

PropEM-L: Radio Propagation Environment Modeling and Learning for Communication-Aware Multi-Robot Exploration

Lillian Clark*, Jeffrey A. Edlund†, Marc Sanchez Net†, Tiago Stegun Vaquero†, and Ali-akbar Agha-mohammadi†

*Ming Hsieh Department of Electrical and Computer Engineering, University of Southern California

†NASA Jet Propulsion Laboratory, California Institute of Technology

Emails: lillianmc@usc.edu, {jeffrey.a.edlund, marc.sanchez.net, tiago.stegun.vaquero, aliakbar.aghahammadi} @jpl.nasa.gov

Abstract—Multi-robot exploration of complex, unknown environments benefits from the collaboration and cooperation offered by inter-robot communication. Accurate radio signal strength prediction enables communication-aware exploration. Models which ignore the effect of the environment on signal propagation or rely on *a priori* maps suffer in unknown, communication-restricted (e.g. subterranean) environments. In this work, we present Propagation Environment Modeling and Learning (PropEM-L), a framework which leverages real-time sensor-derived 3D geometric representations of an environment to extract information about line of sight between radios and attenuating walls/obstacles in order to accurately predict received signal strength (RSS). Our data-driven approach combines the strengths of well-known models of signal propagation phenomena (e.g. shadowing, reflection, diffraction) and machine learning, and can adapt online to new environments. We demonstrate the performance of PropEM-L on a six-robot team in a communication-restricted environment with subway-like, mine-like, and cave-like characteristics, constructed for the 2021 DARPA Subterranean Challenge. Our findings indicate that PropEM-L can improve signal strength prediction accuracy by up to 44% over a log-distance path loss model.

I. INTRODUCTION

Motivated by lunar and planetary exploration [32, 2], we consider the robotic exploration of large-scale and unknown subterranean environments. The increased coverage and redundancy offered by a team of robots can improve exploration performance, relative to a single robot. Multi-robot teams benefit from the collaboration and cooperation offered by inter-robot communication [36]. However, harsh subterranean environments typically have limited communication infrastructure, meaning robots cannot rely on wireless access points for communication. In addition, communication signals see significant degradation due to the scale of the environment, winding passages without line of sight, and obstacles.

Modeling received signal strength (RSS) and how the environment affects it is useful for communication-aware exploration, during which robots autonomously maintain [31], restore [8], and/or improve [35] connectivity. As shown in Fig. 1, we consider the setting of n robots exploring a subterranean environment who must convey their findings to a stationary base station at the entrance of the cave/tunnel, and can relay data through stationary radios which are deployed during

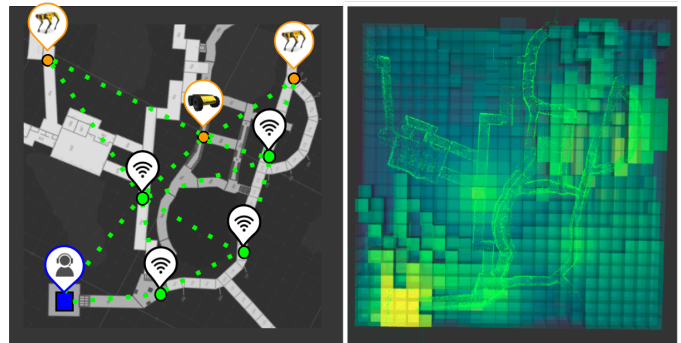


Fig. 1. **Left:** Network of exploring robots, stationary relay radios, and human supervisor base station. The network provides communication during exploration, disaster mitigation, or search and rescue and is visualized over a ground truth map of the 2021 DARPA Subterranean Challenge course. **Right:** 3D map of connectivity predicted by PropEM-L, learned from sparse point cloud data, which can be used to enable communication-aware exploration.

exploration [14]. Predicting whether robot-to-base communication is available (through one or more hops) affects decisions made by the robot, for example when and where to deploy static radios. Additionally, understanding where connectivity is available improves situational awareness for the remote human supervisor and can assist in centralized planning and task allocation [3, 21]. Because the signal strength depends on distance and the environment, accurate prediction can also aid in radio signal source-seeking [12], multi-robot localization [5], and distributed task-planning [33]. However, the mobility of the robots and lack of *a priori* maps presents new challenges for accurate RSS prediction.

A. Related Work

The communication challenges associated with multi-robot exploration have drawn increased attention [31, 26, 27, 34, 23]. These works are largely concerned with maintaining connectivity, which requires models for estimating and/or predicting connectivity. Existing work simplifies the inter-robot communication model to a deterministic communication radius [23], or predicts a probability of connectivity based solely on distance [31]. These simple models break down if, for instance, there is a metal wall within the communication radius which

blocks the signal. For this reason, Miyagusuku *et al.* seek to capture the role of the environment through a learned model of signal attenuation [18, 19]. Their data-driven approach is well-suited for repeated operation in the same environment, but for exploration of an unknown and dynamic environment an online method is advantageous.

The authors of [12] demonstrate an online method for modeling RSS during single-robot exploration. While they present a model which captures the size and location of obstacles (for example, a 1m thick wall 4m from the transmitter), they defer the discussion of estimating these parameters. Quattrini Li *et al.* propose a Gaussian Process (GP)-based method to build a communication map from measurements taken by multiple robots [24]. These works also focus on 2D operations, and give little attention to the challenge of modeling the propagation environment given partial maps in the form of 3D point clouds.

