)}$$

Stress increases with the number of input data and errors in the input distances regarding their Euclidean metric relations (Borg and Groenen, 2005).

Nevertheless, the resulting  $p$ -dimensional point configuration reconstructs the input distances as closely as possible while extracting essential information. Thus, performing MDS with two dimensions on the edge distance matrix of a graph results in an Euclidean visualisation of the graph's inherent connectivity structure. The MDS approach has been used in research before for similar applications, like e.g. reconstructing the spatial layout of a cognitive map Golledge and Stimson (1997).

With regards to the purpose of this work, MDS provides a suitable tool for converting proximity in a graph into geometric proximity (Brandes and Pich, 2011), i.e. Euclidean distances, which can then be visualised. In this way, the graph topology becomes easily comprehensible for a human observer and accessible for interpretation.

It is important to note, that this embedding model requires a significant amount of computational expenses. Firstly, computing the distance matrix for all pairs of  $n$  nodes is quite costly (Brandes and Pich, 2011), even with optimised algorithms for finding the *all-pairs-shortest-paths*. Secondly, finding configuration  $C$  is in  $O(n^2)$  (Pich, 2009) and therefore quickly becomes impractical for large graphs with thousands of nodes. This issue will be addressed in section 3.2.

A consequence of the fact that MDS results are not unique with respect to rotation, reflection, and translation is that it is difficult to directly compare them to a reference layout like the true environment or observed representations like those in Warren *et al.* (2017). Apart from the visualisation, rotation and reflection play a crucial role with regards to shortcut tasks where the angle of direction is a key measurement. In order to be able to analyse simulation results in a meaningful way, it therefore needs to be specified how to display the MDS results accordingly.

### 2.3.2 Procrustes Transformation

One approach is given by adjusting the MDS configuration's structure to the comparison structure without changing its essential information content, i.e. the relations among distances between data points. According to Borg and Groenen (2005), transformations satisfying this constraint are rotation, translation, reflection and scaling - also called similarity transformations. Applying them preserves the invariant shape of the MDS layout including its metric properties (Borg and Groenen, 2005). Performing these transformations is expected to adequately prepare the MDS results for comparison both with respect to scale as well as orientation.

*Procrustes transformation* is a method addressing the problem of fitting one data configuration to another as closely as possible (Borg and Groenen, 2005). It makes use of the idea that in order to compare the shapes of two configurations, they must be optimally "superimposed". This eliminates any superficial differences without consequence to the true underlying structures. To achieve this superposition, one of the configurations is transformed according to the other. Importantly, the method preserves the inherent structure of the transformed configuration by only applying translation, scaling and orthogonal rotation and reflection. As previously defined, these transformations are acceptable according to our purposes. The transformation result  $Z$  can be described as a composition of the input configuration  $Y$  and the computed transformation components:  $Z = bYT + c$ , where  $c$  is the translation component,  $T$  the rotation and reflection component, and  $b$  the scaling factor.

## 2.4 The Present Study

In this thesis, a minimalistic graph structure solely based on visual feature extraction is supplemented with an embedding mechanism based on Multidimensional scaling and Procrustes transformation providing a reference frame. In order to assess the biological plausibility of the model, the embedded graph representation was assessed regarding its ability to reproduce represented global navigational knowledge observed in human behaviour.

Since the graph is constructed on the basis of visual feature recognition and neighbouring place fields, no explicit metric information is included. With MDS as an approach for visualising the graph representation, embedding is solely based on the graph's connectivity and therefore yields a layout displaying the inherent visual based information content of the graph independent from exact coordinates, i.e. without metric information from e.g. path integration. With respect to wormholes, a lack of path integration is in fact expected to pose an advantage in order to avoid an embedding conflict. Still, some kind of reference frame is needed to meet the requirements for modelling human spatial knowledge due to the fact that human navigational behaviour indicates a representation of metric knowledge in some way. Thereby, dif-

ferent mechanisms for acquiring metric information can be employed, for example by path integration, i.e. self-motion, or by perception of visual cues like landmarks.

In order to investigate the role of self-motion in contrast to visual cues as sources for metric information, it is useful to test different reference frames for embedding spatial representations. Firstly, a path integration system solely based on experienced self-motion without updating upon (visual) cues can be utilised as a reference frame for an embedded graph representation. This approach is implemented by applying Procrustes transformation based on a model for self-motion information. On the other hand, including visual cues like object locations in the embedding mechanism as anchor points provides a frame for embedding the graph representation based on fixed positions of those objects in the physical environment. This kind of reference frame is implemented by applying Procrustes transformation based on a model for landmark information.

In case the present model yields spatial knowledge comparable to humans, it can be argued that MDS based on a purely visual graph representation poses a possible model for organising spatial knowledge. Besides that, evaluating the model's performance with both types of reference frames could give insight into the role of different sources for acquiring and updating metric knowledge in spatial representations.

### **Structure of this Thesis**

The Simulation is conducted in a virtual environment, employing an agent guided by the MSN algorithm. In section 3.2, algorithms for processing the constructed graph in preparation for further analysis are described. The graph representation as well as MDS and Procrustes results are evaluated in chapter 4 by means of three types of tasks. Firstly, route knowledge is inquired through finding shortest known paths. Secondly, the goodness of mapping is evaluated by comparing information content in the resulting graph representation with the layout of the actual environment for testing survey knowledge. In a third task, pointing performance is tested giving detailed insight into the relation of embedded objects and possible biases caused by wormholes. All tasks are evaluated with respect to the results from Warren *et al.* (2017), i.e. compared to actual human performance. In subsequent sections, the results are discussed addressing implications, limitations, and possible extensions of the presented model.

## CHAPTER 3

### SIMULATION

#### 3.1 Simulation Environment

The simulation environment for the experiment was constructed in order to replicate the study in Warren *et al.* (2017). Accordingly, a maze in a three-dimensional virtual reality environment was set up using the game engine Unity<sup>®</sup> (B.2.1). The simulation was performed on a HP Spectre x360 - 13-ac033ng with Intel<sup>®</sup> Core™i7-7500U processor and Intel<sup>®</sup> HD-graphics card 620.

The maze layout (see Fig. 3.1a) resembles the one in Warren (see Fig. 1.3), with a centered home location surrounded by eight objects hidden in several corridors. The objects in the environment are free assets available at the Unity Asset Store (for further information see appendix B.2.1). They are not and neither need to be exactly the same kind of objects as used in the original experiment, on the ground that the MSN algorithm does only care about individual gradient features and can not recognize complex objects nor geometry. For that same reason, the decision was made to adapt the visual appearance of the maze environment by adding various nature inspired patterns to the corridor walls (for the specific patterns see appendix B.1), compensating for the human ability to include geometry and abstract concepts in their navigation behaviour. Further, Warren *et al.* (2017) mention that paintings were added to some maze walls in order to aid orientation for the human participants. Therefore, aspiring a useful segmentation, the maze was divided into sections and assigned unique patterns reflecting geometric and contextual concepts like intersecting corridors. See figure 3.1a for the division of the maze with specific patterns assigned. This adaptation ensures the transferability between human performance in the original experiment (Warren *et al.*, 2017) and the performance of the MSN algorithm in the set up environment.

Addressing the core of the experiment, two versions of the maze were set up, reproducing the matched Euclidean and non-Euclidean environments in Warren *et al.* (2017) (see Fig. 1.3). The non-Euclidean version differs only to the effect of newly introduced wormholes, where two corridors previously separated by a 90° turn are merged into one corridor respectively. Because corridors are separated by their wall patterns, the different patterns of the two merged corridors were replaced by only one pattern. Figure 3.1b shows the non-euclidean maze with concerned corridors.

The wormholes were implemented by teleporting the agent from one entry point to the corresponding exit point. Two versions of each wormhole were set up which yields two types of recorded ground truth coordinates. One aligns associated wormhole entry points on a two-dimensional level by teleporting the agent to a rotated version of the maze. Rotated maze versions are stacked on top of each other on a three



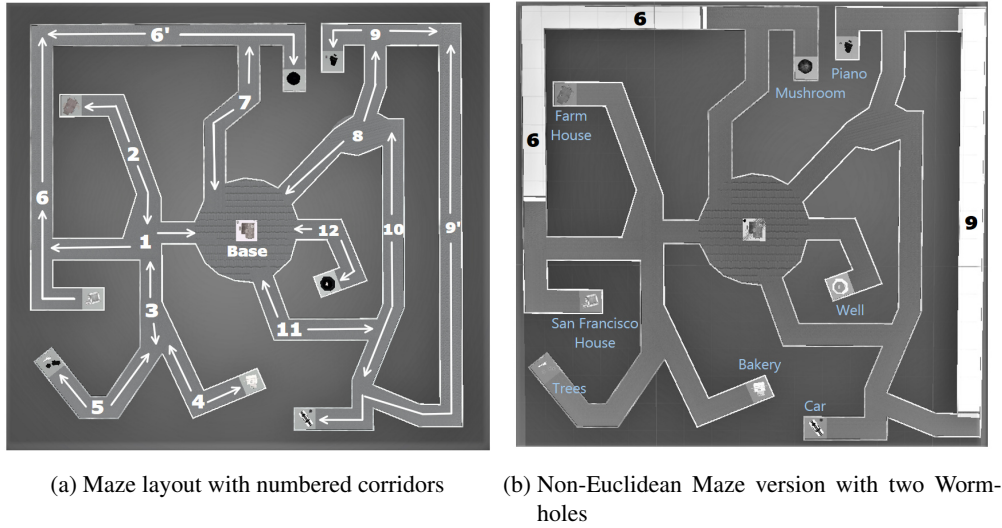


Fig. 3.1: Simulation Environment

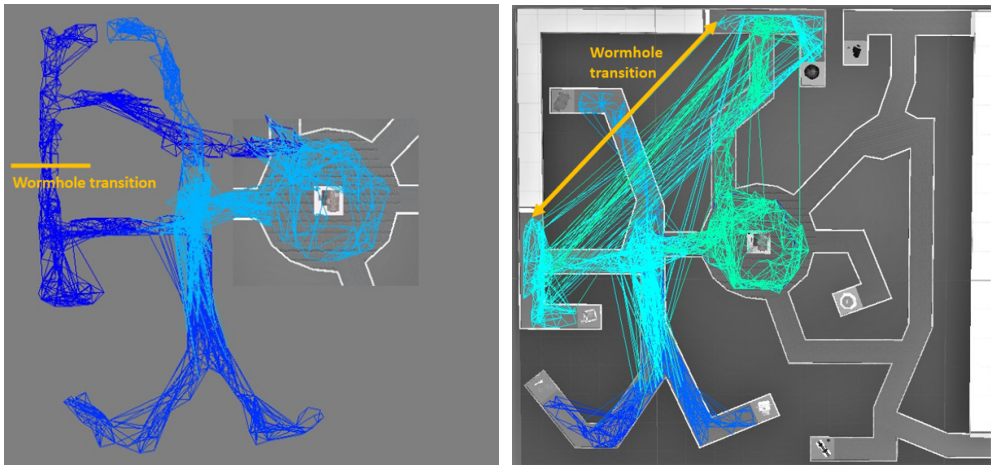
dimensional level, thereby, recorded two dimensional ground truth coordinates conceal the teleportation and do not show a leap (*ground truth coordinates type 1*). The other wormhole version simply teleports the agent within the maze, thus ground truth coordinates capture the wormhole transition by showing a translation and rotation (*ground truth coordinates type 2*). A visualisation of graphs drawn according to each ground truth type is displayed in figure 3.2. Thereby, ground truth type 1 models self-motion experienced by the agent while ground truth type 2 models globally anchored landmark information. The implementation of the wormholes results in a total of eight maze versions, one Euclidean one, one non-Euclidean one, and three rotated versions ( $90^\circ$ ,  $180^\circ$ ,  $270^\circ$ ) of the non-Euclidean maze for each wormhole respectively.

In contrast to the Warren experiment, where wormhole transitions were obscured by making them visually seamless for the participants, this was not implemented in the simulation. That is due to the fact that the MSN algorithm does not experience e.g. optic flow like humans do, therefore tests yielded no significant difference between a visually seamless transition and one with visual inconsistencies due to the wormhole teleportation.

## 3.2 Graph Processing Algorithm

After letting the MSN algorithm explore the virtual reality environment, the generated graph was slightly modified addressing aliases and the size of the graph.

Having the application of Multidimensional Scaling to the constructed graph in mind, the main modification concerned the computational expenses for said MDS application. As mentioned in section 2.3.1, the complexity of MDS increases quadratically with the number of nodes of the graph. Given the limited



(a) Ground truth type 1 modelling self-motion experienced by the agent. (b) Ground truth type 2 modelling globally anchored landmark information.

