

LEGO: Leveraging Experience in Roadmap Generation for Sampling-Based Planning

Rahul Kumar¹, Aditya Mandalika², Sanjiban Choudhury² and Siddhartha S. Srinivasa²

Abstract—We consider the problem of leveraging prior experience to generate roadmaps in sampling-based motion planning. A desirable roadmap is one that is sparse, allowing for fast search, with nodes spread out at key locations such that a low-cost feasible path exists. An increasingly popular approach is to learn a distribution of nodes that would produce such a roadmap. State-of-the-art is to train a conditional variational auto-encoder (CVAE) on the prior dataset with the shortest paths as target input. While this is quite effective on many problems, we show it can fail in the face of complex obstacle configurations or mismatch between training and testing.

We present an algorithm LEGO that addresses these issues by training the CVAE with target samples that satisfy two important criteria. Firstly, these samples belong only to *bottleneck* regions along near-optimal paths that are otherwise difficult-to-sample with a uniform sampler. Secondly, these samples are spread out across *diverse regions* to maximize the likelihood of a feasible path existing. We formally define these properties and prove performance guarantees for LEGO. We extensively evaluate LEGO on a range of planning problems, including robot arm planning, and report significant gains over both heuristics and learned baselines.

I. INTRODUCTION

We examine the problem of leveraging prior experience in sampling-based motion planning. In this framework, the continuous configuration space of a robot is sampled to construct a graph or *roadmap* [1, 2] where vertices represent robot configurations and edges represent potential movements of the robot. A shortest path algorithm [3] is then run to compute a path between any two vertices on the roadmap. The main challenge is to place a *small set of samples in key locations* such that the algorithm can find a high quality path with small computational effort as shown in Fig. 1b.

Typically, low dispersion samplers such as Halton sequences [4] are quite effective in uniformly covering the space and thus bounding the solution quality [5] (Fig. 1a). However, as they decrease dispersion uniformly in C-space, a narrow passage with δ clearance requires $O((\frac{1}{\delta})^d)$ samples to find a path. This motivates the need for biased sampling to *selectively densify* in regions where there might be a narrow passage [6–10]. These techniques are applicable across a wide range of domains and perform quite well in practice.

However, not all narrow passages are relevant to a given query. Biased sampling techniques, which do not have access to the likelihood of the optimal path passing through a region,




Fig. 1: Comparison of roadmaps generated from (a) uniform Halton sequence sampler and (b) from a generative model trained using LEGO. The task is to plan from the shown configuration over the table and obstacle to the other side. The graph is visualized by end effector traces of the edges.

can still end up dropping samples in more regions than necessary. Interestingly, the different environments that a robot operates in share a lot of structural similarity. Hence, we can use information extracted from planning on one such environment for deciding how to sample on another; we can *learn* sampling distributions using tools such as a conditional variational auto-encoder (CVAE). Ichter et al. [11] propose a useful approximation to train a learner to drop samples along the *predicted* shortest path: given a training dataset of worlds, compute shortest paths, and train a model to independently predict nodes belonging to the path. After all, the best a generative model can do, is to sample only along the true shortest path. However, this puts *all of the burden* on the learner. Any amount of prediction error, be it due to approximation or train-test mismatch, results in failure to find a feasible path.

We argue that a sampler, instead of trying to predict the shortest path, needs to only identify key regions to focus sampling at, and let the search algorithm determine the shortest path. Essentially, we ask the following question:

How can we share the responsibility of finding the shortest path between the sampler and search?

Our key insight is for the sampler to predict not the shortest path, but samples that possess two main characteristics. First, we only need to predict samples in bottleneck regions. These are regions containing near-optimal paths, but are difficult for a uniform sampler to reach. Secondly, we need diversity amongst samples. Train-test mismatch is common and to be robust against it we need to sample nodes belonging to a diverse set of alternate paths. The search algorithm can then operate on a sparse graph containing useful but diverse samples to compute the shortest path.

¹Department of Computer Science, Indian Institute of Technology, Kharagpur {vernwalrahul}@iitkgp.ac.in

²Paul G. Allen School of Computer Science and Engineering, University of Washington {adityavk, sanjibac, sidddh}@cs.uw.edu

