AN ABSTRACT OF THE THESIS OF

Manish R. Saroya for the degree of Master of Science in Robotics presented on August 31, 2020.

Title: Topological Learning for Robotic Exploration and Navigation in Uncertain Environments

Abstract approved: _____

Geoffrey A. Hollinger

The effectiveness of robot autonomy is governed by the ability to make decisions based on online sensor measurements and a prior belief of the environment. Uncertainty in the environment introduces challenges to robotic decision making. This thesis address two key robot decision making problems: exploration and navigation. The robotic exploration problem requires a robot to maximally observe an unknown environment. A key challenge of robotic exploration is dealing with prior map uncertainty that arises due to partial knowledge of the world at the beginning of the mission. Since the map is unknown, it is difficult to determine which action gains the most information. The robot navigation problem is to determine a sequence of states in a given environment to move the robot from a start to goal state while avoiding collisions. Sensors mounted on real robots have measurement noise which leads to uncertainty in the mapping of obstacles. This uncertainty in

obstacle location makes it challenging to find collision free paths.

Real-world environments consist of interesting topological structures such as rooms, corridors, connections, intersections, and loops. These topological structures can be exploited to address uncertainty in the environment. Computational topology methods can be utilized to robustly extract environment topology. In this thesis, we resolve robot decision-making challenges due to environment uncertainty by leveraging topological methods that are robust to noise.

We first address prior map uncertainty by predicting unknown regions of subterranean environments. Our method involves a convolutional neural network and a novel loss that combines image inpainting features and topological features via persistent homology. Simulation results on four real mines and a dataset of procedurally-generated networks show our algorithm can exploit subtle topological structural cues in subterranean environments for efficient exploration.

We then address sensor measurement uncertainty by constructing a topologically accurate roadmap from a probabilistic occupancy map. We propose an unsupervised topological learning technique with biased sampling in topologically important regions. Simulation results show a robot can search a collision-free and cost-effective path over the topological roadmap in low runtime.

The proposed topological learning is a significant contribution to the two key problems of robot decision making. Our experimental results show a robot can successfully learn to exploit topological structures for online exploration of subterranean environments. Further we show that through topological biased sampling, a robot can construct a topologically accurate roadmap for robot navigation tasks.

©Copyright by Manish R. Saroya
August 31, 2020
All Rights Reserved

Topological Learning for Robotic Exploration and Navigation in Uncertain Environments

by

Manish R. Saroya

A THESIS

submitted to

Oregon State University

in partial fulfillment of
the requirements for the
degree of

Master of Science

Presented August 31, 2020
Commencement June 2021

Master of Science thesis of Manish R. Saroya presented on August 31, 2020.

APPROVED:

Major Professor, representing Robotics

Director of the Robotics Program

Dean of the Graduate School

I understand that my thesis will become part of the permanent collection of Oregon State University libraries. My signature below authorizes release of my thesis to any reader upon request.

Manish R. Saroya, Author

ACKNOWLEDGEMENTS

I would like to express my deepest appreciation to my supervisor Prof. Geoffrey A. Hollinger for providing direction and guidance to pursue this work. I would also like to extend my deepest gratitude to Dr. Graeme Best for mentoring my decisions and providing practical support.

I would like to thank everyone at Robotic Decision Making laboratory for insightful conversations, especially Bob Debortoli, Seth McCammon, Dylan Jones, Scott Chow, Chris Lee, Emily Scheide, Ian Rankin, Jake Anderson, and Jeff Caley. I am also grateful to thank Abhijeet Agnihotri, Anna Nickelson, Enna Sachdeva, and Dhruv Jawalkar for sharing great ideas. Special thanks to Sumedh Mannar, Sridhar Thiagarajan, and Kartik Gupta for providing a great experience at Oregon State University. I would like to thank Kush Prasad for providing guidance and Gaurav Gupta for assisting with graduate school applications.

I'm extremely grateful to my mother, father, sister Komal, grandmother, and grandfather who have always supported and nurtured me.

Many thanks to Oregon State University for providing Teaching Assistantship that has made this work possible. This research has been funded in part by DARPA under agreement #HR00111820044 and in part by NSF grant IIS-1845227.

TABLE OF CONTENTS

Page

1 Introduction 1
1.1 Exploration with Prior Map Uncertainty 4
1.2 Path Planning with Sensor Measurement Uncertainty 6
1.3 Contributions 8
1.4 Thesis Roadmap 9

2 Related Work 11
2.1 Robot Decision Making 11
2.1.1 Exploration with Prior Map Uncertainty 11
2.1.2 Path Planning with Sensor Measurement Uncertainty 12
2.2 Deep Learning for World Predictions 14
2.3 Unsupervised Topological Learning 14
2.4 Persistent Homology 15
2.5 Summary and Limitations 17

3 Online Exploration of Tunnel Networks Leveraging CNN-based World Predictions 19
3.1 Problem Formulation 20
3.2 Online Exploration Leveraging Topological World Predictions 21
3.2.1 World Prediction with Topological CNNs 22
3.3 Exploration Strategy 29
3.3.1 Explored Map and Mask Generation 29
3.3.2 Frontier Selection Policy and Path Planning 30
3.4 Experimental Results 31
3.4.1 Real-world Tunnel Networks 32
3.4.2 Procedural Generation of Mine-like Environments 32
3.4.3 Experimental Setup 34
3.4.4 Results 35
3.5 Summary 39

TABLE OF CONTENTS (Continued)

Page

4 Roadmap Learning for Probabilistic Maps with Topologically-biased Growing Neural Gas . . . 40
4.1 Problem Formulation . . . 40
4.2 Roadmap Learning for Probabilistic Maps with Topologically-biased Growing Neural Gas . . . 42
4.2.1 Topology of Probabilistic Occupancy Map . . . 43
4.2.2 Topology-biased Sampling Distribution . . . 46
4.2.3 Topologically biased Growing Neural Gas for Roadmap Generation . . . 47
4.2.4 Online Planning . . . 49
4.3 Experiments . . . 49
4.3.1 Scenario . . . 50
4.3.2 Comparison Methods . . . 51
4.3.3 Implementation Details . . . 52
4.3.4 Results . . . 52
4.4 Summary . . . 56

5 Conclusion and Future Work . . . 58
5.1 Summary of Contributions . . . 58
5.2 Future Directions . . . 59

Bibliography . . . 61

LIST OF FIGURES

Figure Page

1.1 An example set of realistic subterranean tunnel networks from our procedurally generated dataset. These environments contain structure in the form of topological features (loops and connected passages). Our proposed method predicts the presence of these features in unexplored space to inform frontier selection strategies for robotic exploration. 5

1.2 Hilbert Map: An example of a probabilistic occupancy map constructed from the Intel Research Lab pointcloud dataset [1]. The bar on right shows the probability spectrum with 0 as free space, 0.5 as unknown space, and 1 as occupied space. Our proposed method extract topological features such as connected component and loops from this map and construct a topologically accurate roadmap for efficient path planning. 7

3.1 Example homology persistence diagram (see Sec. 3.2.1.1). Varying a prediction threshold between 0 and 1 (pixel intensity) results in the birth and death of 0 and 1-dimensional topological structures: connected passages and loops, respectively. Strong topological features persist over a wide range of pixel intensities, while weaker ones exist over a more limited range. The three main loops in the world (see (a)) persist across nearly all threshold levels; therefore, there exists three 1-dim features with $(birth, death)$ near $(0, 1)$, shown as three red dots in the top-left of (b). 23

3.2 U-Net architecture for an image size of 32×32 pixels. The downsampling layers (left) aims to capture spatial context and the upsampling layers (right) accurately localise these features. 26

3.3 The four real-world mine networks used in the experiments, which have been manually snapped to suitable grids while preserving the actual topological structure. The mine entrance and robot start location is shown as a red dot. 33

LIST OF FIGURES (Continued)

Figure | Page

3.4 Example trajectories when for the Edgar Experimental Mine when using our proposed method and the closest first method. Robot (blue) navigates from the start (red) at bottom centre. At each timestep, the robot decides which frontier (orange) to navigate to next. For the top row, the areas predicted to be *open* are shown in green. These lie within the unmasked area of interest (shaded). The percentage of the environment explored is shown underneath. . 36

3.5 Comparison between the exploration strategies on the four real-world mines depicted in Fig. 3.3. 38

4.1 Persistent homology diagram (see Sec. 4.2.1). (a) shows most persistent death locations blue and red corresponding to 0-dim and 1-dim topological structures in (b) respectively. (b) consist of birth and death of topological structures computed via computational topology methods. 45

4.2 2D pointcloud of Intel Research Lab 50

4.3 The same start and goal pair was queried on both (a) and (b). GNG-topological (b) was able to find a path while GNG (a) failed. 53

4.4 PRM-4000 build on Hilbert map. The figure shows PRM can obtain a topologically accurate roadmap when a high number of vertices are selected. However this increases the runtime search requirements. . 54

4.5 PRM-1208 build on Hilbert map. The figure shows PRM-1208 is not able to obtain a topologically accurate roadmap, resulting in a low success rate. 55

LIST OF TABLES

Table | Page

3.1 Layer parameters for the U-Net architecture illustrated in Fig. 3.2, for an environment graph of size 32×32. 28

4.1 Roadmap methods evaluation on Hilbert map for 500 start-goal queries. Average path cost and Dijkstra vertices explored are shown for 431 paths. We compute Success weighted Path Cost (SPC) as described in Sec. 4.3.1 for 500 paths. 57

Chapter 1: Introduction

In this thesis, we investigate mobile robot navigation in the physical world. When navigating in real-world environments, robots must reason about the environment to search for cost-effective paths while avoiding collisions. Many environments have interesting topological structures leading to interesting navigational challenges. The topological structures can include loops, dead-ends, corridors, rooms, and narrow passages. Although these environmental structures bring up navigational challenges, these topological structures can be exploited with the right technique. Two key environments of such types include tunnel networks and structured indoor environments.

Robots are increasingly being tasked with exploring unknown worlds to provide rapid situational awareness to end-users such as mining companies and search and rescue teams. A vitally important domain for robotic exploration is subterranean environments [2, 3, 4, 5], including mining tunnels, urban underground facilities, and natural caves, as exemplified by the ongoing DARPA Subterranean Challenge [6]. These subterranean environments have interesting topological structures such as tunnels, shafts, and intersections. Communication is particularly difficult in these environments, which makes it essential for robots to act autonomously. However, the typically unknown, complex, and sparse structure of underground tunnel networks impedes the ability to make fruitful online decisions. We propose

methods for predicting topological features of partially-known tunnel networks to improve the effectiveness of exploration strategies.

Similar to the increasing application of robots in subterranean environments, robots are increasingly finding application in known environments such as offices and homes to carry out basic navigation tasks. The indoor robot navigation tasks include cleaning [7], helping elderly [8], and delivering items [9]. In these environments, a key part is being able to find a collision-free and cost-effective path to navigate from an initial state to goal state. To successfully complete navigation tasks, robots have to reason about the topological structures such as rooms, corridors, and narrow passages. In this work, we propose a novel method to develop a topological roadmap over which a cost-effective and collision-free path can be effectively planned.

To complete navigation tasks the robot has to make important decisions based on its observations from sensors. Robotic decision making has to be autonomous as communication is often difficult in subterranean environments. Online decision making enables robots to make informed decisions during a particular mission.

In this thesis, we concentrate on two fundamental decision-making problems: robot exploration and path planning. The principal idea behind the robot exploration problem is *given what the robot knows about the environment, what decisions should it make in order to gain maximal information of the environment* [10]. The aim is to select navigation decisions that maximize the fraction of an unknown environment that is visited by the robot in given mission time. The path planning problem relates to computing a sequence of states in a given environment to move

the robot from start to goal state while avoiding collisions. The main objective of the path planning problem is to find a collision-free and cost-effective path that connects the initial state to the goal state.

Both exploration and path planning problems need to deal with uncertainty about the world. In this thesis, we focus on two important sources of uncertainty: prior map uncertainty and sensor measurement uncertainty. Prior map uncertainty arises due to partial knowledge of the world at the beginning of exploration [10, 11]. The robot is uncertain about most of the unknown world. Sensor measurement uncertainty leads to uncertainty in mapping the obstacles in the environment leading to probabilistic occupancy maps. The robot has to reason about the probabilistic free space connections between the narrow passages in probabilistic occupancy maps. Thus it becomes difficult to compute a collision-free path in probabilistic occupancy maps.

In this thesis, we are interested in resolving the challenges of robot decision making by leveraging topological methods. Topological methods are relevant because these environments have interesting topological features that can be exploited. Computational topology based methods such as persistent homology [12] are robust to noise [13] and thus can be used to alleviate uncertainty in robot decision making.

Persistent homology methods can extract topological structures of the world but cannot directly make navigation decisions for the robot. In this thesis, we propose a combination of persistent homology with deep learning and unsupervised topological learning techniques. In Chapter 3, we present a novel technique to

predict unknown environment using a U-net architecture and a novel loss function which combines image inpainting features and topological features via persistent homology. These topological predictions then inform a frontier based exploration strategy substantially increasing the rate of area explored. Later in Chapter 4, we propose a novel technique to create a topologically accurate roadmap using an unsupervised topological learning technique with biased sampling in important regions computed via persistent homology. The aim is to search for a cost-effective and collision-free path over the constructed roadmap.

In the remainder of this section, we further outline the context and contribution of this thesis. In Section 1.1 we introduce exploration with prior map uncertainty and in Section 1.2 we introduce path planning with sensor measurement uncertainty. Section 1.3 states the primary contributions of this thesis. Section 1.4 provides a roadmap of this thesis.

1.1 Exploration with Prior Map Uncertainty

At the core of robotic exploration tasks is a decision making problem where a robot is required to adaptively plan navigation goals. Typical approaches involve maintaining a list of *frontiers*, where these frontiers define boundaries between open and unknown space, and the decision to be made at each time step is which frontier to explore next. Without any knowledge of the unknown environment, the best one can do is distance-based selection, such as selecting the closest frontier [10], and hope that these decisions lead to valuable regions.