Fig. 3.2: Ground truth types: visualisation of example graphs drawn according to each ground truth type.

computation power of a personal computer, it was essential to reduce the graph to a manageable size.

The implementation of a suitable reduction of the graph utilised the fact that the MSN algorithm extracts up to 30 features from each recorded snapshot. Not only does this mean that up to 30 features represent one location, but also does the edge formation mechanism not ensure that all of the features within one snapshot are connected to each other. MDS would unnecessarily disperse these features, although they coincide with one common location. On that basis, the conceptual notion of a node in the graph was changed from representing one visual feature (*feature-node*) to representing a snapshot, or view, (*snapshot-node*). A representation on view level instead of feature level can be seen as an abstraction of the original information content and is backed up by empirical findings of hierarchical representations in humans.

On the implementation level, this means all *feature-nodes* corresponding to one snapshot were merged to a *snapshot-node* respectively. Former *feature-nodes* were clustered into snapshots step by step by selecting all consecutively created nodes with equal coordinate labels. In the process of contracting *feature-nodes* to *snapshot-nodes*, all adjacencies should be preserved and multiple edges between *snapshot-nodes* are supposed to be avoided in the reduced graph. Therefore, the reduction process entails merging the set of all previous edges between different *feature-nodes* of two snapshot clusters to only one edge between the corresponding *snapshot-nodes* in the reduced graph. To preserve specific degrees of connectivity between snapshot clusters, edge weights reflecting the number of edges existing between two snapshot clusters were introduced. For every redundant edge between a given pair of *snapshot-nodes*, the weight of the existing edge between them (starting with an initial value of 1) was reduced by 5%. Thereby, the distance given by the edge weight between *snapshot-nodes* decreased with each found edge between their former snapshot clusters, resulting in a reduced graph reflecting the contracted connectivity by edge

weights. Because edges were still directional and a pair of edges between the same nodes can have different weights due to the merging process, in the end of the process edge weights were made symmetrical by replacing the weights of an edge pair with the average of them.

Addressing the issue of aliases, edges of the original graph were only added to the reduced graph if there existed a minimum number of edges between the respective node neighbourhoods above a certain threshold. The adjustment of edge connections also utilises the two-dimensional-ground-truth Euclidean distance of two nodes. An edge is included in the reduced graph only if the distance between the two nodes of that edge lies below a fixed threshold and is therefore unlikely to be an alias edge. This modification reduces the probability of transferring alias edges between features which could not possibly be the same simply on the basis of their discovery locations.

Considering the present maze environment, this provides additional aid for the MSN algorithm to handle the uniform corridor layout and repeating patterns. A sensitivity for possible alias edges is especially important in the reduced graph. With less edges in total, the resulting distorting influence of aliases on a structural analysis (e.g. MDS) will become proportionally higher.

Of course, there is a trade-off between removing alias edges and preserving true edges which is why the threshold parameters have to be determined exploratively. Due to SURF features being scale invariant, especially long straight corridors could yield valid edges between features with large distances between them. However, removing these edges in the process of reducing the graph does not pose a serious issue because in general concerned corridors are still mapped by shorter edges and edges across a whole corridor would unnecessarily contract this corridor in a MDS analysis. For that reason, alias filtering based on the features' distances cannot be applied to graphs from the wormhole conditions due to leaps or rotated paths respectively. A rudimentary approach was to simply skip the alias filtering based on distances in the non-Euclidean condition which, corresponding graphs still showed a manageable amount of alias edges. The reduced graph is constructed and saved separately. Node and edge labels solely used for navigation (i.e. movement instructions) are not of importance to the MDS computation and were therefore omitted. A side effect of merging individual *feature-nodes* is that the population coding property of the graph disappears and every location is represented by a single *snapshot-node* only. However, because the reduced graph is solely used for the application of MDS - where population coding does not play a role - and the original graph is still available for path planning, this does not affect the present work.

### 3.3 Simulation Procedure

The simulation procedure was structured as follows. First, the MSN algorithm explored the environment through a virtual reality agent. During the exploration, the MSN algorithm generated a graph representing its acquired knowledge of the environment. The graph construction mechanism is described in section 2.1. Parameters for the MSN algorithm were the same as described in Baumann (2019) except for the

ones shown in table 3.1 which were adjusted to provide a more thorough exploration of the environment and therefore a denser connectivity of the graph.

<b>Parameter</b>	<b>Description</b>	<b>Value</b>
<i>kEdgesPerNode</i>	The maximum number of edges that can be added to a node per update step	5 (previously 2)
<i>kGraphUpdateStep</i>	The graph is updated every kGraphUpdateStep frames	10 (previously 12)
<i>Walking Speed</i> (in Unity)	Walking speed of the agent in the virtual reality environment	5 (previously 10)

Table 3.1: Parameter adaptations for the MSN algorithm.

The MSN algorithm performed complete explorations as well as partial explorations of the maze environment. The guideline for every exploration in accordance with Warren *et al.* (2017) was to visit each object and wormhole of interest at least once while exploring corridors and especially intersections densely. The exploration was terminated when the experimenter considered the environment sufficiently explored with regards to the task to be solved in this simulation run. For example, in order to limit the size of the constructed graph, some explorations only included parts of the maze that were relevant to the current task. Because of the specific requirements for each exploration and due to the fact that it would take the agent quite a while to meet these requirements with a random walk exploration, it was decided to manually guide the exploration behaviour of the agent. Even though no exploration routes are exactly the same, the generated graph topologies from two explorations with roughly the same route only show minor differences. Alias edges may vary between explorations but mainly disappear in the reduced graph.

When the exploration phase was completed, the generated graph was reduced as described in section 3.2. Specific parameter settings are shown in table 3.2.

Parameter	Description	Value	Purpose
<i>aliasThreshold</i>	Given an edge $a$ between two nodes, this sets the minimum number of other edges between the nodes' neighbourhoods for edge $a$ to not be considered as aliasing.	20 Determined by balancing the number of preserved aliases against the graph becoming too sparse or even unconnected.	Prevents possible alias edges from being included in the reduced graph.
<i>surroundingDistance</i>	The physical two-dimensional distance two features have to be within for an edge between them to not be deleted.	30 Determined by taking the base (largest open space) as a reference point (with diameter $\sim 20$ units) and adding a tolerance interval of 10 units.	Provides an additional criterion for filtering out alias edges that were not ruled out by the <i>aliasThreshold</i> .

Table 3.2: Parameter settings for the graph reduction algorithm.

Further computation was conducted using MATLAB (B.2.3). A pairwise distance matrix  $\Delta$  was obtained from the reduced graph by calculating the all-pairs-shortest-paths matrix using the MATLAB function `graphallshortestpaths` (B.2.3). This function performs the *Johnson's all pairs shortest paths* algorithm (Johnson, 1977) which is specialised in the application to sparse networks. The used graphs consist of mostly elongated paths along corridors with afterwards contracted nodes with equal coordinates, thus can be described as sparse. With a complexity of  $O(n * \log(n) + n * e)$  ( $n$  being the number of nodes and  $e$  the number of edges), *Johnson's all pairs shortest paths* algorithm reduces the computational costs for solving the all-pairs-shortest-paths problem in sparse weighted graphs compared to other all-pairs-shortest-paths algorithms.

Because edge weights were made symmetrical in the reduction process, the resulting matrix  $\Delta$  is also symmetrical. Therefore, it can be directly passed as input for computing MDS. To perform MDS, two different functions were used, one being the MATLAB function `mdscale` (B.2.3) and MATLAB function `cmdscale` (B.2.3) as the other one. `cmdscale` performs classical MDS, i.e. metric MDS with the assumption of the input data being Euclidean and optionally returns computed Eigenvalues. `mdscale` performs non-metric MDS and additionally returns the remaining stress value. Chosen parameters to `mdscale` are the embedding dimension  $p = 2$  and a specified cost function with metric scaling called '*metricstress*' which normalizes stress by the sum of the squared dissimilarities. By default, non-metric MDS starts with the classical MDS solution as a basis before computing an optimal solution configuration.

As a next step, the respective MDS result was adjusted using the MATLAB function `procrustes` to fit the maze map both in order to visualise and interpret the model's performances. The reference configuration the MDS result is adjusted to by Procrustes transformation is provided by ground truth coordinates for each node which are saved during the graph construction. Since wormholes in the non-Euclidean

maze were implemented in two different ways and thereby yield different kinds of ground truth data, two types of reference structures for the Procrustes transformation are available. *Procrustes type 1* is based on ground truth data from wormholes concealing the teleportation (*ground truth data type 1*). Using this data models self-motion information as a frame of reference for the embedded representation. Accordingly, *Procrustes type 2* is based on *ground truth data type 2* which by capturing the teleportation are aligned with the original maze version and its object locations. This can be seen as a model for using landmark information as anchor points in embedding the spatial representation. Both types of Procrustes were performed and additionally, a Procrustes transformation adjusting each manually marked section of the graph to its corresponding maze section separately (*section-wise Procrustes*) was computed in contrast to transforming the entire graph in a single run.

Finally, tests were performed analysing the acquired spatial representation in order to simulate the tasks described by Warren *et al.* (2017) for both the Euclidean and non-Euclidean maze. Thereby, analysing the Euclidean maze version is used to assess the comparability to the Warren experiment, i.e. whether the computed representation is able to reconstruct the environment's layout in general on a human level of performance (mapping performance). The non-Euclidean maze version on the other hand is analysed with the purpose of investigating how the model deals with physical impossibilities, i.e. wormholes, and whether it shows similar biases like they were caused in humans. Thereby, possible biases have to be evaluated by comparing the model's Euclidean and non-Euclidean representation.

Specific tasks replicating those in Warren *et al.* (2017) were conducted as described in the following. For one thing, the route task was simulated by performing path finding based on the unreduced graph representation as described in section 2.1. Shortcut tasks in the original experiment (Warren *et al.*, 2017) were based on measuring the participant's initial walking direction. Accordingly, this task was geometrically simulated as pointings by drawing a point-to-point line between start and target object in the Procrustes-transformed map calculated by MDS on the reduced graph.

The second point of investigation aimed at simulating experiment two (Rips & Folds) from Warren *et al.* (2017) by also conducting geometric pointing tasks. In the course of this it was especially interesting to examine the structure of the MDS-processed graph for potential distortions similar to those found by Warren *et al.* (2017) in their participant's sketched maps.

## CHAPTER 4

### RESULTS

#### 4.1 Graph Reduction

For an example full exploration of the maze, the reduction process decreased the number of nodes from 13123 to 616 and the number of edges from 167216 to 4288. Figure 4.1 shows the graph topology mapping the environment for an example graph compared to its reduced version respectively. As can be seen in the figure, the overall graph topology is preserved while alias connections are almost completely eliminated without disconnecting the graph.

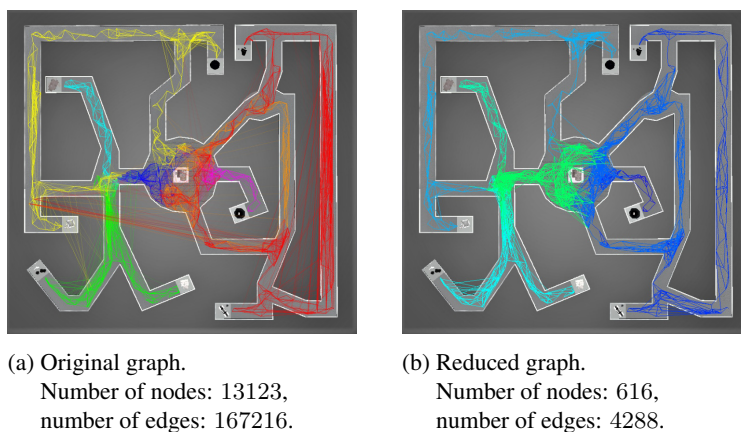


Fig. 4.1: Visualization of the graph reduction process described in section 3.2. The set of colours differs in the reduced graph but still shows the same partitioning into sections.

Effects of the reduction process on the MDS computation are displayed in figure 4.2. As expected, the MDS for the reduced graph shows less distinct points compared to when using the original graph where unnecessary distances between nodes at the same snapshot location occur and add noise to the result. Regarding the goal of reducing computational expenses, applying MDS to the reduced version of the example graph from figure 4.1 yielded an elapsed time of 0.102 seconds for classical MDS and 8.528861 seconds (non-classical MDS) while the original version of the graph took 788 seconds (classical MDS) and respectively more than 80 minutes for non-classical MDS. This illustrates a significant decrease of computing time for a reduced graph compared to its original version.