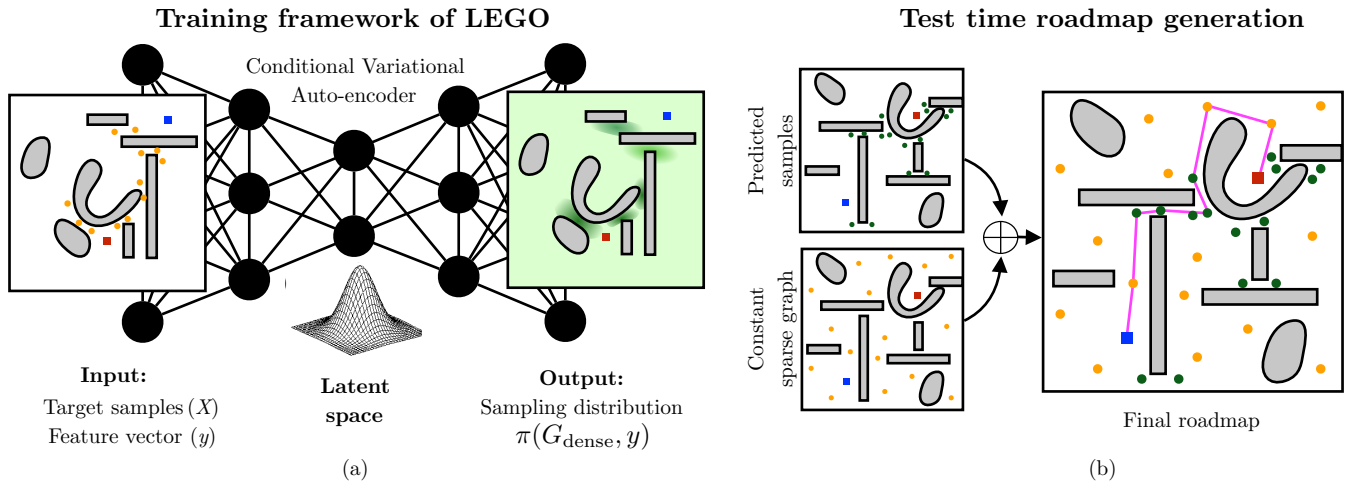


Fig. 2: The LEGO framework for training a CVAE to predict a roadmap. (a) The training process for learning a generative sampling distribution using a CVAE. The input is a pair of candidate samples and feature vector. (b) At test time, the model is sampled to get vertices which are then composed with a constant sparse graph to get a final roadmap.

We present an algorithmic framework, **Leveraging Experience with Graph Oracles (LEGO)** summarized in Fig. 2, for training a CVAE on a prior database of worlds to learn a generative model that can be used for roadmap construction. During training (Fig. 2a), LEGO processes a uniform dense graph to identify a sparse subset of vertices. These vertices are a *diverse* set of *bottleneck* nodes through which a near-optimal path must pass. These are then fed into a CVAE [12] to learn a generative model. At test time (Fig. 2b), the model is sampled to get a set of vertices which is additionally composed with a sparse uniform graph to get a final roadmap. This roadmap is then used by the search algorithm to find the shortest path.

We make the following contributions:

- 1) A framework for training a CVAE to predict a roadmap with different target inputs. We identify two main shortcomings of the state-of-the-art [11] that uses the shortest path as the target input - failures in approximation, and failures due to train-test mismatch (Section IV).
- 2) LEGO, an algorithm that tackles both of these issues. It first generates multiple diverse shortest paths, and then extracts bottleneck nodes along such paths to use as the target input for the CVAE (Section VI).
- 3) We show that LEGO outperforms several learning and heuristic sampling baselines on a set of \mathbb{R}^2 , \mathbb{R}^5 , \mathbb{R}^7 , \mathbb{R}^8 and \mathbb{R}^9 problems. In particular, we show that it is robust to changes in training and test distribution (Section VII).

II. RELATED WORK

The seminal work of Hsu et al. [13] provides a crisp analysis of the shortcomings of uniform sampling techniques in the presence of artifacts such as narrow passages. This has led to a plethora of non-uniform sampling approaches that densify selectively [6–10].

Adaptive sampling in the context of roadmaps aims to exploit structure of the environment to place samples in promising areas. A number of works exploited structure of the workspace

to achieve this. While some of them attempt to sample between regions of collision to identify narrow passages [6, 14–18], others sample near or on the obstacles [19, 20]. There are approaches that divide the configuration space into regions and either select different region-specific planning strategies [21] or use entropy of samples in a particular region to refine sampling [22]. Other methods try to model the free space to speed up planning [23–25]. While these techniques are quite successful in a large set of problems, they can place samples in regions where an optimal path is unlikely to traverse.

A different class of solutions look at adapting sampling distributions online during the planning cycle. This requires a trade-off between exploration of the configuration space and exploitation of the current best solution. Preliminary approaches define a utility function to do so [26, 27] or use online learning [10]; however these are not amenable to using priors. Diankov and Kuffner [28] employs statistical techniques to sample around a search tree. Zucker et al. [29], Kuo et al. [30] formalize sampling as a model-free reinforcement learning problem and learn a parametric distribution. Since these problems are non i.i.d learning problems, they do require interactive learning and do not enjoy the strong guarantees of supervised learning.

There has been a lot of recent effort on finding low dimensional structure in planning [31]. In particular, generative modeling tools like variational autoencoders [32] have been used to great success [33–37]. We base our work on Ichter et al. [11] where a CVAE is trained to learn the shortest path distribution.

III. PROBLEM FORMULATION

Given a database of prior worlds, the overall goal is to learn a policy that predicts a roadmap which in turn is used by a search algorithm to efficiently compute a high quality feasible path. Let \mathcal{X} denote a d -dimensional configuration space. Let \mathcal{X}_{obs} be the portion in collision and $\mathcal{X}_{\text{free}} = \mathcal{X} \setminus \mathcal{X}_{\text{obs}}$ denote the free space. Let a path $\xi : [0, 1] \rightarrow \mathcal{X}$ be a continuous

mapping from index to configurations. A path ξ is said to be collision free if $\xi(\tau) \in \mathcal{X}_{\text{free}}$ for all $\tau \in [0, 1]$. Let $c(\xi)$ be a cost functional that maps ξ to a bounded non-negative cost $[0, c_{\text{max}}]$. Moreover, we set $c(\emptyset) = c_{\text{max}}$. We define a *motion planning problem* $\Lambda = \{x_s, x_g, \mathcal{X}_{\text{free}}\}$ as a tuple of start configuration $x_s \in \mathcal{X}_{\text{free}}$, goal configuration $x_g \in \mathcal{X}_{\text{free}}$ and free space $\mathcal{X}_{\text{free}}$. Given a problem, a path ξ is said to be *feasible* if it is collision free, $\xi(0) = x_s$ and $\xi(1) = x_g$. Let Ξ_{feas} denote the set of all feasible paths. We wish to solve the *optimal* motion planning problem by finding a feasible path ξ^* that minimizes the cost functional $c(\cdot)$, i.e. $c(\xi^*) = \inf_{\xi \in \Xi_{\text{feas}}} c(\xi)$.

