

Efficient Map Prediction via Low-Rank Matrix Completion

Zheng Chen, Shi Bai, Lantao Liu

Abstract—In many autonomous mapping tasks, the maps cannot be accurately constructed due to various reasons such as sparse, noisy, and partial sensor measurements. We propose a novel map prediction method built upon recent success of *Low-Rank Matrix Completion*. The proposed map prediction is able to achieve both map interpolation and extrapolation on raw poor-quality maps with missing or noisy observations. We validate with extensive simulated experiments that the approach can achieve real-time computation for large maps, and the performance is superior to state-of-the-art map prediction approach — Bayesian Hilbert Mapping in terms of mapping accuracy and computation time. Then we demonstrate that with the proposed real-time map prediction framework, the coverage convergence rate (per action step) for a set of representative coverage planning methods commonly used for environmental modeling and monitoring tasks can be significantly improved.

I. INTRODUCTION

Mapping is a critical functionality for an autonomous robot that needs to perform tasks in an initially unknown environment where an underlying map is constructed with onboard sensors. However, this can be challenging if the sensor measurements are inaccurate or if the environment can only be partially measured (sensed), which are commonly seen in autonomous mapping scenarios. For instance, the point clouds obtained from low-end proximity sensors (e.g., LiDAR, IR, Sonar) might be extremely sparse especially in spacious environments. The point clouds might also contain many outliers (noisy points). In addition, oftentimes a mapping robot cannot visit every spot of the environment and the traversed areas may only partially cover the space with only some small sub-regions being measured. Given the sparse, noisy, and partial observations, it is desirable that a complete map can be efficiently predicted to well match the ground-truth environment, which is the objective of this paper.

More formally, the map prediction aims to predict map structure based on already known regions and could be categorized into two classes: interpolation and extrapolation. Map interpolation means given some observations (measurement samples), we want to predict the values *in between* the sampled places; Map extrapolation means given some observation samples (or map structures), we want to predict the map values *beyond* the sampled places.

In this paper, we propose to use Low-Rank Matrix Completion (LRMC) to leverage the ill-conditioned map data to predict the whole map. LRMC exploits some special structures, such as low rank (or linear dependency)



Fig. 1: Urban and residential environments reveal strong structured patterns, such as *Left*: road (or street) network and *Right*: buildings layout. These environments containing linear dependent structures can be modeled as low-rank-matrix maps.

and incoherence hidden in underlying map to perform the prediction. Two examples that possess special patterns in urban/residential environments are shown in Fig. 1. We first illustrate how a structured environmental map satisfies important assumptions required by LRMC model. Then we qualitatively demonstrate with an urban road network map that our proposed LRMC method outperforms state-of-the-art map prediction (or regression) method – Bayesian Hilbert Mapping (BHM) [1] which is a Bayesian extension of Hilbert Mapping (HM) [2], where the map prediction is treated as a kernelized logistic regression problem. BHM eliminates crucial regularization parameter tuning in HM, and outputs continuous maps where the occupancy value of any arbitrary point could be queried.

We conduct exhaustive simulations on a set of maze-like complex maps and quantitatively analyze that our proposed LRMC method is superior to BHM on the maps with varying linear patterns and different rank values. Finally we combine the proposed real-time map prediction with representative coverage planning methods commonly used for environmental mapping and monitoring, and show that the mapping coverage convergence could be significantly improved with high mapping accuracy.

We summarize our contributions as follows:

- This is the first time to examine and report that many complex urban/residential environments possess low-rank and incoherent structure, and to apply Low-Rank Matrix Completion for map prediction based on sparse, noisy, and partially observed maps.
- Our proposed Low-Rank Matrix Completion based map prediction outperforms state-of-the-art map prediction method—Bayesian Hilbert Mapping in terms of mapping accuracy and computation time and is able to perform prediction in real-time.

- We perform extensive simulations and demonstrate the remarkable effectiveness of our proposed method which allows representative coverage planning methods to achieve faster mapping coverage convergence rates.