Fig. 4.2: MDS results compared for the original and the reduced version of a simple graph. The reduced graph (c) shows less distinct points compared to when using the original graph (b) while also reconstructing the ground truth data (a) better.

## 4.2 Route Knowledge

The investigation of route knowledge was conducted on the unreduced graph using the MSN path finding mechanism described in section 2.1. Shortest paths were computed for a pair of objects for each wormhole section respectively. The resulting path bundles are depicted in figure 4.3. Compared with the Euclidean maze condition where paths were found for two different routes, a preference for paths traversing the wormhole can be observed in the non-Euclidean condition. This shows that relative path lengths are already represented in the navigation graph without embedding.

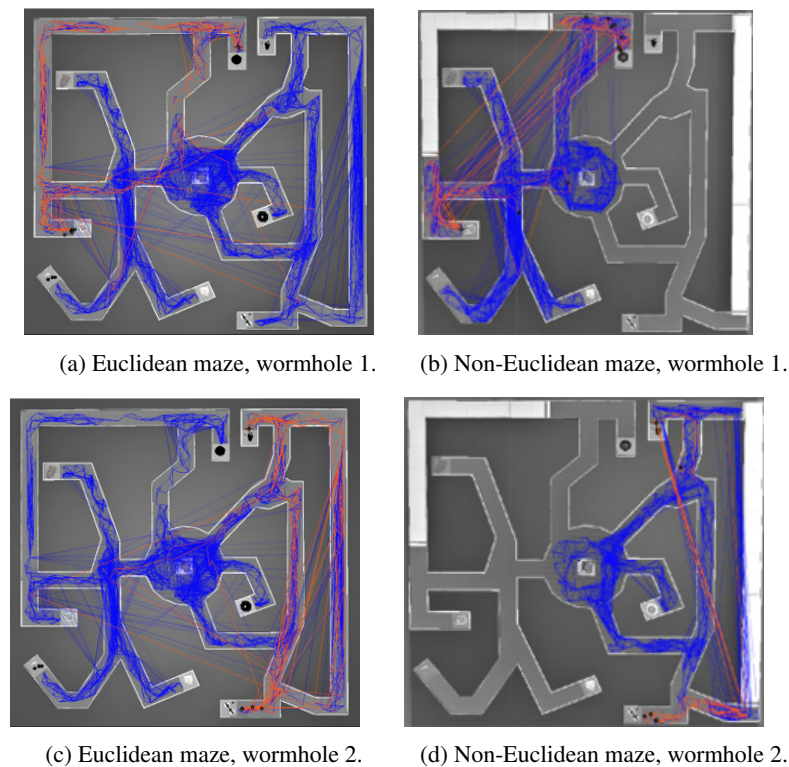


Fig. 4.3: Routes (orange lines) between objects 'Mushroom' and 'San Francisco House' (a, b) and objects 'Piano' and 'Car' (c, d) for the Euclidean maze (a, c) and the non-Euclidean maze (b, d).



### 4.3 Survey Knowledge

Results in the following sections refer to the Procrustes transformed MDS computed on the reduced version of respective graphs. This section only includes selected result figures while also referring to supplementary figures in appendix A.

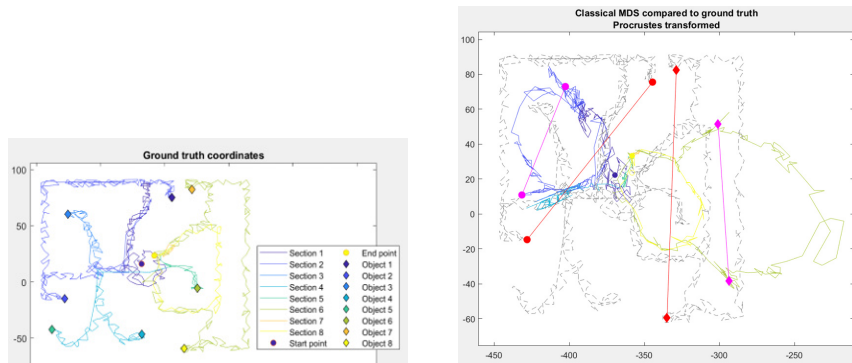
#### 4.3.1 Mapping

The guideline for assessing mapping performances of the MSN algorithm was to investigate how the generated graph represents the environment. For this purpose, structural characteristics that can be found in the graph's layout were compared with characteristics in the environment, i.e. object relations, geometric structures like corridor arrangement, and corridor orientation. These characteristics were also found to be represented in e.g. drawings from participants in Warren, i.e. are included in human spatial representation. The graphs were manually divided into colour coded sections for better comparison, section numbering for each exploration is depicted by the ground truth plot respectively.

##### **Euclidean maze**

Results for a graph for a complete exploration of the Euclidean environment embedded by classical MDS and adjusted to the corresponding ground truth by Procrustes transformation are shown in figure 4.4. In general, a separation of corridors that resembles the one found in the maze can be observed. Graph object positions reflect their ground truth positions relative to the respective corridor, i.e. at the end of each corridor. In particular, intersections of corridors are well represented while turns within one corridor are omitted yielding a loop-shape for section 2 (blue loop) and sections 6/8 (green/yellow loops) or straight shapes for corridors of section 4 (lower left branches with L-shape). The overall orientation of the left and right half of the graph is aligned with the maze while individual corridors branching off the main loops (sections 3, 4, and 5) do not follow their ground truth orientation in general. Further, branches seem to be shrunken which could be an effect of longer edges especially within these branches connecting more distant features. In straight corridors, this is a consequence of the feature extraction algorithm (SURF) being scale invariant. However, when Procrustes transforming the classical MDS result for each section separately, the orientation of all corridor branches can be reconstructed. A similar effect can be achieved with three-dimensional embedding like depicted in figure A.1g. Yet, Eigenvalues indicate that a third dimension does not describe the data with substantial more information (Fig. A.1f).

Section 4 (lower left branches) poses an exception for the separation of corridors by showing corridors which are tied together. These corridor branches also seem to be contracted, i.e. shorter. In the classical MDS representation, section 3 (inner branch in blue loop) is also merged with section 4 but this is resolved in the non-metric as well as in the section-wise transformed classical result. Regarding this aspect,



(a) Visualisation of travelled ground truth route (corridor-wise) and depiction of colour-code for sections. (b) Procrustes transformed classical MDS result superimposed on corresponding ground truth coordinates.

Fig. 4.4: Euclidean maze exploration: sections are colour-coded according to the legend in (a). Pointings: magenta lines correspond to pointings in the MDS representation while red lines show expected pointings based on the respective ground truth coordinates.

non-metric MDS (fig. A.1d) can provide a more accurate representation by differentiating branching corridors but this version also exhibits an additional corridor branching off section 2 (blue loop).


For comparison, figure A.2 shows another exploration of the Euclidean maze with multiple rounds through the entire maze. A common phenomenon between the two different explorations is that lower left corridor branches are contracted and shorter when embedded with classical MDS. Also, non-metric MDS seems to provide better differentiating and also better representation of the shape of these corridors while section-wise Procrustes is able to additionally reconstruct individual branch orientations.

Further observations yielded similar MDS layouts even when computed on a graph with a disconnected loop segment (illustrated in figures A.2f-A.2h).

### Non-Euclidean maze

For investigating the influence of wormholes on the graph representation, the results for the non-Euclidean maze were examined with specific focus on comparing them with the Euclidean results. In general, the maze was divided into its left and right half in order to look at wormhole sections separately.

**Procrustes Type 1.** The graph layout for the non-Euclidean condition (Fig. 4.5 and 4.7) shows a similar separation of corridors and also contracted corridor branches. Furthermore, corridors with a kink are represented as loops as well. A difference to the Euclidean condition can be observed for the wormhole corridors where associated loops are smaller relative to other corridors than the corresponding Euclidean ones. Looking at the separated section for *wormhole 1* (Fig. 4.5), the graph also shows shrunken and contracted corridor branches (displayed as sections 4 and 5) which are better distinguished in the non-metric MDS embedding as can be seen in figure A.3d. The Eigenvalue plot (Fig. A.3g) for this graph interestingly shows a less steep slope between the second and third Eigenvalue than it did for Euclidean



(a) Visualisation of travelled ground truth route (corridor-wise) according to ground truth coordinates type 1 and depiction of colour-code for sections. (b) Procrustes transformed classical MDS result superimposed on corresponding ground truth coordinates. (c) Traversing the wormhole starting from the other end point: resulting pointing does not have the same angle in (b).

Fig. 4.5: Non-Euclidean maze exploration of section *wormhole 1* with ground truth coordinates type 1: one round through the wormhole. Subordinate sections are colour-coded according to the legend in (a). Pointings: magenta lines correspond to pointings in the MDS representation while red lines show expected pointings based on ground truth coordinates of type 1.

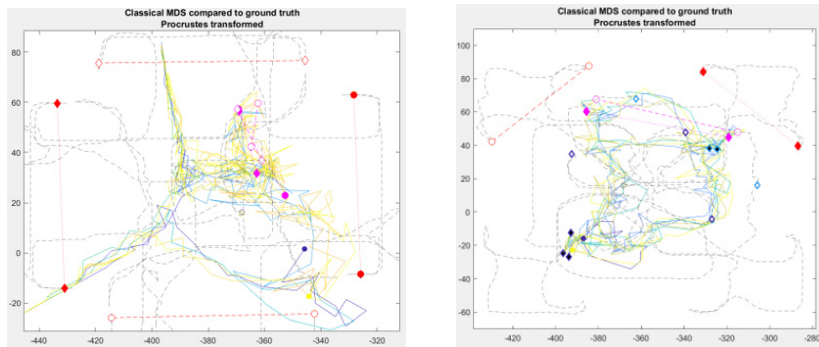
conditions (compare Fig. A.1f). Based on that, it can be inferred that the third dimension could provide additional useful information for describing the graph layout. For example, this can be visually verified when looking at corridors 4 and 5 in a three-dimensional embedding (Fig. A.3h) where they are actually longer than they seem to be in the two-dimensional embedding. An additional length information could not be observed in three-dimensional results for the Euclidean condition. Rotating the three-dimensional embedding also reveals an approximate orthogonality of the main loop (sections 1/2) with section 3 as well as with sections 4/5, i.e. each segment lies in a different plane which could be used as additional directional information.

Figures 4.6a and A.4 show further example explorations for section *wormhole 1* with multiple rounds traversing the wormhole which is of interest because the participants in Warren *et al.* (2017) had to visit each wormhole at least twice. It can be observed that all rounds are embedded into a consistent structure similar to the one resulting from just one dense round like in figure 4.5 despite the ground truth coordinates showing distinct rotated rounds. For two rounds (Fig. A.4a), a clear orientation towards the mean of both rounds is shown (lower left corner) while the entire graph structure is scaled down. This is an effect of the Procrustes transformation which attempts to minimise the squared distance error between each pair of points. Because many edges in the graph connect nodes which were detected in the same corridor but belong to different rounds and therefore have largely distant ground truth coordinates, Procrustes tries to embed these nodes by positioning them in between the competing "round-locations", i.e. having the same distance to both "round-locations". Extending the exploration to four rounds, an extreme down scaling of the overall graph layout can be observed as illustrated in figure A.4b. This can also be traced back to the mean-taking effect of the Procrustes transformation with four competing ground truth rounds which all need to be taken into account when minimising the squared distance errors. With four

rounds no meaningful orientation can be obtained. Despite the loss of orientation, the embedded graph (Fig. 4.6a) shows a consistent grouping of the same objects especially for a non-metric MDS embedding (Fig. A.4d) which substantiates the correct alignment of data from different rounds. Non-metric MDS also shows better distinguished corridor branches but less bundled corridors with more noise compared to the classical MDS result (Fig. 4.6a). A section-wise Procrustes transformation counteracts the correct alignment of multiple rounds into one consistent layout and yielded four differently oriented representations for each round which were pulled apart like illustrated in figure A.4e.

Described observations regarding the object grouping and differences between classical and non-metric MDS can be reproduced with a second exploration over four rounds (see Fig. A.4f and Fig. A.4g) except for a mirrored orientation compared to the first exploration example and more noise included in the non-metric representation. Interestingly, the corresponding Eigenvalue plot (Fig. A.4h) shows a proportional high value for the fourth dimension compared to the third Eigenvalue in this example.