Previous work in the ecology community has considered extracting relevant information from 3D point cloud data. The authors of [6] and [16] use point clouds to determine the location and thickness of trees in a dense forest, which could theoretically feed into the model presented in [12]. Recently, the authors of [10] extend this idea to use learning-based methods on 3D point cloud data to predict the affect of tree canopies on signals propagating between static wireless communication towers. Precise digital terrain models have also been used to inform signal propagation models. [11] extends a network simulator with awareness of the topography, while [29] validates terrain-aware signal strength models with experimental data. Recently, the authors of [17] demonstrate the capability of neural networks to learn from known digital terrain models to improve cellular network design. Our work extends this method to be suitable for a dynamic network of mobile, exploring robots.

B. Contributions and Highlights

To enable communication-aware multi-robot exploration in unknown environments, we propose Propagation Environment Modeling and Learning (PropEM-L), a framework for signal strength prediction which learns the effect of the environment on attenuation. The contributions and highlights of our work are:

- 1) **Propagation Environment Modeling:** We propose PropEM, which leverages sparse 3D geometric representations of the environment (e.g. LiDAR point clouds) to extract features of the physical space which affect signal propagation, including line-of-sight visibility, shadowing due to obstacles, reflection, and diffraction.
- 2) **Learning:** We validate PropEM in conjunction with conventional data-driven approaches to RSS prediction which rely on linear regression. We then propose a neural network-based approach, PropEM-L, which significantly improves prediction accuracy relative to a log-distance path loss model and can estimate RSS within a few decibels.
- 3) **Deployment:** We evaluate the performance of PropEM-L experimentally on a dynamic network of 13

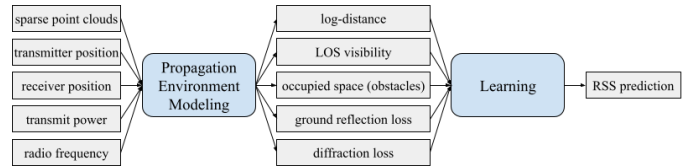


Fig. 2. PropEM-L is a framework for predicting the signal strength between two radios, given their positions, which learns from the geometry of their environment. PropEM acts as an encoding layer from raw sensor data to specific features which are relevant to propagation modeling.

autonomously-deployed static radios, one base station, and six robots exploring an underground environment. Via online learning, we demonstrate that PropEM-L can adapt to challenging new environments and construct signal strength maps.

II. PRELIMINARIES

Radio signal propagation is a multi-scale process where received signal strength is a function of distance between the transmitter and receiver, shadowing due to obstacles, and multi-path phenomena that result from reflections and refractions. Free space path loss is a first-order model which quantifies the expected attenuation in an obstacle-free environment. It is typically modeled as a logarithmic function of distance d given by

$$PL_{dB} = PL(d_0)_{dB} + \eta 10 \log_{10}(d/d_0). \quad (1)$$

$PL(d_0)_{dB}$ is the reference path loss in dB at a known distance d_0 and η is the path loss exponent which captures how quickly the signal falls off and typically takes on different values in different environments: in free space $\eta = 2$, in indoor areas with line-of-sight $1.6 \leq \eta \leq 1.8$, in urban outdoor areas $2.7 \leq \eta \leq 3.5$, and in shadowed urban outdoor areas $3 \leq \eta \leq 5$ [7].

Especially for long-ranges, the effect of shadowing is often captured in this model by the addition of a log-normal random variable with zero mean. Shadowing is complex to model as it depends on the environment itself, and second-order models (capturing path loss and shadowing) are usually determined by fitting samples taken within a specific environment.

Third-order models capture additional information about reflected and diffracted signals. The two-ray model is commonly used to capture the constructive or destructive interference caused by signals which bounce off the ground, and depends on the height of the antennas at the transmitter and receiver (h_{tx} and h_{rx}). The two-ray model [30] is given by

$$PL = \frac{G_{tx} G_{rx} \lambda^2}{(4\pi d)^2} \left(1 + \Gamma \exp\left(\frac{2i\pi(l' - l)}{\lambda}\right) \right)^2 \quad (2)$$

$$l' = \sqrt{(d^2 + (h_{tx} + h_{rx})^2)} \quad (3)$$

$$l = \sqrt{(d^2 + (h_{tx} - h_{rx})^2)} \quad (4)$$

where G_{tx}, G_{rx} are the transmitter and receiver antenna gains. λ is the wavelength of the radio signal. Γ is the ground

reflection coefficient, and we assume perfect reflection ($\Gamma = -1$) for simplicity. $\frac{G_{tx}G_{rx}\lambda^2}{(4\pi d)^2}$ also models the free space path loss, therefore the isolated *reflection loss* can be modeled by

$$RL = \left(1 - \exp\left(\frac{2i\pi(l' - l)}{\lambda}\right)\right)^2 \quad (5)$$

where negative values indicate destructive interference and positive values indicate constructive interference. Note that there exist six-ray and ten-ray models which account for additional reflections, but we limit our attention to the two-ray model to avoid the additional complexity [30].

Given knowledge of the three-dimensional space surrounding the transmitter and receiver, we can also model the Fresnel zone which is a series of ellipsoids defining the area between the transmitter and receiver where obstructions may cause interference. The radius of the first Fresnel zone at a given distance d_1 from the transmitter and d_2 to the receiver can be approximated by

$$r = \sqrt{\frac{d_1 d_2 \lambda}{d_1 + d_2}} \quad (6)$$

making the simplifying assumption that the distance between the radios is much larger than the radius [30]. Even given a clear line-of-sight path, obstacles within the first Fresnel zone can have a significant impact on signal strength [25].