II. RELATED WORK

Map prediction has been studied in occupancy mapping problems. Prevalent methods include Occupancy Grids Mapping (OGM) [3], Gaussian Process Occupancy Mapping (GPOM) [4], [5], Hilbert Mapping (HM) [2] and Bayesian Hilbert Mapping (BHM) [1]. Specifically, OGM belongs to discrete mapping and all grids are independent with each other. To eliminate the requirement of pre-defining discrete grids and the independence assumption in OGM, GPOM is proposed and used in many applications [6], [7]. GPOM naturally captures the neighbouring information thanks to kernelization, but being a non-parametric model, it has a computation complexity of $O(n^3)$, where n is the number of sampled data points. Combining various advantages of GPOM, a faster and simpler parametric model—HM and its Bayesian version, BHM, is proposed later to eliminate the cubical time complexity in GPOM.

Map extrapolation attracts increasing attention in recent years and almost all the current methods for map extrapolation are based on deep neural networks (DNNs) [8], [9], [10], [11], [12], [13]. However, DNNs based methods require huge amount of data for training and hence are still limited in general applications. In this paper, we propose to use Low-Rank Matrix Completion to perform the map prediction. Our proposed LRMC based method falls in between the interpolation and extrapolation methods. Specifically, the LRMC method can not only well interpolate map values, but at the same time can extrapolate the map structure beyond the explored areas.

In the past two decades, the Low-Rank Matrix Completion (LRMC) problem has been well studied. The first theoretically guaranteed exact LRMC algorithm is proposed in [14], where any $n \times n$ incoherent matrices of rank r are proven to be exactly recovered from $Cn^{1.2}r \log n$ uniformly randomly sampled entries with high probability through solving a convex problem of nuclear norm minimization (NNM). Subsequent works [15], [16], [17], [18] refine the provable completion results following the NNM based method. However, since all of the algorithms mentioned above are based on second order methods [19], they can become extremely expensive if the matrix dimension is large [20]. Some first order based methods [20], [21], [22] are developed later. They solve the nuclear norm minimization problem in an iterative manner and rely on Singular Value Decomposition (SVD) of matrices and are suited to large-scale matrix completion problems. In recent years, many other techniques are developed, such as coherent matrix completion [23], [24], [25], non-linear matrix completion [26] and adaptive sampling [27], [28]. A recent survey on LRMC could be found in [29]. In this paper, we mainly employ the first order SVD based iterative method, such as [22] to solve the LRMC problem. The SVD based methods are able to achieve low

prediction error and maintain a low time complexity if certain special problem structure is appropriately leveraged [22].

III. METHODOLOGY

A. Problem Description

Map prediction exhibits the same form as the matrix completion [30] while a 2D map can be treated as a matrix. We propose to use LRMC as the map prediction method. LRMC is defined as completing a partially observed matrix $\tilde{\mathbf{M}} \in \mathbb{R}^{n_1 \times n_2}$ whose part of entries are missing and the corresponding ground-truth matrix \mathbf{M} has a low-rank and incoherent structures. Since the ground-truth rank is usually unknown, a typical way to complete $\tilde{\mathbf{M}}$ is to find a minimal rank matrix \mathbf{X}^* that is consistent with $\tilde{\mathbf{M}}$ on the observed entries. A formal mathematical definition is:

$$\begin{aligned} \mathbf{X}^* = \underset{\mathbf{X}}{\operatorname{argmin}} \quad & \operatorname{rank}(\mathbf{X}) \\ \text{subject to} \quad & X_{ij} = M_{ij}, \quad (i, j) \in \Omega, \end{aligned} \quad (1)$$

where $\Omega \subset \{1, \dots, n_1\} \times \{1, \dots, n_2\}$ is an index set of locations corresponding to the observed entries ($(i, j) \in \Omega$ if M_{ij} is observed).

B. Low-Rank and Incoherent Structures

Two important assumptions for the underlying matrix \mathbf{M} must be satisfied to guarantee the matrix completion results:

- \mathbf{M} should be a low-rank matrix. A low-rank matrix implies its rank value is much smaller than its dimension and this structure enables the possibility to leverage linear dependencies among columns and/or rows of a matrix to predict missing entries [31].
- \mathbf{M} should be an incoherent matrix, which means the coherence of \mathbf{M} is low. Suppose that a *compact* SVD of a rank- r matrix $\mathbf{M} \in \mathbb{R}^{n_1 \times n_2}$ is $\mathbf{M} = \mathbf{U}\Sigma\mathbf{V}^T$, where $\mathbf{U} \in \mathbb{R}^{n_1 \times r}$ and $\mathbf{V} \in \mathbb{R}^{n_2 \times r}$ are stacked left and right singular vectors of \mathbf{M} , while $\Sigma \in \mathbb{R}^{r \times r}$ is the diagonal matrix formed by a number of r singular values of \mathbf{M} . The statistical leverage scores of \mathbf{U} and \mathbf{V} are defined by $\gamma_U = \max_{i \in \{1, \dots, n_1\}} \|\mathbf{U}_{i, :}\|_2^2$ and $\gamma_V = \max_{j \in \{1, \dots, n_2\}} \|\mathbf{V}_{j, :}\|_2^2$, respectively, where $\mathbf{U}_{i, :}$ and $\mathbf{V}_{j, :}$ represent the i^{th} and j^{th} row of \mathbf{U} and \mathbf{V} . Then the coherence of \mathbf{M} is defined as:

$$\gamma_M = \max \{ \gamma_U, \gamma_V \}. \quad (2)$$

The value of γ_M is within $[0, 1]$ and an incoherent structure means the value of γ_M is near 0. A small coherence value implies that the mass of values in the matrix will not concentrate on small portions, but spread out through the whole matrix.

Under the two conditions, the matrix $\tilde{\mathbf{M}}$ has been shown to be provably recoverable [14]. Fig. 2 and Table I illustrate the trends of ranks and coherences of different structured environments which are similar to complex versions of either urban road networks or building layouts. We can observe that the typical structured environments have ranks below 100, which are significantly smaller compared to the original matrix dimensions which could be multiple magnitudes

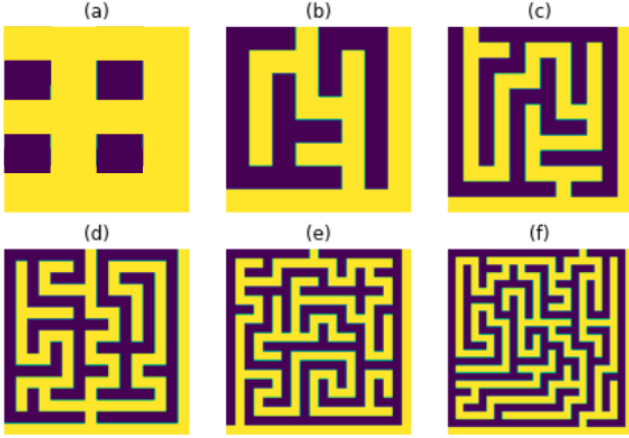


Fig. 2: Maze environments with different ranks.

higher (e.g., 10^4 or above). The trend of ranks also reflects the complexity of the environment, and the higher the rank is, the more complex the environment is. Along with the rank values, we also list the coherence values computed by using Eq. (2). Although the coherence increases proportionally to the rank, they all stay at a low level indicating high incoherence.

TABLE I: Matrix properties of maze environments

	(a)	(b)	(c)	(d)	(e)	(f)
Rank	2	7	11	14	19	22
Coherence	0.0200	0.0400	0.0625	0.0833	0.1000	0.1250

C. Solving LRMC in Map Prediction

To solve the LRMC problem, the rank minimization problem in Eq. (1) is non-convex and NP-hard unfortunately, and hence of little practical use. An alternative way for Eq. (1) is to replace $\text{rank}(\mathbf{X})$ with $\|\mathbf{X}\|_*$, where $\|\mathbf{X}\|_*$ is the nuclear norm of \mathbf{X} , that is, the sum of singular values of \mathbf{X} : $\|\mathbf{X}\|_* = \sum_{k=1}^{\min\{n_1, n_2\}} \sigma_k(\mathbf{X})$, where $\sigma_k(\mathbf{X})$ is the k^{th} largest singular values of \mathbf{X} . The new formulation is:

$$\begin{aligned} \mathbf{X}^* = \underset{\mathbf{X}}{\operatorname{argmin}} \quad & \|\mathbf{X}\|_* \\ \text{subject to} \quad & X_{ij} = M_{ij}, \quad (i, j) \in \Omega. \end{aligned} \quad (3)$$

The nuclear norm $\|\mathbf{X}\|_*$ is an effective convex relaxation to the rank objective and the optimization in Eq. (3) could be solved by semi-definite programming [32].