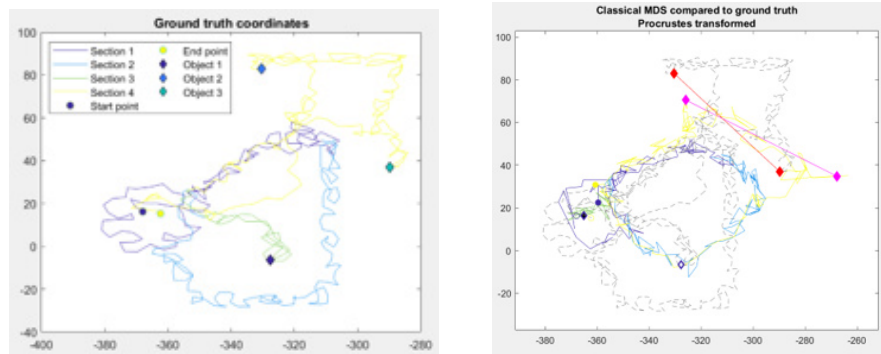
The isolated *wormhole 2* section shows for one round (Fig. 4.7) a correct alignment of the lower seg-



(a) Section *wormhole 1*: classical MDS result enlarged by a factor of 35 and superimposed on ground truth coordinates. Pointing: pair-wise pointings between the same objects belonging to different rounds.  
 (b) Section *wormhole 2*: classical MDS result enlarged by a factor of 10, superimposed on ground truth coordinates. Pointings: two example pointings of the same object pair.

Fig. 4.6: Non-Euclidean explorations: multiple rounds through each wormhole section.

ment of the inner loop (light blue) with the wormhole path (yellow) which in contrast do not lie on top of each other in the ground truth coordinates. This yields a much smaller wormhole loop compared to the Euclidean condition. As it was observed in section *wormhole 1*, the inner corridor (green) branching off the base does not reproduce the original ground truth orientation in the classical MDS as well as the non-metric MDS representation (Fig. A.5d). The Eigenvalues for classical MDS indicate no important descriptive role of a third dimension as illustrated in figure A.5h. Nevertheless, three-dimensional embedding (Fig. A.5g) reveals an almost orthogonal orientation of the wormhole loop (yellow) with respect to the inner loop (dark and light blue). A second exploration with four rounds through *wormhole 2* yielded a similar layout compared to the one-round exploration. The overall representation was shrunk which is why figure 4.6b shows the MDS result enlarged by a factor of 10. In this example, section-wise



(a) Visualisation of travelled ground truth route (corridor-wise) according to ground truth coordinates type 1 and depiction of colour-code for sections. (b) Procrustes transformed classical MDS result superimposed on corresponding ground truth coordinates.

Fig. 4.7: Non-Euclidean explorations of section *wormhole 2* with ground truth coordinates type 1. Pointings: magenta lines correspond to pointings in the MDS representation while red lines show expected pointings based on ground truth coordinates of type 1.

Procrustes transformation also resulted in distinct represented rounds pulled apart. The Eigenvalues imply no particular important role for a third dimension.

**Procrustes Type 2.** A complete exploration of the maze transformed by Procrustes with type 2 ground truth coordinates shows the familiar loop structure for both the left half and the right half of the maze (Fig. 4.8a and for more details Fig. A.6a - A.6d). In this example, the loop of the left half, i.e. *wormhole 1* section, was mirrored compared to the ground truth reference during MDS which was not resolved by applying Procrustes transformation. When isolating this section in the same exploration (Fig. 4.8b, A.6f, and A.6g) or when using Procrustes section-wise like illustrated in figure A.6d, the mirroring disappears. However, another exploration only covering section *wormhole 1* and conducted in a corridor-wise instead of round-wise manner (see Fig. A.7b) also yielded a mirrored wormhole loop. The MDS representation especially for section *wormhole 1* was shrunken in general but figure A.7f (isolated wormhole corridors) and section-wise Procrustes (Fig. A.6g) yielded the least down scaling. In contrast to type 1 Procrustes, multiple rounds do not increase the down scaling factor since the ground truth coordinates do not change between rounds. Another difference is that instead of contracting the two lower branching corridors of section *wormhole 1* with one another, the left one of those (dark blue object markers) was merged with the wormhole loop like illustrated in figure 4.8b. This could be a consequence of more alias edges between the loop and the concerned branch in the constructed graph. Similar to previous observations, the separation of those corridors is better in the non-metric MDS result (see Fig. A.6f) and even more exhibited in the section-wise Procrustes representation (Fig. A.6d). Taking a closer look at the second exploration of section *wormhole 1* (Fig. A.7b) with respect to this point, it can be observed that here lower corridor branches (displayed in green and ocker in figure A.7b and A.7c) are neither merged with one another nor with the loop but are well separated, especially in the non-metric MDS result (Fig. A.7d)

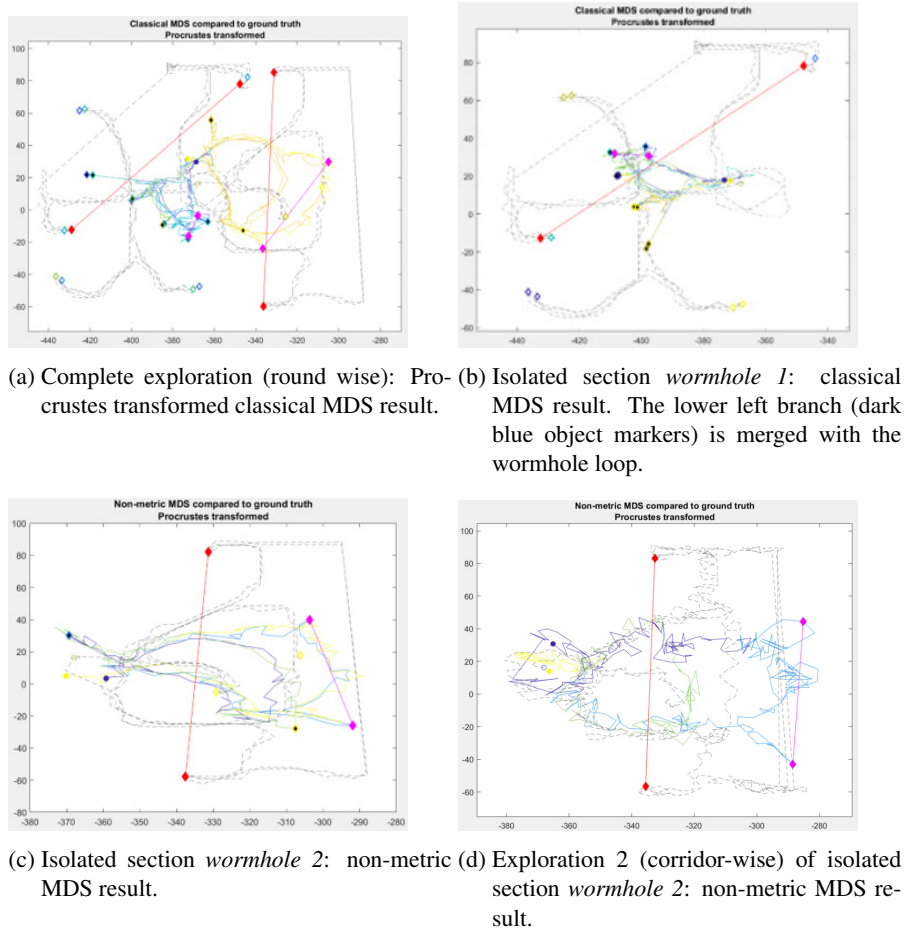


Fig. 4.8: Non-Euclidean explorations with ground truth coordinates type 2. Pointings: magenta lines correspond to pointings in the MDS representation while red lines show expected pointings based on ground truth coordinates of type 2.

and the section-wise Procrustes transformed result (Fig. A.7e).

### 4.3.2 Pointing

To evaluate pointing performances in the MDS representation, both the expected pointing direction from respective ground truth coordinates as well as pointing results from Warren need to be considered. Thereby, non-Euclidean conditions are interpreted with respect to the results from the Euclidean condition.

#### Euclidean maze

The MDS representation in figure 4.4b is able to reproduce expected pointings based on the ground truth which is shown by approximate parallel lines of computed pointing (magenta lines) and ground truth pointing (red lines). Different explorations yield slightly different angles for respective pointings,

nevertheless, this is in range of tolerance when compared to the high variance for pointing direction in Warren *et al.* (2017). As it is shown in figure A.1d, non-metric MDS for exploration 1 computes pointings which are even more parallel to reference pointings than section-wise transformed MDS (Fig. A.1g).

In contrast, exploration 2 shows greater deviations for pointings, e.g. section *wormhole 2* predicts a pointing with an orientation of  $45^\circ$  compared to the ground truth pointing (Fig. A.2a). Here, non-metric MDS can improve pointing predictions for *wormhole 1* as shown in figure A.2b but for *wormhole 2* only section-wise Procrustes can resemble the expected pointing direction better while also showing an even more accurate pointing for *wormhole 1* (see Fig. A.2c).

### Non-Euclidean maze

In the non-Euclidean condition, the MDS representation was tested for wormhole biases comparable to those found by Warren *et al.* (2017). Since no particular effect could be observed for a non-metric MDS computation compared to classical MDS, the analysis of pointing performance for non-Euclidean conditions focused on classical MDS results.

**Procrustes Type 1.** With ground truth coordinates of type 1, reference pointings to which the MDS result is compared represent expected pointings toward the wormhole location of the target object. When only considering one round through the respective wormhole section, the computed MDS pointing is close to the ground truth pointing (which in this case represents the wormhole-position of the target object) demonstrating a bias of slightly less (wormhole 1, Fig. 4.5) and slightly more (wormhole 2, Fig. 4.7) than  $45^\circ$ . In addition, pointings between the same pair of objects but computed on two different representations for entering the wormhole from each end point respectively (compare 4.5b and 4.5c) yields non-opposite pointings as expected with regards to results from Warren *et al.* (2017). However, the angle between non-opposite pointings deviates from the observed  $75.27^\circ$  in the Warren experiment.

When interpreting results for four rounds through *wormhole 1*, it needs to be considered that the overall orientation of the MDS representation is not aligned with the original maze. Therefore, pointing directions can only be characterised relatively to the expected ground truth pointings for the respective round. Figure 4.6a displays pair-wise pointings between the same objects belonging to different rounds. It can be observed that all four MDS pointings show a diagonal orientation, i.e. a bias close to  $45^\circ$ , with regards to their reference pointings although it has to be kept in mind that the Euclidean condition also showed slightly non-parallel pointings. In figure 4.6b, two example pointings are shown for section *wormhole 2* with diagonal MDS pointings as well, although a comparison to the expected pointing is difficult to interpret here, too.

**Procrustes Type 2.** Investigating section *wormhole 1* with Procrustes transformation type 2 yields an angle difference of approximately  $45^\circ$  between ground truth and MDS pointing, predicting a bias towards the wormhole location of the target object and away from the actual location (shown by ground truth co-

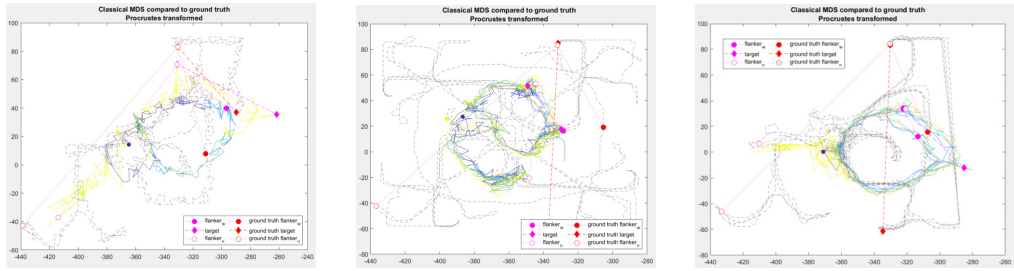
ordinates) as displayed in figures 4.8b and A.6e-A.6g. The bias for section *wormhole 1* is less prominent when looking at the MDS result for the complete maze (Fig. 4.8a) and even disappearing in the second example with mirrored orientation (Fig. A.7, except for figure A.7f displaying isolated wormhole corridors). In contrast, section *wormhole 2* shows a less prominent bias when considering the isolated section (Fig. 4.8c and A.6h-A.6j) compared to the whole maze representation (Fig. 4.8a, A.6c) where a bias of  $45^\circ$  can be observed except for the section-wise Procrustes representation in figure A.6d. Furthermore, pointings between the same pair of objects do not show a bias dependent on which end point was used to enter the wormhole.

### 4.3.3 Folds in spatial Representations

Checking the MDS representation for possible folds like defined in Warren *et al.* (2017) can give insight into how individual parts of the maze are embedded in relation to each other. The evaluation concerning folds was only conducted on results for section *wormhole 2* because of distorted orientations of the within-loop corridor branch in examples for section *wormhole 1* and due to the fact that in Warren *et al.* (2017) there only exists a plot of found foldings in section *wormhole 2* for direct comparison.

