

SePaint: Semantic Map Inpainting via Multinomial Diffusion

Zheng Chen, Deepak Duggirala, David Crandall, Lei Jiang, Lantao Liu

Abstract—Prediction beyond partial observations is crucial for robots to navigate in unknown environments because it can provide extra information regarding the surroundings beyond the current sensing range or resolution. In this work, we consider the inpainting of semantic Bird’s-Eye-View maps. We propose SePaint, an inpainting model for semantic data based on generative multinomial diffusion. To maintain semantic consistency, we need to condition the prediction for the missing regions on the known regions. We propose a novel and efficient condition strategy, *Look-Back Condition (LB-Con)*, which performs one-step look-back operations during the reverse diffusion process. By doing so, we are able to strengthen the harmonization between unknown and known parts, leading to better completion performance. We have conducted extensive experiments on different datasets, showing our proposed model outperforms commonly used interpolation methods in various robotic applications.

I. INTRODUCTION

A *dense* semantic map of the surrounding environments can be extremely helpful for autonomous robots to perform tasks such as goal-oriented navigation and autonomous exploration. However, in many cases the robot only has partial observations due to constraints like object occlusions, restricted energy, limited sensor resolution or field of view, etc. Incomplete semantic maps containing missing labels can cause significant confusion for downstream planning modules. In this work, we consider the completion problem for Bird’s-Eye-View (BEV) semantic maps.

One scenario with partially observed semantic BEV maps is the task of coverage mapping, where the goal of the robot is to map an entire environment. An example of this scenario can be seen in Fig. 1-(a), in which an Unmanned Aerial Vehicle (UAV) is assigned to map a city street using a downward-facing sensor. Due to the limited field of view of the sensor, the UAV can only map a small region at each time step, which may lead to an incomplete map if the overall time or energy budget is constrained. An example of an incomplete map is in Fig. 1-(b), where the dark color represents regions that are not mapped.

The problem of semantic BEV map completion is similar to image inpainting and completion [1]–[3], which aims at filling missing regions within a natural RGB image. The core of both problems is the same — learning the distribution of semantic maps or natural RGB images and predicting the absent part based on the known part using the learned distribution. However, an important difference lies in the continuity of the data: values in BEV semantic map are



Fig. 1. (a) Semantic coverage mapping is performed by a UAV in an urban environment using a downward-facing sensor with a limited field of view. (b) Only a partially revealed map can be obtained if time or energy is tightly constrained (dark color represents missing regions. See Fig. 6 for other colors). (c) Our proposed SePaint is able to reason about the underlying semantic structures. (d) Ground truth of the environment.

label IDs and thus are discrete, while values in natural RGB images represent colors and are continuous. Therefore, we can regard the semantics completion as a discrete image inpainting problem, where semantic labels are sampled from categorical distributions.

In this work, we propose **SePaint**: a diffusion-based **Semantic inPainting** framework for categorical data. Our proposed method can be solely based on an unconditionally trained multinomial Denoising Diffusion Probabilistic Model (DDPM). To achieve the functionality of inpainting, we condition the inference process on the known parts and propose an efficient condition strategy, *Look-Back Condition (LB-Con)*. LB-Con applies noise to both the missing region and the known region, performing a one-step *look-back* operation when reversing. This operation achieves a bidirectional merging of known content and generated content, both in the forward and reverse process. In addition, our method is able to eliminate the requirement of paired input and corresponding ground truth data. This does not introduce any bias for specific missing patterns and thus allows the model to generalize any masks.

We validate our proposed framework in different experimental settings with different datasets. First, we show that SePaint is able to inpaint 2D perspective segmentation images with various masks. Then we apply our approach to semantic maps with different sparsities obtained by a depth-based unprojection. Finally, we demonstrate that our method can also produce multiple samples, and this provides a useful way to quantify the inpainting uncertainty. In all settings, we show our proposed method outperforms popular interpolation methods such as nearest-neighbor, linear, and cubic interpolation. In summary, our contributions include:

- We apply a diffusion model to the problem of inpainting for discrete data, and we show our proposed model has very useful potential in various robotic applications.

All authors are with the Luddy School of Informatics, Computing, and Engineering, Indiana University, Bloomington, IN 47408, USA.
Corresponding author email: zc11@iu.edu

- We propose a novel and efficient condition strategy to harmonize the prediction results with the known regions.
- We conduct extensive experiments to validate the advantages of our proposed method over other baselines. We will make the code of our work public. Link: <https://github.com/ArlenCHEN/SePaint>

II. RELATED WORK

Map Inpainting aims to predict missing information in a partially observed map, so that information beyond the robot’s perception range can be estimated. This is important to improve the performance and efficiency of a robot’s behaviors. There is much existing work attempting to inpaint partial maps. For example, Variational Autoencoders (VAEs) have been used to predict unknown map regions beyond frontiers in a 2D occupancy grid map [4]. An occupancy anticipation model is proposed in [5] to learn to extrapolate to unseen regions from RGB and depth images. A Conditional Neural Process (CNP) [6] is used as the backbone to predict unknown regions based on the context of observed parts [7]. Occupancy Prediction Network (OPNet) [8], a faster 3D fully-convolutional network has been proposed to predict 3D occupancy maps, and can be trained using self-supervision. All the above methods belong to data-driven methods and require paired data (observation-gt) to train the map inpainting model. To eliminate the dependency on data, a Low-Rank Matrix Completion (LRMC) based method is proposed to complete a binary map [9]. LRMC is able to achieve both map interpolation and extrapolation on raw poor-quality maps with missing or noisy observations. Although existing work has shown excellent performance, it all targets predicting binary occupancy maps and lacks the capability to predict semantic information.