For a single object, like a hill or boulder, the attenuation caused by diffraction can be estimated by treating the obstruction as a diffracting knife edge where the additional *diffraction loss* is given by

$$DL_{dB} = 20 \log_{10}(|F(v)|). \quad (7)$$

$F(v)$ is the Fresnel integral, a function of the Fresnel-Kirchoff diffraction parameter v , defined as

$$v = h \sqrt{\frac{2(d_1 + d_2)}{\lambda d_1 d_2}} \quad (8)$$

where h is the relative height of the obstruction (i.e. 0 if exactly in the line of sight). An approximate solution for Eq. 7 is given in [11]. As demonstrated by Filiposka and Trajanov, the combination of the two-ray model and the knife-edge model can predict values within a few dB of the ground truth and can identify isolated weak-reception areas well [11].

III. PROPAGATION ENVIRONMENT MODELING

In this section we introduce PropEM, as depicted in Fig. 2, which serves as an encoding layer from raw point cloud data to physically-relevant context of the propagation environment between two radios. In spacious environments with multiple robots equipped with 3D LiDAR taking frequent measurements (e.g. 10Hz), the aggregated point cloud capturing the physical space around them is constantly changing in size and can quickly exceed millions of points. Learning on the raw data would be challenging; the reduction in feature space offered by PropEM is advantageous.

Occupancy grid: PropEM creates a discretized representation of the world via OpenVDB, an open source library for

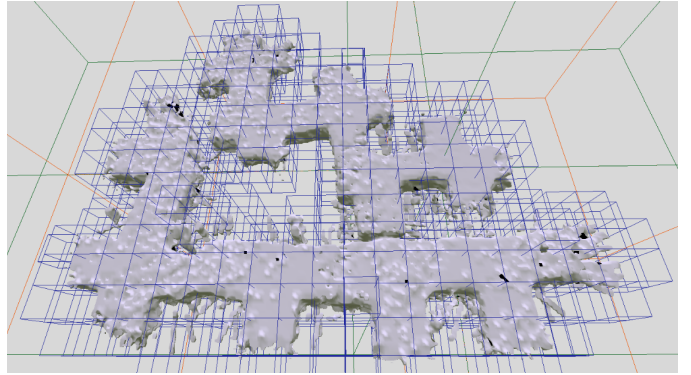


Fig. 3. 3D occupancy grid constructed for predicting line of sight.

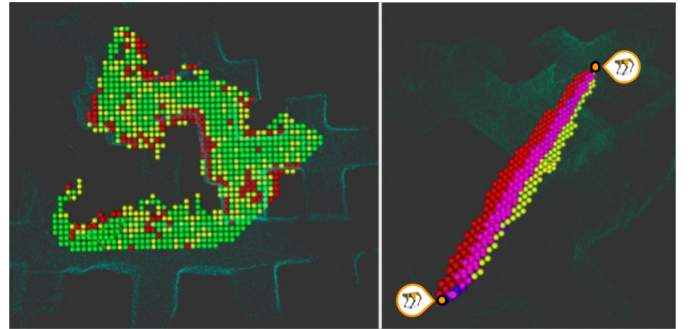


Fig. 4. **Left:** Occupied (red), maybe occupied (yellow) and free (green) voxels intersected during raytracing. **Right:** Voxels which comprise the first Fresnel zone between two radios (black), where colors help illustrate perspective.

the efficient storage and manipulation of sparse volumetric data in three-dimensional grids [20]. Each voxel in this grid is associated with the probability that the space it represents is occupied. We store probabilities as log-likelihoods so that performing Bayesian updates on a voxel amounts to addition and subtraction.

At each time t that a robot takes a LiDAR scan, we associate the robot's estimated pose, $x_r(t)$, and the scan, $s_r(t)$. Localization is performed via a pose-graph SLAM algorithm called LAMP, designed for the large-scale exploration of perceptually-degraded environments [9, 22]. PropEM performs raytracing from the robot's estimated position to each point in the point cloud representation of the scan. Voxels along these rays have an increased likelihood of being unoccupied (as the light ray did not intersect an obstacle) while voxels at the points have an increased likelihood of being occupied. Following the approach presented in [13], PropEM performs this Bayesian update for all robots and all scans in a stationary, shared frame. The result is a 3D occupancy grid as shown in Fig 3. Likelihoods continue to update as robots re-visit an area, such that we can account for dynamic obstacles.

Line of sight: Given an arbitrary pair of positions representing a transmitter and receiver, we can use this occupancy grid to determine the propagation environment which would affect the signal passing between them. To trace the LOS

ray, PropEM uses digital differential analyzers (DDAs), a technique from computer graphics used to interpolate between two points. We traverse the voxels between the transmitter and receiver, and classify them as free, occupied, maybe occupied, or unknown. This classification is done by thresholding the log-likelihoods associated with each voxel (see Fig. 4). The number of total and occupied voxels informs our calculation of free space path loss and shadowing.

Reflection and diffraction: To consider third-order effects like reflection and diffraction, we follow the approach presented in [11] to consider a model of the terrain (in our case the occupancy grid representing the partially known environment) as it causes reflections from the ground and as it intersects with the first Fresnel zone. We calculate the reflection loss RL given by Eq. 5 as a function of the transmitter and receiver locations. Then, incrementally along the line-of-sight path, we traverse voxels horizontally and vertically in the robot’s frame of reference to check the first Fresnel zone clearance (see Fig. 4). Along each direction we calculate the Fresnel-Kirchoff diffraction parameter v determined from the occupancy grid, and use the worst case v to approximate the diffraction loss DL given by Eq. 7.