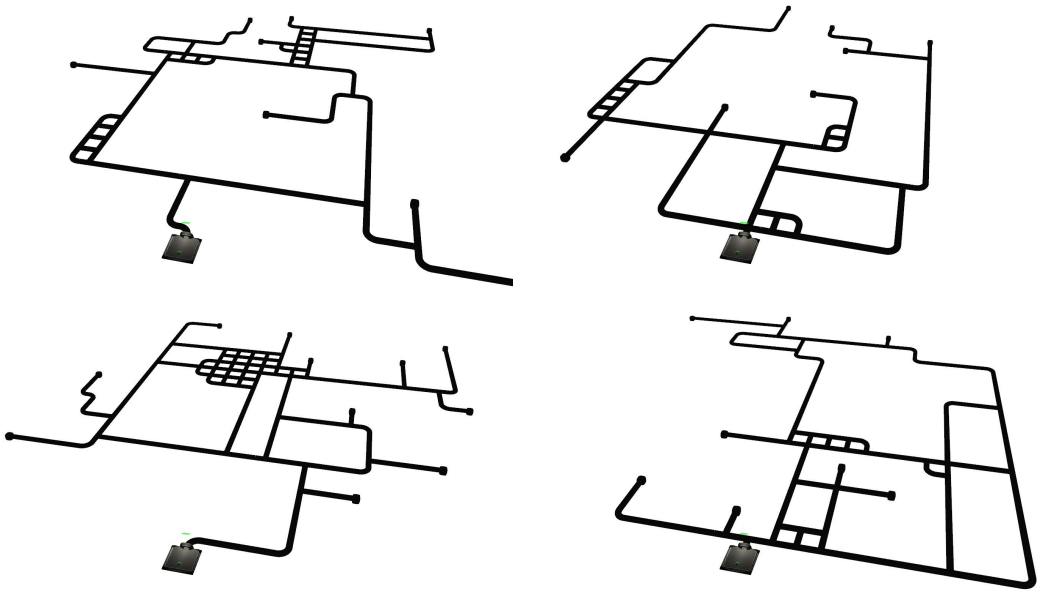

Figure 1.1: An example set of realistic subterranean tunnel networks from our procedurally generated dataset. These environments contain structure in the form of topological features (loops and connected passages). Our proposed method predicts the presence of these features in unexplored space to inform frontier selection strategies for robotic exploration.

However, even in unknown worlds, there is often knowledge available regarding the typical structure of environments similar to the current environment of interest. This is particularly evident in human-populated environments, such as buildings and cities, which exhibit common features between environments that can be exploited for improved robotic decision making [14, 15, 16, 17, 18]. Subterranean environments are often less structured, due to for example, complex interactions between ventilation requirements [19] and non-uniform distributions of mineral deposits [20]. However, we hypothesise that subtle structural cues do exist in these environments and can be exploited for improved exploration efficiency.

A key differentiating characteristic of subterranean environments is the relatively large proportion of inaccessible areas and resulting sparse topology. Prediction methods that focus on 'completing' holes in the known world [17] are thus not appropriate for tunnel networks. Instead, it is advantageous to reason over topological features, such as the connectivity and relationships between loops and dead-ends, as evident in the example environments of Fig. 1.1. Knowing or estimating these structural features is instrumental to effective decision making. To achieve this, we investigate and extend recently proposed topological loss functions designed for learning topological features of images [21, 13].

1.2 Path Planning with Sensor Measurement Uncertainty

One of the problems in real-world robot navigation is planning a collision-free and low-cost path from an initial state to a goal state. Since the path planning problem is NP-complete [22], approximate methods are used to solve the problem effectively. Sampling-based planning methods such as the Probabilistic roadmap (PRM) algorithm [23] are proven to be effective for planning in high dimensional space. PRM deals with constructing a high level graphical representation of the robot workspace, by sampling the entire space and reasoning about the local connections with the collision detection module. This high-level representation is called roadmap, which is a type of topological map of the robot free space. PRMs can learn a topologically accurate roadmap when a large number of samples are provided. Paths can be searched by planning algorithms such as Dijkstra over the

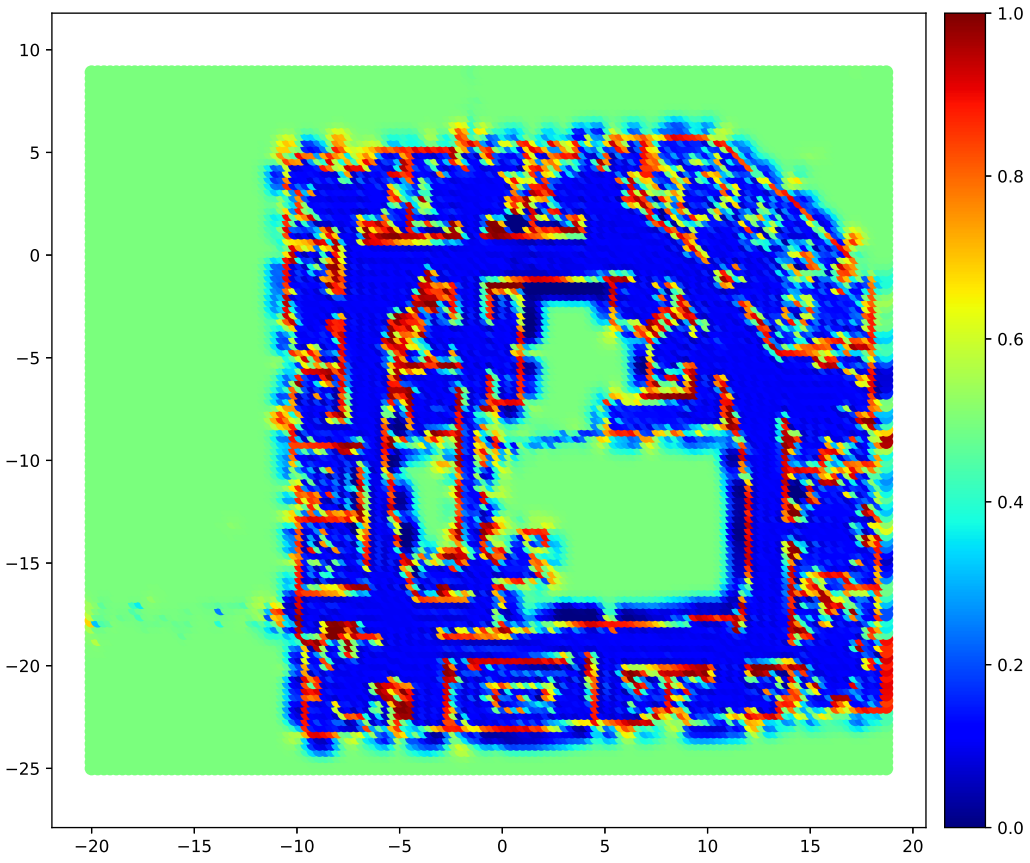

Figure 1.2: Hilbert Map: An example of a probabilistic occupancy map constructed from the Intel Research Lab pointcloud dataset [1]. The bar on right shows the probability spectrum with 0 as free space, 0.5 as unknown space, and 1 as occupied space. Our proposed method extract topological features such as connected component and loops from this map and construct a topologically accurate roadmap for efficient path planning.

developed roadmap in order to connect the start state and the goal state. However, one of the limitations of PRM is they cannot reason about collision when the environment is represented by a continuous occupancy map. Therefore, in this work, we present a method to create a roadmap that reasons for collision space using the topology of the environment represented by a continuous occupancy map.

Occupancy grid maps represent the environment in evenly spaced cells, where each cell is either occupied or free space [24]. However the real-world environment is continuous in nature, therefore in this work, we employ recently developed Hilbert Map [25] a continuous occupancy map representation of the environment. Hilbert map can be queried at any location to get its occupancy probability. Hilbert map also eliminates the independent cell assumption from the occupancy grid map making it a better representation of the environment as shown in the Fig. 1.2.

Since continuous occupancy maps do not have a clear boundary between the occupied cell and free cell, the vanilla PRM algorithm cannot construct a roadmap representing the robot's workspace. Thus in order to solve the path planning problem we need an algorithm that can develop a topologically accurate roadmap from a continuous occupancy map.

1.3 Contributions

Following the context of this thesis, we outline primary contributions of this thesis as follows:

1. A deep learning framework for online exploration of tunnel networks that

exploits subtle topological structural cues in subterranean environment for efficient exploration.

2. A procedural method for generating a database of realistic random subterranean environment graphs that contain inherent structure that can be learned.

3. Simulation results showing a robot can successfully leverage world predictions to effectively explore four real-world mines.

4. A framework to compute a topologically accurate roadmap of probabilistic occupancy map using an unsupervised topological learning technique with biased sampling in topologically important regions.

5. Simulation results showing a robot can search a collision-free and cost-effective path over the topological roadmap in low runtime.

The work from this thesis has partially appeared in our previous publication. Most of Chapter 3 is based on Saroya et al. [26].

1.4 Thesis Roadmap

The remainder of this thesis is structured as follows.

Chapter 2 surveys related work in robot exploration, robot navigation, deep learning for environment prediction, unsupervised topological learning, and persistent homology.

Chapter 3 presents a deep learning framework that exploits structural topology for online exploration of subterranean environments.

Chapter 4 presents an unsupervised learning-based method to construct a topologically accurate roadmap for robot navigation tasks.

Chapter 5 concludes this thesis and discusses important areas for future work.

Chapter 2: Related Work

In this chapter, we introduce necessary related work and their limitations for the two key problems addressed in this thesis exploration and path planning. Later we relate the literature of predictions with deep learning, unsupervised topological learning, and persistent homology to exploration and navigation tasks.

2.1 Robot Decision Making

2.1.1 Exploration with Prior Map Uncertainty

Exploration of an unknown environment is one of the most prominent robotic information gathering problems found in the literature. Robot exploration requires adaptively selecting next goals that result in the rapid discovery and mapping of an unknown environment.

Most approaches to exploration does not reason for unexplored world. The seminal work of Yamauchi [10] introduced the concept of frontiers, which describes boundaries between open and unknown space. Frontier-based exploration strategies then typically involve selecting the next frontier according to a distance-based measure, such as closest-first [10] and related variants [27]. Information-theoretic approaches additionally reason over the expected observation value at candidate frontier viewpoints by, for example, evaluating the sensor area coverage or Shannon

information gain for the map [28, 29, 30, 31]. However, these approaches typically only greedily reason over the information value of *single* observations at frontier locations; in tunnels of consistent width, all single observations may provide similar information gain. What these methods fail to capture is the information provided by *subsequent* observations as a result of moving to specific regions.

Predicting the value of future observations beyond the current frontiers requires making further assumptions about the unknown. Fortunately, in many real world environments, there is inherent structure available to be exploited for decision making. These structural features may be difficult to predict, but even weak predictions help guide exploration strategies. Promising methods have been proposed for cluttered and well-structured environments, such as matching the known world to a given database of building maps [14, 18], imitating the actions of agents in a similar known world [32], and CNN-based prediction methods for frontier selection [17]. Related algorithms that exploit predictions have also been proposed for navigating to a known [16] or unknown [33] goal, or set of goals [34]. As argued above, tunnel networks and partially-structured subterranean environments are topologically different to the environments considered in these studies, thus requiring new prediction techniques.

2.1.2 Path Planning with Sensor Measurement Uncertainty

The path planning problem relates to computing a sequence of states in a given environment to move the robot from a start to goal state while avoiding colli-

sions. The principal objective is to find a collision-free and cost-effective path that connects the initial state to the goal state.

Most methods in path planning assume a deterministic representation of the world. Sampling-based path planning methods are extensively used while planning path on deterministic maps. Probabilistic Roadmap [23], rapidly-exploring random tree [35] and visibility graph [36] are some of the variants of sampling-based planning over deterministic occupancy maps. Several approaches have been developed to increase the performance of sampling-based techniques by biasing the samples to important regions [37] such as narrow passages [38]. Placing samples adaptively according to previously found solution [39]. Recent work, Critical PRMs [40] involves deep learning-based methods to learn sampling in the critical regions like narrow passages. All the above approaches work well on deterministic occupancy maps but cannot scale on probabilistic occupancy maps which is a more realistic representation of the environment.

Another set of work deals with adjusting the cost of edges according to the occupancy probability at the edge's location. This provides a penalty to the path passing through high occupancy probability regions [41, 42]. However, this approach does not provide an optimal solution as a concession of a short path can overcome the penalty of high occupancy probability regions leading the robot to move through high occupancy probability regions.

2.2 Deep Learning for World Predictions

Convolutional neural networks (CNNs) and related variants are widely used for image processing problems such as object segmentation and classification [43]. While most of this large body of work has focused on analysing *complete* images, there has been recent work in the problem of predicting or 'completing' partially-known images, known as *image inpainting*, which is analogous to the world prediction problem we consider here.

The inpainting technique by Liu et al. [44] features a variant of the popular U-Net architecture [45] with pixel masking, which directly inspired our approach, albeit for a different problem domain.

Closely related techniques have been proposed for robotics problems. CNN-based prediction methods for frontier selection [17, 46]; unlike our work, [17] focuses on filling holes in a mostly-known world and [46] considers predictions of ocean models for marine monitoring.

2.3 Unsupervised Topological Learning

Unsupervised learning based methods such as k-mean [47] clustering are widely used to cluster the high dimensional input data to a low dimensional subspace. The k-means based methods often heavily depend on the initial starting conditions and require an initial guess of the number of vertices to be used. Also, the vertices turn to have an isotropic form i.e. circular or spherical clusters, not suitable for topological learning of complex environments.

Another branch of unsupervised learning methods addresses the problem of topological learning of the high dimensional input data. The self-organizing Kohonen maps [48] learn the topology of the data but requires an initial number of vertices and network structure. Because of the predefined number of vertices and structure of the network, some of the important topological information is lost by these methods. Other approaches competitive Hebbian learning (CHL) [49] and Neural gas (NG) [50] also require prior knowledge of roadmap size. Another developed algorithm is Growing Neural Gas (GNG) [51] which is able to continue topological learning by adding new vertices and connections as required, thus eliminating the prior decision of network size.