Diffusion Models [10], [11] or score-based models [12]–[14] are recently emerging powerful alternatives for generative modeling. Essentially, diffusion models are trained to generate an image by denoising randomly sampled noise in a step-by-step manner. Thanks to principled probabilistic modeling, the objective of diffusion models is surprisingly simple, and the training process is more stable than adversarial training in GANs [15]. In addition to the more explicit modeling of the generation process and the more efficient training, diffusion models also show powerful ability on many tasks, such as image translation [16], [17], image editing [18], video synthesis [19], and language-guided image generation [17], [20], and achieve state-of-the-art performance on image synthesis [21] and image inpainting [2]. However, diffusion models suffer from slow inference since they need iterations over many diffusion steps in a chain. Much work has tried to improve the sampling process during inference, including variances during learning [22], non-Markovian diffusion processes [23], and high-order fast ODE-solvers [24]. Based on prior work, a diffusion-based image inpainting model is proposed in [2], which is close to our work and achieves state-of-the-art inpainting performance compared with GANs. However, that paper only

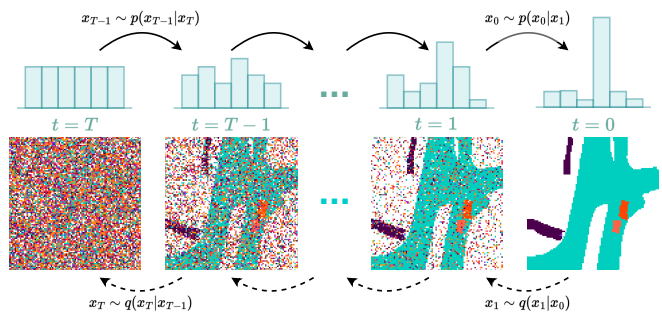


Fig. 2. Diffusion processes on discrete semantic map data. Top arrows: reverse process. Bottom arrows: forward process.

works on ordinal data such as natural RGB images, thus is not suitable for semantic inpainting.

A framework for discrete diffusion is first described in [23], but no evaluation is conducted. Concurrent to [23], another independent work [25] also proposes two methods for learning discrete data, Argmax Flows and Multinomial Diffusion, and conducts extensive experiments and shows remarkable generation capability on both segmentation images and natural language. There is other work modeling discrete data [26]–[31] but all that work only evaluates their methods on text modeling. In this paper, we use the multinomial diffusion model in [25] as our backbone for semantic inpainting.

III. METHODOLOGY

A. Discrete Diffusion Models

The multinomial diffusion model [25] is proposed to generate discrete data, e.g., segmentation images and text. The multinomial diffusion model adopts a categorical distribution as the distribution for transiting among intermediate samples. In this section, we review key steps in the multinomial diffusion model.

Using the categorical distribution as the transition, the forward diffusion process is:

$$q(\mathbf{x}_t|\mathbf{x}_{t-1}) = \mathcal{C}(\mathbf{x}_t; \mathbf{p} = (1 - \beta_t)\mathbf{x}_{t-1} + \beta_t/K), \quad (1)$$

where β_t is determined by a pre-defined variance schedule, and K is the number of classes considered in generation.

We can use \mathbf{x}_0 to predict any intermediate sample \mathbf{x}_t :

$$q(\mathbf{x}_t|\mathbf{x}_0) = \mathcal{C}(\mathbf{x}_t; \mathbf{p} = \bar{\alpha}_t\mathbf{x}_0 + (1 - \bar{\alpha}_t)/K), \quad (2)$$

where $\bar{\alpha}_t = \prod_{i=1}^t \alpha_i$, and $\alpha_i = 1 - \beta_i$.