We now embed the problem on a graph $G = (V, E)$ such that each vertex $v \in V$ is an element of $v \in \mathcal{X}$. The graph follows a connectivity rule expressed as an indicator function $\text{Link} : \mathcal{X} \times \mathcal{X} \rightarrow \{0, 1\}$ to denote if two configurations should have an edge¹. The weight of an edge $c(u, v)$ is the cost of traversing the edge. We reuse ξ to denote a path on the graph.

We introduce a couple graph operations. Let $|G|$ denote the cardinality of the graph, i.e. the size of $|V|$ ². We use the notation $G \stackrel{\pm}{\leftarrow} X$ to compactly denote insertion of a new set of vertices X , i.e. $V \leftarrow V \cup X$, $E \leftarrow E \cup \{(u, v) \mid u \in X, \text{Link}(u, v) = 1\}$.

A graph search algorithm ALG is given a graph G and a planning problem Λ . First, it adds the start goal pair to the graph, i.e. $G' = G \stackrel{\pm}{\leftarrow} \{x_s, x_g\}$. It then collision checks edges against $\mathcal{X}_{\text{free}}$ till it finds the shortest feasible path ξ^* which is then returned. Hence, the cost of such a path can be found by evaluating $c(\text{ALG}(G, \Lambda))$. If ALG is unable to find any feasible path, it returns \emptyset which corresponds to c_{max} .

Definition 1 (Dense Graph). *We assume we have a dense graph $G_{\text{dense}} = (V_{\text{dense}}, E_{\text{dense}})$ that is sufficiently large to connect the space i.e. for any plausible planning problem, it contains a sufficiently low cost feasible path.*

Henceforth, we care about competing with G_{dense} . We reiterate that searching this graph, $\text{ALG}(G_{\text{dense}}, \Lambda)$, is too computationally expensive to perform online.

We wish to learn a mapping from features extracted from the problem to a sparse subgraph of G_{dense} . Let $y \in \mathbb{R}^m$ be a feature representation of the planning problem. Let $\pi(G_{\text{dense}}, y)$ be a *subgraph predictor oracle* that maps the feature vector to a subgraph $G \subset G_{\text{dense}}$, such that $|G| \leq N$. We wish to solve the following optimization problem:

Problem 1 (Optimal Subgraph Prediction). *Given a joint distribution $P(\Lambda, y)$ of features and problems, and a dense graph G_{dense} , compute a subgraph predictor oracle π that minimizes the ratio of the cost of the shortest feasible path in the subgraph to the dense graph:*

$$\pi^* = \arg \min_{\pi \in \Pi} \mathbb{E}_{(\Lambda, y) \sim P(\Lambda, y)} \left[\frac{c(\text{ALG}(\pi(G_{\text{dense}}, y), \Lambda))}{c(\text{ALG}(G_{\text{dense}}, \Lambda))} \right] \quad (1)$$

¹Note this does not involve collision checking. We consider undirected graphs for simplicity. However, it easily extends to directed graphs.

²Alternatively we can also use the size of $|E|$

IV. FRAMEWORK FOR PREDICTING ROADMAPS

We now present a framework for training graph predicting oracles as illustrated in Fig. 2(a). This is a generalization of the approach presented in [11]. The framework applies three main approximations. First, instead of predicting a subgraph $G \subset G_{\text{dense}}$, we learn a mapping that directly predicts states $x \in \mathcal{X}$ in continuous space.³ Secondly, instead of solving a structured prediction problem, we learn an i.i.d sampler that will be invoked repeatedly to get a set of vertices. These vertices are then connected according to an underlying connection rule, such as k -NN, to create a graph. Thirdly, we compose the sampled graph with a *constant sparse graph* $G_{\text{sparse}} \subset G_{\text{dense}}$, $|G_{\text{sparse}}| \leq N$. This ensures that the final predicted graph has some minimal coverage.⁴

The core component of the framework is a Conditional Variational Auto-encoder (CVAE) [38] which is used for approximating the desired sample distribution. CVAE is an extension of a traditional variational auto-encoder [12] which is a directed graphical model with low-dimensional Gaussian latent variables. CVAE is a *conditional* graphical model which makes it relevant for our application where conditioning variables are features of the planning problem. We provide a high level description for brevity, and refer the reader to [32] for a comprehensive tutorial.

Here $x \in \mathcal{X}$ is the output random variable, $z \in \mathbb{R}^L$ is the latent random variable and $y \in \mathbb{R}^m$ is the conditioning variable. We wish to learn two *deterministic mappings* - an *encoder* and a *decoder*. An encoder maps (x, y) to a mean and variance of a Gaussian $q_\phi(z|x, y)$ in latent space, such that it is “close” to an isotropic Gaussian $\mathcal{N}(0, I)$. The decoder maps this Gaussian and y to a distribution in the output space $p_\theta(x|z, y)$. This is achieved by maximizing the following objective $\mathcal{L}(x, y; \theta, \phi)$:

$$-D_{KL}(q_\phi(z|x, y) \parallel \mathcal{N}(0, I)) + \frac{1}{L} \sum_{l=1}^L \log p_\theta(x|y, z^{(l)}) \quad (2)$$

Note that the encoder is needed only for training. At test time, only the decoder is used to map samples from an isotropic Gaussian in the latent space to samples in the output space.