Recent work has employed self-organizing maps (another class of unsupervised topological learning) to the robot decision problems. However, these methods consider deterministic environments and do not rationale for uncertainties. Xia et al. [52] proposed learning sampling distribution for motion planning using a self-organizing incremental neural network. Zhong et al. used a self-organizing neural network for robot navigation to learn a route map. Best et al. [34] proposed a decentralized variant of self-organizing maps for planning a sequence of states for a team of robots deployed for an inspection task.

2.4 Persistent Homology

Persistent homology is a widely used of topological data analysis (TDA) to extract qualitative features from the data. Persistent homology has been used to analyse

data in various fields like sensor-network coverage [53], 3-D structure of DNA in nucleus [54], signals in images [55].

Persistent homology can be used with CNNs to capture topological structure of a scene. For exploration application, topological information, such as the connectivity and relationship between loops and deadends, is vital for computing accurate paths across known and unknown space. To capture topological structure into CNNs, we borrow recent ideas for learning topological features of images. More specifically, Brüel-Gabrielsson et al. [13] presented a persistent homology based regulariser for denoising, and Hu et al. [21] generalised this approach by learning topological features from ground truth segmentations of similar images, and applied it to medical image processing. In Chapter 3 we employ the loss function in [21] with a modified feature matching technique as a component of our world prediction method.

Persistent homology can be used to infer topology of the uncertain environment. This is particularly useful for the context of path planning over uncertain environments. Standard sampling based method do no explicitly consider topology. Some methods that involves persistent homology include Bhattacharya et. al. [56] to find optimal trajectories that are robust to uncertainty in the environment. Florian et. al [57] used persistent homology to determine a collection of homotopy inequivalent collision-free trajectories. The clustering of topological trajectory was presented in [58] with relative persistent homology. In this work, we use persistent homology in a different context of building a roadmap for multi-query path planning.

In Chapter 4, we propose a GNG algorithm to learn a roadmap from proba-

bilistic occupancy map. In order to speed up the topological leaning of GNG, we borrow techniques from persistent homology theory [12] to compute the regions of topological importance and bias the sampling distribution towards topologically important regions.

2.5 Summary and Limitations

In this chapter, we surveyed the literature on robot decision making in uncertain environments and the use of topology to assist decision making. We started by discussing robotic exploration with and without world prediction methods. Followed by path planning literature in deterministic and probabilistic environments.

In this thesis, we build on three methods: (1) CNNs for image inpainting, (2) unsupervised topological learning through Growing Neural Gas, and (3) topology extraction by persistent homology. By extending and adapting these methods, we propose new algorithms for addressing the challenges of robot exploration and navigation.

CNNs are great for predicting images but have not been used for predicting the unknown part of the world for robot exploration. Also, the topological features have not been used for improving exploration strategies. Chapter 3 of this thesis introduces a novel method to leverage topological prediction from a deep neural network to increase exploration efficiency.

Unsupervised topological learning has been used heavily for data processing but not to build roadmaps from probabilistic maps. In Chapter 4 we combine the

Growing Neural Gas algorithm with persistent homology to bias samples towards topologically-important regions and build a topologically accurate roadmap.

Chapter 3: Online Exploration of Tunnel Networks Leveraging CNN-based World Predictions

Robotic exploration requires adaptively selecting navigation goals that result in the rapid discovery and mapping of an unknown world. In many real-world environments, subtle structural cues can provide insight about the unexplored world, which may be exploited by a decision maker to improve the speed of exploration. In sparse subterranean tunnel networks, these cues come in the form of topological features, such as loops or dead-ends, that are often common across similar environments. In this chapter we propose a method for learning these topological features using techniques borrowed from topological image segmentation and image inpainting to learn from a database of worlds. These world predictions then inform a frontier-based exploration policy. Our simulated experiments with a set of real-world mine environments and a database of procedurally-generated artificial tunnel networks demonstrate a substantial increase in the rate of area explored compared to techniques that do not attempt to predict and exploit topological features of the unexplored world.

3.1 Problem Formulation

We address the problem of exploring an unknown environment by a mobile robot. We are particularly interested in exploring environments that feature sparsely connected corridors, such as subterranean tunnel networks. The aim is to select navigation decisions that maximise the fraction of the unknown environment that is visited by the robot within a given mission time.

The environment is represented by a navigation graph $\mathcal{G}$ with vertices $\mathcal{V}$ and edges $\mathcal{E}$. We assume the vertices $\mathcal{V}$ are formulated as a uniformly-spaced grid of locations over a rectangular area. Each vertex $v \in \mathcal{V}$ is either OPEN or CLOSED, where OPEN implies the location specified by v is open space and to be visited by the robot, and CLOSED implies a location is unreachable by the robot. Edges are connected as a 4-connected grid, with an edge $e \in \mathcal{E}$ between neighbouring vertices $v_i, v_j \in \mathcal{V}$ existing if and only if both v_i and v_j are OPEN.

At time $t = 0$, a robot starts at a predefined location specified by a vertex $s(0) = v_0 \in \mathcal{V}$. The robot follows a path along connected edges in the graph, such that at each timestep $t \in (1, 2, ..., T)$, the robot is located at a vertex $s(t) \in \mathcal{V}$ with each such vertex being OPEN. Each single edge traversal takes 1 timestep to navigate. The robot has a maximum mission time denoted T.

The graph $\mathcal{G}$ is initially unknown to the robot, but is revealed as it explores the world. At time t, the known environment is denoted $\mathcal{G}_t$, where $\mathcal{G}_t$ has the same vertices as $\mathcal{G}$, except the vertices with unknown traversability are labeled UNKNOWN. When the robot is located at a vertex v, the robot observes each of

the 4 adjacent vertices as being either OPEN or CLOSED, and this information is incorporated into the known graph $\mathcal{G}_t$.

Given these environment, motion, and observation models, the problem we aim to solve is stated as follows. At each timestep $t \in (1, 2, ..., T)$ with the robot located at $s(t)$, move the robot to an adjacent OPEN vertex $s(t+1)$, such that the number of unique vertices visited during the mission is maximised.

3.2 Online Exploration Leveraging Topological World Predictions

We propose a decision-making algorithm for online graph exploration involving frontier selection that is informed by value predictions of unexplored space. These value predictions are achieved using a novel world prediction method that uses topological features to infer the structure of unexplored space using a U-Net architecture.

The key steps of the proposed approach are enumerated as follows:

1. In an offline phase, train a neural network with a database of example world maps while using a loss function that measures the prediction error of topological features.

2. When performing online exploration, iteratively:

   (a) Update an environment map by incorporating observations from the current location.

   (b) Predict the unknown parts of the environment by inputting the cur-

rently explored environment map into the trained neural network.

(c) Estimate the exploration value of each frontier.

(d) Select the best frontier, and plan and execute a path over the known graph to the selected frontier.

These steps are described further as follows, beginning by describing the proposed world prediction model, then the frontier exploration strategy that leverages these predictions.

3.2.1 World Prediction with Topological CNNs

We propose a novel world prediction method that estimates the unknown parts of the environment graph $\mathcal{G}$. Our primary motivation for these predictions is to leverage this information in robotic exploration algorithms (discussed later in Sec. 3.3). For robotic exploration, the predictions do not necessarily need to be perfect point-to-point matches to the ground truth, but rather should provide a meaningful estimate of the connectivity of the unknown parts of the world. Thus, central to our prediction method is a CNN loss function that incorporates topological features of the world and spatial features inspired by image inpainting techniques.

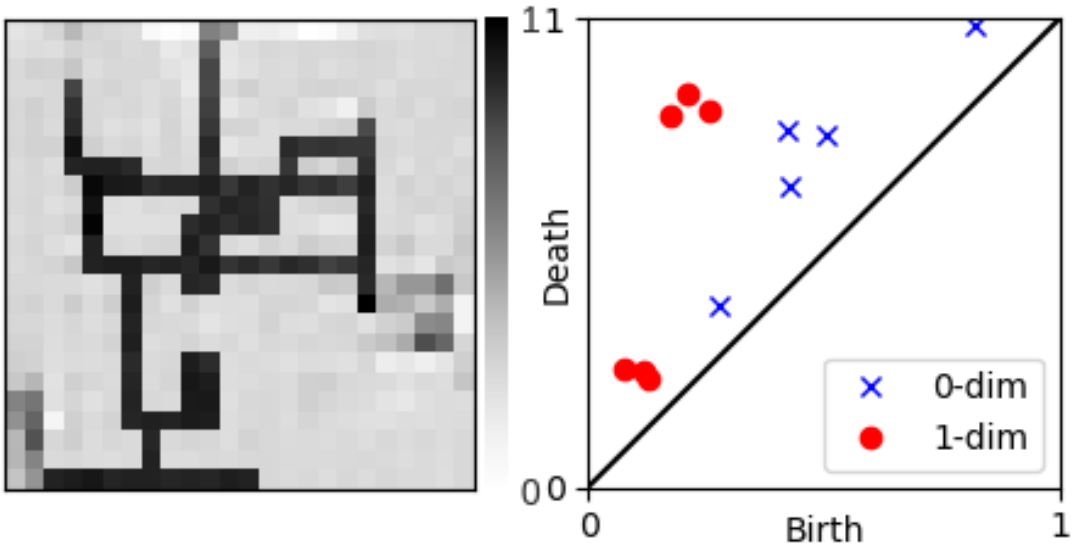

(a) Example world prediction. (b) Homology persistence diagram.

Figure 3.1: Example homology persistence diagram (see Sec. 3.2.1.1). Varying a prediction threshold between 0 and 1 (pixel intensity) results in the birth and death of 0 and 1-dimensional topological structures: connected passages and loops, respectively. Strong topological features persist over a wide range of pixel intensities, while weaker ones exist over a more limited range. The three main loops in the world (see (a)) persist across nearly all threshold levels; therefore, there exists three 1-dim features with $(birth, death)$ near $(0, 1)$, shown as three red dots in the top-left of (b).

3.2.1.1 Topological Loss

The intuition beyond the topological loss function is to match topological structures between a predicted environment graph $\mathcal{G}_{pred}$ and the ground truth environment graph $\mathcal{G}$. The 0-dimensional structures of graph $\mathcal{G}$ are defined as connected components, the 1-dimensional structures are loops, and higher dimensional structures are not applicable in our case. We compute these structures using the library in [13], which is based on persistent homology theory [12].

In image processing contexts, one approach to matching these structures between pairs of greyscale images would be to threshold the pixels at a fixed intensity and match topological structures on the resultant images. However, this approach yields loss functions that are not continuous or differentiable, making them ill-suited for learning algorithms such as neural networks. Instead, we employ the theory of persistent homology to track topological structures over a spectrum of intensity thresholds.

In particular, we generate *persistent homology diagrams* [12], which consist of the *birth* and *death* of every topological structure in a greyscale image. The *birth* of a structure is defined as the lowest threshold level of an image where this structure exists; similarly, the *death* is defined as the highest threshold where this structure exists. By definition, a structure is guaranteed to exist for all thresholds between its *birth* and *death*.

An example persistent homology diagram is presented in Fig. 3.1. Three clear 1-dim structures (loops) are present in this greyscale world prediction (Fig. 3.1a), along with several noisy structures. Each dot in Fig. 3.1b represents the *birth* and *death* of a single structure in the image. The further a structure is from the $birth = death$ line, the more persistent that structure is in the image.

Correspondences of structures in a pair of images are computed by first computing correspondence costs as the Euclidean distance between the $(birth, death)$ coordinate of each structure in one image with the $(birth, death)$ coordinates in the other image. The 0-dim structures can only correspond with other 0-dim structures, and similarly for the 1-dim structures. Correspondence is achieved using

the Hungarian assignment method [59], which connects mutually exclusive pairs of structures such that the sum of correspondence costs is minimised. Structures that are not assigned a corresponding structure are penalised with a cost equal to the distance to the $birth = death$ line, i.e. $\frac{birth-death}{2}$, which encourages such structures to be removed. We empirically observed that this approach produces more intuitive correspondences than the nearest-neighbour approach presented in [21].

Finally, given these correspondences, we define the topological loss $L_{toploss}$ as the resulting sum of correspondence costs plus the penalties for structures that do not have a correspondence.

3.2.1.2 Inpainting Loss

The topological loss function defined above encourages the existence of matching topological structures, but does not consider the size or spatial relationships between these structures in the environment. To address these issues, we also employ a set of spatial loss functions that have been shown to be useful for image inpainting [44]. We briefly describe these loss functions as follows; full details are provided in [44, Sec. 3.3].