The conditional posterior $q(\mathbf{x}_{t-1}|\mathbf{x}_t, \mathbf{x}_0)$ can be computed in a closed form:

$$q(\mathbf{x}_{t-1}|\mathbf{x}_t, \mathbf{x}_0) = \mathcal{C}(\mathbf{x}_{t-1}; \mathbf{p} = \mathbf{p}_{\text{post}}(\mathbf{x}_t, \mathbf{x}_0)), \quad (3)$$

where $\mathbf{p}_{\text{post}}(\mathbf{x}_t, \mathbf{x}_0) = \mathbf{p} / \sum_{k=1}^K p_k$ and

$$\mathbf{p} = [\alpha_t\mathbf{x}_t + (1 - \alpha_t)/K] \odot [\bar{\alpha}_{t-1}\mathbf{x}_0 + (1 - \bar{\alpha}_{t-1})/K]. \quad (4)$$

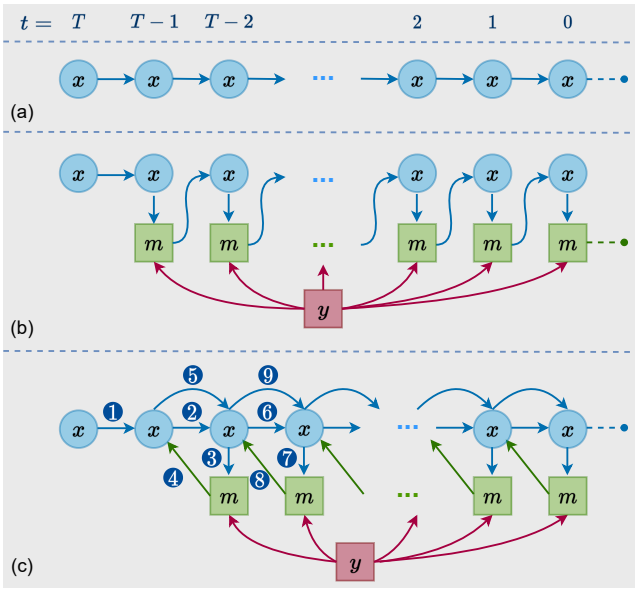


Fig. 3. Different diffusion-based inference strategies. (a) A standard reverse process. (b) A reverse process with Seq-Con. (c) A reverse process with LB-Con.

We assume that $p_\theta(\mathbf{x}_{t-1}|\mathbf{x}_t)$ has the same form as $q(\mathbf{x}_{t-1}|\mathbf{x}_t, \mathbf{x}_0)$. Thus we have:

$$p_\theta(\mathbf{x}_{t-1}|\mathbf{x}_t) = \mathcal{C}(\mathbf{x}_{t-1}; \mathbf{p} = \mathbf{p}_{\text{post}}(\mathbf{x}_t, \hat{\mathbf{x}}_0)), \quad (5)$$

where $\hat{\mathbf{x}}_0$ is the prediction from the neural network $\hat{\mathbf{x}}_0 = NN_\theta(\mathbf{x}_t, t)$. We then are able to obtain the loss for training the multinomial diffusion model:

$$\begin{aligned} L_{t-1} &= \mathcal{D}_{KL}(q(\mathbf{x}_{t-1}|\mathbf{x}_t, \mathbf{x}_0) || p_\theta(\mathbf{x}_{t-1}|\mathbf{x}_t)) \\ &= \mathcal{D}_{KL}(\mathcal{C}(\mathbf{x}_{t-1}; \mathbf{p}_{\text{post}}(\mathbf{x}_t, \mathbf{x}_0)) || \mathcal{C}(\mathbf{x}_{t-1}; \mathbf{p}_{\text{post}}(\mathbf{x}_t, \hat{\mathbf{x}}_0))). \end{aligned} \quad (6)$$

An illustration of the multinomial diffusion model can be seen in Fig. 2.

B. Semantic Inpainting

We denote the results generated during the reverse diffusion process as \mathbf{x} , the input semantic map which contains missing regions as \mathbf{y} , and the binary mask indicating the known (valued as 1) and unknown (valued as 0) as \mathbf{M} . The generated content in the missing regions can be extracted by $\mathbf{M} \odot \mathbf{x}$ while the known regions in \mathbf{y} can be obtained by $(\mathbf{1} - \mathbf{M}) \odot \mathbf{y}$. We denote the mixture of \mathbf{x} and \mathbf{y} based on the mask \mathbf{M} as \mathbf{m} ,

$$\begin{aligned} \mathbf{m} &= g(\mathbf{x}, \mathbf{y}) \\ &= \mathbf{M} \odot \mathbf{x} + (\mathbf{1} - \mathbf{M}) \odot \mathbf{y}. \end{aligned} \quad (7)$$

Our proposed SePaint depends on an unconditionally pre-trained multinomial diffusion model. The inpainting only occurs during inference such that no mask bias is introduced into the model. Our proposed SePaint is only applied to the inference process, where we treat the trained diffusion model as a generative prior, and is used to predict the data sample at any intermediate time step. In Fig. 3-(a), we show a standard diffusion-based inference (the reverse process) from a pure