By maintaining and updating the occupancy grid representation internally, PropEM can determine for an arbitrary pair of radio positions the distance, visibility, and obstacles between them as well as the ground reflection loss and diffraction loss from the cave/tunnel walls. In the next section, we discuss how these features are used in RSS prediction.

IV. LEARNING

PropEM outputs the physically-relevant features of the propagation environment, and we use these features to perform data-driven radio signal strength prediction.

A. Field Test Data

We performed extensive field testing in an underground limestone mine in Kentucky and captured about 80 minutes of autonomous exploration with 1 base station radio, 6 ground robots (3 legged Boston Dynamics Spot robots and 3 wheeled Clearpath Robotics Husky robots), and 13 communication relay radios which were autonomously deployed by the ground robots when signal-to-noise ratio (SNR) fell below a desired threshold. All radios were MIMO StreamCaster 4000-series radios from Silvus Technologies [15, 28]. Exploration spanned about 300 meters in each of the x and y directions, and we collected 1.3 million datapoints capturing transmitter position, receiver position, frequency, noise, transmitter power, RSS, SNR, loss rate, etc [1].

Although we can only collect signal strength data between pairs of radios which have connectivity, we have location estimate data for all radios. For pairs of radios which have certainly lost connectivity, meaning neither of the robots report RSS measurements at a given location for a two-minute period, we augment the dataset following the approach presented in [18], where the received signal strength is inferred to be the

TABLE I
CORRELATION BETWEEN PROPEM OUTPUT FEATURES AND RSS.

Feature	Correlation with RSS (dB)
Log distance (m)	-0.84
Line of sight	0.39
Not-free voxels (#)	-0.78
Reflection loss (dB)	0.58
Diffraction loss (dB)	0.34

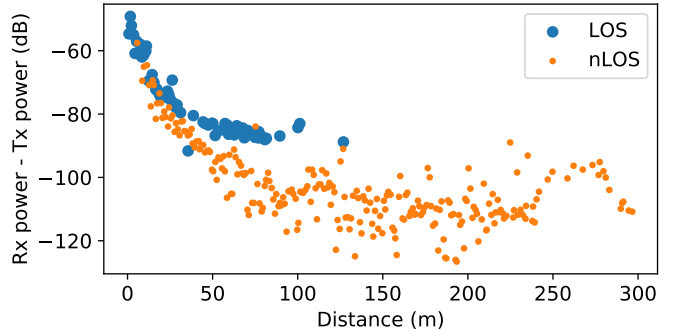


Fig. 5. Signal attenuation plotted against distance. Measurements at each distance are averaged. Line-of-sight (LOS) links see less degradation than non-line-of-sight (nLOS) links.

noise floor of the radios (here -94dB). To ensure synchronization between the location estimate and the signal strength measurement, we remove any data for which the two were observed more than one second apart. Additionally, we observe a limited number of outliers caused by localization failure which are easily removed for static radios whose location estimates should never vary more than a few meters.

As a preliminary step in validating PropEM, we note the correlation between the features PropEM outputs and the RSS measurements. Table I captures the Pearson correlation coefficient of data collected during field testing, which illustrates that signal attenuation is correlated with distance and the number of *not-free* voxels (occupied, maybe occupied, or unknown), followed by the additional loss predicted by the two-ray ground reflection model.

B. Conventional Propagation Models

To demonstrate PropEM as part of an RSS prediction framework, and as baselines for comparison, we implement the standard log-distance path loss prediction model and several extensions of this model. Table II presents a summary of the accuracy of each model, and Fig. 6 offers a snapshot of their performance.

Simple Path Loss: Using linear regression, we fit the parameters of the simple path loss model given by Eq. 1. We abbreviate $PL(d_0)_{dB}$ at $d_0 = 1\text{m}$ to PL_{d_0} in this and the following sections. This regression gives $PL_{d_0} = 14.84$ and $\eta = 4.73$.

Visibility: Motivated by the distinctly different curves observed in the data, as demonstrated in Fig. 5, we implement a model which incorporates visibility. We use linear regression to fit two distinct lines to Eq. 1 depending on whether there

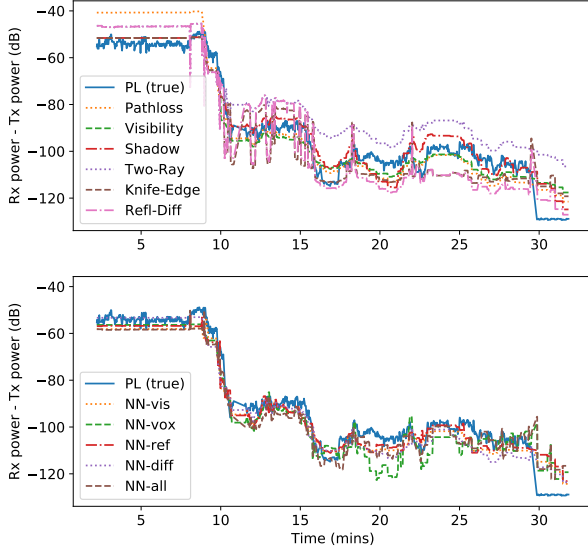


Fig. 6. Example performance of the learning models visualized against ground-truth attenuation (PL). **Top**: conventional data-driven propagation models. **Bottom**: neural-network based prediction. These plots depict the link between a single exploring robot and the base station.

is line of sight, which gives $PL_{d0}^{LOS} = 36.5$, $\eta^{LOS} = 2.75$, $PL_{d0}^{nLOS} = 13.72$ and $\eta^{nLOS} = 4.81$.