Nevertheless, the optimizations in Eq. (1) and Eq. (3) aim to exactly recover the partially observed matrix, and this way was claimed as too rigid and may result in over-fitting [22]. A more robust way is to add a regularization parameter in Eq. (3):

$$\begin{aligned} \mathbf{X}^* = \underset{\mathbf{X}}{\operatorname{argmin}} \quad & \|\mathbf{X}\|_* \\ \text{subject to} \quad & \sum_{(i,j) \in \Omega} (X_{ij} - M_{ij})^2 \leq \delta, \end{aligned} \quad (4)$$

where $\delta \geq 0$ is a parameter regulating the training error tolerance. Equivalently Eq. (4) could be reformulated in a

Lagrange form:

$$\mathbf{X}^* = \underset{\mathbf{X}}{\operatorname{argmin}} \sum_{(i,j) \in \Omega} (X_{ij} - M_{ij})^2 + \lambda \|\mathbf{X}\|_*, \quad (5)$$

where λ is a regularization parameter controlling the nuclear norm. A SVD based iterative method proposed in [22] can be used to solve the Eq. (5).

D. Mapping with Non-myopic Planner

We are intrigued if the proposed map prediction can facilitate and improve existing environmental mapping and coverage methods. Popular planning paradigms include lawnmower mapping [33], [34], myopic (short-horizon) planning [35], as well as non-myopic (long-horizon) planning [36]. Lawnmower planning pursues complete spatial coverage, and myopic methods are oftentimes greedy schemes. The non-myopic approaches have been favored by many researchers for environmental mapping and monitoring [37], [38]. Since non-myopic planners optimize paths over some horizon, we adopt and adapt the standard *Traveling Salesman Problem* (TSP), which will be used in our evaluation in the experimental section.

A TSP can be defined as that, given the positions of a collection of way points, what is the shortest possible path of reaching each way point only once and eventually returning to the starting point. To efficiently solve the TSP, we extend our recently developed adaptive k -opt method [39] which exploits and modifies the adaptive k -swap mechanism originally designed for multi-agent task allocation. The core idea is to approximate the NP-hard problem with an opportunistic and localized tour improvement heuristic by taking account of tour constraints adaptively and flexibly, so that real-time solution can be obtained with small sacrifice on solution accuracy.

To determine the number of the way points for navigation guidance in TSP, previous work [14] has theoretically proven that a partially observed low-rank matrix could be perfectly recovered if the number of sampled entries m obeys:

$$m \geq Cn^{1.2}r \log n, \quad (6)$$

where C is a positive numerical coefficient, $n = \max\{n_1, n_2\}$, and r is the rank of the underlying matrix \mathbf{M} [14]. Based on Eq. (6), the placements of way points can be defined as:

$$N'_s = \left\lceil \frac{C \cdot n^{1.2} r \log n}{N} \right\rceil \quad (7)$$

where N'_s represents the number of way points and N is the number of observed map samples within each way point's local sensing range. Note that, in coverage planning scenarios, the robot can take observations in between any pair of successive sampling way points and may introduce additional observations. This implies we may want to reduce the number of sampled way points, and we do this by using a scale factor ε , so that $N_s = \lfloor \varepsilon N'_s \rfloor$, where N_s is the number of sampled way points in TSP planning. We name this TSP based planning with different ε values as TSP_ε . Once a number of N_s points are sampled, a fast and near-shortest




Fig. 3: *Left*: road network displays linear dependent structures. *Right*: The road network is extracted where the dark color represents the road while others are off-road places.




Fig. 4: Comparison between LRMC and BHM under NO (first row) and PO (second row) patterns, respectively. (a)&(e) Ground-truth maps. (b)&(f) Partially observed maps for NO and PO patterns. The missing parts are indicated in green patches. (c)&(g) Predicted maps from LRMC model. (d)&(h) Predicted maps from BHM.

path will be formed using the adaptive k -opt TSP solver and thereafter our robot could start the coverage mission by following the path.

IV. EXPERIMENTS

To validate the proposed framework, we perform extensive simulation evaluations with a variety of complex maps.

A. Experimental Setups and Metrics Design

To account for the sensor limitations as in robotics applications, we use two measurement patterns, *Noisy Observation* (NO) and *Partial Observation* (PO), for constructing the Ω . NO means the observations are imperfect due to random noises (e.g., due to factors of environments/weathers such as snows, or sensor hardware limitations such as sparse and noisy measurements); PO means missed or uncovered regions which are typical in environmental mapping and monitoring scenarios.