noise \mathbf{x}_T to a data sample \mathbf{x}_0 . In semantic inpainting, to generate missing content, we need to condition the reverse process on the known regions. One straightforward way to achieve this is to convert the known region to a noisy version at each time step and merge the noisy known regions with the generated sample according to the binary mask. We call this simple condition Sequential Condition (**Seq-Con**), and we show an illustration of an inference process with Seq-Con in Fig. 3-(b). Seq-Con can achieve inpainting and harmonize the content in known regions and unknown regions because the generation of samples \mathbf{x}_t , $t \in [0, T-2]$ at each time step depends on \mathbf{m}_{t+1} , a mixture of content from both regions. More information on known regions from \mathbf{y} is merged as the inference process goes deeper. However, we notice the Seq-Con is not sufficient to harmonize different regions because the merging only happens in one direction. To address this, we propose a bidirectional merging strategy called Look-Back Condition (**LB-Con**), an illustration of which can be seen in Fig. 3-(c). We mark the first 9 steps of the inference to improve readability. In contrast to Seq-Con, LB-Con adds noise to the mixture \mathbf{m} , performing a one-step forward process, meaning the inference process looks back one step when two reverse steps are made. LB-Con performs merging in both the forward process and the reverse process, and thus the content in different regions can be better harmonized. Note that sampling from a categorical distribution can be performed using the Gumbel-Max trick [32], [33], which achieves sampling by adding i.i.d. Gumbel noise to the unnormalized log probabilities and selecting the index with the maximum value.

We now provide more details for implementing LB-Con and present the complete inference pseudo-code for our proposed SePaint. We first describe the merging process at time step t , $\mathbf{m}_t = g(\mathbf{x}_t, \mathbf{y}_t)$, where g is the merging function in Eq. (7); \mathbf{x}_t and \mathbf{y}_t can be obtained based on Eq. (5) and Eq. (2):

$$\begin{aligned} \mathbf{x}_t &\sim p_\theta(\mathbf{x}_t|\mathbf{x}_{t+1}) \\ &\sim \mathcal{C}(\mathbf{x}_t; \mathbf{p} = \mathbf{p}_{\text{post}}(\mathbf{x}_{t+1}, \hat{\mathbf{x}}_0)), \end{aligned} \quad (8)$$

where $\mathbf{x}_0 = NN_\theta(\mathbf{x}_{t+1}, t+1)$.

$$\begin{aligned} \mathbf{y}_t &\sim q(\mathbf{y}_t|\mathbf{y}_0) \\ &\sim \mathcal{C}(\mathbf{y}_t; \mathbf{p} = \bar{\alpha}_t \mathbf{y}_0 + (1 - \bar{\alpha}_t)/K), \end{aligned} \quad (9)$$

where \mathbf{y}_0 is the input partial semantic map.

To perform the *look-back* operation in LB-Con, e.g., the step ④ in Fig. 3-(c), we base the computation on Eq. (1) and have:

$$\begin{aligned} \mathbf{x}_{T-1} &\sim q(\mathbf{x}_{T-1}|\mathbf{m}_{T-2}) \\ &\sim \mathcal{C}(\mathbf{x}_{T-1}; \mathbf{p} = (1 - \beta_{T-2})\mathbf{m}_{T-2} + \beta_{T-2}/K). \end{aligned} \quad (10)$$

We use the cosine schedule for β_t proposed in [22]. The complete pseudo-code of our proposed LB-Con is shown in Algorithm 1.

IV. EXPERIMENTS

A. Settings

We perform extensive experiments on different datasets to demonstrate the advantages of our proposed SePaint. We

Algorithm 1: LB-Con for SePaint

Input: Partially observed semantic map \mathbf{y}_0 ; Binary mask \mathbf{M}
Output: Dense semantic map \mathbf{x}_0