We train the CVAE by passing in a dataset $\mathcal{D} = \{X_i, y_i\}_{i=1}^D$. y_i is the feature vector (conditioning variable) extracted from the planning problem Λ_i . X_i is the desired set of nodes extracted from the dense graph G_{dense} that we want our learner to predict. Hence we train the model by maximizing the following objective.

$$\mathcal{R}(\mathcal{D}; \theta, \phi) = \frac{1}{|\mathcal{D}|} \sum_{i=1}^{|\mathcal{D}|} \sum_{j=1}^{|X_i|} \mathcal{L}(x_j, y_i; \theta, \phi) \quad (3)$$

³For cases where a subgraph is preferred, e.g. G_{dense} lies on a constraint manifold, one can design a projection operator $\mathcal{P} : \mathcal{X} \rightarrow V_{\text{dense}}$

⁴Since G_{dense} is a Halton graph, we use the first N Halton sequences.




Fig. 3: (a) The training input generated by SHORTESTPATH. (b) Failure of SHORTESTPATH model to route through gaps (c) Another failure of the model due to unexpected obstacles (d) The training input generated by LEGO. Diverse shortest paths are generated followed by extraction of bottleneck nodes.

A. General Train and Test Procedure

To summarize, the overall training framework is as follows:

- 1) Load a database of planning problems Λ_i and corresponding feature vectors y_i .
- 2) For each Λ_i , extract relevant nodes from the dense graph G_{dense} by invoking $X_i = \text{EXTRACTNODES}(\Lambda_i, G_{\text{dense}})$.
- 3) Feed dataset $\mathcal{D} = \{X_i, y_i\}_{i=1}^D$ as input to CVAE.
- 4) Train CVAE and return learned decoder $p_\theta(x|y, z)$.

At test time, given a planning problem Λ , the graph predicting oracle $\pi(G_{\text{dense}}, y)$ performs the following set of steps:

- 1) Extract feature vector y from planning problem Λ .
- 2) Sample N nodes using decoder $p_\theta(x|y, z)$.
- 3) Connect nodes to create a graph G . Compose sampled graph with a constant sparse graph $G \leftarrow G \oplus G_{\text{sparse}}$.

The focus of this work is on examining variants of the node extraction function $X = \text{EXTRACTNODES}(\Lambda, G_{\text{dense}})$. While the parameters of the CVAE are certainly relevant (discussed in Appendix A), in this paper we ask the question:

What is a good input X to provide to the CVAE?

To that end, we explore the following schemes:

- 1) SHORTESTPATH: Extract nodes X_{sp} belonging to the shortest path. This is the baseline approach (Section V).
- 2) BOTTLENECKNODE: Extract nodes X_{bn} that correspond to bottlenecks along the shortest path (Section VI-A).
- 3) DIVERSEPATHSET: Extract nodes X_{div} belonging to multiple diverse shortest paths (Section VI-B).
- 4) LEGO: Extract nodes X_{lego} that correspond to bottlenecks along multiple diverse shortest paths. This is our proposed approach (Section VI-C).

V. THE SHORTESTPATH (ICHTER ET AL. [11]) PROCEDURE

We examine the scheme applied in [11] of using nodes belonging to the shortest path on the dense graph as input for training the CVAE. The rationality for this scheme is that the distribution of states belonging to the shortest path might lie on a manifold that can be captured by the latent

space of the CVAE. This hypothesis is validated across many high-dimensional planning domains.

We argue that the presented results should not be entirely surprising. The intrinsic difficulty of a planning problem stems from having to search in multiple potential homotopy classes to find a feasible high quality solution. This often manifests in problems involving mazes, bugtraps or narrow passages where the search has to explore and backtrack frequently. Simply increasing the dimension of the problem does not necessarily render it difficult. On the contrary, since the volume of free space increases substantially, there is often an abundance of feasible paths. The challenge, of course, is to find a manifold on which such paths lie with high probability. This is where we found the CVAE to be critical - it learns to *interpolate* between the start and goal along a low dimensional manifold.

However, we are interested in more *difficult* problems where such interpolations would break down. Based on extensive evaluations of this SHORTESTPATH scheme, we were able to identify two concrete vulnerabilities:

- 1) *Failure to route through gaps*: Fig. 3(b) shows the output of the CVAE when there is a gap through which the search has to route to get to the goal. The model gets stuck in a poor local minimum between linearly interpolating start-goal and routing through the gap since the network is not expressive enough to map the feature vector to such a path. This is tantamount to burdening the sampler to solve the planning problem.
- 2) *Presence of unexpected obstacles in test data*: Fig. 3(c) shows the output of the CVAE when there are small, unexpected obstacles in the test data which were not present in the training data. The learned distribution samples over this obstacle as it only predicts what it thinks is the shortest path. Even if we were to have such examples in the training data, unless the feature extractor detects such obstacles, the problem remains.

VI. APPROACH

In this section, we present LEGO (Leveraging Experience with Graph Oracles), an algorithm to train a CVAE to predict sparse yet high quality roadmaps. We do so by tackling head-on the challenges identified in Section IV. Firstly, we recognize that the learner does not have to directly predict the shortest path. Instead, we train it to predict only *bottleneck*

Algorithm 1: BOTTLENECKNODE

Input : Planning problem Λ , Bottleneck tolerance ϵ
Dense path ξ_{dense}^* , Sparse graph G_{sparse}
Output : Bottleneck nodes V_{bn}

- 1 $G_{\text{inf}} \leftarrow G_{\text{sparse}} \oplus \xi_{\text{dense}}^*$; \triangleright Add to sparse graph
- 2 $\eta \leftarrow 1$;
- 3 **while** $c(\text{ALG}(G_{\text{inf}}, \Lambda)) \leq (1 + \epsilon)c(\xi_{\text{dense}}^*)$ **do**
- 4 $\eta \leftarrow \eta + \delta\eta$; \triangleright Increase inflation
- 5 **for** $(u, v) \in E_{\text{inf}} \setminus E_{\text{sparse}}$ **do**
- 6 $c(u, v) \leftarrow \eta c(u, v)$; \triangleright Inflate added edges
- 7 $\xi_{\text{inf}}^* \leftarrow \text{ALG}(G_{\text{inf}}, \Lambda)$; \triangleright Shortest inflated path
- 8 $V_{\text{bn}} \leftarrow \xi_{\text{dense}}^* \cap \xi_{\text{inf}}^*$; \triangleright Bottleneck nodes
- 9 **return** V_{bn} ;

nodes that can assist the underlying search in finding a near-optimal solution. Secondly, the roadmap must be robust to prediction errors of the learner. We safeguard against this by training the learner to predict a *diverse set of paths* with the hope that at-least one of them is feasible.