In particular, we use location-specific losses to improve per-pixel accuracy of predictions, defined as the pixel-wise L1 loss for masked (L_{mask}) and non-masked ($L_{\neg mask}$) pixels, as used in [21, 44]. In our context, the masked and non-masked regions define explored regions and unknown regions of interest, respectively, as defined later in Sec. 3.3.1. The purpose of L_{mask} is to heavily penalise changes to

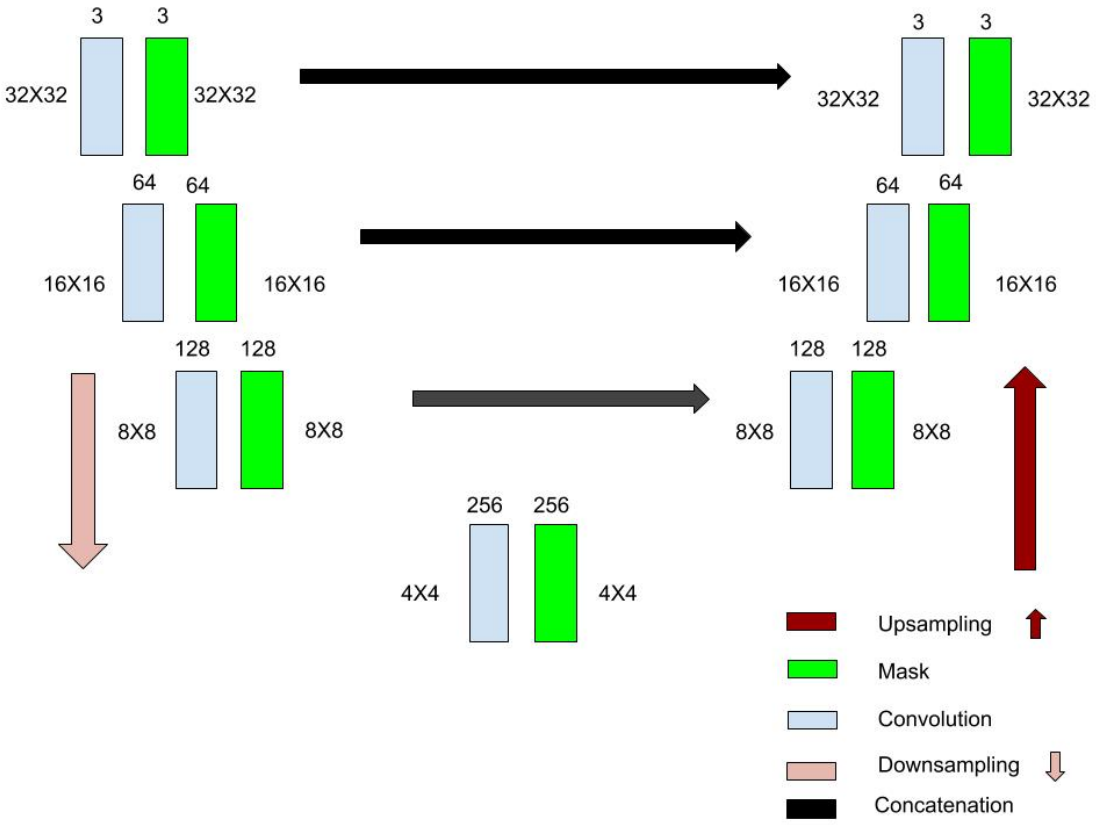

Figure 3.2: U-Net architecture for an image size of 32×32 pixels. The downsampling layers (left) aims to capture spatial context and the upsampling layers (right) accurately localise these features.

the known pixels of the image.

However, using pixel-wise loss alone is often a poor learning signal since it does not capture contextual knowledge; thus we also use style (L_{style}) and perceptual ($L_{perceptual}$) loss functions, which encourage the matching of local spatial features. L_{style} and $L_{perceptual}$ project the pair of images into the feature space of the VGG-16 network then compute distances between the images in this space. The VGG-16 network, which we use without modification, is a widely-used ImageNet-pretrained network for image classification tasks [60].

Finally, we also use the total variation (smoothing) loss L_{tv}, which encourages continuity between neighbouring pixels [44].

3.2.1.3 Combined loss function

Our loss function L_{total} combines the topological, style and perceptual losses as

$$L_{total} = 0.8L_{toploss} + 150L_{mask} + 10.5L_{\neg mask} + 0.05L_{perceptual} + 120L_{style} + 2L_{tv}. \tag{3.1}$$

This is based on the loss function proposed in [44], except with the additional $L_{toploss}$ term, and minor adjustments to the weightings based on our preliminary empirical results.

Table 3.1: Layer parameters for the U-Net architecture illustrated in Fig. 3.2, for an environment graph of size 32×32.

Type	Kernel	Stride	Input dimensions	Output dimensions
down	7	2	32×32×3	16×16×64
down	5	2	16×16×64	8×8×128
down	5	2	8×8×128	4×4×256
down	3	2	4×4×256	2×2×512
up	3	1	4×4×(256+512)	4×4×256
up	3	1	8×8×(256+128)	8×8×128
up	3	1	16×16×(128+64)	16×16×64
up	3	1	32×32×(64+3)	32×32×3

3.2.1.4 U-Net architecture

To make predictions of the world, the combined loss function defined in (3.1) is used to train a CNN from a database of example worlds. More specifically, we use a CNN with a U-Net architecture, which has been shown to outperform other CNN architectures in terms of speed and generalisability for image processing contexts [45]. The U-Net architecture consists of a contraction path for capturing spatial context followed by an expansion path for precise localisation of features.

We select the parameters of the U-Net architecture layers in a similar, but reduced, way to what has been proposed for image inpainting [44]. More specifically, we define the relevant parameters as presented in Fig. 3.2 and Table 3.1. We replicate pixel values across the 3 colour channels since our graphs correspond to greyscale images (see Sec. 3.3.1). The greyscale output predictions are then thresholded to be suitable for path planning.

3.3 Exploration Strategy

In this section, we detail our proposed exploration strategy. The key steps are to input the currently known environment graph $\mathcal{G}_t$ and an associated mask into the world prediction model proposed above; then, the predictions are used to estimate the value of each frontier, defined as a decaying sum of unexplored OPEN vertices that are reachable from each frontier. We define these concepts further as follows.

3.3.1 Explored Map and Mask Generation

To generate a prediction $\mathcal{G}_{pred}$ of the unknown world, the currently known environment graph $\mathcal{G}_t$ is converted to an image such that each vertex corresponds to a pixel in an image. All OPEN vertices are given a pixel value of 1, and CLOSED and UNKNOWN vertices are given 0. The prediction network also takes as input a mask that describes which vertices in the world are known. All OPEN and CLOSED vertices in the known graph $\mathcal{G}_t$ are included in the mask. We note that, by this definition, frontiers (OPEN vertices that have at least 1 adjacent UNKNOWN vertex) are included in the mask. Additionally, to concentrate on predicting the regions of primary interest to the current decision, we also add to the mask all UNKNOWN vertices beyond a specified number of vertices distance away from current frontiers. Due to the L_{mask} loss in Eqn. (3.1), predictions are heavily penalized if they are made in the masked region; this is to incentivize prediction in the unmasked area of interest.

3.3.2 Frontier Selection Policy and Path Planning

Exploration is performed by maintaining a set $\mathcal{U}$ of frontiers u, with $u \in \mathcal{U} \subset \mathcal{V}$, where each frontier u is an OPEN vertex that has at least 1 adjacent UNKNOWN vertex. First, the shortest path distance over the known graph $\mathcal{G}_t$ from the current position $s(t)$ to each vertex u is computed, denoted as $d(s(t) \rightarrow u)$.

Second, the information value $I(u)$ of each frontier u is estimated by reasoning over the number of OPEN vertices reachable from u that exist in the unexplored regions of the world. This value is computed as a weighted sum of predicted vertices, with the weights decreasing as a function of path distance from u. More specifically,

$$I(u) = \sum_{v \in \mathcal{V}_{pred}} \frac{1}{d(u \rightarrow v)}, \tag{3.2}$$

where $\mathcal{V}_{pred}$ is the set of predicted OPEN vertices v in the unexplored regions of the world, according to the world predictions, and $d(u \rightarrow v)$ is the shortest path distance over the union of the known graph $\mathcal{G}_t$ and predicted graph $\mathcal{G}_{pred}$. We emphasise that here we use path distance rather than Euclidean distance so that the predicted topological information is taken into account in these value estimates.

The current goal frontier u^* is selected as the frontier u that maximises the value divided by the distance, i.e.,

$$u^* = \underset{u \in \mathcal{U}}{\operatorname{argmax}} \frac{I(u)}{d(s(t) \rightarrow u)}. \tag{3.3}$$

The intuition behind this policy is to select frontiers that are both nearby the

current location and are predicted to lead towards unexplored OPEN vertices. Selecting nearby frontiers greedily minimises the amount of time spent backtracking over already-explored vertices, which is a desirable strategy when the problem is formulated as visiting as many vertices as possible in a given time T.

Finally, the action taken by the robot at time t is chosen as the next vertex along the shortest path in the known graph $\mathcal{G}_t$ from $s(t)$ to u^*. The robot cycles between making an observation, making a new prediction of the unexplored regions based on the observations, selecting a frontier by following the policy in (3.3), and moving to the next vertex. This cycle continues until the mission time T is exceeded or the entire world is explored.

3.4 Experimental Results

We present a set of simulated experimental results that demonstrate the behaviour and viability of our proposed approach. We begin by describing our method for procedurally generating a library of random worlds that aim to represent a distribution of tunnel-like environments, such as mines and natural caves. The model learned from this library is then used to make world predictions during exploration of both a database of 100 artificial tunnels and four real-world mine networks. Our experimental results demonstrate the benefits of predicting topological features for robotic exploration.

3.4.1 Real-world Tunnel Networks

We perform experiments on the maps of four real-world mines illustrated in Fig. 3.3. The first three mines were recently used in the ongoing DARPA Subterranean Challenge [6]: the Edgar Experimental Mine[1] in Idaho Springs, CO, and the NIOSH Safety Research Mine and NIOSH Experimental Mine[2] both in Bruceton, PA. The fourth environment, referred to as the Chilean Mine, is located near Santiago, Chile, and is being used for ongoing robotics research [61]. The original maps of these four mines were converted into a format suitable for our world prediction method by manually snapping the ground-truth maps to a 24×24 grid while preserving the topological features of the environments. These four environments have realistic and varied topologies that are representative of underground mines.

3.4.2 Procedural Generation of Mine-like Environments

Our proposed world prediction method learns a model from a given database of environment graphs. For these experiments, we developed a procedural method for generating a database of realistic random environment graphs that contain inherent structure that may be learned. We emphasise that our proposed method is not limited to this dataset, and if datasets of tunnel networks are available then these could be used instead. The following experiments test on both artificial and real-world tunnel networks, which demonstrate the models learned from our

[1] `https://mining.mines.edu/edgar-experimental-mine/`
[2] `https://bitbucket.org/subtchallenge/tunnel_ground_truth/src/master/`

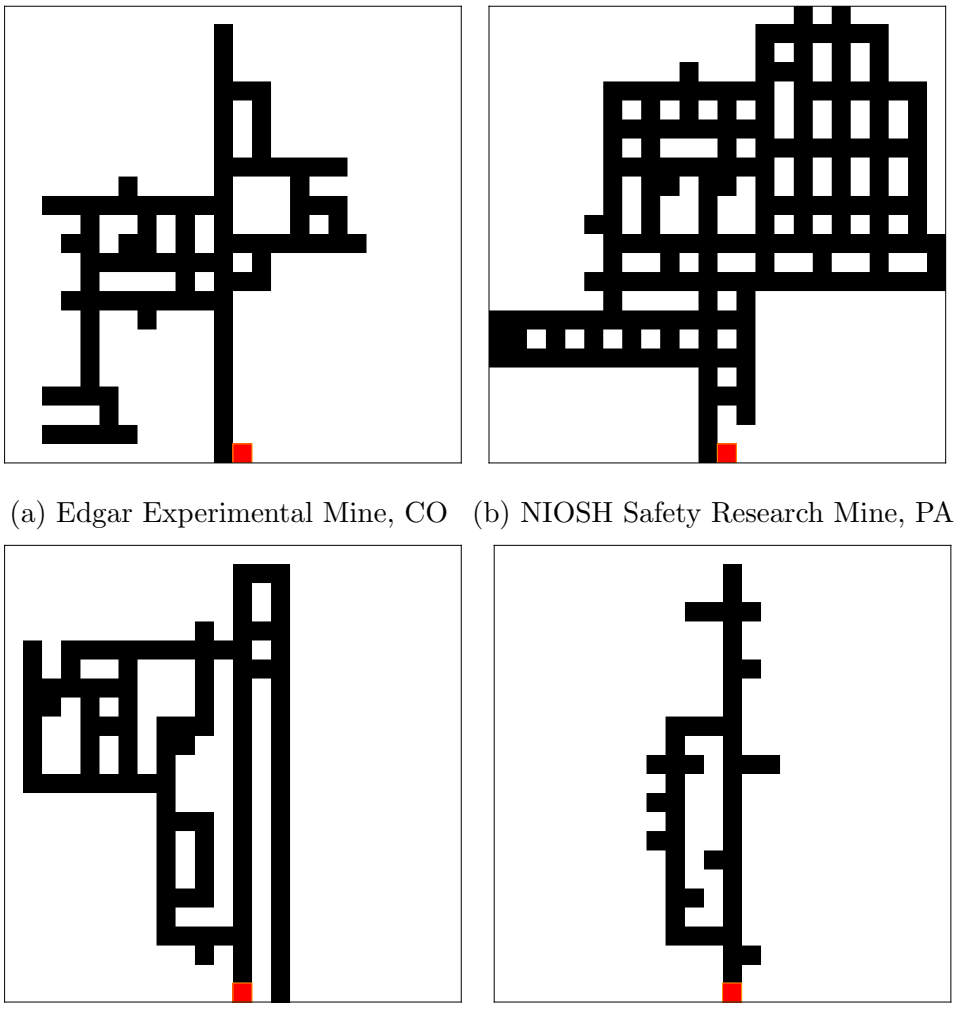

(a) Edgar Experimental Mine, CO (b) NIOSH Safety Research Mine, PA

(c) NIOSH Experimental Mine, PA (d) Chilean Mine [61]

Figure 3.3: The four real-world mine networks used in the experiments, which have been manually snapped to suitable grids while preserving the actual topological structure. The mine entrance and robot start location is shown as a red dot.

procedurally-generated dataset transfer to the real-world environments introduced above.

An example set of generated maps is illustrated earlier in Fig. 1.1. Our generation procedure begins by selecting a random set of 40 vertices as keypoints. The keypoints are drawn from a random Gaussian mixture model (GMM) with 5 mixture components at uniform-random locations and variance of 5 units. This GMM model results in non-uniform clusters of keypoints, which may represent, for example, a subterranean tunnel network for mining a non-uniform distribution of minerals. This non-uniformity results in subtle structures in the environments that may be learned by a prediction model. The start location of the robot is also marked as a keypoint. The keypoints are randomly sorted with the start location appearing first. For each keypoint, a random number (1 with probability 0.4, 2 with probability 0.6) of closest keypoints that appear earlier in the sorted list are selected. The vertices along the shortest paths between and including these pairs of vertices are labeled as OPEN. All other vertices are CLOSED. The result of this procedure is a set of world maps that feature a large variation of random topological features, as illustrated in Fig. 1.1.