- 1 Import the Gumbel-Max function from Algorithm 1.
- 2 $\mathbf{x}_T \leftarrow \text{Gumbel-Max}(\mathcal{C}(\mathbf{x}_T; \mathbf{p} = \mathbf{1}/K))$
- 3 $\{\varepsilon_i\}_{i=1}^K, \varepsilon_i \sim U[0, 1]$
- 4 $\hat{\mathbf{x}}_0 \leftarrow \text{NN}_\theta(\mathbf{x}_T, T)$
- 5 // Step ①; Eq. (8)
- 6 $\mathbf{x}_{T-1} \leftarrow \text{Gumbel-Max}(\mathcal{C}(\mathbf{x}_{T-1}; \mathbf{p} = \mathbf{p}_{\text{post}}(\mathbf{x}_T, \hat{\mathbf{x}}_0)), \{\varepsilon_i\}_{i=1}^K)$
- 7 **for** $t = T-2, \dots, 0$ **do**
- 8 $\{\varepsilon_i\}_{i=1}^K, \varepsilon_i \sim U[0, 1]$ *if* $t > 1$, *else* $\{\varepsilon_i\}_{i=1}^K = \mathbf{0}$
- 9 $\hat{\mathbf{x}}_0 \leftarrow \text{NN}_\theta(\mathbf{x}_{t+1}, t+1)$
- 10 // Step ②; Eq. (8)
- 11 $\mathbf{x}_t \leftarrow \text{Gumbel-Max}(\mathcal{C}(\mathbf{x}_t; \mathbf{p} = \mathbf{p}_{\text{post}}(\mathbf{x}_{t+1}, \hat{\mathbf{x}}_0)), \{\varepsilon_i\}_{i=1}^K)$
- 12 $\{\varepsilon_i\}_{i=1}^K, \varepsilon_i \sim U[0, 1]$ *if* $t > 1$, *else* $\{\varepsilon_i\}_{i=1}^K = \mathbf{0}$
- 13 // Eq. (9)
- 14 $\mathbf{y}_t \leftarrow \text{Gumbel-Max}(\mathcal{C}(\mathbf{y}_t; \mathbf{p} = \bar{\alpha}_t \mathbf{y}_0 + (1 - \bar{\alpha}_t)/K), \{\varepsilon_i\}_{i=1}^K)$
- 15 // Step ③; Eq. (7)
- 16 $\mathbf{m}_t = \mathbf{M} \odot \mathbf{y}_t + (\mathbf{1} - \mathbf{M}) \odot \mathbf{x}_t$
- 17 $\{\varepsilon_i\}_{i=1}^K, \varepsilon_i \sim U[0, 1]$ *if* $t > 1$, *else* $\{\varepsilon_i\}_{i=1}^K = \mathbf{0}$
- 18 // Step ④; Eq. (10)
- 19 $\mathbf{x}_{t+1} \leftarrow \text{Gumbel-Max}(\mathcal{C}(\mathbf{x}_{t+1}; \mathbf{p} = (1 - \beta_t) \mathbf{m}_t + \beta_t/K), \{\varepsilon_i\}_{i=1}^K)$
- 20 $\hat{\mathbf{x}}_0 \leftarrow \text{NN}_\theta(\mathbf{x}_{t+1}, t+1)$
- 21 $\{\varepsilon_i\}_{i=1}^K, \varepsilon_i \sim U[0, 1]$ *if* $t > 1$, *else* $\{\varepsilon_i\}_{i=1}^K = \mathbf{0}$
- 22 // Step ⑤; Eq. (8)
- 23 $\mathbf{x}_t \leftarrow \text{Gumbel-Max}(\mathcal{C}(\mathbf{x}_t; \mathbf{p} = \mathbf{p}_{\text{post}}(\mathbf{x}_{t+1}, \hat{\mathbf{x}}_0)), \{\varepsilon_i\}_{i=1}^K)$
- 24 **end**
- 25 **return** \mathbf{x}_0

show that our proposed method can be applied to both perspective data and various forms of BEV data.

We use two popular datasets for evaluation: Cityscapes [34] and nuScenes [35].

Cityscapes is a standard 2D semantic segmentation dataset with pixel-wise labels for 19 common categories (e.g., road, sidewalk, tree, person, car, etc.) for images from different urban environments. To reduce the computational burden, we downsample the images used for training the multinomial diffusion model to 128×256 using nearest-neighbor interpolation. We use all categories as prediction targets. The official training set (2,975 images) is split into 2,500 images for training and 475 images for validation. The original testing set is not used. All experimental results reported for Cityscapes are based on the validation images.

nuScenes is a large-scale dataset for autonomous driving research consisting of high-resolution 3D sensor data collected from a fleet of autonomous vehicles driven in the cities of Boston and Singapore. The dataset contains 1,000 short video sequences, each frame of which has rich semantic BEV annotations with 14 classes (e.g., drivable surface, sidewalk, car, truck, bus, etc.). We extract a total of 28,508 BEV semantic annotations and split those into a training set (28,008 images) and a testing set (500 images). We use a downsampled size of 100×100 during training and testing.

nuScenes also includes semantic annotations for all sensing points from a 32-beam LiDAR. To approximate unprojected BEV maps from cameras, we project the semantically annotated LiDAR points onto the front camera’s image plane,




Fig. 4. Qualitative evaluation on Cityscapes. We use white to represent the missing regions (1st column).




Fig. 5. Quantitative evaluations for different methods on Cityscapes.

giving us both semantic and depth images. We use different interpolation methods to densify the projected images: nearest neighbor for semantic images and linear interpolation for depth images. The densification gives us dense unprojected BEV maps. For our sparse unprojected BEV maps, we directly project the raw annotated LiDAR points to BEV space.

We use the model proposed in [25] as our unconditional multinomial diffusion model. During training, we use a learning rate of 1×10^{-4} , batch size of 4 for Cityscapes and 16 for nuScenes, the total number of diffusion steps is set as 4,000, and we use random horizontal flips as data augmentation. Hand-crafted masks from [2] are used to approximate the missing regions during inference. For baselines we use nearest-neighbor interpolation, linear interpolation, and cubic interpolation. Two metrics are used to quantitatively evaluate methods: Accuracy and mean Intersection over Union (mIoU), where $\text{Acc} = \frac{n_{tp} + n_{tn}}{n_{tp} + n_{tn} + n_{fp} + n_{fn}}$ and $\text{IoU} = \frac{n_{tp}}{n_{tp} + n_{fp} + n_{fn}}$, with n_{tp}, n_{tn}, n_{fp} and n_{fn} as true positives, true negatives,




Fig. 6. Qualitative evaluation on nuScenes. We use black to represent the missing regions (1st column).