A. Bottleneck Nodes

We begin by noting that G_{sparse} has a uniform coverage over the entire configuration space. Hence, the learner only has to contribute a critical set of nodes that allow G_{sparse} to represent paths that are near-optimal with respect to the path in G_{dense} . We call these *bottleneck nodes* as they correspond to regions that are difficult for a uniform sampler to cover. We define $X_{\text{bn}} = \text{BOTTLENECKNODE}(\Lambda, G_{\text{dense}})$ as:

Definition 2 (Bottleneck Nodes). *Given a dense graph G_{dense} , find the smallest set of nodes which in conjunction with a sparse subgraph G_{sparse} contains a near-optimal path, i.e.*

$$\begin{aligned} \arg \min_{V \subset V_{\text{dense}}} & |V| \\ \text{s.t.} & \frac{c(\text{ALG}(G_{\text{sparse}} \oplus V, \Lambda))}{c(\text{ALG}(G_{\text{dense}}, \Lambda))} \leq 1 + \epsilon \end{aligned} \quad (4)$$

Here $G \oplus V'$ represents a merge operation, i.e. $V \leftarrow V \cup V'$, $E \leftarrow E \cup \{(u, v) \mid u \in V', (u, v) \in E_{\text{dense}}\}$.

The optimization Section 4 is combinatorially hard. We present an approximate solution in Algorithm 1. We use the optimal path ξ_{dense}^* on the dense graph and create an *inflated graph* $G_{\text{inf}}(\eta)$ by composing $G_{\text{sparse}} \oplus \xi_{\text{dense}}^*$ and inflating weights of newly added edges by η (Line 6). The idea is to disincentivize the search from using any of the newly added edges. This inflation factor is increased till a near-optimal path can no longer be found (Lines 3-6). At this point, the additional vertices that the shortest path on this inflated path pass through are essential to achieve near-optimality. This is formalized by the following guarantee:

Proposition 1 (Bounded bottleneck edge weights). *Let $E_{\text{bn}} \leftarrow \xi_{\text{inf}}^* \setminus E_{\text{sparse}}$ be the chosen bottleneck edges, E_{bn}^* be the optimal bottleneck edges and ξ_{dense}^* be the optimal path on G_{dense} .*

$$\sum_{e_i \in E_{\text{bn}}} c(e_i) \leq \sum_{e_i \in E_{\text{bn}}^*} c(e_i) + \frac{(1 + \epsilon)c(\xi_{\text{dense}}^*)}{\eta} \quad (5)$$

Proof: (Sketch) Let V_{bn}^* be the optimal bottleneck nodes and E_{bn}^* be the optimal bottleneck edges. Let ξ_{bn}^* be the path returned by $\text{ALG}(G_{\text{sparse}} \oplus V_{\text{bn}}^*, \Lambda)$. From Definition 2, the following holds:

$$\begin{aligned} \sum_{e_i \in E_{\text{bn}}^*} c(e_i) + \sum_{e_i \in \xi_{\text{bn}}^* \setminus E_{\text{bn}}^*} c(e_i) &\leq (1 + \epsilon)c(\xi_{\text{dense}}^*) \\ \sum_{e_i \in \xi_{\text{bn}}^* \setminus E_{\text{bn}}^*} c(e_i) &\leq (1 + \epsilon)c(\xi_{\text{dense}}^*) \end{aligned}$$

Since ξ_{inf}^* is the shortest path on the inflated graph, we have:

$$\begin{aligned} \sum_{e_i \in E_{\text{bn}}} c(e_i) + \eta \sum_{e_i \in \xi_{\text{inf}}^* \setminus E_{\text{bn}}} c(e_i) \\ \leq \eta \sum_{e_i \in E_{\text{bn}}^*} c(e_i) + \sum_{e_i \in \xi_{\text{bn}}^* \setminus E_{\text{bn}}^*} c(e_i) \end{aligned}$$

Putting the two inequalities together we have:

$$\begin{aligned} \eta \sum_{e_i \in \xi_{\text{inf}}^* \setminus E_{\text{bn}}} c(e_i) &\leq \eta \sum_{e_i \in E_{\text{bn}}^*} c(e_i) + \sum_{e_i \in \xi_{\text{bn}}^* \setminus E_{\text{bn}}^*} c(e_i) \\ \sum_{e_i \in \xi_{\text{inf}}^* \setminus E_{\text{bn}}} c(e_i) &\leq \sum_{e_i \in E_{\text{bn}}^*} c(e_i) + \frac{\sum_{e_i \in \xi_{\text{bn}}^* \setminus E_{\text{bn}}^*} c(e_i)}{\eta} \\ \sum_{e_i \in \xi_{\text{inf}}^* \setminus E_{\text{bn}}} c(e_i) &\leq \sum_{e_i \in E_{\text{bn}}^*} c(e_i) + \frac{(1 + \epsilon)c(\xi_{\text{dense}}^*)}{\eta} \end{aligned} \quad \blacksquare$$

Fig. 4 illustrates the samples generated by (a) SHORTESTPATH and (b) LEGO trained with samples from BOTTLENECKNODE; and the successful routing through narrow passages using samples from LEGO.