3.4.3 Experimental Setup

In our experiments, we compare five strategies:

(a) our proposed method with the complete loss function in Eqn. (3.1),

(b) an oracle planner that uses the same frontier selection strategy but has access

to the ground truth map instead of predicted maps,

(c) our method except without the topological loss $L_{toploss}$, but still including the other terms of (3.1),

(d) our method except with only the pixel-wise losses L_{mask} and $L_{\neg mask}$, and

(e) a standard closest-first strategy that does not consider environment predictions.

All methods were tested on a set of 100 random environments of size 24×24. The closest 10% to 20% of each map was initially revealed to the robot in all cases so that some information was available to make predictions. Each decision took on the order of milliseconds of computation time to compute and does not require GPU. In all cases, the prediction network models were trained with 40,000 procedurally generated maps, with each method training for 400 epochs and taking approximately 12 hours on a standard desktop computer with an Nvidia GeForce GTX 1080 GPU.

3.4.4 Results

3.4.4.1 Illustrated example

In Fig. 3.4, we illustrate example exploration trajectories through the Edgar Experimental Mine. When performing exploration with our proposed method, reasonable predictions of the ground truth world are achieved at all time steps; although the

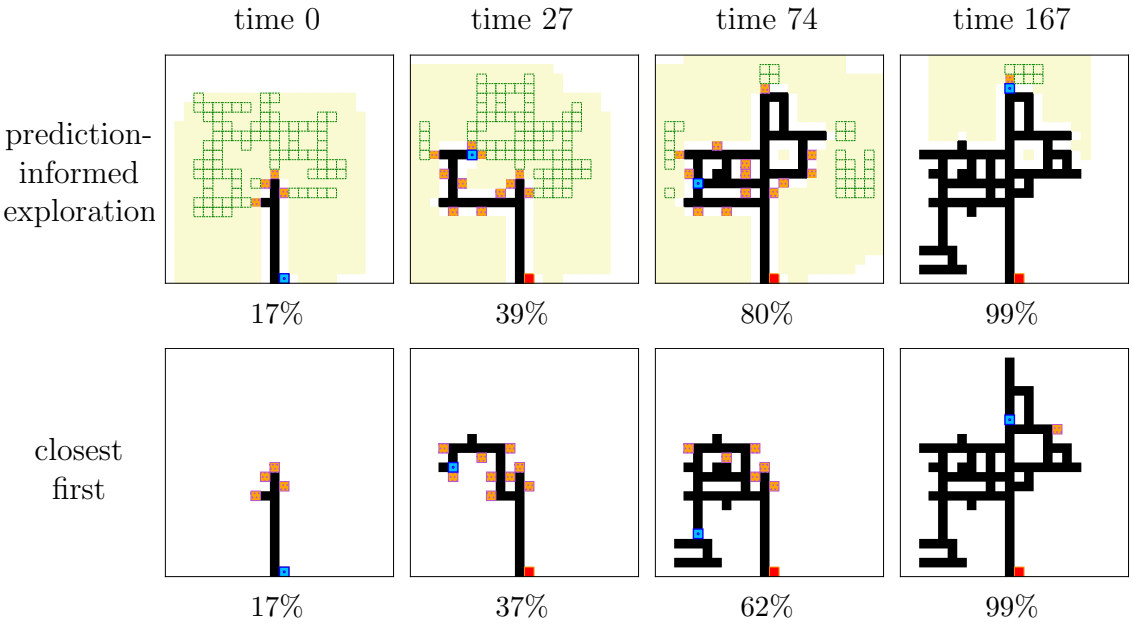

Figure 3.4: Example trajectories when for the Edgar Experimental Mine when using our proposed method and the closest first method. Robot (blue) navigates from the start (red) at bottom centre. At each timestep, the robot decides which frontier (orange) to navigate to next. For the top row, the areas predicted to be *open* are shown in green. These lie within the unmasked area of interest (shaded). The percentage of the environment explored is shown underneath.

predictions are not a perfect pixel-to-pixel match, it correctly identifies the regions of high value. This drives our prediction-informed exploration strategy towards the high-value regions in the top part of the world first. On the other hand, the closest-first strategy arbitrarily picks between frontiers adjacent to the current location; in this case the robot begins exploring towards the dead-ends in the lower left, which results in backtracking early on and therefore poorer exploration performance.

3.4.4.2 Exploration of real-world mines

We performed experiments with the four real-world mine environments illustrated earlier in Fig. 3.3. The results for each individual mine are presented in Fig. 3.5. In all four mines, our proposed approach outperformed all of the comparison methods and was competitive with the oracle solution. In particular, our prediction-informed strategy explored 60% of each mine faster than the closest-first strategy by a factor of 30% for the Edgar Mine, 11% for Safety Research Mine, 16% for Experimental Mine, and 28% for the Chilean Mine. These results show that our method results in more effective exploration since it focuses on exploring the predicted high-value regions first. Our proposed method also outperformed the comparison methods that did not use our full prediction model; these differences were less substantial, but this result still demonstrates the benefit of using our combined loss function that includes the topological features in conjunction with the spatial features. In most cases, the time to full completion is similar between all of the methods since the prediction-informed methods often skipped less valu-

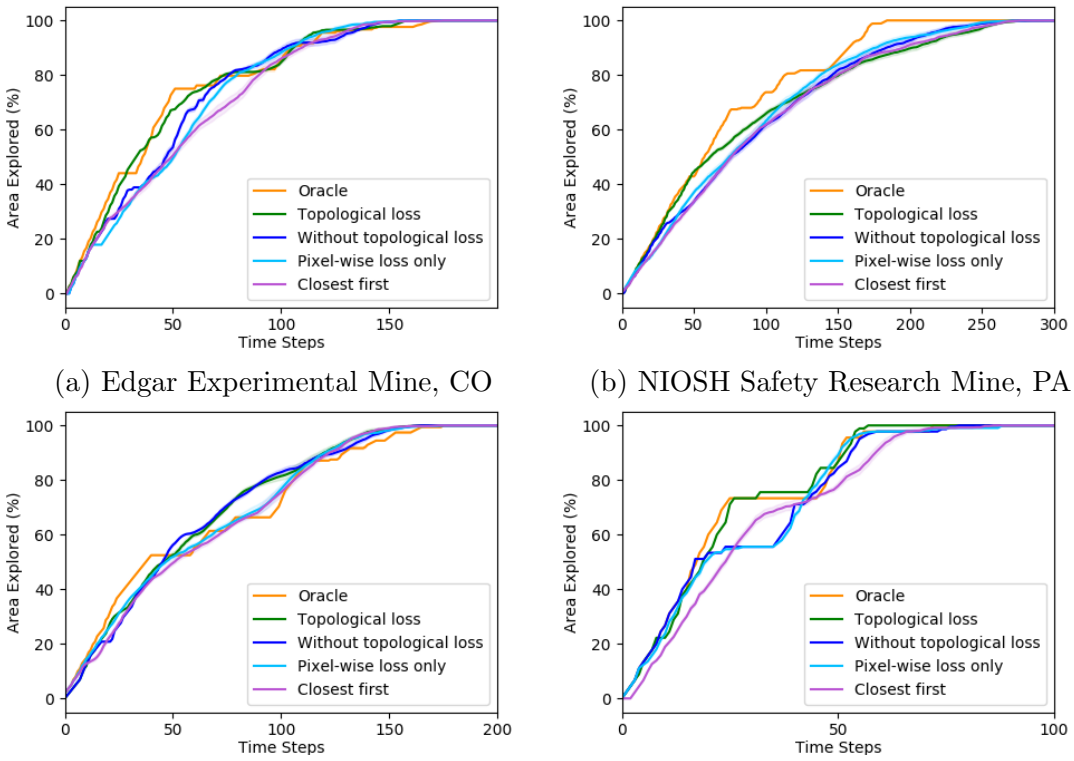

(a) Edgar Experimental Mine, CO

(b) NIOSH Safety Research Mine, PA

(c) NIOSH Experimental Mine, PA

(d) Chilean Mine [61]

Figure 3.5: Comparison between the exploration strategies on the four real-world mines depicted in Fig. 3.3.

able regions early on that need to be navigated to later if full coverage is desired. We note that the oracle strategy performed relatively poorly on the Experimental Mine since the exploration heuristic favours taking the robot towards the dead-end on the right-side of the tunnel network (see Fig. 3.3c) early in the mission due to the concentration of tunnels in the top right.

3.4.4.3 Testing on procedurally-generated database

For further validation, we also performed experiments on 100 procedurally-generated environments. An average of 172 timesteps was required to complete a full coverage, which was consistent across the algorithms. However, the value of exploiting environment predictions is more evident earlier in the mission, where the prediction-informed strategies were able to target the most promising areas of the world first. On average, our proposed method took 27% fewer timesteps to explore the first third of the environment, and 10% fewer timesteps to explore the first half, compared to the closest-first method. The *oracle*, which had access to the ground truth map, clearly outperformed all methods, which reinforces the hypothesis that having access to accurate predictions can be exploited by exploration strategies.

3.5 Summary

In this chapter, we showed topology alleviates prior map uncertainty. We presented a new exploration strategy that leverages a novel world prediction model designed for subterranean environments. We employ persistent homology to exploit subtle structural cues in tunnel networks for efficient exploration. The key findings of our results show our approach yields a substantially faster time to complete a partial coverage of the unexplored world. In the next chapter, we show persistent homology alleviates sensor measurement uncertainty and improves navigation task efficiency in uncertain environments.

Chapter 4: Roadmap Learning for Probabilistic Maps with Topologically-biased Growing Neural Gas

The main objective of navigation is to find a collision-free and cost-effective path from a start state to goal state. Sampling-based approaches are extensively used to compute a roadmap to meet this objective in deterministic environments. Sensors mounted on real robots have sensor measurement uncertainty which leads to uncertainty in mapping. In this chapter, we propose a topologically-biased Growing Neural Gas algorithm to learn a topologically accurate roadmap of the probabilistic map such as Hilbert map. Our simulation results show topologically-biased Growing Neural Gas algorithm surpasses Probabilistic Roadmaps (PRMs) in the success weighted path cost metric and provides fast online planning.

4.1 Problem Formulation

We address the problem of finding cost-effective and collision-free paths for multiple start and goal queries in a probabilistic occupancy map created by a robot with sensor measurement uncertainty. We are interested in indoor environments with topological structures such as rooms, corridors, and narrow passages. We exploit these topological structures present in the environment to alleviate sensor measurement uncertainty and obtain a topologically-accurate roadmap for multi-

query planning. Given a probabilistic map of indoor environments, we aim to develop an algorithm that can find collision-free minimum cost paths for multi-query start and goal pairs, with low run-time computational requirements and a high success rate.

The environment is represented by a 2-dimensional probabilistic occupancy map that consists of a probability of occupancy at every location. One of the examples of probabilistic occupancy map is the Hilbert map [25], built from lidar-range observation as shown in Fig. 1.2. Let ζ be a query location, then Hilbert map $\mathcal{M}$ provides occupancy probability as $p(\zeta)$.

Let $\mathcal{G}$ be a roadmap graph with vertices $\mathcal{V}$ and edges $\mathcal{E}$. Each vertex $v \in \mathcal{V}$ refers to a location vector $w_v \in R^n$ in the occupancy map. Each edge $e \in \mathcal{E}$ connects two vertices $v_i, v_j \in \mathcal{V}$ when a collision-free path exists between them. Each edge e has an associated cost c_e represented by the Euclidean distance between two vertices.

Given a start and goal location represented by $v_s, v_g \in \mathcal{V}$, a path X corresponds to a sequence of vertices $X = (v_1, v_2, v_3, ..., v_n)$ such that $\forall i, (v_i, v_{i+1}) \in \mathcal{E}$, $v_1 = v_s$ and $v_n = v_g$. The overall path cost c_X is given by the sum of all the c_e as in the path

$$c_X = \sum_{i=1}^{n-1} c_{(v_i, v_{i+1})}. \tag{4.1}$$

We address the problem of finding collision-free and cost-effective paths for multiple start ζ_s and goal ζ_s position queries using the Dijkstra algorithm. In order to solve a multi-query minimum path cost problem on a probabilistic occupancy map, we opt to learn a sparse topologically accurate roadmap of the environment.

The roadmap should meet the below three criteria:

1. Define X' as the shortest path over the roadmap from ζ_s to ζ_g. Define X^* as the shortest path over the ground-truth map from ζ_s to ζ_g. We aim to find a roadmap such that $c_{X'} \approx c_{X^*}$ for all (ζ_s, ζ_g).

2. Provide high success rate S for a set of start-goal pairs (ζ_s, ζ_g). The success rate is defined as the number of feasible solutions found for all start-goal pairs to the total number of start-goal queries.

3. Provides fast online solution; i.e., minimise the number of Dijkstra vertices expanded to find a solution.

4.2 Roadmap Learning for Probabilistic Maps with Topologically-biased Growing Neural Gas

We propose an algorithm to perform navigation tasks in an environment representation containing sensor measurement noise. As robots are extensively being used in indoor offices and homes we are interested to perform navigational tasks in office-like environments. We leverage unsupervised learning-based method Growing Neural Gas (GNG) in order to learn a topologically accurate roadmap from a probabilistic map. The indoor environment consists of important topological connections between rooms and corridors. We introduce biased sampling towards important/persistent topological structures to obtain topologically accurate roadmap via persistent homology. Collision-free and cost-effective paths are planned over

the constructed roadmap for multiple start and goal pairs.

4.2.1 Topology of Probabilistic Occupancy Map

We propose a roadmap algorithm that accurately represents the topology of the probabilistic occupancy map. To achieve such a roadmap we compute topological structures of the environment such as rooms, corridors, and narrow passages via persistent homology theory [12].