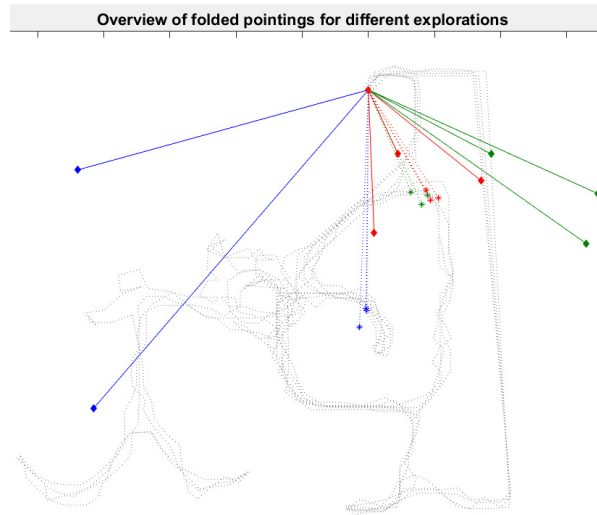
Individual results from the investigation of possible foldings in the spatial representation are depicted in figure 4.9. With type 2 Procrustes (Fig. 4.9c), a folding similar to that found in Warren *et al.* (2017) can be observed by switched pointings to the target object and the wormhole flanker. Compared to the actual location of those objects displayed by the ground truth, the target object is also ripped away from the neutral flanker. MDS pointings for Procrustes type 1 with only one round shows similar rips and foldings results and can be used to illustrate that these results resemble the expected pointings according to the object's wormhole location depicted as the ground truth for Procrustes type 1 (see Fig. 4.9a). When looking at an exploration with four rounds for Procrustes type 1 (Fig. 4.9b) foldings do not occur. Still, target and wormhole flanker are represented almost at the same location, i.e. moved closer together while in contrast to previous results target pointings are also not ripped away from pointings towards neutral flankers. Figure 4.10 shows an overview for observed pointings where Procrustes type 2 and Procrustes type 1 with one round are illustrated with solid lines while Procrustes type 1 with multiple rounds is depicted by dashed lines. One example for Procrustes type 2 was conducted without a neutral flanker therefore the plot contains one more pair of target and wormhole flanker pointings than there are neutral flanker pointings. As can be seen when comparing the results (4.10a) to found pointings from Warren *et al.* (2017) (4.10b), target pointings (green) are in general ripped away from their neutral flanker (blue) while being folded over the respective wormhole flanker (red). Procrustes type 1 with multiple rounds poses an exception to this, only showing closer target and wormhole flanker pointings but without foldings.






(a) Procrustes type 1: pointing according to the wormhole location of objects. (b) Procrustes type 1 with ground truth pointing according to the actual location of objects in the maze. (c) Procrustes type 2: ground truth pointing according to the actual location of objects in the maze.

Fig. 4.9: Experiment 2: Folded pointings.



(a) Procrustes type 1 (one round) and Procrustes type 2: solid lines. Procrustes type 1 (multiple rounds): dashed lines.



(b) Visualisation of rip and folds results from Warren *et al.* (2017).

Fig. 4.10: Overview of folded pointing results. Blue lines indicate neutral flanker pointings, green lines target pointings and red lines wormhole flanker pointings. In accordance with the results from Warren, an ordinal reversal of target pointings with respect to wormhole flanker pointings can be observed for MDS embedded representations. This reversal is revoked when using Procrustes transformation type 1 on multiple rounds through the wormhole.

# CHAPTER 5

## DISCUSSION

### 5.1 Evaluation of Results

The present study simulated experiments on humans conducted by Warren *et al.* (2017) with a graph approach based on visual features. In the course of the simulation, graph representations of non-Euclidean as well as matching Euclidean environments were built and embedded with Multidimensional Scaling providing a two dimensional layout and Procrustes transformation serving as a reference frame for orientation. The respective resulting representation of the environment was analysed with respect to findings in the Warren experiment including shortcut (here pointing) performances.

#### 5.1.1 Summary of Results

Both results from finding paths in the original graph as well as the MDS embedding show that relative path lengths are represented by the graph. For the non-Euclidean maze in particular, shorter wormhole routes are chosen by the navigation algorithm and the MDS representation yields smaller wormhole loops.

As for mapping performances, the overall layout of the maze, round-walks and intersecting corridors are captured. This poses a comparable performance to human behaviour as observed in Warren *et al.* (2017) where participants were able to reproduce the overall maze layout, object configuration as well as intersections. Besides that, the model was not able to reconstruct curves and corners due to the lack of angle information and represented distances better than orientations. One reason could be the limited spatial resolution and maybe performance with respect to this aspect could be improved by a very dense exploration. Also in agreement with the results from Warren, wormholes did not pose major problems for the model.

Non-metric MDS yields a better mapping performance with respect to distinguishing branching corridors in general. While it also shows an improvement for pointing performances compared to classical MDS in the Euclidean condition, this effect could not be found in non-Euclidean conditions where non-metric MDS rather added more noise to the result.

The model predicts pointings as described in Warren both for the Euclidean condition (also with variation) as well as for non-Euclidean conditions where a similar bias of 45° could be observed. However, this bias was less prominent in results for e.g. Procrustes transformation based on *type 2 ground truth data* while this approach could also not reproduce different pointings for the same object pair dependent

on which one is the start object. Procrustes type 1 on the other hand yielded better pointing predictions in general. However, the most apparent difference to type 2 is that the embedding of multiple rounds through the maze lack a meaningful orientation of the overall maze representation which makes pointings difficult to interpret. Nevertheless, a tendency towards replicating found biases from Warren can be observed in non-Euclidean conditions as shown for example in figures 4.5b and 4.7b. Especially when transforming MDS results section-wise with Procrustes, the majority of resulting representations show pointing biases very close to those expected in Warren. Also, using this type of reference for orientating MDS results improves mapping performances with respect to the orientation of individual corridors. The only weakness of section-wise Procrustes was observed for multiple rounds through the maze with Procrustes type 1, where it pulls identical corridors apart even though they were aligned during MDS embedding.

The second part of the pointing analysis investigated possible folds in the representation of the environment. Results show that pointings to the target object were folded towards its wormhole location indicating an ordinal reversal of objects (Fig. 4.10). Since Procrustes does not change the relation of points in the MDS representation, foldings were observed for both Procrustes types. In agreement with observations for pointing tasks, it can also be inferred that a Procrustes transformation of type 1 is not applicable for multiple rounds through the wormhole. In this representation, no foldings were identified and therefore this type of reference frame does not replicate results from the Warren experiment.

Thus, it can be inferred that MDS alone as a model for embedding is able to replicate found distortions in the form of rips and foldings caused by wormholes in humans. When additionally applying Procrustes transformation, pointings best resemble the results from Warren *et al.* (2017) with ground truth data modelling self-motion as a frame of reference for orientating the MDS result. On the other hand, results also indicate that self-motion without updating (i.e. Procrustes type 1 on multiple rounds) counteracts the folding effect as can be seen in figure 4.9b.

### 5.1.2 Interpretation of Results

Results from the present study have the potential to substantiate the following hypotheses about the organisation of spatial knowledge. Firstly, the replication of results from experiments on humans with a graph-based representation supports the hypothesis of a labelled graph as the structure for spatial knowledge. In particular, it was shown that biases and foldings caused by non-Euclidean environmental characteristics like wormholes can be reproduced. Based on that, it is inferred that a MDS-based approach is able to compute an embedding of topological feature graphs which resemble the layout of spatial knowledge in humans. Furthermore, it was also shown that even when using a graph approach, an embedding strongly anchored by globally consistent landmark information (modelled by *type 2 ground truth data*) cannot reconstruct human navigational performances. This poses further evidence against the Euclidean

map hypothesis which was already found to be less plausible in the experiment from Warren *et al.* (2017). Besides that, the plausibility of assuming that globally consistent landmark information is available in the first place is questionable.

Secondly, the presented MSN algorithm provides a labelled graph based on individual salient points in the environment while only the graph's connectivity was used as input to the MDS computation. Therefore, it can further be argued that results substantiate the theory that vision alone provides topological knowledge (given by the graph representation) while metric relations like path lengths and junction angles are added to the representation by embedding (here modelled by MDS) and by path integration (modelled by Procrustes type 1). Thereby, the replication of foldings in human spatial representation indicates that an embedding based on visual cues alone already shows relative positions of objects like they were found to be represented in humans.

Furthermore, results substantiate the approach to assign explicit metric information only a supporting role while strengthening the role of visual information for embedding and updating during the acquisition of spatial knowledge. As can be observed from results of Procrustes type 1, an internal reference frame based on self-motion information (e.g. path integration) is able to replicate pointing biases. In particular, the most promising approach seems to be orientating the representation according to locally experienced headings for each segment of the environment (i.e. corridors in the maze) which is modelled by section-wise Procrustes. This type of reference frame can be substantiated by findings of head direction cells (Knierim *et al.*, 1995) in rodents which are believed to encode allocentric heading directions. However, it was also found that local metric information can only play a supporting role and cannot be relied on, since wormholes would pose a problem otherwise. In fact, the embedding based on visual feature recognition (modelled by MDS) needs to be prioritised over orientating the representation based on experienced self-motion. This was observed in explorations with multiple rounds through a wormhole yielding competing experienced headings for the same corridor. One approach to integrate both the visual-based embedding and self-motion cues is to reset the path integration system providing these cues upon rediscovery of places. The results from Warren *et al.* (2017) indicate that the updating process even conceals major inconsistencies caused by wormholes because participants do not report any structural inconsistencies which would be noticed with conflicting angles acquired by path integration. From that, it can be inferred that that the overall experienced layout according to visual cues is more important in reconstructing the entire environment than exact metric information. This is also substantiated by reproducing folded pointings, i.e. relative object locations, based on the graph representation solely embedded by MDS.

## 5.2 Limitations and possible Extensions

Firstly, shown results are based on a rather small sample of different explorations, therefore one needs to be careful when trying to generalise observed effects. As shown in these examples, results depend among other things on the nature of the respective exploration behaviour (e.g. round-wise versus corridor-wise) which might limit comparability and at least makes it possible to find other effects in different example explorations. Therefore, a longer and more thorough investigation of the given environment has to be conducted to see if effects can be replicated. Furthermore, alternative environments should be investigated which differ from the present maze structure, i.e. which do not solely involve clear decision points given by corridor intersections. This could also give more data on how to choose parameter values for the employed algorithms which are currently strongly adapted to the present environment with e.g. its specific wall patterns. Another aspect putting the results into perspective is that variations in observed performances are rather high both in the experiment on humans (Warren *et al.*, 2017) as well as for the presented model. As Warren *et al.* (2017) already stated, it is difficult and rather a question of relative degree of error to infer a bias versus dismissing deviations from expected pointings as general inaccuracies.

Another limitation is that the application of MDS is only practicable for graphs up to a certain size. Even with a reduction of the graph like proposed in this work, huge graphs could result in larger environments. In order to process graphs of larger sizes, some speed-up methods for MDS can be considered. Examples for a more efficient computation of MDS like Pivot or Landmark MDS are described in Pich (2009). These approaches focus only on a (strategically selected) representative subset of nodes in the graph and therefore reduce the size of input data passed to the MDS algorithm. In addition, prior knowledge about the environment could be used to choose nodes representing the layout or even an initial configuration could be passed to the MDS algorithm starting from which it builds the optimal embedding solution.

On the other hand, computed stress by applying MDS could be helpful to gain more knowledge about possible alias connections in the graph. A node with a proportional high point-wise stress value is considered to have a relatively uncertain location which either indicates that this node is better represented with an additional dimension (Borg and Groenen, 2005) or, considering the present context, this could mean that edges connected to this node are alias connections and should therefore have a weaker influence on the entire embedding.

Another aspect which has to be noted is that different types of Procrustes transformations were performed on different explorations since respective ground truth coordinates are used. This limits the comparability of effects between both types and in order to obtain reliable differences, each exploration would need to be tested with both types of Procrustes transformation.

A further limitation relating to the evaluation of pointing performances is that quite some assumptions

based on Euclidean principles are involved in the analysis. First, classical MDS tries to reconstruct assumed Euclidean input data. Nevertheless, non-metric MDS yielded similar results at least for non-Euclidean conditions which implies that this assumption has no important influence on the obtained representations. However, also the pointing procedure, i.e. drawing a line between objects in a two dimensional layout, was based on Euclidean geometry. Hence, it needs to be considered that this procedure could not be adequate for analysing the specific layout and object configuration within the MDS representation. In general, one has to keep in mind that the MDS representation itself was visualised by adjusting it to the respective ground truth resembling the original maze. Therefore, structures found during the analysis of mapping performances could be projected onto the MDS representation from prior knowledge about the maze layout.