Shadowing Heuristic: Considering not just visibility but the amount of occupied space, we use linear regression to approximate the dB of attenuation per meter of not-free voxels. Similar to the approach presented in [4], this model adds an additional attenuation of 0.16 dB/m to the output of the path loss model with LOS parameters.

Two-Ray: This ground reflection model subtracts RL_{dB} , given by Eq. 5, from the simple path loss model using LOS parameters.

Knife-Edge: This diffraction model subtracts DL_{dB} , given by Eq. 7, from the simple path loss model using LOS parameters.

Reflection-Diffraction: As presented in [11], this model combines the effects of two-ray ground reflection and knife-edge diffraction. For a maximum $v \leq -0.8$, the Fresnel zone is considered clear and this model reduces to the two-ray model. Otherwise, the diffraction loss is subtracted from the two-ray model with the modification that any constructive reflective interference is ignored.

Of these models, the visibility model achieves the most accurate prediction, as shown in Fig. 6 (top). The two-ray, knife-edge, and reflection-diffraction models are very sensitive to small differences in antenna height, distance between transmitter and receiver, and precise location of neighboring walls. Because PropEM relies on sparse point clouds, uncertain location measurements, and discretized space (voxels are up to 1.5m wide), we see rapid changes in the predicted diffraction loss as new voxels are explored or the robot’s position changes even slightly. This sensitivity is likely the reason the more parsimonious models actually perform better.

TABLE II
MEAN ABSOLUTE ERROR (DB) FOR EACH MODEL (OFFLINE)

Learning Model	Field Test	Day 1	Day 2	Day 3
Simple path loss	6.73	15.87	15.37	13.0
Visibility	6.5	13.93	14.14	12.57
Shadowing Heuristic [4]	6.61	14.12	14.55	12.88
Two-Ray	12.07	17.59	17.86	16.57
Knife-edge	10.18	13.54	14.69	13.36
Reflection-Diffraction [11]	10.15	13.93	15.45	14.11
NN-vis	6.11	12.63	13.42	12.03
NN-vox	3.87	11.87	12.70	11.21
NN-ref	6.14	13.00	13.56	12.58
NN-diff	6.36	12.64	14.99	12.43
NN-all	3.73	11.42	12.95	11.70

C. Learned Propagation Models (Neural Networks)

As implemented, the conventional models make a few simplifying assumptions: the simple path loss model assumes no obstacles, the visibility and shadowing heuristic models assume all obstacles have the same material properties, the two-ray model assumes a flat reflective terrain, and the knife-edge model assumes the worst-case from obstructions. While these models have the advantage of being interpretable and grounded in the physics of how radio signals propagate, non-linear combinations of these features, learned via stochastic gradient descent, can capture subtleties in the environmental dependence which are inadequately captured in the simpler models.

We introduce five fully-connected neural networks (NN), each with the same simple structure: an input layer, a sixteen-unit hidden layer with restricted linear activation function, and an output layer which predicts attenuation PL_{dB} . Each NN learns to minimize the L2 loss between the predicted attenuation and the measured attenuation, using the Adam optimizer to train for 100 epochs or until convergence. The field test data from the limestone mine is normalized such that all features have zero-mean and unit standard deviation. 30% of the data is withheld from training, and the remainder is divided randomly into batches of size 2048.

Each NN takes a subset of the features which PropEM outputs, as described below:

- **NN-vis** uses logarithmic distance and LOS, which is broken into two boolean input features: *strictly visible* indicates all the voxels between the transmitter and receiver are free; *strictly not visible* indicates there is at least one occupied voxel.
- **NN-vox** uses logarithmic distance and the number of occupied, maybe occupied, free, and unknown voxels.
- **NN-ref** uses logarithmic distance and reflection loss.
- **NN-diff** uses logarithmic distance and diffraction loss.
- **NN-all** combines the features of NN-vox, -vis, and -diff.

Table II compares the performance of each NN model, and Fig. 7 depicts the probability density function of errors observed using each model. NN-all saw the best performance of any method with an average error of only 3.72 dB, and is less likely to produce errors above 10dB. To contextualize this performance, the standard deviation of PL_{dB} for a pair of

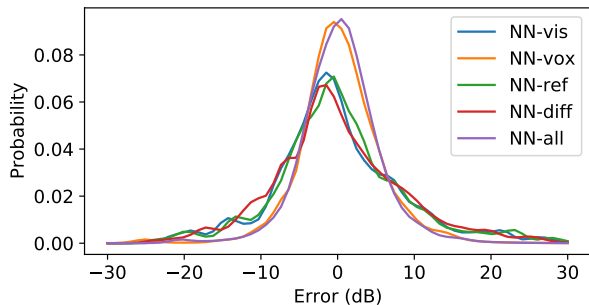


Fig. 7. Error probability density functions for each NN model.

static radios (at a constant distance) observed in the limestone mine was up to 3.43 dB, indicating that errors in this range may be attributed to inevitable channel noise.