Persistent homology theory is a branch of computational topology that robustly extract topological features from noisy representations. The 0-dimensional topological structures of an environment correspond to connected components such as free space connections between rooms, the 1-dimensional topological structures correspond to free space loops around obstacles. Higher n-dimensional structures are not applicable for 2-dimensional maps, similar to our described in section 3.2.1.1.

In image processing contexts, one approach to compute these structures from a greyscale image would be to threshold the pixels at a fixed intensity and compute free space connected component between rooms and free space loops around an obstacle. However, this approach would not yield persistence/importance of topological features making them ill-suited for topological learning [62]. Instead, we employ the theory of persistent homology to track topological structures over a spectrum of intensity thresholds.

Similar to 3.2.1.1, we generate *persistent homology diagrams* [12] shown in Fig. 4.1, which consist of the *birth* and *death* of every topological structure in a

greyscale image. The *birth* of a structure is defined as the lowest threshold level of an image where this structure exists; similarly, the *death* is defined as the highest threshold where this structure exists. By definition, a structure is guaranteed to exist for all thresholds between its *birth* and *death*. Longer lifespan means the structure is more persistent and important.

The *birth* and *death* threshold value correspond to *birth* and *death* pixels (representing locations) in the probabilistic occupancy map. In Fig. 4.1a we show *death* pixels representing death locations (blue stars) corresponding to the 0-dimensional features from 4.1b. Similarly we also show *death* pixels 1-dimensional features (red stars) corresponding to the 1-dimensional features from 4.1b. We use these death locations to bias sampling distribution in Sec. 4.2.2.

To compute *persistent homology diagram* of an intensity image the computational topology methods requires the input image to be converted into a simplicial complex. The probabilistic map in Fig. 1.2 consists of intensity values corresponding to probability of occupancy at a location. We first convert probabilistic map Fig. 1.2 into a simplicial complex using Freudenthal triangulation [63]. Later we compute persistent homology diagram [12] using Guddi Library [64] with Freudenthal simplicial complex as input. Fig. 4.1b shows persistent homology diagram that consist of most persistent 0-dim and 1-dim features. For simplicity of each dimension we only consider 25 strong features that persist over a wide range of pixel intensities.

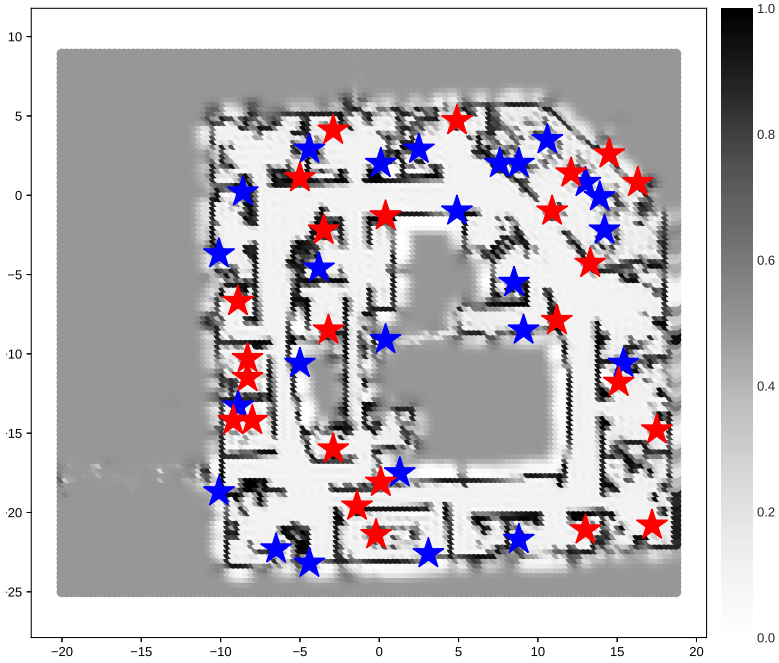

(a) Most persistent structures

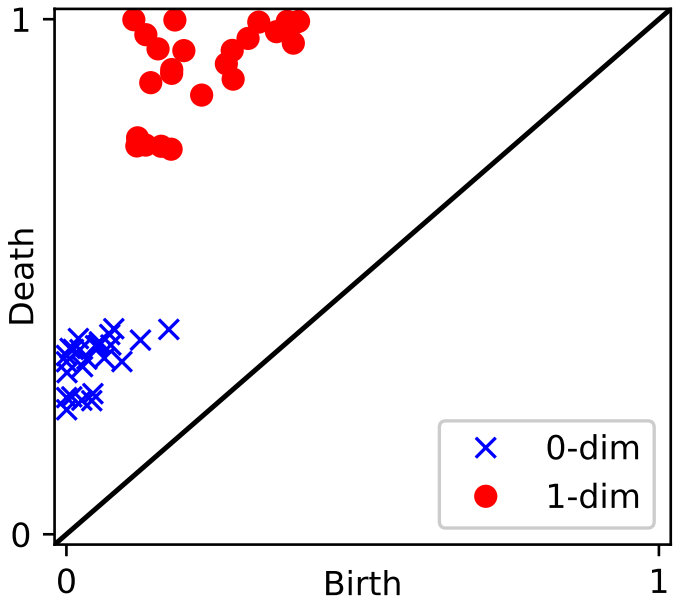

(b) Persistent homology diagram.

Figure 4.1: Persistent homology diagram (see Sec. 4.2.1). (a) shows most persistent death locations blue and red corresponding to 0-dim and 1-dim topological structures in (b) respectively. (b) consist of birth and death of topological structures computed via computational topology methods.

4.2.2 Topology-biased Sampling Distribution

In this section, we introduce two topologically biased sampling distributions that are used to learn a topologically accurate roadmap. We leverage topological structures present in the environment to develop these sampling distributions. Firstly we create map probability distribution and later develop a Gaussian mixture distribution using structure locations computed via persistent homology.

The map probability distribution $\mathcal{P}(\zeta)$ is obtained by scaling probabilistic occupancy map $\mathcal{M}$ shown in Fig. 1.2 as $e^{\gamma p(\zeta)}$ (where γ is a constant parameter) and later applying normalization.

A Gaussian mixture model(GMM) $\mathcal{Z}(\zeta)$ is created with 2 mixture components with standard deviation of 1.5 units placed at locations randomly selected from important topological regions. 25 most persistent 0-dimensional *death* pixels and 25 most persistent 1-dimensional *death* pixels from the environment are shown in the figure Fig. 4.1a. At every iteration of Growing Neural gas see Sec. 4.2.3 2 Gaussian mixture component means are randomly selected from these 50 pixel locations.

We create a topologically biased sampling distribution $\mathcal{T}(\zeta)$ by normalizing $\mathcal{P}(\zeta)\mathcal{Z}(\zeta)$. Considering map probability distribution $\mathcal{P}(\zeta)$ in topologically biased sampling distribution $\mathcal{T}(\zeta)$ eliminates/reduces the chances of sampling from high occupancy regions.

4.2.3 Topologically biased Growing Neural Gas for Roadmap Generation

We generate a roadmap through unsupervised learning based Growing Neural Gas (GNG) algorithm [51]. The growing neural gas can learn the topology of n-dimensional input signals drawn from an unknown probability distribution. To construct a topologically accurate roadmap we leverage the GNG algorithm with our topologically biased sampling distribution described in Sec. 4.2.2.

We sample 2-d points ζ_i from the map probability distribution $\mathcal{P}(\zeta)$ proportional to its free space probability to learn a roadmap representing the topology of the space. Below are the steps to generate a roadmap.

Input: Set of points ζ drawn from normalized occupancy map distribution $\mathcal{P}(\zeta)$

Output: Roadmap graph $\mathcal{G}(\mathcal{V}, \mathcal{E})$

1. Initialize $\mathcal{G}(\mathcal{V} = \{\}, \mathcal{E} = \{\})$

2. Sample (ζ_1, ζ_2) from $\mathcal{P}(\zeta)$ and initialize two vertices $\mathcal{V} = \{v_1, v_2\}$ with location vectors $w_{v_1} = \zeta_1 \in R^n$ and $w_{v_2} = \zeta_2 \in R^n$.

3. Generate input signal ζ from a combined probability distribution of $\mathcal{P}(\zeta)$ and $\mathcal{T}(\zeta)$; find winning vertices s_1, s_2 first and second nearest neighbor to ζ. Compute N_{s_1} set of adjacent vertices of s_1.

4. Create an edge $e = (s_1, s_2)$ with age zero if there is no connection between

s_1, s_2; if there is already an edge make its age zero;

$$\mathcal{E} = \mathcal{E} \cup e(s_1, s_2) \tag{4.2}$$

5. Increment the error of winner vertex s_1

$$error_{s_1} = error_{s_1} + \|w_{s_1} - \zeta\| \tag{4.3}$$

6. Move s_1 and its neighbor vertices towards ζ by a fraction of ϵ_b and ϵ_n respectively.

$$w_{s_1} = w_{s_1} + \epsilon_b(w_{s_1} - \zeta) \tag{4.4}$$

$$w_n = w_n + \epsilon_n(w_{s_1} - \zeta), \forall n \in N_{s_1} \tag{4.5}$$

7. Increment the age of all edges connected to s_1 by one; remove all the edges with age greater than a_{max} and if this leads to no emanating edge for vertices, remove those vertices.

8. Add new vertex after every λ iterations.

   (a) location the vertex q with max accumulated error and its neighbor f with largest accumulated error.

   (b) decay the error of q and f with a parameter α

   (c) insert vertex r between q and its neighbor f with largest accumulated

error; set the error of r to zero.

$$\mathcal{V} = \mathcal{V} \cup r; w_r = (w_q + w_f)/2 \tag{4.6}$$

(d) remove old edge between q and f; insert two new edge $e(q, r)$ and $e(r, f)$ with age zero.

9. Decay error of all the vertices by β

10. If stopping criteria is not satisfied go to step 3.

4.2.4 Online Planning

During online planning phase, a start and goal state pair (ζ_s, ζ_g) is provided. Each state is connected with the roadmap graph $\mathcal{G}$. Following this, Dijkstra search is performed to compute a cost-effective path from ζ_s to ζ_g.

4.3 Experiments

In this section, we evaluate our built roadmaps in an indoor office environment. We begin by describing our indoor test scenario followed by a description of comparison probabilistic roadmap planners. Our simulation results demonstrate the benefits of constructing roadmap using unsupervised topological learning approach (GNG) and biasing samples towards important topological regions to generate a topologically accurate roadmap.

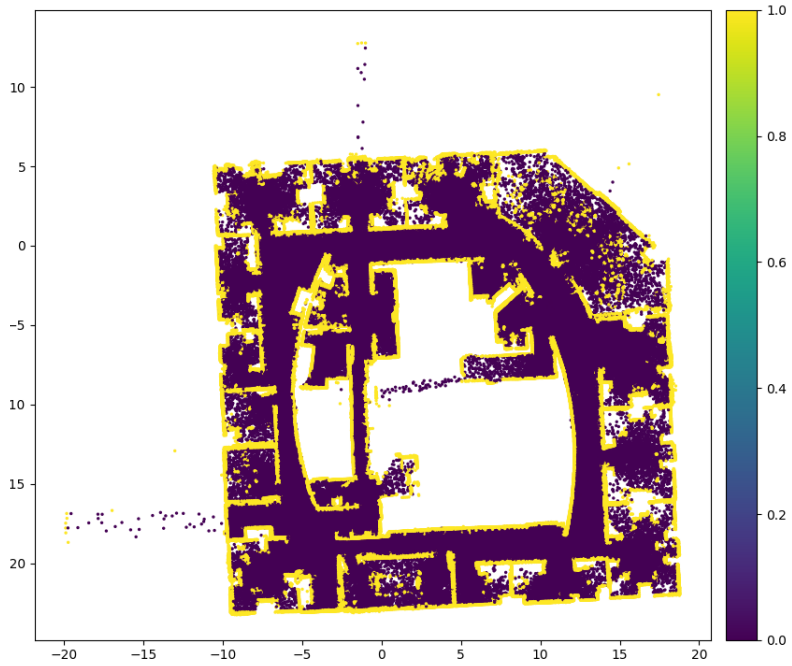

Figure 4.2: 2D pointcloud of Intel Research Lab

4.3.1 Scenario

We perform our experiments on a 2D pointcloud of Intel Research Lab obtained from Robotics Data Set Repository [1]. The dataset has manually labeled laser-range observations to void errors due to matching algorithms see Fig. 4.2. The Intel research lab consists of interesting topological structures such as rooms, corridors, and narrow passages. We exploit these topological features to construct a topologically accurate roadmap that surpasses multi-query probabilistic roadmap planners in success rate, path length, and computational runtime requirements.

We develop a Hilbert map from the laser-range observations obtained from the Intel Research Lab dataset as shown in Fig. 1.2. Hilbert map reasons for

the probability of occupancy that makes it a more realistic representation of the environment. Therefore all our experiments are based on Hilbert map.

We evaluate our approach based on success rate S, path cost c_X (or length), runtime Dijkstra vertices explored, and Success weighted Path Cost (SPC) metric. All methods are tested for 500 start-goal pairs (ζ_s, ζ_g). The Success rate is defined as the number of feasible solutions found for all start-goal pairs to the total number of start and goal pairs. Path cost is defined as the path length in meters. Runtime Dijkstra vertices explored provides a number of vertices explored while finding a path to goal on the roadmap. Success weighted Path Cost is a function of success rate and path cost [65] and is defined by

$$SPC = \frac{1}{N} \sum_i S_i \frac{c^*_{X_i}}{max(c^*_{X_i}, c_{X_i})}. \tag{4.7}$$

N denotes a total number of start-goal pairs. $c^*_{X_i}$ and c_{X_i} denotes optimal path cost and path cost obtained form the roadmap respectively.

4.3.2 Comparison Methods

In this section, we illustrate the comparison methods used in our experiments.

(a) GNG-topological method as proposed in this chapter;

(b) GNG method as proposed in this chapter except without topological biasing.