This also affects the proposed application of Procrustes transformation which is based on recorded ground truth coordinates. These coordinates reflect the experienced maze structure very accurately but are actually not included in the graph representation itself. To obtain a more plausible model for e.g. underlying reference orientations, ground truth coordinates could be replaced by a path integration system which would be also used during the generation of the graph representation. This path integration system could make use of the internal compass (a model for reference direction, comparable to an allocentric path integrator (Baumann, 2019)) employed in the MSN algorithm and would additionally track the agent's self-motion.

Just like Baumann (2019) mentions that the internal compass mostly relies on updating from feature positions, path integration with respect to travelled self-motion also depends on correction upon visual cues Hübner and Mallot (2007). This was also shown in the presented results in this thesis. a path integration system could support the embedding mechanism by interacting with it. Besides providing an internal reference frame for orientating the MDS embedding by Procrustes transformation, a first embedded representation of the environment currently under exploration could in turn support the updating step of the path integrator. Acquired local knowledge about path length and angles is often noisy and unreliable, therefore, it could be useful to update this local knowledge upon rediscovery of places according to that first embedded representation. Thereby, it is not necessary and probably less optimal to utilise a global representation of the entire environment but rather compare path integration results to embedded relations of current local segments. In order to enable such an interaction between path integration and embedding, the embedding step needs to be performed online during the exploration. For example, a MDS step could occur every time a loop closure is detected and stored path integrator information would be updated according to computed relations by the MDS step. Thereby, local measurements can be adjusted according to the overall experienced geometry but are also used as a reference frame for orientating the MDS result. For a method integrating multiple iterative solutions of an MDS embedding, see Borg and Groenen (2005). Of course, it needs to be determined what weighted influence each factor has, i.e. cer-

tain thresholds need to be set for both relying on path integrator data as well as relying on MDS relations. This involves a trade-off between avoiding inconsistencies caused by wormholes and losing too much direction information.

A path integration system could also be useful for the reduction of the graph. Summarising *feature nodes* to *snapshot nodes* can be made independent from ground truth coordinates by employing computed path integration position indicating features at the same location. Another extension could involve not only restraining the abstraction to snapshots (i.e. features at exactly the same location) but also to whole regions like corridors or squares. A study by Martinet *et al.* (2011) implies that hierarchical abstractions are employed in the prefrontal cortex ultimately yielding specific sets of neurons representing structural properties of the environment like corridors. Therefore, such abstractions could be plausible in terms of modelling human spatial knowledge.

Especially hierarchies pose an interesting approach for organising spatial representations. Although wormholes did not pose major problems for the model, section *wormhole 1* for example showed a proportional high Eigenvalue for the third dimension in some explorations. This indicates that a higher-dimensional representation could be more adequate, e.g. for representing the inner-loop corridor which spatially overlaps with the loop in two dimensional ground truth data. Such higher-dimensional representations could be implemented by means of hierarchies where relations of points in a third dimension are encoded by their relation on a higher level of the hierarchy. In addition, representing e.g. corridors as their own coherent structures on a respective level of the hierarchy enables the assignment of individual orientations to each corridor. This could substantiate the presented approach to separately orientate corridors which was found to yield best performances for the model. Evidence for a hierarchical structure in humans was found for example by McNamara and Diwadkar (1997) where participants' distance estimates between landmarks are influenced by the landmark's level of hierarchy. Additionally, Meilinger *et al.* (2016) observed that only the number of spatial subdivisions between target and participant influenced reaction time independently from the true Euclidean distance. Based on this, it may be argued that metric knowledge is in fact only included on a higher level of the hierarchy (Hübner and Mallot, 2007) which can be aligned with presented results indicating only a supporting role for metric information instead of posing a structure spatial knowledge is based on. Furthermore, it could also be investigated how a hierarchy with multiple levels of metric knowledge could model human behaviour. For example, local and global metrics could be represented on different levels of the hierarchy where for example local metric knowledge is used when navigating from one place to another. When combining individual local metric representations, global inconsistencies can arise which was shown for example in Warren *et al.* (2017) where shortcuts between the same pair of objects were non-opposite to each other which could also be reproduced in the present study with an embedding based on self-motion cues. With a graph-approach, local metric information does not need to be globally consistent. In addition, a globally

adjusted metric representation could also be stored on a different level of the hierarchy which enables a coherent reproduction of the experienced environment. This can be substantiated for example with the fact that participants in Warren *et al.* (2017) did not notice any anomalies in the overall wormhole environment while also being able to draw a consistent map of the experienced environment. With respect to the presented model, MDS is able to both provide a global embedding as well as a local one for individual segments while Procrustes transformation could add metric knowledge either on a global scale or only locally using respective frames of references. Considering the previous proposal to employ an interaction between MDS embedding and path integration during the generation of the graph, a hierarchical organisation could avoid the mentioned trade-off by simply encoding both representations on different levels. In addition, it would be interesting to investigate how MDS would perform on a higher level of a hierarchical representation since in the present study the graph is based on individual intensity features extracted from the environment. In order to model human spatial knowledge, it would be more plausible to consider geometric shapes and objects as "features" and see how a graph based on this could be embedded using the presented model. A study by Becker (2019) showed that hierarchically organised spatial abstractions like corridors can already be inferred from graphs generated by the MSN algorithm when applying a method called spectral clustering. These findings can also be connected to the reduction of the graph where extracted clusters could provide further abstraction from individual features to entire regions.

### 5.3 Conclusion

The present study simulated wormhole experiments on humans conducted by Warren *et al.* (2017). Thereby, a feature-based navigation algorithm was employed to construct a graph representation of matching Euclidean and non-Euclidean environments. Furthermore, a model is proposed for embedding constructed feature-graphs by Multidimensional Scaling and Procrustes transformation. Embedding the graph structure with this approach shows a comparable representation of the environment as observed in humans when being confronted with non-Euclidean environmental characteristics. In particular, wormholes do not pose major problems for the representation while relative object positions and distortions thereof can be reproduced solely applying Multidimensional Scaling to topological graphs. Based on that it can be inferred that structural information containing relative positions of objects can be acquired solely based on visual features without relying on explicit metric information.

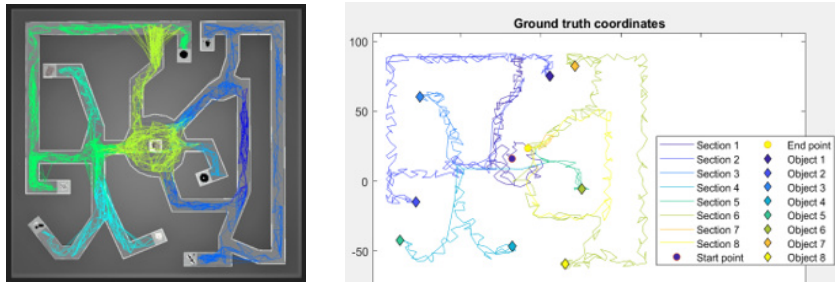
In adding Procrustes transformation as a frame of reference for orientation to the embedding procedure, directional shortcut biases can be replicated as well. From the investigation of different reference frames it can be inferred that self-motion seems to provide helpful supplementary information to feature-graphs for adding metric knowledge. In doing so, embedded graph representations reproduce spatial knowledge as it can be observed in humans. On the other hand, the results also emphasize that relying too strongly




on accurate self-motion cues does not provide a plausible spatial representation. It is therefore inferred that self-motion provided for example by a path integration system should only play a supporting role and therefore could be subordinated to visual-based embedding provided e.g. by Multidimensional Scaling. Future investigations could address effects of alternative frames of reference for the Procrustes transformation for example with varying degree of influence from self-motion and external global cues. Besides providing these ideas for further development of the Microsnapshot Navigation algorithm, the presented replication of experiments on humans with a graph-based approach substantiates the model's biological plausibility as well as provides further evidence for strengthening the labelled graph hypothesis in regards to the question about the organisation of spatial knowledge in humans.

# APPENDIX A

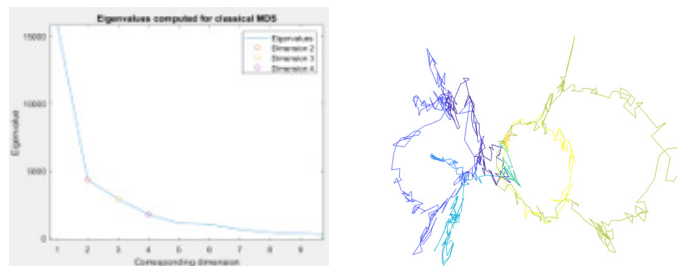
## FIGURES RESULTS



(a) Visualisation of the constructed graph. (b) Visualisation of travelled ground truth route (corridor-wise) and depiction of colour-code for sections.




(c) Procrustes transformed classical MDS result superimposed on corresponding ground truth coordinates. (d) Procrustes transformed non-metric MDS result. (e) Section-wise Procrustes transformation of the classical MDS result.




(f) Eigenvalue plot for the classical MDS result. (g) Manually rotated three dimensional embedding.

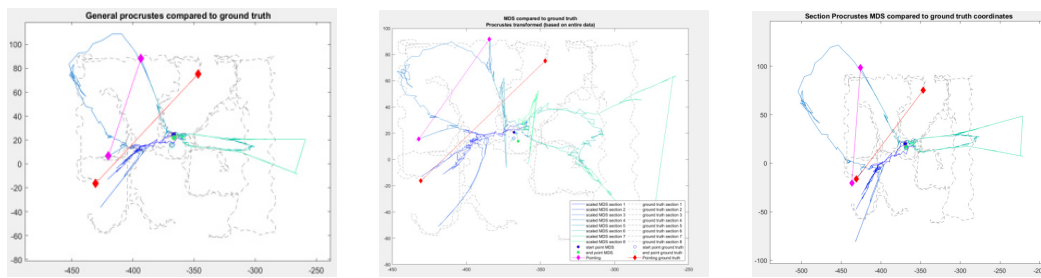
Fig. A.1: Euclidean maze exploration: sections are colour-coded according to the legend in (b) - except for the graph visualisation (a) where the same partitioning is shown with deviating colours. Pointings: magenta lines correspond to pointings in the MDS representation while red lines show expected pointings based on the respective ground truth coordinates.



(a) Exploration 2: Procrustes transformed classical MDS result superimposed on corresponding ground truth coordinates. (b) Exploration 2: Procrustes transformed non-metric MDS result. (c) Exploration 2: Classical MDS result, section-wise Procrustes transformed.

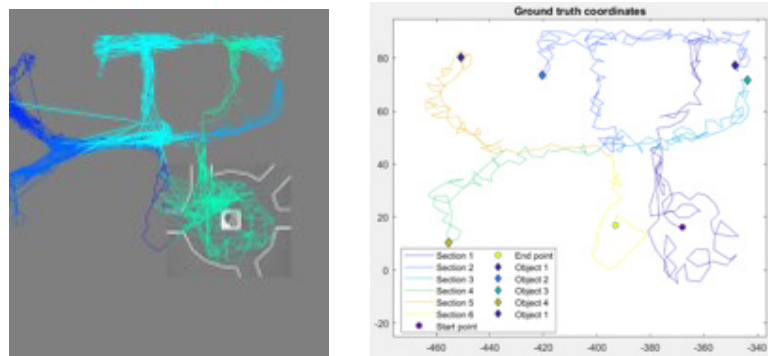


(d) Exploration 2 (multiple rounds through the entire maze). (e) Exploration 3: disconnected branch in section *wormhole 2* (turquoise).




(f) Exploration 3: classical MDS result. (g) Exploration 3: non-metric MDS result. (h) Exploration 3: classical MDS result, section-wise Procrustes transformed.


Fig. A.2: Further explorations of the Euclidean maze. Pointings: magenta lines correspond to pointings in the MDS representation while red lines show expected pointings based on the respective ground truth coordinates.



(a) Visualisation of the constructed graph. (b) Visualisation of travelled ground truth route (corridor-wise) according to ground truth coordinates type 1 and depiction of colour-code for sections.