Fig. 7. Quantitative evaluations for different methods on nuScenes.




Fig. 8. Completion for unprojected BEV maps. Unprojected BEV map using Top: depth image; and Bottom: LiDAR points.

false positives, and false negatives, respectively.

B. Perspective Semantics Completion

We first validate the effectiveness and advantages of our proposed SePaint on 2D perspective semantic images from Cityscapes. A set of qualitative evaluations can be seen in Fig. 4. It can be seen that SePaint outperforms all other baselines in terms of prediction accuracy and realistic inference about the scene structure. The reason for this advantage is that SePaint is based on a powerful diffusion generative model and is able to learn the categorical distribution of semantic labels. In contrast, interpolation-based methods only make rigid predictions based on values in known regions and lack prior knowledge of urban scenarios. Among different interpolation-based methods, Nearest-Neighbor has the best performance since semantic labels are discrete. Linear interpolation has the worst performance among all listed methods in Fig. 4. Cubic interpolation behaves even worse than Linear interpolation (not shown due to space limitations). Linear/Cubic interpolation can generate new classes out-of-distribution for the current image because there is no causal relation among labels. For example, interpolating between the class 8 and class 10 leads to a class of 9, which is not necessarily the class that is missed and is physically close to the class 8 and 10. The true underlying class close to the class 8 and 10 can be another possible labels, say 15. This problem exists because all labels assigned to classes are randomly set — they are fully independent of each other. Quantitative comparisons can be seen in Fig. 5, where values of mIoU and Acc for different methods are reported. We run semantic inpainting for all images in the validation set and show the mean of the metric values.

C. BEV Semantics Completion

In this section, we show comparisons of different methods for completing city semantic BEV maps from nuScenes. We first show qualitative and quantitative results for scenarios of semantic coverage mapping. Then we validate that SePaint is also able to complete unprojected BEV maps where extreme sparsity can happen and strong extrapolations are required.

1) *Completion for Coverage Mapping:* We approximate the partially observed semantic BEV map by applying various binary masks to the semantic BEV data in nuScenes.

TABLE I
COMPARISON OF DIFFERENT METHODS FOR COMPLETING THE
UNPROJECTED BEV MAP.

Input map type	Method	mIoU	Acc
Dense unprojected map	SePaint	14.55	53.54
	Nearest	9.77	37.16
	Linear	10.51	42.63
	Cubic	7.21	42.59
Sparse unprojected map	SePaint	12.98	59.21
	Nearest	7.86	21.27
	Linear	8.57	42.45
	Cubic	7.66	47.87




Fig. 9. Different samples from SePaint.

A set of qualitative evaluations can be seen in Fig. 6, where the dark regions in the 1st column are missed. Similar to the results for Cityscapes, our proposed method can predict the most semantically-consistent BEV structures and achieve the best performance among all the methods under comparison. Quantitative comparisons can be seen in Fig. 7.

2) *Completion for Unprojected BEV Map*: One of the commonly used ways for having a BEV semantic map from vision is to unproject a 2D segmentation image using a corresponding depth image or LiDAR points. However, extreme sparsity can be brought down due to the limited resolution of depth images or LiDAR points. We show examples of completing unprojected BEV maps in Fig. 8. We can see from the 1st column that the dark color (missing regions) takes most of the pixels, especially in the LiDAR-based projected map (the bottom row). Even with the presence of extreme sparsity, our method can still *imagine* a highly reasonable structure (see the 3rd column of Fig. 8). In contrast, the representative of interpolation methods for discrete data, Nearest-Neighbor, can only predict missing regions in an overly rigid and simplified manner, leading to unacceptable densified maps. More quantitative details of comparing different methods for the unprojected BEV maps can be found in Table I.

D. Multiple Samples

In this section, we show another unique advantage of our proposed method that can be applied to many different robotic applications such as coverage mapping, goal-oriented navigation, and autonomous exploration: the capability to

TABLE II
ABLATION STUDY FOR LB-CON AND SEQ-CON

		City	nuScenes	Dense	Sparse	Samples
mIoU	LB-Con	31.88	37.21	14.55	12.98	35.82
	Seq-Con	27.65	33.11	12.34	8.56	30.98
Acc	LB-Con	89.02	92.33	53.54	59.21	89.25
	Seq-Con	85.88	88.39	51.14	54.27	83.66

generate multiple samples, each of which can be consistent and harmonized well with the known regions. We list some examples in Fig. 9, where 5 different samples for each partial map are listed in a row-wise manner. The ability to generate multiple samples can not only provide different possibilities for missing regions but also offer a way to quantify the uncertainty of predictions for the missing parts. The uncertainty can be defined as high if the predictions across different samples are highly varying, and as low if the predictions across different samples are almost the same. The estimated uncertainty can be very useful to robot planning when decision-making on partial maps is needed.