(c) Probabilistic roadmaps [37] with 7 nearest neighbor node connections in 5m radius. We show three variants PRM-4000, PRM-2000, PRM-1208 with 4000,

2000, 1208 nodes respectively. All of the PRMs are obtained from the Hilbert map with a 0.4 intensity threshold. For a fair comparison, we choose PRM-1208 as GNG uses the same number of nodes.

4.3.3 Implementation Details

We construct roadmap from continuous occupancy maps using GNG algorithm with topologically biased sampling. Some of the parameters used in construction of roadmap are as follows: map preprocessing scaling factor $\gamma = 30$, $\epsilon_b = 0.2$, $\epsilon_n = 0.005$, $a_{max} = 20$, $\lambda = 100$, $\alpha = 0.5$, $\beta = 0.995$. For computing Hilbert map from the pointcloud shown in Fig. 4.2 we use the implementation of [25]. To compute the persistent homology diagram we use Guddi Library [64]. All experiments were performed on Intel i7 8th generation core processor.

4.3.4 Results

We illustrate the roadmaps constructed with GNG-topological, GNG, PRM-4000, and PRM-1208 method. A path with the same start and goal pair was queried on all of the roadmaps. GNG-topological and PRM-4000 successfully find a path for the queried start and goal pair as shown in Fig. 4.3b and Fig. 4.4 respectively. However, the GNG method was not able to find this path due to lack of connections to the room where the goal was present see Fig 4.3a. GNG-topological was able to find the path to the start-goal query because the room consisting of the goal

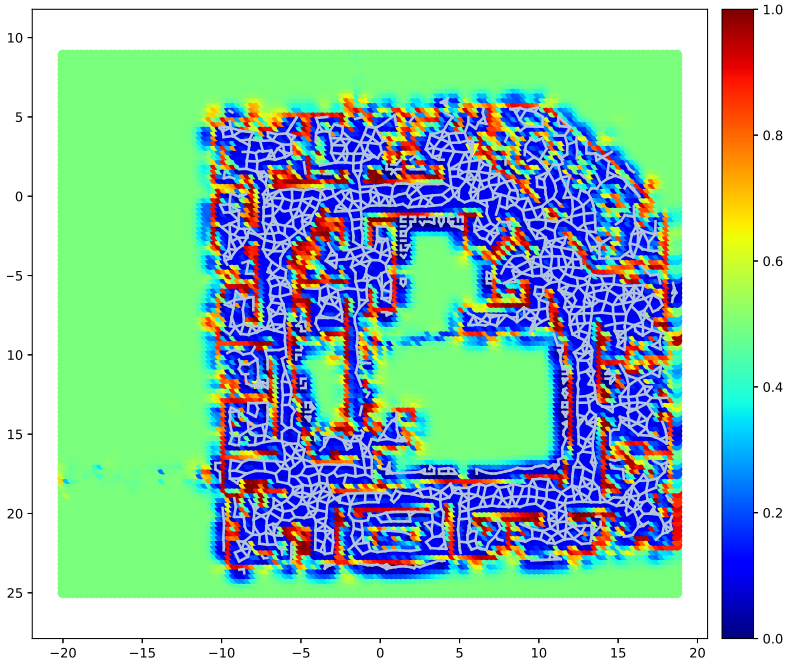

(a) Growing Neural Gas Roadmap.

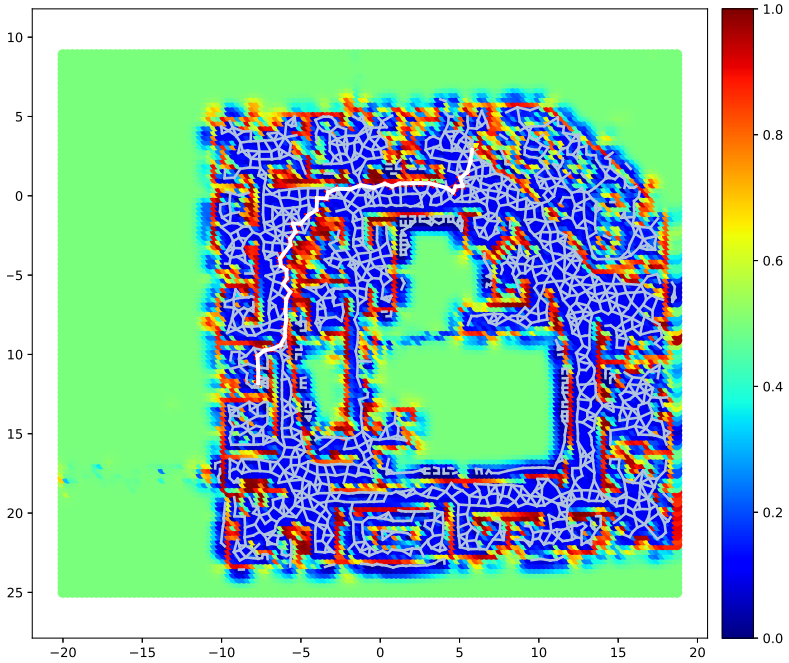

(b) Topologically biased Growing Neural Gas Roadmap

Figure 4.3: The same start and goal pair was queried on both (a) and (b). GNG-topological (b) was able to find a path while GNG (a) failed.

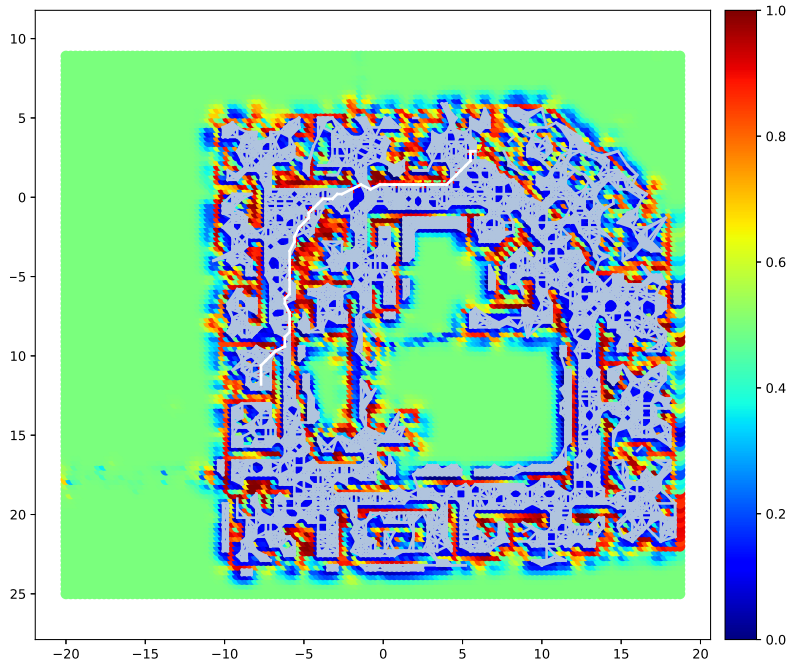

Figure 4.4: PRM-4000 build on Hilbert map. The figure shows PRM can obtain a topologically accurate roadmap when a high number of vertices are selected. However this increases the runtime search requirements.

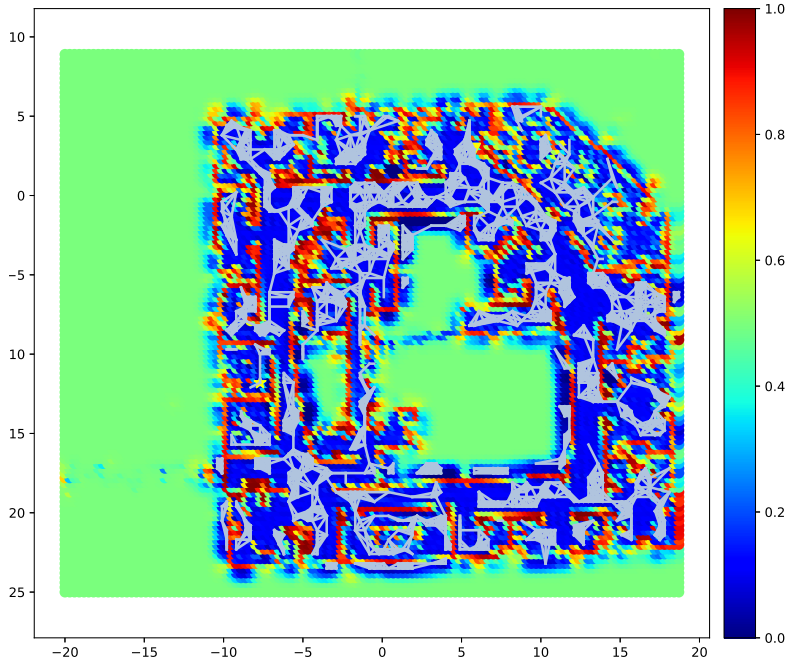

Figure 4.5: PRM-1208 build on Hilbert map. The figure shows PRM-1208 is not able to obtain a topologically accurate roadmap, resulting in a low success rate.

had two 1-dim topological features see Fig. 4.1a. These two topological features assisted in learning a topologically accurate roadmap nearby the room using our proposed topologically biased sampling. Lastly, we illustrate roadmap generated by PRM-1208 in Fig. 4.5. The graph generated by PRM-1208 is quite disconnected thus not able to find a path for the queried start and goal pair.

We evaluate our approach based on success rate, path cost, runtime Dijkstra vertices explored and SPC metric. Our results are shown in Table 4.1. All of the methods are tested on Hilbert maps. The results on 500 start and goal pair queries suggest that our GNG-topological method has the highest success rate with requiring significantly less runtime Dijkstra vertices explored. The mean and standard deviation of path cost and explored Dijkstra vertices are shown for 431 paths where GNG-topological, PRM-4000, and GNG successfully found a path altogether. The GNG method has a lower path cost and runtime Dijkstra vertices explored than GNG-topological. However, GNG-topological has a higher success weighted path cost because of the high success rate difference between GNG-topological and GNG.

4.4 Summary

In this chapter, we show topology alleviates sensor measurement uncertainty. We have proposed a new approach for navigation tasks in probabilistic environments such as Hilbert map. A navigation objective is to find a collision-free and cost-effective path for multiple start and goal pairs. We show unsupervised topology

Table 4.1: Roadmap methods evaluation on Hilbert map for 500 start-goal queries. Average path cost and Dijkstra vertices explored are shown for 431 paths. We compute Success weighted Path Cost (SPC) as described in Sec. 4.3.1 for 500 paths.

Roadmap	Success rate (%)	Path cost (m)	Dijkstra vertices explored	SPC
GNG Topological (1204 nodes)	94.8	23.34 ± 11.57	750 ± 443	0.850
PRM (4000 nodes)	93.2	21.53 ± 10.78	1601 ± 934	0.896
GNG (1208 nodes)	90.6	22.78 ± 11.28	730 ± 427	0.837
PRM (2000 nodes)	79.4	n/a	n/a	0.728
PRM (1208 nodes)	13.4	n/a	n/a	0.126

learning methods such as Growing Neural Gas can generate a roadmap on probabilistic maps to meet the navigation objective. We employ persistent homology based topological sampling bias to exploit the structural information present in indoor environments. We show that roadmap generated by GNG-topological methods yields better performance than GNG. Chapter 5 of this thesis discusses important extension areas of this work.

Chapter 5: Conclusion and Future Work

Topological learning has great potential in improving the performance of robot decision making in uncertain environments. In this thesis, we addressed two key real-world robot decision-making problems: exploration and navigation. We showed that topological reasoning can alleviate map prior uncertainty and sensor measurement uncertainty in robotic exploration and navigation problems respectively. We resolved real-world robot decision-making challenges due to environmental uncertainty by exploiting topological structures extracted via persistent homology.

5.1 Summary of Contributions

In Chapter 3, we presented a novel exploration strategy of tunnel networks that exploits topological structural cues from the environment for efficient exploration. Our simulation results shows that a robot can successfully leverage topological predictions to effectively explore four real-world mines. This work has partially appeared in our previous publication [26].

In Chapter 4, we presented a framework to compute a topologically accurate roadmap from a probabilistic occupancy map using Growing Neural Gas with biased sampling in topologically important regions. Simulation results show that GNG-topological generates topologically accurate roadmap meeting collision-free

and cost-effective navigation objectives.

5.2 Future Directions

In future work, we are interested in generalising the topological exploration approach for different problems, such as non-uniform coverage objectives and motion costs, or larger and 3D environments. Other learning methods for prediction would also be worth pursuing, such as extracting probabilistic inferences from CNNs [46] and learning directly in the space of topological maps rather than image space [15]. We would also like to explore using datasets of real tunnel networks for training our model, if made available, or consider other methods for procedural world generation that exhibit different topological characteristics. Multi-robot generalisations would also be interesting, where the proposed world prediction model could inform generalised frontier selection strategies [31, 34, 66, 67].

We would like to extend our topological roadmap approach to exploration tasks. Instead of generating a roadmap for a fixed map, it would be worth pursuing a self-organizing incremental neural network for generating roadmaps in exploration tasks [52]. It would also be interesting to see how the GNG-topological method scale in 3D environments.

In this thesis, we investigate topology for map prior uncertainty and sensor measurement uncertainty. However, robots encounter other types of uncertainty as well such as communication and actuation uncertainty. We are currently investigating ways of incorporating our proposed world prediction method into communi-

cation deployment strategies [68]. It would be interesting to integrating topological methods in motion planning problems for reducing actuation uncertainty. Further, topological methods can also be useful in reducing covariance of the position estimates in Simultaneous Localization and Mapping.

Bibliography

[1] A. Howard and N. Roy, "The robotics data set repository (radish)," 2003. [Online]. Available: http://radish.sourceforge.net/

[2] D. Tardioli, L. Riazuelo, D. Sicignano, C. Rizzo, F. Lera, J. L. Villarroel, and L. Montano, "Ground robotics in tunnels: Keys and lessons learned after 10 years of research and experiments," *Journal of Field Robotics*, vol. 36, no. 6, pp. 1074–1101, 2019.