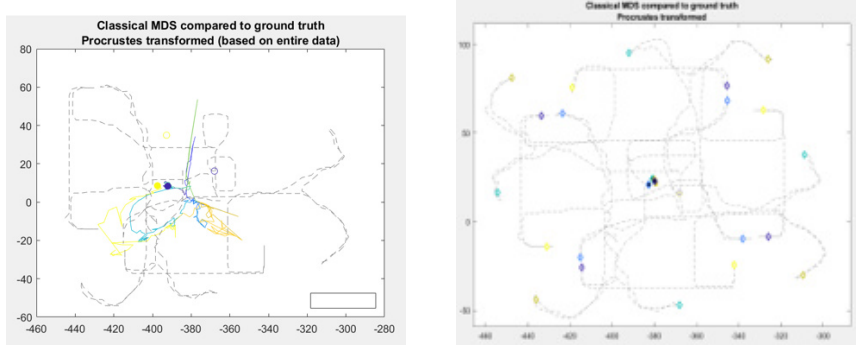
(c) Procrustes transformed classical MDS result superimposed on corresponding ground truth coordinates. (d) Procrustes transformed non-metric MDS result. (e) Classical MDS result, section-wise Procrustes transformed.




(f) Traversing the wormhole starting from the other end point: resulting pointing does not have the same angle as in comparison pointings (shown in e-g). (g) Eigenvalue plot the classical MDS result. (h) Manually rotated three dimensional embedding: sections 3 (middle blue), 4/5 (green and orange) and 1/2 (darker blues) are orthogonal to each other.

Fig. A.3: Non-Euclidean maze exploration of section *wormhole 1* with ground truth coordinates type 1: one round through the wormhole. Subordinate sections are colour-coded according to the legend in (b) - except for the graph visualisation (a) where the same partitioning is shown with deviating colours.


Pointings: magenta lines correspond to pointings in the MDS representation while red lines show expected pointings based on ground truth coordinates of type 1.



(a) Two rounds through wormhole 1: Pro- (b) Four rounds through wormhole 1: ground crustes transformed classical MDS result. truth and original size of the Procrustes transformed MDS result enlarged in c.

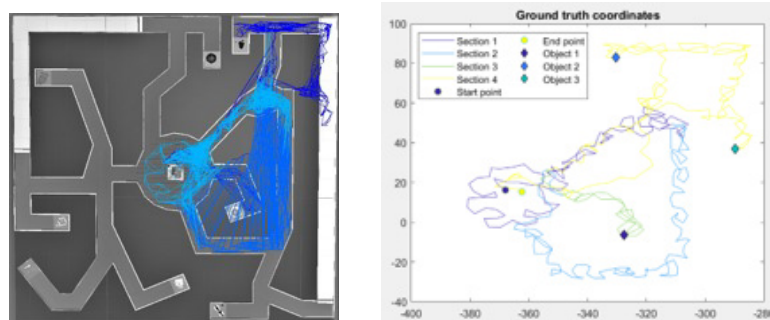


(c) Classical MDS result enlarged by a factor of 35 and superimposed on ground truth coordinates. Pointing: pair-wise pointings between the same objects belonging to different rounds. (d) Non-metric MDS version of A.4c. Markers correspond to all objects across the four rounds, identical objects are displayed in the same colour. (e) Section-wise Procrustes transformation of four rounds through wormhole 1 with one example pointing.

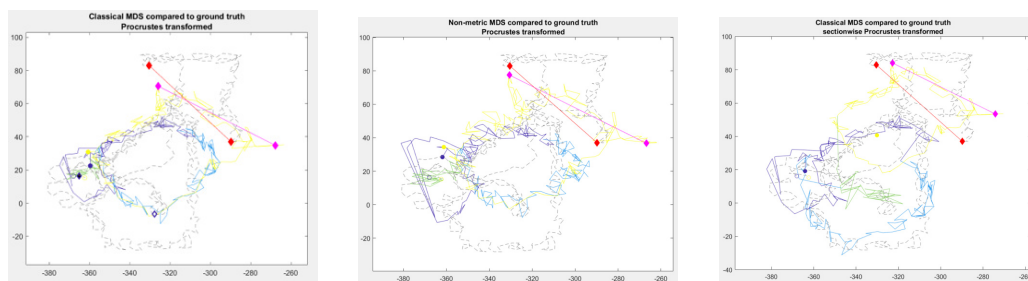


(f) Exploration 2 with four rounds: more densely explored (g) Exploration 2 with four rounds: non-metric MDS result. (h) Exploration 2 with four rounds: Eigenvalue plot with proportional high value for a fourth dimension.


Fig. A.4: Non-Euclidean explorations of section *wormhole 1* with ground truth coordinates type 1: multiple rounds through the wormhole. Pointings: magenta lines correspond to pointings in the MDS representation while red lines show expected pointings based on ground truth coordinates of type 1.



(a) Visualisation of the constructed graph. (b) Visualisation of travelled ground truth route (corridor-wise) according to ground truth coordinates type 1 and depiction of colour-code for sections.



(c) Procrustes transformed classical MDS result superimposed on corresponding ground truth coordinates. (d) Procrustes transformed non-metric MDS result. (e) Section-wise Procrustes transformed classical MDS result.






(f) Four rounds through the wormhole: classical MDS result enlarged by a factor of 10, superimposed on ground truth coordinates. Pointings: two example pointings of the same object pair. (g) Manually rotated three dimensional embedding: the section containing the wormhole (yellow) is bent towards the base (dark blue) and inner loop (middle blue). (h) Eigenvalue plot for the exploration shown in a.




Fig. A.5: Non-Euclidean explorations of section *wormhole 2* with ground truth coordinates type 1. Pointings: magenta lines correspond to pointings in the MDS representation while red lines show expected pointings based on ground truth coordinates of type 1.




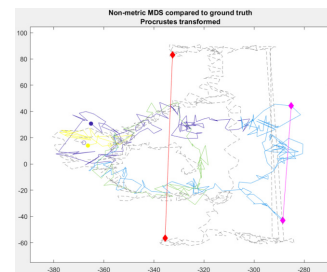

(a) Visualisation of the constructed graph.



(b) Complete exploration (round wise): Procrustes transformed classical MDS result. (c) Complete exploration: Procrustes transformed non-metric MDS result. (d) Complete exploration: section-wise Procrustes transformed classical MDS result.




(e) Isolated section *wormhole 1*: classical MDS result. The lower left branch (dark blue object markers) is merged with the wormhole loop. (f) Isolated section *wormhole 1*: non-metric MDS result. (g) Isolated section *wormhole 1*: section-wise Procrustes transformed classical MDS result.




(h) Isolated section *wormhole 2*: non-metric MDS result. (i) Exploration 2 (corridor-wise) of isolated section *wormhole 2*: non-metric MDS result. (j) Exploration 2 (corridor-wise) of isolated section *wormhole 2*: section-wise Procrustes transformed classical MDS result.

Fig. A.6: Non-Euclidean explorations with ground truth coordinates type 2. Pointings: magenta lines correspond to pointings in the MDS representation while red lines show expected pointings based on ground truth coordinates of type 2.



(a) Visualisation of the constructed graph. (b) Procrustes transformed classical MDS result. (c) Close-up on the classical MDS result with colour-coded sections for the exploration shown in a.



(d) Procrustes transformed non-metric MDS result. (e) Section-wise Procrustes transformed classical MDS result. (f) Isolated wormhole loop without branching corridors.

Fig. A.7: Additional non-Euclidean exploration of section *wormhole 1* with ground truth coordinates type 2. Pointings: magenta lines correspond to pointings in the MDS representation while red lines show expected pointings based on ground truth coordinates of type 2.



# APPENDIX B

## SUPPLEMENTARY INFORMATION

### B.1 Abbreviations

MSN	Microsnapshot Navigation (Baumann, 2019)
MDS	Multidimensional Scaling

### B.2 Software references

#### B.2.1 Unity

Unity<sup>®</sup> 2019.1.0f2, Unity Technologies, 2019, <https://unity.com/>.

##### Asset store

<https://assetstore.unity.com/>

#### B.2.2 Visual Studio

Microsoft Visual Studio 16.4.5, Microsoft, 2019, <https://visualstudio.microsoft.com/de/vs/>.

#### B.2.3 MATLAB

MATLAB 9.6 (R2019a), The MathWorks Inc., Natick Massachusetts.

##### All shortest paths algorithm

graphallshortestpaths, Johnson (1977), <https://de.mathworks.com/help/bioinfo/ref/graphallshortestpaths.html>

##### Multidimensional Scaling

mdscale, <https://de.mathworks.com/help/stats/mdscale.html>

cmdscale, <https://de.mathworks.com/help/stats/cmdscale.html>

##### Procrustes

procrustes, <https://de.mathworks.com/help/stats/procrustes.html>

## B.2.4 Code Libraries

### LEMON

LEMON Graph Library 1.3.1, A. Jüttner, B. Dezso and P. Kovacs, 2014, <https://lemon.cs.elte.hu/trac/lemon>.

### OpenCV

OpenCV 3.4.0 (2014), Bradski, G.. The OpenCV Library. Dr. Dobb's Journal of Software Tools, 2000, <https://opencv.org/>.

### Boost

Boost 1.66.0, Beman Dawes, David Abrahams (1998-2005), Rene Rivera (2004-2005), <https://www.boost.org/>.

## B.3 Image references

1. <https://i.pinimg.com/736x/2e/23/26/2e23269940c17d5878e730e0a75e2361.jpg>
2. <http://planetolog.com/maps/map-metro/big/stuttgart-metro-map.jpg>

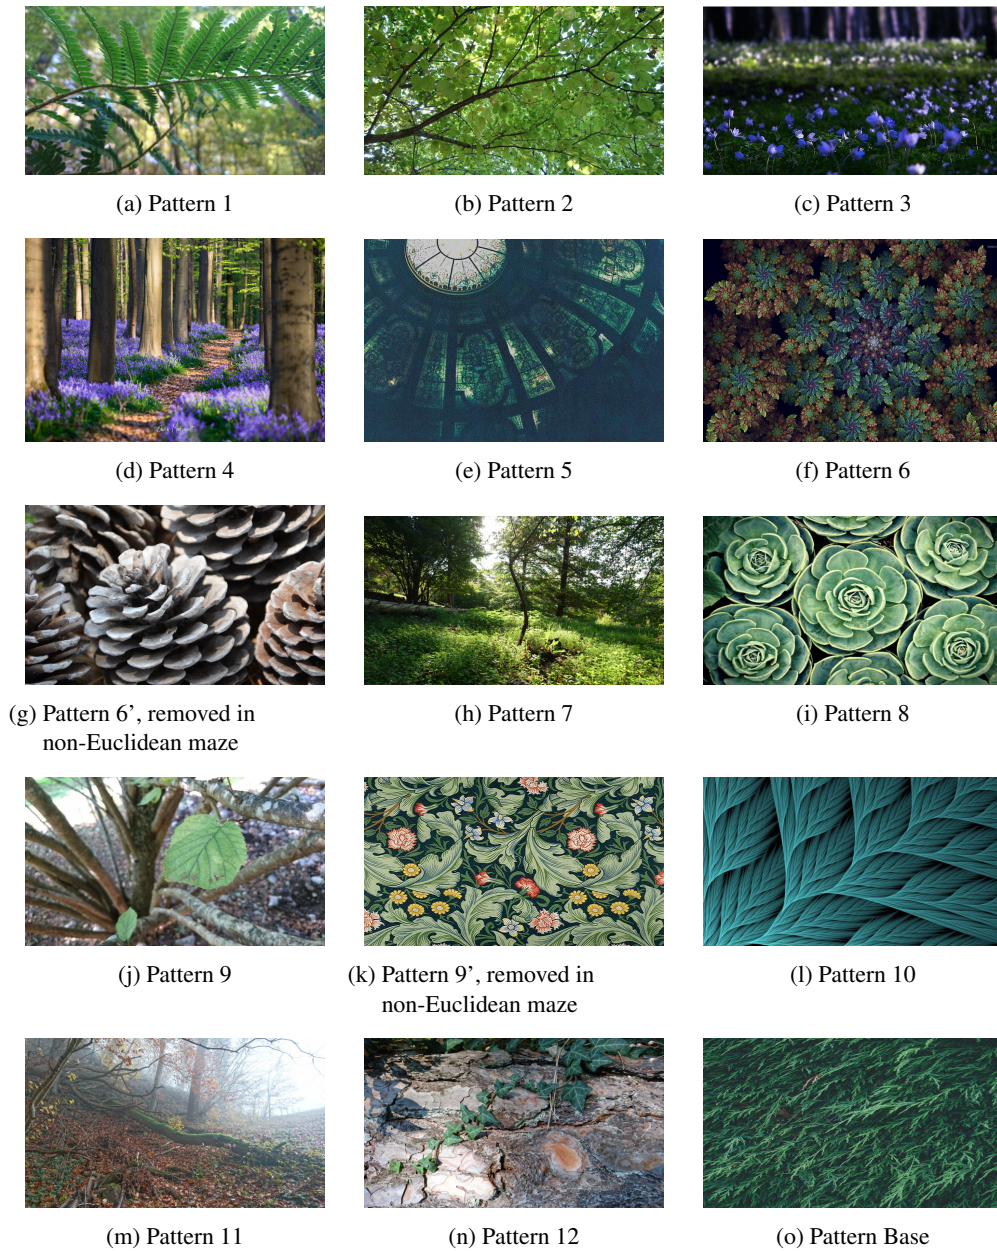


Fig. B.1: **Listing of patterns:** used to distinguish corresponding maze sections in figure 3.1.

