

Online Exploration of Tunnel Networks Leveraging Topological CNN-based World Predictions

Manish Saroya, Graeme Best, and Geoffrey A. Hollinger

Abstract—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. 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.

I. INTRODUCTION

Robots are increasingly being tasked with exploring unknown worlds to provide rapid situational awareness to end users. A vitally important domain for robotic exploration is subterranean environments [1], [2], [3], [4], including mining tunnels, urban underground facilities, and natural caves, as exemplified by the ongoing DARPA Subterranean Challenge [5]. 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.

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




Fig. 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.

the closest frontier [6], and hope that these decisions lead to valuable regions.

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 [7], [8], [9], [10], [11]. Subterranean environments are often less structured, due to, for example, complex interactions between ventilation requirements [12] and non-uniform distributions of mineral deposits [13]. 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 [10] 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. 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 [14], [15].

In this paper, we propose a new method for robotic exploration of sparse tunnel networks that aims to learn and exploit subtle structural cues present in these types of environments. Our core contribution is a novel world predic-

*Approved for public release; distribution is unlimited. This work was in part sponsored by DARPA under agreement #HR00111820044 and in part by NSF grant IIS-1845227. Any opinions, findings and conclusions or recommendations expressed in this material are those of the authors and do not necessarily reflect those of the sponsor.

*The authors are with the Collaborative Robotics and Intelligent Systems (CoRIS) Institute, Oregon State University, Corvallis OR, USA. {saroyam,bestg,geoff.hollinger}@oregonstate.edu

tion technique that predicts the unknown environment. This technique uses a U-Net architecture [16] and a novel loss function that combines pixel-to-pixel and topological feature matching via persistent homology [14], the Hungarian assignment method [17] and image inpainting features [18]. Our primary motivation for these predictions is to leverage this information to inform robotic exploration algorithms; we incorporate this information into a frontier selection policy to rapidly explore the most promising regions of the world. We have released our code publicly on our repository¹.

We present results for simulated experiments where our network is trained with a library of procedurally-generated realistic tunnel networks that exhibit random topological structures. Using this learned model, we show a robot can successfully leverage these predictions to effectively explore four real-world mines, including three used in the DARPA Subterranean Challenge [5], as well as a large database of artificial mine networks. Our key findings indicate that our proposed method of combining topological predictions with frontier selection improves the exploration efficiency over standard methods that do not attempt to explicitly learn topological features. Additionally, our results demonstrate that models learned from our artificial dataset of tunnel networks successfully transfer to real-world environments. Overall, our proposed methodology has the potential to substantially increase the information discovery efficiency of robotic exploration systems in practice.

II. RELATED WORK

Exploration of an unknown environment is one of the most prominent robotic information gathering problems found in the literature. The seminal work of Yamauchi [6] 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 [6] and related variants [19]. 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 [20], [21], [22], [23]. 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 [7], [11], imitating the actions of agents in a similar known world [24], and CNN-based prediction methods for frontier selection [10]. Related algorithms that exploit predictions have also been proposed for navigating to a known [9] or unknown [25] goal, or set of goals [26]. 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.

Convolutional neural networks (CNNs) and related variants are widely used for image processing problems such as object segmentation and classification [27]. 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. [18] features a variant of the popular U-Net architecture [16] with pixel masking, which directly inspired our approach, albeit for a different problem domain. Closely related techniques have been proposed for robotics problems [10], [28]; unlike our work, [10] focuses on filling holes in a mostly-known world and [28] considers predictions of ocean models for marine monitoring.

Standard CNNs with location-specific loss functions, such as pixel-wise cross entropy, are appropriate for many problems; however, they generally fail to capture the topological structure of a scene. For our 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 achieve this, we borrow recent ideas for learning topological features of images. More specifically, Br  l-Gabrielsson et al. [15] presented a topology-based regulariser for denoising, and Hu et al. [14] generalised this approach by learning topological features from ground truth segmentations of similar images, and applied it to medical image processing. We employ the loss function in [14] with a modified feature matching technique as a component of our world prediction method.

III. 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

¹https://github.com/manishsaroya/map_inpainting

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 $t \in (1, 2, \dots, T)$, 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 $t \in (1, 2, \dots, T)$ 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.

IV. 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 currently 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.

A. 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. IV-B). For robotic exploration, the predictions do not




Fig. 2. Example homology persistence diagram (see Sec. IV-A.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).

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.

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 [15], which is based on persistent homology theory [29].

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* [29], 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. 2. Three clear 1-dim structures (loops) are present in this greyscale world prediction (Fig. 2a), along with several noisy structures. Each dot in Fig. 2b 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




Fig. 3. U-Net architecture for an image size of 32×32 pixels. The down-sampling layers (left) aims to capture spatial context and the upsampling layers (right) accurately localise these features.