[3] A. Morris, D. Ferguson, Z. Omohundro, D. Bradley, D. Silver, C. Baker, S. Thayer, C. Whittaker, and W. Whittaker, "Recent developments in subterranean robotics," *Journal of Field Robotics*, vol. 23, no. 1, pp. 35–57, 2006.

[4] R. Zlot and M. Bosse, "Efficient large-scale three-dimensional mobile mapping for underground mines," *Journal of Field Robotics*, vol. 31, no. 5, pp. 758–779, 2014.

[5] W. Tabib, M. Corah, N. Michael, and R. Whittaker, "Computationally efficient information-theoretic exploration of pits and caves," in *Proc. of IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, 2016, pp. 3722–3727.

[6] DARPA. (2019) Subterranean challenge. [Online]. Available: www.subtchallenge.com

[7] Z. Zhao, W. Chen, C. C. Peter, and X. Wu, "A novel navigation system for indoor cleaning robot," in *Proc. of IEEE International Conference on Robotics and Biomimetics (ROBIO)*, 2016, pp. 2159–2164.

[8] Z. Yin, D. Wang, and J. Liu, "A method of constructing robotics service platform for assisting handicapped or elderly people," *Journal of Robotics*, 2020.

[9] Y. Sun, L. Guan, Z. Chang, C. Li, and Y. Gao, "Design of a low-cost indoor navigation system for food delivery robot based on multi-sensor information fusion," *Sensors*, vol. 19, no. 22, p. 4980, 2019.

[10] B. Yamauchi, "A frontier-based approach for autonomous exploration," *Proc. of IEEE Computational Intelligence in Robotics and Automation (CIRA)*, pp. 146–151, 1997.

[11] Y. Mei, Y.-H. Lu, C. S. G. Lee, and Y. C. Hu, "Energy-efficient mobile robot exploration," in *IEEE International Conference on Robotics and Automation (ICRA)*, 2006, pp. 505–511.

[12] N. Otter, M. A. Porter, U. Tillmann, P. Grindrod, and H. A. Harrington, "A roadmap for the computation of persistent homology," *EPJ Data Science*, vol. 6, no. 1, p. 17, 2017.

[13] R. Brüel-Gabrielsson, B. J. Nelson, A. Dwaraknath, P. Skraba, L. J. Guibas, and G. Carlsson, "A topology layer for machine learning," *arXiv CoRR*, 2019.

[14] A. J. Smith and G. A. Hollinger, "Distributed inference-based multi-robot exploration," *Autonomous Robots*, vol. 42, no. 8, pp. 1651–1668, 2018.

[15] M. Luperto and F. Amigoni, "Predicting the global structure of indoor environments: A constructive machine learning approach," *Autonomous Robots*, vol. 43, no. 4, pp. 813–835, 2019.

[16] R. A. MacDonald and S. L. Smith, "Active sensing for motion planning in uncertain environments via mutual information policies," *The International Journal of Robotics Research (IJRR)*, vol. 38, no. 2-3, pp. 146–161, 2019.

[17] R. Shrestha, F. Tian, W. Feng, P. Tan, and R. Vaughan, "Learned map prediction for enhanced mobile robot exploration," in *Proc. of IEEE International Conference on Robotics and Automation (ICRA)*, 2019, pp. 1197–1204.

[18] D. P. Ström, F. Nenci, and C. Stachniss, "Predictive exploration considering previously mapped environments," in *Proc. of IEEE International Conference on Robotics and Automation (ICRA)*, 2015, pp. 2761–2766.

[19] H. L. Hartman, J. M. Mutmansky, R. V. Ramani, and Y. J. Wang, *Mine ventilation and air conditioning*. John Wiley & Sons, 2012.

[20] J. Horne, J. Ferm, F. Caruccio, and B. Baganz, "Depositional models in coal exploration and mine planning in appalachian region," *AAPG bulletin*, vol. 62, no. 12, pp. 2379–2411, 1978.

[21] X. Hu, L. Fuxin, D. Samaras, and C. Chen, "Topology-preserving deep image segmentation," *arXiv CoRR*, 2019.

[22] S. M. LaValle, *Planning algorithms*. Cambridge university press, 2006.

[23] L. E. Kavraki, P. Svestka, J.-C. Latombe, and M. H. Overmars, "Probabilistic roadmaps for path planning in high-dimensional configuration spaces," *IEEE Transactions on Robotics and Automation*, vol. 12, no. 4, pp. 566–580, 1996.

[24] A. Elfes, "Using occupancy grids for mobile robot perception and navigation," *Computer*, vol. 22, no. 6, pp. 46–57, 1989.

[25] R. Senanayake and F. Ramos, "Bayesian Hilbert maps for dynamic continuous occupancy mapping," in *Conference on Robot Learning*, 2017, pp. 458–471.

[26] M. Saroya, G. Best, and G. A. Hollinger, "Online exploration of tunnel networks leveraging topological CNN-based world predictions," in *Proc. of IEEE/RSJ Int. Conf. Intelligent Robots and Systems (IROS)*, 2020.

[27] N. Megow, K. Mehlhorn, and P. Schweitzer, "Online graph exploration: New results on old and new algorithms," *Theoretical Computer Science*, vol. 463, pp. 62–72, 2012.

[28] F. Bourgault, A. A. Makarenko, S. B. Williams, B. Grocholsky, and H. F. Durrant-Whyte, "Information based adaptive robotic exploration," in *Proc. of IEEE/RSJ Int. Conf. Intelligent Robots and Systems (IROS)*, vol. 1, 2002, pp. 540–545.

[29] W. Burgard, M. Moors, C. Stachniss, and F. Schneider, "Coordinated multi-robot exploration," *IEEE Transactions Robotics*, vol. 21, no. 3, pp. 376–386, 2005.

[30] M. Lauri and R. Ritala, "Planning for robotic exploration based on forward simulation," *Robotics and Autonomous Systems*, vol. 83, pp. 15–31, 2016.

[31] M. Corah and N. Michael, "Distributed matroid-constrained submodular maximization for multi-robot exploration: theory and practice," *Autonomous Robots*, vol. 43, no. 2, pp. 485–501, 2019.

[32] S. Choudhury, M. Bhardwaj, S. Arora, A. Kapoor, G. Ranade, S. Scherer, and D. Dey, "Data-driven planning via imitation learning," *The International Journal of Robotics Research (IJRR)*, vol. 37, no. 13-14, pp. 1632–1672, 2018.

[33] J. A. Caley, N. R. J. Lawrance, and G. A. Hollinger, "Deep learning of structured environments for robot search," *Autonomous Robots*, vol. 43, no. 7, pp. 1695–1714, 2019.

[34] G. Best and G. A. Hollinger, "Decentralised self-organising maps for multi-robot information gathering," in *Proc. of IEEE/RSJ Int. Conf. Intelligent Robots and Systems (IROS)*, 2020.

[35] S. M. LaValle and J. J. Kuffner, "Rapidly-exploring random trees: Progress and prospects," *Algorithmic and computational robotics: new directions*, no. 5, pp. 293–308, 2001.

[36] T. Siméon, J.-P. Laumond, and C. Nissoux, "Visibility-based probabilistic roadmaps for motion planning," *Advanced Robotics*, vol. 14, no. 6, pp. 477–493, 2000.

[37] D. Hsu, J.-C. Latombe, and H. Kurniawati, "On the probabilistic foundations of probabilistic roadmap planning," *The International Journal of Robotics Research*, vol. 25, no. 7, pp. 627–643, 2006.

[38] D. Hsu, L. E. Kavraki, J.-C. Latombe, R. Motwani, S. Sorkin *et al.*, "On finding narrow passages with probabilistic roadmap planners," in *Robotics: the algorithmic perspective: 1998 workshop on the algorithmic foundations of robotics*, 1998, pp. 141–154.

[39] J. D. Gammell, S. S. Srinivasa, and T. D. Barfoot, "Batch informed trees (bit*): Sampling-based optimal planning via the heuristically guided search of implicit random geometric graphs," in *Proc. of IEEE international conference on robotics and automation (ICRA)*, 2015, pp. 3067–3074.

[40] B. Ichter, E. Schmerling, T.-W. E. Lee, and A. Faust, "Learned critical probabilistic roadmaps for robotic motion planning," *arXiv preprint arXiv:1910.03701*, 2019.

[41] B. Cohen, S. Chitta, and M. Likhachev, "Search-based planning for dual-arm manipulation with upright orientation constraints," in *Proc. of IEEE International Conference on Robotics and Automation (ICRA)*, 2012, pp. 3784–3790.

[42] P. Vernaza, M. Likhachev, S. Bhattacharya, S. Chitta, A. Kushleyev, and D. D. Lee, "Search-based planning for a legged robot over rough terrain," in

Proc. of IEEE International Conference on Robotics and Automation (ICRA), 2009, pp. 2380–2387.

[43] A. Krizhevsky, I. Sutskever, and G. E. Hinton, "ImageNet classification with deep convolutional neural networks," in *Advances in Neural Information Processing Systems 25*, F. Pereira, C. J. C. Burges, L. Bottou, and K. Q. Weinberger, Eds. Curran Associates, Inc., 2012, pp. 1097–1105.

[44] G. Liu, F. A. Reda, K. J. Shih, T.-C. Wang, A. Tao, and B. Catanzaro, "Image inpainting for irregular holes using partial convolutions," in *Proc. of European Conference on Computer Vision*, 2018.

[45] O. Ronneberger, P. Fischer, and T. Brox, "U-Net: Convolutional networks for biomedical image segmentation," in *Proc. of International Conference on Medical Image Computing and Computer Assisted Intervention (MICCAI)*, 2015, pp. 234–241.

[46] J. A. Caley, "Deep learning for robotic exploration," Ph.D. dissertation, Oregon State University, 2019.

[47] S. Lloyd, "Least squares quantization in pcm," *IEEE transactions on information theory*, vol. 28, no. 2, pp. 129–137, 1982.

[48] T. Kohonen, "Self-organized formation of topologically correct feature maps," *Biological cybernetics*, vol. 43, no. 1, pp. 59–69, 1982.

[49] T. Martinetz, "Competitive hebbian learning rule forms perfectly topology preserving maps," in *International conference on artificial neural networks*. Springer, 1993, pp. 427–434.

[50] T. Martinetz and K. Schulten, "A" neural-gas" network learns topologies," in *International conference on artificial neural networks*, 1991, pp. 397–402.

[51] B. Fritzke, "A growing neural gas network learns topologies," in *Advances in neural information processing systems*, 1995, pp. 625–632.

[52] C. Xia, Y. Zhang, and I.-M. Chen, "Learning sampling distribution for motion planning with local reconstruction-based self-organizing incremental neural network," *Neural Computing and Applications*, vol. 31, no. 12, pp. 9185–9205, 2019.

[53] V. De Silva and R. Ghrist, "Coverage in sensor networks via persistent homology," *Algebraic & Geometric Topology*, vol. 7, no. 1, pp. 339–358, 2007.

[54] K. Emmett, B. Schweinhart, and R. Rabadan, "Multiscale topology of chromatin folding," *arXiv preprint arXiv:1511.01426*, 2015.

[55] M. K. Chung, P. Bubenik, and P. T. Kim, "Persistence diagrams of cortical surface data," in *International Conference on Information Processing in Medical Imaging*. Springer, 2009, pp. 386–397.

[56] S. Bhattacharya, R. Ghrist, and V. Kumar, "Persistent homology for path planning in uncertain environments," *IEEE Transactions on Robotics*, vol. 31, no. 3, pp. 578–590, 2015.

[57] F. T. Pokorny and D. Kragic, "Data-driven topological motion planning with persistent cohomology." in *Robotics: Science and Systems*, 2015.

[58] F. T. Pokorny, K. Goldberg, and D. Kragic, "Topological trajectory clustering with relative persistent homology," in *Proc. of IEEE International Conference on Robotics and Automation (ICRA)*, 2016, pp. 16–23.

[59] H. W. Kuhn, "The Hungarian method for the assignment problem," *Naval research logistics quarterly*, vol. 2, no. 1-2, pp. 83–97, 1955.

[60] K. Simonyan and A. Zisserman, "Very deep convolutional networks for large-scale image recognition," *arXiv cs.CV*, 2014.

[61] K. Leung, D. Lühr, H. Houshiar, F. Inostroza, D. Borrmann, M. Adams, A. Nüchter, and J. R. del Solar, "Chilean underground mine dataset," *Int. J. Robotics Research*, vol. 36, no. 1, pp. 16–23, 2017.

[62] S. McCammon and G. A. Hollinger, "Topological hotspot identification for informative path planning with a marine robot," in *Proc. of IEEE International Conference on Robotics and Automation (ICRA)*, 2018, pp. 4865–4872.

[63] J. Bey, "Simplicial grid refinement: on freudenthal's algorithm and the optimal number of congruence classes," *Numerische Mathematik*, vol. 85, no. 1, pp. 1–29, 2000.

[64] The GUDHI Project, *GUDHI User and Reference Manual*, 3.3.0 ed. GUDHI Editorial Board, 2020. [Online]. Available: https://gudhi.inria.fr/doc/3.3.0/

[65] P. Anderson, A. Chang, D. S. Chaplot, A. Dosovitskiy, S. Gupta, V. Koltun, J. Kosecka, J. Malik, R. Mottaghi, M. Savva *et al.*, "On evaluation of embodied navigation agents," *arXiv preprint arXiv:1807.06757*, 2018.

[66] J. Faigl and M. Kulich, "On benchmarking of frontier-based multi-robot exploration strategies," in *Proc. of European Conference on Mobile Robots (ECMR)*, 2015.

[67] G. Best, O. Cliff, T. Patten, R. R. Mettu, and R. Fitch, "Dec-MCTS: Decentralized planning for multi-robot active perception," *International Journal of Robotics Research*, vol. 38, no. 2-3, pp. 316–337, 2019.

[68] M. Tatum, "Communications coverage in unknown underground environments," Master's thesis, Carnegie Mellon University, 2020.