To quantitatively compare the two methods, we use two evaluation metrics, *total prediction accuracy* (TPA), and *computational time* (CT). TPA indicates the correct entry matches (including the observed entries) between the entire predicted map and the ground-truth map while CT measures the required time for map prediction computation. In addition, we also evaluate the *mapping convergence rate* during the coverage planning.

B. Demonstration with Road Network Mapping

We first demonstrate that our proposed LRMC method is able to perform map interpolation and extrapolation on a road network map extracted from an urban environmental map, as shown in Fig. 3. The road surface can be detected by proximity sensors such as IR. We assume the robot is an aerial vehicle so that it can spot regions locally below it.

The comparison of the two mapping methods on the urban road map under NO and PO patterns can be seen in Fig. 4. With NO pattern, both methods can generally predict the road network pattern (topology), although BHM is more sensitive to the noises and thus fails to capture the exact road shape. On the other hand, with PO pattern, LRMC significantly outperforms BHM in the presence of multiple missing patches. Different from NO pattern, PO requires that the map prediction not only has interpolation capability, but also can perform slight extrapolation to reason the map structure within missing patches. LRMC leverages the linear dependency among the rows/columns of the map matrix and is able to predict the structure even multiple missing patches occur. In contrast, although BHM can capture local neighbouring spatial relationship by using kernelization, it is unable to capture the linear dependency brought by the low-rank structure, and thus performs worse than LRMC.

C. Quantitative Comparison between LRMC and BHM

We first generate 20 differing complex maze-like layouts/environments by using the Daedalus tool [40], and consider the TPA and CT statistical comparisons under NO and PO patterns for different values of C . The results are shown in Fig. 5 and Fig. 7. The trend of TPA and CT in different rank-valued maps are shown in Fig. 6 for NO and Fig. 8 for PO, respectively.

We first analyze the results for NO pattern. In Fig. 5-*Left*, it can be observed that LRMC could achieve a higher accuracy level. From Fig. 5-*Left*, we want to determine a proper scale of the coefficient C . We can think the coefficient C is actually reflecting the number of observed map samples according to Eq. (6). In Fig. 5-*Left*, we can see that TPA could reach a high level and stay stable when $C \geq 1.5$. We then compare the CT performance as shown in Fig. 5-*Right*. The result clearly distinguishes the two methods. The required computation time for LRMC is much lower than that of BHM. Given the size of the entire map (matrix), the computational complexity of the LRMC algorithm (SoftImpute [22]) we use is $O((m+n)r^2)$, where m and n are the dimensions and r is the rank of the underlying ground-truth matrix. This means the LRMC model has a constant time complexity regardless of the number of observations. BHM is a logistic regression based method and has linear time complexity in terms of the number of observations. Note however, from Fig. 5-*Right* we can see that the LRMC model has a better CT performance than constant time! The reason for this is the LRMC algorithm we use relies on SVD of the matrix. SVD is the most efficient for dense matrices but not that efficient for sparse ones. In our case, at first the number of observations can be small so the input matrix is a sparse




Fig. 5: *Left*: Statistical TPA performance and *Right*: CT performance per value of C for LRMC and BHM under NO pattern from 20 different mazes (with different linear dependencies but the same rank).




Fig. 6: *Left*: TPA performance and *Right*: CT performance per value of C of both methods under NO pattern on different maps with varying values of rank.

one. The runtime benefit will actually increase with more observations received, which is a significant feature for real-time robot mapping applications.

In Fig. 6, we compare the TPA and CT for environments with different ranks. It is clear to see in Fig. 6-*Left* that when the environment becomes more complex (as the rank value increases), the accuracy of BHM decreases while LRMC maintains the same high-level prediction (almost 100%). From the first row of Fig. 4, we can find BHM generally performs well but many minor discrepancies can be found along the road edges. When the environment becomes complex, the number of edges increases and thus the minor errors accumulate and this results in the increase of the total




Fig. 7: *Left*: Statistical TPA performance and *Right*: CT performance per value of C for LRMC and BHM under PO pattern from 20 different mazes (with different linear dependencies but the same rank).




Fig. 8: *Left*: TPA performance and *Right*: CT performance per value of C of both methods under PO pattern on different maps with varying values of rank.

map prediction error. The time performances (Fig. 6-*Right*) of the two methods for different rank-valued maps show the same trend as in Fig. 5-*Right*.