E. Ablation Study

We perform an ablation study to demonstrate the advantages of our proposed LB-Con over the simple Seq-Con. In this study, we evaluate diffusion-based inpaintings with different condition strategies on five experimental settings discussed in the above sections. The settings include the completion for **Cityscapes**, **nuScenes**, **Dense** unprojected BEV maps, **Sparse** unprojected BEV maps, and Generation of multiple **Samples**. Quantitative evaluations in terms of mIoU and Acc are reported in Table II, from which we can see our proposed LB-Con has systematic improvements in all tested data compared with the simple Seq-Con.

V. CONCLUSION AND FUTURE WORK

We present an inpainting model, SePaint, for discrete semantic data based on multinomial diffusion. Our proposed SePaint is a conditional generative model where the generation for missing regions is conditional on the known regions. To harmonize the content of the two regions, we propose a new condition strategy, Look-back Condition (LB-Con), which performs one-step look-back operations during the reverse diffusion process. We have reported extensive experimental results and validated that our proposed model outperforms commonly used interpolation methods for completing data. To the best of our knowledge, there are no other generative models, such as VAEs or GANs, working on the generation of discrete segmentation images/maps.

One bottleneck for the current model is the speed during inference. The current generation is still far from being useful in real-time scenarios, which are commonly required in robotic applications. In the future, we will investigate how to accelerate the discrete multinomial diffusion model and refine our proposed LB-Con such that the inpainting inference speed can be real-time or close to be real-time. We will also combine real robotic data with our model and validate our method in real-world navigation experiments.