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 [17], 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 [14].

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.

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 [18]. We briefly describe these loss functions as follows; full details are provided in [18, 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 [14], [18]. In our context, the masked and non-masked regions define explored regions and unknown regions of interest, respectively, as defined later in Sec. IV-B.1. The purpose of L_{mask} is to heavily penalise changes to the known pixels of the image.

TABLE I

Layer parameters for the U-Net architecture illustrated in Fig. 3, for an environment graph of size 32×32 .

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

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 [30].

Finally, we also use the total variation (smoothing) loss L_{tv} , which encourages continuity between neighbouring pixels [18].

3) *Combined loss function*: Our loss function L_{total} combines the topological, style and perceptual losses as

$$L_{total} = 0.8L_{toploss} + 150L_{mask} + 10.5L_{\neg mask} + 0.05L_{perceptual} + 120L_{style} + 2L_{tv}. \quad (1)$$


This is based on the loss function proposed in [18], except with the additional $L_{toploss}$ term, and minor adjustments to the weightings based on our preliminary empirical results.

4) *U-Net architecture*: To make predictions of the world, the combined loss function defined in (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 [16]. 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 [18]. More specifically, we define the relevant parameters as presented in Fig. 3 and Table I. We replicate pixel values across the 3 colour channels since our graphs correspond to greyscale images (see Sec. IV-B.1). The greyscale output predictions are then thresholded to be suitable for path planning.

B. 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



(a) Edgar Experimental Mine, CO (b) NIOSH Safety Research Mine, PA (c) NIOSH Experimental Mine, PA (d) Chilean Mine [31]

Fig. 4. 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.

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.

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. (1), predictions are heavily penalized if they are made in the masked region; this is to incentivize prediction in the unmasked area of interest.

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,

$$I(u) = \sum_{v \in \mathcal{V}_{pred}} \frac{1}{d(u \rightarrow v)}, \quad (2)$$

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.,

$$u^* = \operatorname{argmax}_{u \in \mathcal{U}} \frac{I(u)}{d(s(t) \rightarrow u)}. \quad (3)$$

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), and moving to the next vertex. This cycle continues until the mission time T is exceeded or the entire world is explored.

V. 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.

A. Real-world tunnel networks

We perform experiments on the maps of four real-world mines illustrated in Fig. 4. The first three mines were recently used in the ongoing DARPA Subterranean Challenge [5]: the Edgar Experimental Mine² in Idaho Springs, CO, and

²<https://mining.mines.edu/edgar-experimental-mine/>




Fig. 5. 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.

the NIOSH Safety Research Mine and NIOSH Experimental Mine³ 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 [31]. 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.

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

An example set of generated maps is illustrated earlier in Fig. 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.

C. Experimental Setup

In our experiments, we compare five strategies:

- our proposed method with the complete loss function in Eqn. (1),
- an oracle planner that uses the same frontier selection strategy but has access to the ground truth map instead of predicted maps,
- our method except without the topological loss $L_{toploss}$, but still including the other terms of (1),
- our method except with only the pixel-wise losses L_{mask} and $L_{\neg mask}$, and
- 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

³https://bitbucket.org/subtchallenge/tunnel_ground_truth/src/master/




Fig. 6. Comparison between the exploration strategies on the four real-world mines depicted in Fig. 4.

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.

D. Results

1) *Illustrated example*: In Fig. 5, 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 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.

2) *Exploration of real-world mines*: We performed experiments with the four real-world mine environments illustrated earlier in Fig. 4. The results for each individual mine are presented in Fig. 6. 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 valuable 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. 4c) early in the mission due to the concentration of tunnels in the top right.

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.

VI. CONCLUSIONS

We have presented a new exploration strategy that leverages a novel world prediction model designed for subterranean environments. The key findings of our results show our approach yields a substantially faster time to complete a partial coverage of the unexplored world. In future work, we are interested in generalising the 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 [28] and learning directly in the space of topological maps rather than image space [8]. 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. We are currently investigating ways of incorporating our proposed world prediction method into communication deployment strategies [32]. Multi-robot generalisations would also be interesting, where the proposed world prediction model could inform generalised frontier selection strategies [23], [26], [33], [34].