V. ADAPTING TO NEW ENVIRONMENTS

One of the most important aspects of a prediction framework is the ability to generalize and transfer across different, complex environments. We collected data from two 30-minute (days 1 and 2) and one 60-minute (day 3) analog exploration missions in the course constructed for the DARPA Subterranean Challenge [1]. This environment, shown in Fig. 1 was characterized by narrow, winding passageways and had three distinct subsections: an urban environment similar to a subway station, a mine-like environment, and a subterranean cave-like environment. We had no *a priori* knowledge of the course or the materials from which it was constructed.

Table II summarizes the results of each model for each day of the competition. The average error for each model was more than 10dB, which could be due to a number of things. The scale of this environment was drastically different from that of the field test, with tunnels of 2m width rather than 20m width. Additionally, the environment was constructed of different materials (e.g. thin metal panel dividers in place of thick limestone columns). These materials attenuated signal more significantly than expected, as illustrated in Fig. 8. As a result, a simple path loss model was unable to accurately predict signal strength, with more than 15dB of error on average during each 30-minute experiment.

A. Online Learning

To adapt to drastically different environments, PropEM-L can perform online learning to leverage the partial 3D geometric representations available during exploration. Note that for simplicity we assume in this section RSS measurements from all radios are available at the base station immediately, ignoring any networking latencies.

Linear regression: Using the visibility model and re-fitting the parameters $\eta^{\text{LOS}}, PL_{d0}^{\text{LOS}}, \eta^{\text{nLOS}}, PL_{d0}^{\text{nLOS}}$ every k minutes, we can achieve a mean absolute error of 10.99dB (21% improvement over the offline-learned model) on day 1, 12.47dB (12% improvement) on day 2, and 11.89dB (5% improvement) on day 3 for $k = 1$. These improvements are due to the fact that the model updates the attenuation it associates with a

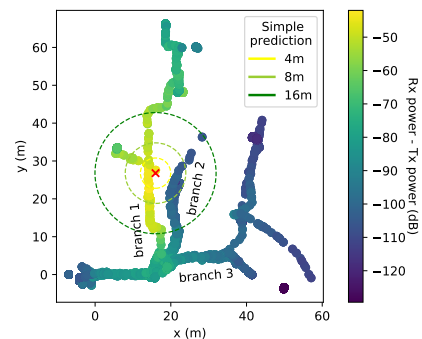


Fig. 8. Spatial plot of signal strength from a static transmitter (red) in the environment depicted in Fig. 1. A simple path loss model predicts the signal will attenuate with distance (circles around transmitter). However, branch 2 sees more degradation than predicted due to barriers between branches.

TABLE III
MEAN ABSOLUTE ERROR (dB) FOR EACH NN MODEL TRAINED ONLINE.

Model	Day 1	Day 2	Day 3
NN-vis	11.20	11.77	10.48
NN-vox	8.65	9.60	7.28
NN-ref	10.55	11.81	10.70
NN-diff	9.57	11.69	9.51
NN-all 5	8.76	9.70	7.61

given number of occupied voxels between a transmitter and receiver. This reflects a change in the model’s belief about the RF properties under identical geometries.

Perhaps more notably, the final path loss parameters learned via linear regression can shed insight on the environment itself. The learned η^{LOS} in the new environment was between 1.59-1.94, which indicates indoor areas and is significantly lower than the initial estimate of 2.75 which is typical for outdoor urban environments. Without any other information about the map, we can conclude the competition environment was more confined. Similarly, re-fitting the heuristic model increased our approximation of attenuation over not-free space from 0.16 to 0.43-0.66 dB/m. This lets us infer that the competition environment had material which blocked radio waves on average up to 4 times more than expected. This interpretability is an advantage of the conventional linear regression models over the NN models.

Neural Networks: To train the neural network component of PropEM-L online during exploration, every 60 seconds we train for 10 epochs (selected empirically) on data collected in the last minute. The number of samples used to retrain is on average 2644 (with a minimum of 498 and a maximum of 5414). Starting with the model trained offline prevents overfitting to initial measurements, while over time the new environment is learned. Fig. 9 shows how the mean absolute error decreases and converges over time. Over the entire duration of the competition run (day 3), NN-vox trained online can achieve a mean absolute error of 7.28 dB, improving performance over the simple path loss model trained offline almost two-fold.

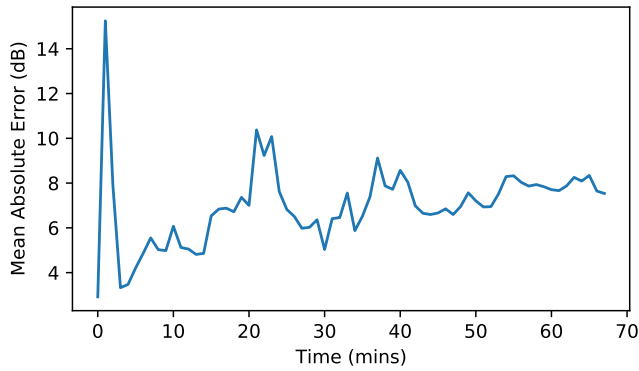


Fig. 9. Performance of NN-vox trained online in the competition environment on day 3. By the final minute of exploration, mean absolute error is 7.54 dB.

B. Signal Strength Mapping

PropEM-L can predict the strength of a signal received by a robot at any arbitrary location from one or multiple nearby radios. Given the location of static radios and their transmit powers, we can use our framework to construct a signal strength map, as shown in Fig. 1. PropEM-L can update this map as areas are explored or the network topology changes. Connectivity maps have a number of uses in communication-aware exploration: robots can autonomously identify areas to deploy additional relay radios [28], prioritize exploring new areas which are expected to offer connectivity, or navigate to areas with strong signal to prioritize data transfer and improve situational awareness for the human supervisor [8].