We then consider the results under PO pattern. Similar to what we have in the NO pattern, the TPA and CT results are shown in Fig. 7 and Fig. 8, respectively. The comparison results follow similar trends to the NO pattern.

D. Coverage Planning with Real-Time Map Prediction

We investigate if the proposed map prediction can facilitate and improve existing environmental mapping and coverage planning methods. Our evaluations are based on three most prevalent strategies. The first planning method is the lawnmower (LM) planning: a simple but the most widely used coverage planning method in environmental surveying and monitoring. Specifically, the traversal patterns and resolutions are pre-determined such that the whole area could be swept incrementally by the robot. The second planning is the myopic planning (MP) which is typically based on greedy choices (or best-first actions). Here at each action step, a set of points are randomly sampled in all unexplored areas and the closest one to the robot is selected as the local goal. The third planning is the non-myopic customized $TSP-\epsilon$ as described in Section III-D, where $\epsilon \in \{0.25, 0.5, 0.75\}$.

The mapping process with non-myopic planning on the urban road network map is demonstrated in Fig. 9, where the real observed (sensed) map and corresponding predicted map from LRMC at four different action steps are presented. The results are consistent with that shown in Fig. 4.

To statistically evaluate the methods, we first consider the coverage planning with only raw observations but no map prediction. (As mentioned earlier, real-time predicted complete maps might not be available due to expensive computation cost in large environments.) The results are shown in Fig. 10-*Left*, in which none of the listed methods could complete the map coverage within the given number of steps. This also implies the necessity of real-time map prediction, and the results with map prediction are shown in Fig. 10-*Right*. We can see that, within the same or less number of steps, nearly all of the planning methods could obtain a complete (almost 100%) mapping coverage with real-time map prediction. In addition, the rate for coverage




Fig. 9: Environmental mapping of an urban road map using TSP_0.5. The path is calculated by using an adaptive k -opt TSP planner. In each of (a), (b), (c) and (d), the *Left* figure denotes the observed map while *Right* is the predicted map using LRMC model, respectively.




Fig. 10: Coverage convergence per action step for different coverage planning methods *Left*: without real-time map prediction and *Right*: with real-time map prediction.

convergence is also remarkably improved for all of the listed coverage planning methods.

According to Fig. 10-*Right*, among all the coverage planning methods, TSP_ ϵ methods significantly outperform other methods in terms of coverage convergence, although TSP based method can reach different completeness level as the value of ϵ varies. Empirically a larger value of ϵ can lead to more observations, a longer path (and action steps), but also a better level of completeness (i.e., TSP_0.75). A proper choice of ϵ value can give us a satisfactory trade-off between completeness and required action steps. Even with extremely restricted number of action steps (i.e., less than 200), with a proper ϵ , the TSP based methods can still achieve around 90% coverage. The mapping results at different action steps by TSP_0.5 are shown in Fig. 9-(a), (b), (c) and (d), where we can see that our proposed method is able to perform map interpolation and extrapolation simultaneously on the fly during the execution of environmental mapping.

We also list the required number of action steps of all

TABLE II: Required number of action steps for varying levels of coverage completeness for different methods. The format for the entries indicates # *without real-time map prediction* / # *with real-time map prediction*, where ‘-’ means the corresponding method fails to finish the respective completeness level.

	20%	50%	90%	100%
LM	153 / 69	- / 176	- / 320	- / 356 (99%)
MP	238 / 38	- / 155	- / 566	- / -
TSP_0.25	124 / 37	- / 70	- / 151	- / -
TSP_0.5	127 / 19	- / 101	- / 195	- / 251 (98%)
TSP_0.75	120 / 16	- / 61	- / 240	- / 309

planning methods for different levels of completeness in Table II. It is clear to see that with the proposed real-time LRMC map prediction all coverage planning methods can obtain a substantial improvement of coverage quality, and the required action steps at all listed completeness quantiles are remarkably reduced compared to the results without real-time map prediction.