REFERENCES

- [1] 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," *J. Field Robotics*, vol. 36, no. 6, pp. 1074–1101, 2019.
- [2] A. Morris, D. Ferguson, Z. Omohundro, D. Bradley, D. Silver, C. Baker, S. Thayer, C. Whittaker, and W. Whittaker, "Recent developments in subterranean robotics," *J. Field Robotics*, vol. 23, no. 1, pp. 35–57, 2006.
- [3] R. Zlot and M. Bosse, "Efficient large-scale three-dimensional mobile mapping for underground mines," *J. Field Robotics*, vol. 31, no. 5, pp. 758–779, 2014.
- [4] W. Tabib, M. Corah, N. Michael, and R. Whittaker, "Computationally efficient information-theoretic exploration of pits and caves," in *Proc. of IEEE/RSJ IROS*, 2016, pp. 3722–3727.
- [5] DARPA. (2019) Subterranean challenge. [Online]. Available: www.subtchallenge.com
- [6] B. Yamauchi, "A frontier-based approach for autonomous exploration," *Proc. of IEEE CIRA*, pp. 146–151, 1997.
- [7] A. J. Smith and G. A. Hollinger, "Distributed inference-based multi-robot exploration," *Autonomous Robots*, vol. 42, no. 8, pp. 1651–1668, 2018.
- [8] 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.
- [9] R. A. MacDonald and S. L. Smith, "Active sensing for motion planning in uncertain environments via mutual information policies," *Int. J. Robotics Research*, vol. 38, no. 2-3, pp. 146–161, 2019.
- [10] R. Shrestha, F. Tian, W. Feng, P. Tan, and R. Vaughan, "Learned map prediction for enhanced mobile robot exploration," in *Proc. of IEEE ICRA*, 2019, pp. 1197–1204.
- [11] D. P. Ström, F. Nenci, and C. Stachniss, "Predictive exploration considering previously mapped environments," in *Proc. of IEEE ICRA*, 2015, pp. 2761–2766.
- [12] H. L. Hartman, J. M. Mutmansky, R. V. Ramani, and Y. J. Wang, *Mine ventilation and air conditioning*. John Wiley & Sons, 2012.
- [13] 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.
- [14] X. Hu, L. Fuxin, D. Samaras, and C. Chen, "Topology-preserving deep image segmentation," *arXiv CoRR*, 2019.
- [15] R. Brül-Gabrielsson, B. J. Nelson, A. Dwaraknath, P. Skraba, L. J. Guibas, and G. Carlsson, "A topology layer for machine learning," *arXiv CoRR*, 2019.
- [16] O. Ronneberger, P. Fischer, and T. Brox, "U-Net: Convolutional networks for biomedical image segmentation," in *Proc. of MICCAI*, 2015, pp. 234–241.
- [17] H. W. Kuhn, "The Hungarian method for the assignment problem," *Naval research logistics quarterly*, vol. 2, no. 1-2, pp. 83–97, 1955.
- [18] 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 ECCV*, 2018.
- [19] 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.
- [20] 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 IROS*, vol. 1, 2002, pp. 540–545.
- [21] W. Burgard, M. Moors, C. Stachniss, and F. Schneider, "Coordinated multi-robot exploration," *IEEE Trans. Robotics*, vol. 21, no. 3, pp. 376–386, 2005.
- [22] M. Lauri and R. Ritala, "Planning for robotic exploration based on forward simulation," *Robotics and Autonomous Systems*, vol. 83, pp. 15–31, 2016.
- [23] 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.
- [24] S. Choudhury, M. Bhardwaj, S. Arora, A. Kapoor, G. Ranade, S. Scherer, and D. Dey, "Data-driven planning via imitation learning," *Int. J. Robotics Research*, vol. 37, no. 13-14, pp. 1632–1672, 2018.
- [25] 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.
- [26] G. Best and G. A. Hollinger, "Decentralised self-organising maps for multi-robot information gathering," in *Proc. of IEEE/RSJ IROS*, 2020.
- [27] 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.
- [28] J. A. Caley, "Deep learning for robotic exploration," Ph.D. dissertation, Oregon State University, 2019.
- [29] 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.
- [30] K. Simonyan and A. Zisserman, "Very deep convolutional networks for large-scale image recognition," *arXiv cs.CV*, 2014.
- [31] K. Leung, D. Lüth, 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.
- [32] M. Tatum, "Communications coverage in unknown underground environments," Master's thesis, Carnegie Mellon University, 2020.
- [33] J. Faigl and M. Kulich, "On benchmarking of frontier-based multi-robot exploration strategies," in *Proc. of EECMR*, 2015.
- [34] G. Best, O. Cliff, T. Patten, R. R. Mettu, and R. Fitch, "Dec-MCTS: Decentralized planning for multi-robot active perception," *Int. J. Robotics Research*, vol. 38, no. 2-3, pp. 316–337, 2019.