VI. CONCLUSION

In this work, we develop Propagation Environment Modeling and Learning (PropEM-L), a framework for signal strength prediction which learns the effect of complex, communication-restricted environments on attenuation. Our approach leverages sparse sensor-derived 3D geometric representations to model features of the propagation environment between a transmitter and receiver, including line-of-sight visibility, shadowing due to obstacles, reflection, and diffraction. We compare our neural network-based online learning with conventional approaches to RSS prediction, and demonstrate PropEM-L on a dynamic network of exploring robots and stationary radios in multiple communication-restricted, underground environments. Our findings indicate that PropEM-L can significantly improve RSS prediction and adapt to new environments, which will be an important aspect of transferring communication-aware exploration strategies from analog missions to real exploration missions.

Through online learning we can also infer certain things about the environment itself, which can give human supervisors (e.g. scientists and first responders) a better understanding of its scale and material properties. Formalizing this study is an interesting direction for future work. Additionally, future work could consider other learning methods (e.g. Gaussian

Processes), other representations (e.g. time-series RSS measurements), or other methods of encoding raw sensor data (e.g. autoencoders).

ACKNOWLEDGEMENTS

We gratefully acknowledge all members of Team CoSTAR. We also thank Belal Wang and Silvus Technologies for their hardware and technical expertise, and Kentucky Underground Storage for access to their facilities. This research was carried out at the Jet Propulsion Laboratory, California Institute of Technology, under a contract with the National Aeronautics and Space Administration (80NM0018D0004). This work was supported in part by NASA Space Technology Research Fellowship Grant No. 80NSSC19K1189.

REFERENCES

- [1] AA Agha-mhammedi, K Otsu, B Morrell, DD Fan, R Thakker, A Santamaria-Navarro, et al. Nebula: Quest for Robotic Autonomy in Challenging Environments; Team CoSTAR at the DARPA Subterranean Challenge. *arXiv preprint arXiv:2103.11470*, 2021.
- [2] AA Agha-mohammadi, KL Mitchell, and PJ Boston. Robotic Exploration of Planetary Subsurface Voids in Search for Life. In *AGU Fall Meeting Abstracts*, pages P41C–3463, 2019.
- [3] J Banfi, A Quattrini Li, N Basilico, I Rekleitis, and F Amigoni. Asynchronous Multirobot Exploration Under Recurrent Connectivity Constraints. In *2016 IEEE ICRA*, pages 5491–5498. IEEE, 2016.
- [4] B Bruggemann, A Tiderko, and M Stilkerieg. Adaptive Signal Strength Prediction Based on Radio Propagation Models for Improving Multi-Robot Navigation Strategies. In *2009 Second Int. Conference on Robot Communication and Coordination*, pages 1–6. IEEE, 2009.
- [5] Y Cao, M Li, I Švogor, S Wei, and G Beltrame. Dynamic Range-Only Localization for Multi-Robot Systems. *IEEE Access*, 6:46527–46537, 2018.
- [6] JC Carr and JB Slyder. Individual Tree Segmentation from a Leaf-Off Photogrammetric Point Cloud. *Int. Journal of Remote Sensing*, 39(15-16):5195–5210, 2018.
- [7] FJ Casadevall Palacio, R Agustí Comes, J Pérez Romero, M López Benítez, S Grimoud, B Sayrac, I Dagres, A Polydoros, J Riihijärvi, JN Nasreddine, et al. Radio Environmental Maps: Information Models and Reference Model. Document number D4. 1, 2011.
- [8] L Clark, J Galante, B Krishnamachari, and K Psounis. A Queue-Stabilizing Framework for Networked Multi-Robot Exploration. *IEEE Robotics and Automation Letters*, 6(2):2091–2098, 2021.
- [9] K Ebadi, Y Chang, M Palieri, A Stephens, A Hatteland, E Heiden, A Thakur, N Funabiki, B Morrell, S Wood, et al. LAMP: Large-Scale Autonomous Mapping and Positioning for Exploration of Perceptually-Degraded Subterranean Environments. In *2020 IEEE ICRA*, pages 80–86. IEEE, 2020.