V. CONCLUSIONS

We consider the map prediction problem where the measurements are noisy, sparse and partially observed. We first show that many maps possess low-rank and incoherent structures. Then we propose to use LRMC to perform the map prediction. Our extensive experiments validate that the LRMC method can effectively perform map prediction (interpolation and extrapolation), and the LRMC method outperforms state-of-the-art Bayesian Hilbert Mapping in terms of mapping accuracy and computation time and could update the entire map in real-time. Lastly we show that the proposed real-time map prediction method can be easily combined with popular coverage planning methods and can significantly improve their coverage convergence rates.

REFERENCES

- [1] Ransalu Senanayake and Fabio Ramos. Bayesian hilbert maps for dynamic continuous occupancy mapping. In *Conference on Robot Learning*, pages 458–471, 2017.
- [2] Fabio Ramos and Lionel Ott. Hilbert maps: scalable continuous occupancy mapping with stochastic gradient descent. *The International Journal of Robotics Research*, 35(14):1717–1730, 2016.
- [3] Alberto Elfes. Using occupancy grids for mobile robot perception and navigation. *Computer*, 22(6):46–57, 1989.
- [4] Simon T O’Callaghan and Fabio T Ramos. Gaussian process occupancy maps. *The International Journal of Robotics Research*, 31(1):42–62, 2012.
- [5] Jinkun Wang and Brendan Englot. Fast, accurate gaussian process occupancy maps via test-data outrees and nested bayesian fusion. In *2016 IEEE International Conference on Robotics and Automation (ICRA)*, pages 1003–1010. IEEE, 2016.
- [6] Maani Ghaffari Jadidi, Jaime Valls Miró, Rafael Valencia, and Juan Andrade-Cetto. Exploration on continuous gaussian process frontier maps. In *2014 IEEE International Conference on Robotics and Automation (ICRA)*, pages 6077–6082. IEEE, 2014.
- [7] Maani Ghaffari Jadidi, Jaime Valls Miro, and Gamini Dissanayake. Mutual information-based exploration on continuous occupancy maps. In *2015 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, pages 6086–6092. IEEE, 2015.
- [8] Jeffrey A Caley, Nicholas RJ Lawrance, and Geoffrey A Hollinger. Deep learning of structured environments for robot search. *Autonomous Robots*, 43(7):1695–1714, 2019.
- [9] Andrzej Pronobis and Rajesh PN Rao. Learning deep generative spatial models for mobile robots. In *2017 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, pages 755–762. IEEE, 2017.
- [10] Kapil Katyal, Katie Popek, Chris Paxton, Joseph Moore, Kevin Wolfe, Philippe Burlina, and Gregory D Hager. Occupancy map prediction using generative and fully convolutional networks for vehicle navigation. *arXiv preprint arXiv:1803.02007*, 2018.
- [11] Rakesh Shrestha, Fei-Peng Tian, Wei Feng, Ping Tan, and Richard Vaughan. Learned map prediction for enhanced mobile robot exploration. In *2019 International Conference on Robotics and Automation (ICRA)*, pages 1197–1204. IEEE, 2019.
- [12] Manish Saroya, Graeme Best, and Geoffrey A Hollinger. Online exploration of tunnel networks leveraging topological cnn-based world predictions. In *Proc. of IEEE/RSJ IROS*, 2020.
- [13] Kapil Katyal, Katie Popek, Chris Paxton, Phil Burlina, and Gregory D Hager. Uncertainty-aware occupancy map prediction using generative networks for robot navigation. In *2019 International Conference on Robotics and Automation (ICRA)*, pages 5453–5459. IEEE, 2019.
- [14] Emmanuel J Candès and Benjamin Recht. Exact matrix completion via convex optimization. *Foundations of Computational mathematics*, 9(6):717, 2009.
- [15] Emmanuel J Candès and Terence Tao. The power of convex relaxation: Near-optimal matrix completion. *IEEE Transactions on Information Theory*, 56(5):2053–2080, 2010.
- [16] Yudong Chen. Incoherence-optimal matrix completion. *IEEE Transactions on Information Theory*, 61(5):2909–2923, 2015.
- [17] David Gross. Recovering low-rank matrices from few coefficients in any basis. *IEEE Transactions on Information Theory*, 57(3):1548–1566, 2011.