B. Diverse PathSet

In this training scheme, we try to ensure the roadmap is *robust* to errors introduced by the learner. One antidote to this process is *diversity* of samples. Specifically, we want the roadmap to have enough diversity such that if the predicted shortest path is in fact infeasible, there are low cost alternates.

We set this up as a two player game between a planner and an adversary. The role of the adversary is to invalidate as many shortest paths on the dense graph G_{dense} as possible with a fixed budget of edges that it is allowed to invalidate. The role of the planner is to find the shortest feasible path on the invalidated graph and add this to the set of diverse paths Ξ_{div} . The function $X_{\text{div}} = \text{DIVERSEPATHSET}(\Lambda, G_{\text{dense}})$ then returns nodes belonging Ξ_{div} . We formalize this as:

Definition 3 (Diverse PathSet). *We begin with a graph $G^0 = G_{\text{dense}}$. At each round i of the game, the adversary chooses a set of edges to invalidate:*

$$E_i^* = \arg \max_{E_i \subset E, |E_i| \leq \ell} c(\text{ALG}(G^{i-1} \ominus E_i, \Lambda)) \quad (6)$$

Algorithm 2: DIVERSEPATHSET

Input : Planning problem Λ , Pathset size k ,
Dense graph G_{dense} , Sparse graph G_{sparse}

Output : Diverse pathset Ξ_{div}

```

1  $G \leftarrow G_{\text{dense}}$ ;
2  $\Xi_{\text{div}} \leftarrow \emptyset$ ;
3 for  $i = 1, \dots, k$  do
4    $\Xi \leftarrow \text{ALG}^L(G, \Lambda)$ ;  $\triangleright$  L-shortest paths
5    $E_i \leftarrow \emptyset, \Xi_{\text{inv}} \leftarrow \emptyset$ ;
6   while  $|E_i| < \ell$  do
7      $\Xi \leftarrow \{\xi \mid \xi \in \Xi, \xi \cap E_i = \emptyset\}$ ;
8     for  $j = |E_i|, \dots, \ell$  do
9        $e_j \leftarrow \arg \max_{e \in E} \min_{\xi \in \Xi, e \notin \xi} c(\xi)$ ;  $\triangleright$  Greedy
10       $E_i \leftarrow E_i \cup \{e_j\}$ ;  $\triangleright$  Add edge to set
11       $\Xi_j \leftarrow \{\xi \in \Xi \mid e_j \in \xi\}$ ;
12       $\Xi_{\text{inv}} \leftarrow \Xi_{\text{inv}} \cup \Xi_j$ ;  $\triangleright$  Invalidate paths
13       $\Xi \leftarrow \Xi \setminus \Xi_j$ ;
14       $E_{\text{sc}} \leftarrow \text{SETCOVER}(\Xi_{\text{inv}})$ ;  $\triangleright$  Greedy cover
15      if  $|E_{\text{sc}}| \leq |E_i|$  then
16         $E_i \leftarrow E_{\text{sc}}$ ;  $\triangleright$  Use better cover
17       $G \leftarrow G \ominus E_i$ ;  $\triangleright$  Remove edges
18       $\xi_i \leftarrow \text{ALG}(G, \Lambda)$ ;  $\triangleright$  New shortest path
19       $\Xi_{\text{div}} \leftarrow \Xi_{\text{div}} \cup \xi_i$ ;  $\triangleright$  Add to diverse pathset
20 return  $\Xi_{\text{div}}$ ;

```

and the graph is updated $G^i = G^{i-1} \ominus E_i^*$. The planner choose the shortest path $\xi_i = \text{ALG}(G^i, \Lambda)$ which is added to the set of diverse paths Ξ_{div} .

The optimization problem (6) is similar to a set cover problem (NP-Hard [39]) where the goal is to select edges to cover as many paths as possible. If we knew the exact set of paths to cover, it is well known that a greedy algorithm will choose a near-optimal set of edges [39]. We have the inverse problem - we do not know how many consecutive shortest paths can be covered with a budget of ℓ edges.

Algorithm 2 describes the procedure. We greedily choose a set of edges to invalidate as many consecutive shortest paths till we exhaust our budget (Lines 8-13). We then apply greedy set cover (Line 14). If it leads to a better solution, we continue repeating the process. At termination, we ensure:

Proposition 2 (Near-optimal Invalidated EdgeSet). *Let Ξ_{inv} be the contiguous set of shortest paths invalidated by Algorithm 2 using a budget of ℓ . Let ℓ^* be the size of the optimal set of edges that could have invalidated Ξ_{inv} .*

$$\ell \leq (1 + \log |\Xi_{\text{inv}}|) \ell^* \quad (7)$$

Proof: (Sketch) We briefly explain the equivalence to a set cover problem. Each path in Ξ_{inv} corresponds to an element that has to be covered. Each edge $e \in E_{\text{dense}}$ corresponds to a set of paths in Ξ_{inv} , where each path in the set contains the edge. Invalidating the edge invalidates all paths in the set.