Freely accessible resources: patterns 3, 4, 5, 6, 6', 8, 9', 10, Base.

Personal photographs taken by the author: 1, 2, 7, 9, 11, 12.

## LIST OF TABLES

1.1	Extracted from Warren <i>et al.</i> (2017): authors' predictions for the topological graph, the labelled graph and the Euclidean map hypothesis. . . . .	8
3.1	Parameter adaptations for the MSN algorithm. . . . .	23
3.2	Parameter settings for the graph reduction algorithm. . . . .	24

## LIST OF FIGURES

1.1	Maps: two types of spatial representations . . . . .	3
1.2	Models of spatial knowledge. Figure and description directly adopted from Warren <i>et al.</i> (2017). (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. . . . .	3
1.3	Experimental environment, directly adopted from Warren <i>et al.</i> (2017). Red lines correspond to paths taken in the route task. . . . .	8
1.4	Results of shortcut tasks from Warren <i>et al.</i> (2017). . . . .	9
1.5	Results for experiment two and maze sketches both adopted from Warren <i>et al.</i> (2017). . . . .	10
2.1	Example graph generated by the MSN algorithm, image from Baumann (2019). Blue lines are edges, long blue lines across non-traversable space of the environment are alias edges. Orange lines show path bundles found by the navigation mechanism. . . . .	14
3.1	Simulation Environment . . . . .	20
3.2	Ground truth types: visualisation of example graphs drawn according to each ground truth type. . . . .	21
4.1	Visualization of the graph reduction process described in section 3.2. The set of colours differs in the reduced graph but still shows the same partitioning into sections. . . . .	26
4.2	MDS results compared for the original and the reduced version of a simple graph. The reduced graph (c) shows less distinct points compared to when using the original graph (b) while also reconstructing the ground truth data (a) better. . . . .	27
4.3	Routes (orange lines) between objects 'Mushroom' and 'San Francisco House' (a, b) and objects 'Piano' and 'Car' (c, d) for the Euclidean maze (a, c) and the non-Euclidean maze (b, d). . . . .	27
4.4	Euclidean maze exploration: sections are colour-coded according to the legend in (a). Pointings: magenta lines correspond to pointings in the MDS representation while red lines show expected pointings based on the respective ground truth coordinates. . . . .	29
4.5	Non-Euclidean maze exploration of section <i>wormhole 1</i> with ground truth coordinates type 1: one round through the wormhole. Subordinate sections are colour-coded according to the legend in (a). Pointings: magenta lines correspond to pointings in the MDS representation while red lines show expected pointings based on ground truth coordinates of type 1. . . . .	30

4.6 Non-Euclidean explorations: multiple rounds through each wormhole section. . . . . 31

4.7 Non-Euclidean explorations of section *wormhole 2* with ground truth coordinates type 1. Pointings: magenta lines correspond to pointings in the MDS representation while red lines show expected pointings based on ground truth coordinates of type 1. . . . . 32

4.8 Non-Euclidean explorations with ground truth coordinates type 2. Pointings: magenta lines correspond to pointings in the MDS representation while red lines show expected pointings based on ground truth coordinates of type 2. . . . . 33

4.9 Experiment 2: Folded pointings. . . . . 36

4.10 Overview of folded pointing results. Blue lines indicate neutral flanker pointings, green lines target pointings and red lines wormhole flanker pointings. In accordance with the results from Warren, an ordinal reversal of target pointings with respect to wormhole flanker pointings can be observed for MDS embedded representations. This reversal is revoked when using Procrustes transformation type 1 on multiple rounds through the wormhole. . . . . 36

A.1 Euclidean maze exploration: sections are colour-coded according to the legend in (b) - except for the graph visualisation (a) where the same partitioning is shown with deviating colours. Pointings: magenta lines correspond to pointings in the MDS representation while red lines show expected pointings based on the respective ground truth coordinates. 45

A.2 Further explorations of the Euclidean maze. Pointings: magenta lines correspond to pointings in the MDS representation while red lines show expected pointings based on the respective ground truth coordinates. . . . . 46

A.3 Non-Euclidean maze exploration of section *wormhole 1* with ground truth coordinates type 1: one round through the wormhole. Subordinate sections are colour-coded according to the legend in (b) - except for the graph visualisation (a) where the same partitioning is shown with deviating colours. Pointings: magenta lines correspond to pointings in the MDS representation while red lines show expected pointings based on ground truth coordinates of type 1. . . . . 47

A.4 Non-Euclidean explorations of section *wormhole 1* with ground truth coordinates type 1: multiple rounds through the wormhole. Pointings: magenta lines correspond to pointings in the MDS representation while red lines show expected pointings based on ground truth coordinates of type 1. . . . . 48

A.5 Non-Euclidean explorations of section *wormhole 2* with ground truth coordinates type 1. Pointings: magenta lines correspond to pointings in the MDS representation while red lines show expected pointings based on ground truth coordinates of type 1. . . . . 49

A.6	Non-Euclidean explorations with ground truth coordinates type 2. Pointings: magenta lines correspond to pointings in the MDS representation while red lines show expected pointings based on ground truth coordinates of type 2. . . . .	50
A.7	Additional non-Euclidean exploration of section <i>wormhole 1</i> with ground truth coordinates type 2. Pointings: magenta lines correspond to pointings in the MDS representation while red lines show expected pointings based on ground truth coordinates of type 2. . .	51
B.1	<b>Listing of patterns:</b> used to distinguish corresponding maze sections in figure 3.1. Freely accessible resources: patterns 3, 4, 5, 6, 6', 8, 9', 10, Base. Personal photographs taken by the author: 1, 2, 7, 9, 11, 12. . . . .	54

## REFERENCES

- Baumann, T.** (2019). *Microsnapshot Navigation*. Master's thesis, Cognitive Neuroscience, Dept. of Biology, University of Tübingen.
- Becker, F.** (2019). *Identifying hierarchical structures in feature graphs by using spectral clustering*. Bachelor's thesis, Cognitive Neuroscience, Dept. of Biology, University of Tübingen.
- Borg, I.** and **P. Groenen** (2005). *Modern Multidimensional Scaling*. Springer.
- Brandes, U.** and **C. Pich** (2011). Fast Multidimensional Scaling of Network Distances and its Application to Large Bibliographic Networks.
- Cartwright, B.** and **T. S. Collett** (1983). Landmark learning in bees. *Journal of Comparative Physiology*, **4**, 521–543. ISSN 151.
- Chrastil, E. R.** and **W. H. Warren** (2013). Active and passive spatial learning in human navigation: acquisition of survey knowledge. *J. Exp. Psychol. Learn. Mem. Cogn.*, 1520. ISSN 39.
- Derdikman, D.** and **E. I. Moser** (2010). A manifold of spatial maps in the brain. *Trends in Cognitive Sciences*, **12**, 561–569. ISSN 14.
- Dijkstra, E. W.** (1959). A note on two problems in connexion with graphs. *Numerische Mathematik*, **1**, 269–271. ISSN 1.
- Duckett, T., S. Marsland,** and **J. Shapiro** (2002). Fast, On-Line Learning of Globally Consistent Maps. *Autonomous Robots*, 287–300. ISSN 12, doi:10.1023/A:1015269615729.
- Foo, P., W. H. Warren, A. Duchon,** and **M. Tarr** (2005). Do humans integrate routes into a cognitive map? Map- vs. landmark-based navigation of novel shortcuts. *J. Exp. Psychol. Learn. Mem. Cogn.*, 195–215. ISSN 31.
- Franz, M. O., B. Schölkopf, H. A. Mallot,** and **H. H. Bülthoff** (1998). Learning view graphs for robot navigation. *Autonomous robots*, 111–125. ISSN 5.
- Fyhn, M., S. Molden, M. P. Witter, E. I. Moser,** and **M.-B. Moser** (2004). Spatial representation in the entorhinal cortex. *Science*, **305**, 1258–1264. ISSN 5688, doi:10.1126/science.1099901.
- Gallistel, C. R.** (1990). Learning, development, and conceptual change. *The Organization of Learning*. Cambridge, MA: The MIT Press.
- Golledge, R.** and **R. Stimson** (1997). *Spatial behaviour: a geographic perspective*. Guilford, San Francisco.
- Hübner, W.** and **H. A. Mallot** (2007). Metric embedding of view-graphs. *Autonomous Robots*, 183–196. ISSN 23, doi:10.1007/s10514-007-9040-0.
- Johnson, D. B.** (1977). Efficient algorithms for shortest paths in sparse networks. *Journal of the ACM*, 1–13. ISSN 24.
- Knierim, J., H. Kudrimoti,** and **B. McNaughton** (1995). Place cells, head direction cells, and the learning of landmark stability. *J. Neurosci.*, **15**, 1648–1659. ISSN 11.
- Lu, F.** and **E. Milios** (1997). Globally consistent range scan alignment for environment mapping. *Autonomous Robots*, **4**, 333–349.
- Mallot, H. A.** (2018). Analysing Data with MATLAB - Mathematical Background. *Lecture notes*.
- Martinet, L.-E., D. Sheynikhovich, K. Benchenane,** and **A. Arleo** (2011). Spatial Learning and Action Planning in a Prefrontal Cortical Network Model. *PLoS Computational Biology*, **5**, e1002045. ISSN 7, doi:10.1371/journal.pcbi.1002045.



- McNamara, T. P.** (1986). Mental representations of spatial relations. *Cognitive Psychology*, 87–121. ISSN 87.
- McNamara, T. P.** and **V. A. Diwadkar** (1997). Symmetry and asymmetry of human spatial memory. *Cognitive Psychology*, 2, 160–190. ISSN 34.
- Meilinger, T., M. Strickrodt,** and **H. H. Bühlhoff** (2016). Qualitative differences in memory for vista and environmental spaces are caused by opaque borders, not movement or successive presentation. *Cognition*, 77–95. ISSN 155.
- Milford, M. J.** and **G. F. Wyeth** (2008). Mapping a Suburb With a Single Camera Using a Biologically Inspired SLAM System. *IEEE Transactions on Robotics*, 24, 1038–1053, doi:10.1109/TRO.2008.2004520.
- Moar, I.** and **G. H. Bower** (1983). Inconsistency in spatial knowledge. *Memory & Cognition*, 2, 107–113. ISSN 11.
- O’Keefe, J.** and **J. Dostrovsky** (1971). The hippocampus as a spatial map. Preliminary evidence from unit activity in the freely-moving rat. *Brain Research*, 171–175. ISSN 34.
- Pich, C.** (2009). *Applications of Multidimensional Scaling to Graph Drawing*. Doctoral thesis, Department of Computer & Information Science, University of Konstanz, Germany, Konstanzer Online-Publikations-System (KOPS). URL <http://nbn-resolving.de/urn:nbn:de:bsz:352-opus-83992>.
- Rothkegel, R.** (2000). *Die mentale Repräsentation von Distanzen*. Doctoral thesis, Department of Psychology, University Trier, Germany. URL <http://ub-dok.uni-trier.de/diss/diss11/20001108/20001108.htm>.
- Sadalla, E. K.** and **L. J. Staplin** (1980). The perception of traversed distance: Intersections. *Environ. Behav.*, 167–182. ISSN 12.
- Strickrodt, M.** (2019). *The impossible puzzle: No global embedding in environmental space memory*. Doctoral thesis, Department of Mathematics and Natural Sciences, University of Tübingen, Germany. URL <http://dx.doi.org/10.15496/publikation-31622>.
- Thrun, S., W. Burgard,** and **D. Fox** (2005). *Probabilistic robotics*. MIT Press, Cambridge.
- Warren, W. H.** (2019). Non-Euclidean navigation. *Journal of Experimental Biology*, 152–163. ISSN 222, doi:doi:10.1242/jeb.187971.
- Warren, W. H., D. B. Rothman, B. H. Schnapp,** and **J. D. Ericson** (2017). Wormholes in virtual space: From cognitive maps to cognitive graphs. *Cognition*, 152–163. ISSN 166, doi:http://dx.doi.org/10.1016/j.cognition.2017.05.020.
- Wilson, M. A.** and **B. L. McNaughton** (1993). Dynamics of the hippocampal ensemble code for space. *Science*, 1055–1058. ISSN 261.