- [18] Benjamin Recht. A simpler approach to matrix completion. *Journal of Machine Learning Research*, 12(12), 2011.
- [19] Zhang Liu and Lieven Vandenbergh. Interior-point method for nuclear norm approximation with application to system identification. *SIAM Journal on Matrix Analysis and Applications*, 31(3):1235–1256, 2010.
- [20] Jian-Feng Cai, Emmanuel J Candès, and Zuowei Shen. A singular value thresholding algorithm for matrix completion. *SIAM Journal on optimization*, 20(4):1956–1982, 2010.
- [21] Shuiwang Ji and Jieping Ye. An accelerated gradient method for trace norm minimization. In *Proceedings of the 26th annual international conference on machine learning*, pages 457–464, 2009.
- [22] Rahul Mazumder, Trevor Hastie, and Robert Tibshirani. Spectral regularization algorithms for learning large incomplete matrices. *The Journal of Machine Learning Research*, 11:2287–2322, 2010.
- [23] Guangcan Liu, Qingshan Liu, and Xiaotong Yuan. A new theory for matrix completion. In *Advances in Neural Information Processing Systems*, pages 785–794, 2017.
- [24] Srinadh Bhojanapalli and Prateek Jain. Universal matrix completion. *arXiv preprint arXiv:1402.2324*, 2014.
- [25] Yudong Chen, Srinadh Bhojanapalli, Sujay Sanghavi, and Rachel Ward. Coherent matrix completion. In *International Conference on Machine Learning*, pages 674–682, 2014.
- [26] Brian Eriksson, Laura Balzano, and Robert Nowak. High-rank matrix completion. In *Artificial Intelligence and Statistics*, pages 373–381, 2012.
- [27] Yudong Chen, Srinadh Bhojanapalli, Sujay Sanghavi, and Rachel Ward. Completing any low-rank matrix, provably. *The Journal of Machine Learning Research*, 16(1):2999–3034, 2015.
- [28] Armin Eftekhari, Michael B Wakin, and Rachel A Ward. Mc2: A two-phase algorithm for leveraged matrix completion. *Information and Inference: A Journal of the IMA*, 7(3):581–604, 2018.
- [29] Luong Trung Nguyen, Junhan Kim, and Byonghyo Shim. Low-rank matrix completion: A contemporary survey. *IEEE Access*, 7:94215–94237, 2019.
- [30] Benjamin Recht. A simpler approach to matrix completion. 2009.
- [31] Greg Ongie, Daniel Pimentel-Alarcón, Laura Balzano, Rebecca Willett, and Robert D Nowak. Tensor methods for nonlinear matrix completion. *arXiv preprint arXiv:1804.10266*, 2018.
- [32] Monique Laurent and Antonios Varvitsiotis. Positive semidefinite matrix completion, universal rigidity and the strong arnold property. *Linear Algebra and its Applications*, 452:292–317, 2014.
- [33] Teng Li, Chaoqun Wang, Max Q-H Meng, and Clarence W de Silva. Coverage sampling planner for uav-enabled environmental exploration and field mapping. *arXiv preprint arXiv:1907.05910*, 2019.
- [34] Junnan Song and Shalabh Gupta. An online coverage path planning algorithm. *IEEE Transactions on Robotics*, 34(2):526–533, 2018.
- [35] Fahad Ahmad Khan, Saad Ahmad Khan, Damla Turgut, and Ladislau Bölöni. Greedy path planning for maximizing value of information in underwater sensor networks. In *39th Annual IEEE Conference on Local Computer Networks Workshops*, pages 610–615. IEEE, 2014.
- [36] Yiannis Kantaros, Brent Schlotfeldt, Nikolay Atanasov, and George J Pappas. Asymptotically optimal planning for non-myopic multi-robot information gathering. In *Robotics: Science and Systems*, 2019.
- [37] Kai-Chieh Ma, Lantao Liu, and Gaurav S Sukhatme. An information-driven and disturbance-aware planning method for long-term ocean monitoring. In *2016 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, pages 2102–2108. IEEE, 2016.
- [38] Weizhe Chen and Lantao Liu. Pareto monte carlo tree search for multi-objective informative planning. In *Robotics: Science and Systems*, 2019.
- [39] Zhibei Ma, Lantao Liu, and Gaurav S Sukhatme. An adaptive k-opt method for solving traveling salesman problem. In *2016 IEEE 55th Conference on Decision and Control (Cdc)*, pages 6537–6543. IEEE, 2016.
- [40] Walter D. Pullen. Daedalus program. Available at <http://www.astrolog.org/labyrnth/daedalus.htm>.