- [10] Y Egi and CE Otero. Machine-Learning and 3D Point-Cloud Based Signal Power Path Loss Model for the Deployment of Wireless Communication Systems. *IEEE Access*, 7:42507–42517, 2019.
- [11] S Filiposka and D Trajanov. Terrain-Aware Three-Dimensional Radio-Propagation Model Extension for NS-2. *Simulation*, 87(1-2):7–23, 2011.
- [12] J Fink and V Kumar. Online Methods for Radio Signal Mapping with Mobile Robots. In *2010 IEEE ICRA*, pages 1940–1945. IEEE, 2010.
- [13] R Garg. Planning for exploration, 2020. URL <https://www.youtube.com/watch?v=vCzQDLjywwU>.
- [14] MF Ginting, T Touma, J Edlund, A Buscicchio, K Otsu, and AA Agha-mohammadi. Deployable Mesh Network for Enabling Reliable Communication from within Sub-surface Voids to the Planetary Surface. In *AGU Fall Meeting Abstracts*, pages P055–0001, 2020.
- [15] MF Ginting, K Otsu, JA Edlund, J Gao, and AA Agha-Mohammadi. CHORD: Distributed Data-sharing via Hybrid ROS 1 and 2 for Multi-robot Exploration of Large-scale Complex Environments. *IEEE Robotics and Automation Letters*, 6(3):5064–5071, 2021.
- [16] V Kubelka, P Dandurand, P Babin, P Giguère, and F Pomerleau. Radio Propagation Models for Differential GNSS based on Dense Point Clouds. *Journal of Field Robotics*, 37(8):1347–1362, 2020.
- [17] U Masood, H Farooq, and A Imran. A Machine Learning Based 3D Propagation Model for Intelligent Future Cellular Networks. In *2019 IEEE Global Communications Conference*, pages 1–6. IEEE, 2019.
- [18] R Miyagusuku, A Yamashita, and H Asama. Improving Gaussian Processes Based Mapping of Wireless Signals using Path Loss Models. In *2016 IEEE/RSJ IROS*, pages 4610–4615. IEEE, 2016.
- [19] R Miyagusuku, A Yamashita, and H Asama. Precise and Accurate Wireless Signal Strength Mappings using Gaussian Processes and Path Loss Models. *Robotics and Autonomous Systems*, 103:134–150, 2018.
- [20] K Museth. VDB: High-resolution Sparse Volumes with Dynamic Topology. *ACM Transactions on Graphics (TOG)*, 32(3):1–22, 2013.
- [21] K Otsu, S Tepsuporn, R Thakker, TS Vaquero, JA Edlund, W Walsh, G Miles, T Heywood, MT Wolf, and AA Agha-Mohammadi. Supervised Autonomy for Communication-degraded Subterranean Exploration by a Robot Team. In *2020 IEEE Aerospace Conference*, pages 1–9. IEEE, 2020.
- [22] M Palieri, B Morrell, A Thakur, K Ebadi, J Nash, A Chatterjee, C Kanellakis, L Carlone, C Guaragnella, and AA Agha-Mohammadi. Locus: A Multi-Sensor LiDAR-centric Solution for High-precision Odometry and 3D Mapping in Real-time. *IEEE Robotics and Automation Letters*, 6(2):421–428, 2020.
- [23] Y Pei, MW Mutka, and N Xi. Connectivity and Bandwidth-aware Real-time Exploration in Mobile Robot Networks. *Wireless Communications and Mobile Computing*, 13(9):847–863, 2013.
- [24] A Quattrini Li, PK Penumarthy, J Banfi, N Basilico, JM O’Kane, I Rekleitis, S Nelakuditi, and F Amigoni. Multi-robot Online Sensing Strategies for the Construction of Communication Maps. *Autonomous Robots*, 44(3):299–319, 2020.
- [25] TS Rappaport et al. *Wireless Communications: Principles and Practice*, volume 2. Prentice Hall PTR New Jersey, 1996.
- [26] P Robuffo Giordano, A Franchi, C Secchi, and HH Bühlhoff. A Passivity-based Decentralized Strategy for Generalized Connectivity Maintenance. *Int. Journal Robotics Research*, 32(3):299–323, 2013.
- [27] MN Rooker and A Birk. Multi-Robot Exploration under the Constraints of Wireless Networking. *Control Engineering Practice*, 15(4):435–445, 2007.
- [28] M Saboia, L Clark, V Thangavelu, JA Edlunch, K Otsu, GJ Correa, V Shankar Varadharajan, A Santamaria-Navarro, et al. ACHORD: Communication-Aware Multi-Robot Coordination with Intermittent Connectivity. *Under Review*, 2022.
- [29] S Santra, LB Paet, M Laine, K Yoshida, and E Staudinger. Experimental Validation of Deterministic Radio Propagation Model developed for Communication-aware Path Planning. In *2021 IEEE 17th Int. CASE*, pages 1241–1246. IEEE, 2021.
- [30] T Schwengler. Wireless & Cellular Communications, Version 3.9. *Telecommunication Systems Laboratory, Colorado, USA.*, 2, 2016.
- [31] E Stump, A Jadbabaie, and V Kumar. Connectivity Management in Mobile Robot Teams. In *2008 IEEE ICRA*, pages 1525–1530. IEEE, 2008.
- [32] TN Titus, JJ Wynne, MJ Malaska, AA Agha-Mohammadi, PB Buhler, EC Alexander, JW Ashley, A Azua-Bustos, PJ Boston, DL Buczkowski, et al. A Roadmap for Planetary Caves Science and Exploration. *Nature Astronomy*, 5(6):524–525, 2021.
- [33] J Vander Hook, TS Vaquero, F Rossi, M Troesch, MS Net, J Schoolcraft, JP de la Croix, and S Chien. Mars On-site Shared Analytics Information and Computing. In *Proceedings of ICAPS*, volume 29, pages 707–715, 2019.
- [34] TS Vaquero, M Troesch, and S Chien. An Approach for Autonomous Multi-Rover Collaboration for Mars Cave Exploration: Preliminary Results. In *Int. Symposium on Artificial Intelligence, Robotics, and Automation in Space (i-SAIRAS 2018)*. Also appears at the *ICAPS PlanRob*, 2018.
- [35] TS Vaquero, M Saboia, K Otsu, M Kaufmann, JA Edlund, and AA Agha-mohammadi. Traversability-aware Signal Coverage Planning for Communication Node Deployment in Planetary Cave Exploration. In *The Int. Symposium on Artificial Intelligence, Robotics and Automation in Space (i-SAIRAS)*, 2020.
- [36] Z Yan, N Jouandeau, and AA Cherif. A Survey and Analysis of Multi-robot Coordination. *Int. Journal Advanced Robotic Systems*, 10(12):399, 2013.