REFERENCES

- [1] Wenbo Li, Zhe Lin, Kun Zhou, Lu Qi, Yi Wang, and Jiaya Jia. Mat: Mask-aware transformer for large hole image inpainting. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 10758–10768, 2022.
- [2] Andreas Lugmayr, Martin Danelljan, Andres Romero, Fisher Yu, Radu Timofte, and Luc Van Gool. Repaint: Inpainting using denoising diffusion probabilistic models. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 11461–11471, 2022.
- [3] Wentao Wang, Li Niu, Jianfu Zhang, Xue Yang, and Liqing Zhang. Dual-path image inpainting with auxiliary gan inversion. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 11421–11430, 2022.
- [4] 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.
- [5] Santhosh K Ramakrishnan, Ziad Al-Halah, and Kristen Grauman. Occupancy anticipation for efficient exploration and navigation. In *Computer Vision–ECCV 2020: 16th European Conference, Glasgow, UK, August 23–28, 2020, Proceedings, Part V 16*, pages 400–418. Springer, 2020.
- [6] Marta Garnelo, Dan Rosenbaum, Christopher Maddison, Tiago Ramalho, David Saxton, Murray Shanahan, Yee Whye Teh, Danilo Rezende, and SM Ali Eslami. Conditional neural processes. In *International conference on machine learning*, pages 1704–1713. PMLR, 2018.
- [7] Amine Elhafsi, Boris Ivanovic, Lucas Janson, and Marco Pavone. Map-predictive motion planning in unknown environments. In *2020 IEEE International Conference on Robotics and Automation (ICRA)*, pages 8552–8558. IEEE, 2020.
- [8] Lizi Wang, Hongkai Ye, Qianhao Wang, Yuman Gao, Chao Xu, and Fei Gao. Learning-based 3d occupancy prediction for autonomous navigation in occluded environments. In *2021 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, pages 4509–4516. IEEE, 2021.
- [9] Zheng Chen, Shi Bai, and Lantao Liu. Efficient map prediction via low-rank matrix completion. In *2021 IEEE International Conference on Robotics and Automation (ICRA)*, pages 13953–13959. IEEE, 2021.
- [10] Jascha Sohl-Dickstein, Eric Weiss, Niru Maheswaranathan, and Surya Ganguli. Deep unsupervised learning using nonequilibrium thermodynamics. In *International Conference on Machine Learning*, pages 2256–2265. PMLR, 2015.
- [11] Jonathan Ho, Ajay Jain, and Pieter Abbeel. Denoising diffusion probabilistic models. *Advances in Neural Information Processing Systems*, 33:6840–6851, 2020.
- [12] Yang Song and Stefano Ermon. Generative modeling by estimating gradients of the data distribution. *Advances in neural information processing systems*, 32, 2019.
- [13] Yang Song and Stefano Ermon. Improved techniques for training score-based generative models. *Advances in neural information processing systems*, 33:12438–12448, 2020.
- [14] Yang Song, Jascha Sohl-Dickstein, Diederik P Kingma, Abhishek Kumar, Stefano Ermon, and Ben Poole. Score-based generative modeling through stochastic differential equations. *arXiv preprint arXiv:2011.13456*, 2020.
- [15] Ian Goodfellow, Jean Pouget-Abadie, Mehdi Mirza, Bing Xu, David Warde-Farley, Sherjil Ozair, Aaron Courville, and Yoshua Bengio. Generative adversarial nets. In Z. Ghahramani, M. Welling, C. Cortes, N. Lawrence, and K.Q. Weinberger, editors, *Advances in Neural Information Processing Systems*, volume 27. Curran Associates, Inc., 2014.
- [16] Xuan Su, Jiaming Song, Chenlin Meng, and Stefano Ermon. Dual diffusion implicit bridges for image-to-image translation. In *International Conference on Learning Representations*, 2022.
- [17] Robin Rombach, Andreas Blattmann, Dominik Lorenz, Patrick Esser, and Björn Ommer. High-resolution image synthesis with latent diffusion models. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 10684–10695, 2022.
- [18] Chenlin Meng, Yutong He, Yang Song, Jiaming Song, Jiajun Wu, Jun-Yan Zhu, and Stefano Ermon. Sdedit: Guided image synthesis and editing with stochastic differential equations. In *International Conference on Learning Representations*, 2021.
- [19] Jonathan Ho, Tim Salimans, Alexey Gritsenko, William Chan, Mohammad Norouzi, and David J Fleet. Video diffusion models. *arXiv preprint arXiv:2204.03458*, 2022.
- [20] Gwanghyun Kim, Taesung Kwon, and Jong Chul Ye. Diffusionclip: Text-guided diffusion models for robust image manipulation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 2426–2435, 2022.
- [21] Prafulla Dhariwal and Alexander Nichol. Diffusion models beat gans on image synthesis. *Advances in Neural Information Processing Systems*, 34:8780–8794, 2021.
- [22] Alexander Quinn Nichol and Prafulla Dhariwal. Improved denoising diffusion probabilistic models. In *International Conference on Machine Learning*, pages 8162–8171. PMLR, 2021.
- [23] Jiaming Song, Chenlin Meng, and Stefano Ermon. Denoising diffusion implicit models. *arXiv preprint arXiv:2010.02502*, 2020.
- [24] Cheng Lu, Yuhao Zhou, Fan Bao, Jianfei Chen, Chongxuan Li, and Jun Zhu. Dpm-solver: A fast ode solver for diffusion probabilistic model sampling in around 10 steps. *arXiv preprint arXiv:2206.00927*, 2022.
- [25] Emiel Hoogeboom, Didrik Nielsen, Priyank Jaini, Patrick Forré, and Max Welling. Argmax flows and multinomial diffusion: Learning categorical distributions. *Advances in Neural Information Processing Systems*, 34:12454–12465, 2021.
- [26] Jacob Austin, Daniel D Johnson, Jonathan Ho, Daniel Tarlow, and Rianne van den Berg. Structured denoising diffusion models in discrete state-spaces. *Advances in Neural Information Processing Systems*, 34:17981–17993, 2021.
- [27] Andrew Campbell, Joe Benton, Valentin De Bortoli, Tom Rainforth, George Deligiannidis, and Arnaud Doucet. A continuous time framework for discrete denoising models. *arXiv preprint arXiv:2205.14987*, 2022.
- [28] Sander Dieleman, Laurent Sartran, Arman Roshannai, Nikolay Savinov, Yaroslav Ganin, Pierre H Richemond, Arnaud Doucet, Robin Strudel, Chris Dyer, Conor Durkan, et al. Continuous diffusion for categorical data. *arXiv preprint arXiv:2211.15089*, 2022.
- [29] Ting Chen, Ruixiang Zhang, and Geoffrey Hinton. Analog bits: Generating discrete data using diffusion models with self-conditioning. *arXiv preprint arXiv:2208.04202*, 2022.
- [30] Haoran Sun, Lijun Yu, Bo Dai, Dale Schuurmans, and Hanjun Dai. Score-based continuous-time discrete diffusion models. *arXiv preprint arXiv:2211.16750*, 2022.
- [31] Robin Strudel, Corentin Tallec, Florent Althé, Yilun Du, Yaroslav Ganin, Arthur Mensch, Will Grathwohl, Nikolay Savinov, Sander Dieleman, Laurent Sifre, et al. Self-conditioned embedding diffusion for text generation. *arXiv preprint arXiv:2211.04236*, 2022.
- [32] Emil Julius Gumbel. *Statistical theory of extreme values and some practical applications: a series of lectures*, volume 33. US Government Printing Office, 1954.
- [33] Chris J Maddison, Daniel Tarlow, and Tom Minka. A* sampling. *Advances in neural information processing systems*, 27, 2014.
- [34] Marius Cordts, Mohamed Omran, Sebastian Ramos, Timo Rehfeld, Markus Enzweiler, Rodrigo Benenson, Uwe Franke, Stefan Roth, and Bernt Schiele. The cityscapes dataset for semantic urban scene understanding. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 3213–3223, 2016.
- [35] Holger Caesar, Varun Bankiti, Alex H Lang, Sourabh Vora, Venice Erin Liong, Qiang Xu, Anush Krishnan, Yu Pan, Giancarlo Baldan, and Oscar Beijbom. nuscenes: A multimodal dataset for autonomous driving. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 11621–11631, 2020.