Line 14 invokes a greedy set cover algorithm which at every iteration chooses the edge which covers the largest number of uncovered paths. Let ℓ_{greedy} be the number of edges selected by the greedy algorithm, and ℓ^* be the optimal. From [39],

Algorithm 3: LEGO

Input : Planning problem Λ , Bottleneck tolerance ϵ ,
Pathset
size k , Dense graph G_{dense} , Sparse graph G_{sparse}

Output : LEGO nodes V_{lego}

```

1  $\Xi_{\text{div}} \leftarrow \text{DIVERSEPATHSET}(\Lambda, k, G_{\text{dense}}, G_{\text{sparse}})$ ;
2  $V_{\text{lego}} \leftarrow \emptyset$ ;
3 while  $\Xi_{\text{div}} \neq \emptyset$  do
4    $\xi^* \leftarrow \arg \min_{\xi \in \Xi_{\text{div}}} c(\xi)$ ;  $\triangleright$  Pick diverse path
5    $\Xi_{\text{sparse}} \leftarrow \text{ALG}^L(G_{\text{sparse}}, \Lambda)$ ;
6    $\Xi_{\text{sparse}} \leftarrow \{\xi \mid \xi \in \Xi_{\text{sparse}}, c(\xi) \leq (1 + \epsilon)c(\xi^*)\}$ ;
7    $E_{\text{inv}} \leftarrow \text{SETCOVER}(\Xi_{\text{sparse}})$ ;  $\triangleright$  Edges to remove
8    $G_{\text{sparse}} \leftarrow G_{\text{sparse}} \ominus E_{\text{inv}}$ ;  $\triangleright$  Invalidate paths
9    $V_{\text{bn}} \leftarrow \text{BOTTLENECKNODE}(\Lambda, \epsilon, \xi^*, G_{\text{sparse}})$ ;
10   $V_{\text{lego}} \leftarrow V_{\text{lego}} \cup V_{\text{bn}}$ ;  $\triangleright$  Add to LEGO nodes
11 return  $V_{\text{lego}}$ ;

```




Fig. 4: Comparison of samples (distribution illustrated as heatmap) generated by SHORTESTPATH (a, c) and DIVERSEPATHSET (b, d) in different environments. In the first environment (left), LEGO (b) is trained with samples from BOTTLENECKNODE. In the second environment (right), LEGO (d) is trained with samples from DIVERSEPATHSET. In both instances, SHORTESTPATH fails to find a solution.

we have the following near-optimality guarantee:

$$\ell_{\text{greedy}} \leq (1 + \log |\Xi_{\text{inv}}|) \ell^*$$

If $\ell_{\text{greedy}} \leq \ell$, i.e. we have budget remaining, we continue adding edges that can only invalidate more paths in Lines 8-13. This continues till the budget is exhausted. \blacksquare

Fig. 4 illustrates the samples generated by (a) SHORTESTPATH and (b) LEGO trained with samples from DIVERSEPATHSET; and the robustness to unexpected obstacles exhibited by LEGO.

C. Combining Diversity with Bottleneck Extraction

We present LEGO in Algorithm 3 which combines the characteristics of BOTTLENECKNODE and DIVERSEPATHSET to extract a set of diverse bottleneck nodes. We first find a set of diverse paths on the dense graph (Line 1). We then iterate over each path, and adversarially invalidate edges of the sparse graph to ensure it does contain a feasible shorter path (Line 4-8). The bottleneck nodes for this path are extracted and added to the set of nodes to be returned (Line 9).

VII. EXPERIMENTAL RESULTS

In this section we evaluate the performance of LEGO on various problem domains and compare it against other samplers. We consider samplers that do not assume offline computation or learning such as Medial-Axis PRM (MAPRM) [6, 14], Randomized Bridge Sampler (RBB) [15], Workspace Importance Sampler (WIS) [16], a Gaussian sampler, GAUSSIAN [20], and a uniform Halton sequence sampler, HALTON [4]. Additionally, we also compare our framework against the state-of-the-art learned sampler SHORTESTPATH [35] upon which our work is based.

a) Evaluation Procedure: For a given sampler and a planning problem, we invoke the sampler to generate a fixed number of samples. We then evaluate the performance of the samplers on three metrics: a) sampling time b) success rate in solving shortest path problem and c) the quality of the solution obtained, on the graph constructed with the generated samples.

b) Problem Domains: To evaluate the samplers, we consider a spectrum of problem domains. The \mathbb{R}^2 problems have random rectilinear walls with random narrow passages (Fig. 6(a)). These passages can be small, medium or large in width. The n -link arms are a set of n line-segments fixed to a base moving in an uniform obstacle field (Fig. 6(b)). The n -link snakes are arms with a free base moving through random rectilinear walls with passages (Fig. 6(c)). Finally, the manipulator problem has a 7DoF robot arm [40] manipulating a stick in an environment with varying clutter (Fig. 6(d)). Two variants are considered - constrained (\mathbb{R}^7), when the stick is welded to the hand, and unconstrained, when the stick can slide along the hand (\mathbb{R}^8).

c) Experiment Details: For the learned samplers SHORTESTPATH and LEGO, we use 4000 training worlds and 100 test worlds. Dense graph is an r -disc Halton graph [5]: 2000 vertices in \mathbb{R}^2 to 30,000 vertices in \mathbb{R}^8 . The CVAE was implemented in TensorFlow [41] with 2 dense layers of 512 units each. Input to the CVAE is a vector encoding source and target locations and an occupancy grid. Training time over 4000 examples ranged from 20 minutes in \mathbb{R}^2 to 60 minutes in \mathbb{R}^8 problems. At test time, we time-out samplers after 5 sec. The code is open sourced⁵ with more details in [42].

A. Performance Analysis

a) Sampling time: Table I reports the average time each sampler takes for 200 samples across 100 test instances. SHORTESTPATH and LEGO are the fastest. MAPRM and RBB both rely on heavy computation with multiple collision checking steps. WIS, by tetrahedralizing the workspace and identifying narrow passages, is relatively faster but slower than the learners. Unfortunately, some of the baselines time-out on manipulator planning problem due to expense of collision checking.

⁵<https://github.com/vernalrahul/LearningRoadmaps>

b) Success Rate: Table II reports the success rates (95% confidence intervals) on 100 test instances when sampling 500 vertices. Success rate is the fraction of problems for which the search found a feasible solution. LEGO has the highest success rate. The baselines are competitive in \mathbb{R}^2 , but suffer for higher dimensional problems.

c) Normalized Path Cost: This is the ratio of cost of the computed solution w.r.t. the cost of the solution on the dense graph. Fig. 6 shows the normalized cost for HALTON, SHORTESTPATH and LEGO- these were the only baselines that consistently had bounded 95% confidence intervals (i.e. when success rate is $\geq 60\%$). SHORTESTPATH has the lowest cost, however LEGO is within 10% bound of the optimal.

B. Observations

We report on some key observations from Table II and Fig. 6.

○ 1. LEGO consistently outperforms all baselines

As shown in Table II, LEGO has the best success rate (for 500 samples) on all datasets. The second row in Fig. 6 shows that LEGO is within 10% bound of the optimal path.

○ 2. LEGO places samples only in regions where the optimal path may pass.

Fig. 5 shows samples generated by various baseline algorithms on a 2D problem. The heuristic baselines use various strategies to identify important regions - MAPRM finds medial axes, RBB finds bridge points, GAUSSIAN samples around obstacles, WIS divides up space non-uniformly and samples accordingly. However, these methods places samples everywhere irrespective of the query. SHORTESTPATH takes the query into account but fails to find the gaps. LEGO does a combination of both - it finds the right gaps.

○ 3. LEGO has a higher performance gain on harder problems (narrow passages) as it focuses on bottlenecks.

Table II shows how success rates vary in 2D problems with small / medium / large gaps. As the gaps gets narrower, LEGO outperforms more dominantly. The BOTTLENECKNODE component in LEGO seeks the bottleneck regions (Fig. 4b).

For manipulator planning \mathbb{R}^8 problems, when stick is unconstrained, LEGO and SHORTESTPATH are almost identical. We attribute this to such problems being easier, i.e. the shortest path simply slides the stick out of the way and plans to the goal. When the stick is constrained, LEGO does far better. Fig. 6(d) shows that LEGO is able to sample around the table while SHORTESTPATH cannot find this path.

○ 4. LEGO is robust to a certain degree of train-test mismatch as it encourages diversity.

Fig. 7 shows the success rate of learners on a 2D test environment that has been corrupted. Environment 1 is less corrupted than environment 2. Fig. 7(a) shows that on environment 1, LEGO is still the best sampler. SHORTESTPATH (Fig. 7(c)) ignores the corruption in the environment and fails. LEGO (Fig. 7(d)) still finds the correct bottleneck.

TABLE I: Average time (sec.) by sampling algorithms to generate 200 samples over 100 planning problems

	Non-Learned Samplers				Learned Samplers		
	HALTON	MAPRM	RBB	GAUSSIAN	WIS	SHORTESTPATH	LEGO
Point Robot (2D)	0.0036	0.53	0.22	0.02	0.25	0.006	0.006
N-link Arm (3D)	0.0058	—	23.96	1.95	0.36	0.016	0.016
N-link Arm (7D)	0.0071	—	37.24	3.77	1.12	0.017	0.017
Snake Robot (5D)	0.0069	39.56	142.21	3.43	0.54	0.013	0.013
Snake Robot (9D)	0.0074	40.01	180.43	8.71	2.11	0.017	0.017
Manipulator (8D)	0.01	—	—	3.33	—	0.018	0.018

TABLE II: Success Rates of different algorithms on 100 trials over different datasets (reported with a 95% C.I.)

	Non-Learned Samplers				Learned Samplers		
	HALTON	MAPRM	RBB	GAUSSIAN	WIS	SHORTESTPATH	LEGO
2D Point Robot Planning							
2D Large (easy)	0.73 ± 0.08	0.73 ± 0.08	0.74 ± 0.09	0.65 ± 0.09	0.78 ± 0.08	0.86 ± 0.07	0.97 ± 0.03
2D Medium	0.48 ± 0.08	0.63 ± 0.09	0.61 ± 0.09	0.48 ± 0.10	0.63 ± 0.09	0.69 ± 0.09	0.89 ± 0.06
2D Small (hard)	0.36 ± 0.09	0.53 ± 0.09	0.48 ± 0.08	0.32 ± 0.09	0.52 ± 0.09	0.59 ± 0.09	0.83 ± 0.07
N-Link Arm							
3D	0.39 ± 0.09	—	0.54 ± 0.09	0.46 ± 0.10	0.52 ± 0.10	0.61 ± 0.09	0.74 ± 0.08
7D	0.29 ± 0.09	—	0.46 ± 0.09	0.41 ± 0.09	0.46 ± 0.09	0.57 ± 0.10	0.71 ± 0.08
N-Link Snake Robot							
5D	0.41 ± 0.09	0.42 ± 0.09	0.48 ± 0.10	0.41 ± 0.09	0.50 ± 0.10	0.77 ± 0.08	0.84 ± 0.07
9D	0.49 ± 0.09	0.45 ± 0.09	0.52 ± 0.10	0.51 ± 0.10	0.53 ± 0.09	0.82 ± 0.07	0.86 ± 0.07
Manipulator Arm Planning							
Unconstrained (8D)	0.24 ± 0.09	—	—	—	—	0.81 ± 0.08	0.82 ± 0.07
Constrained (7D)	0.09 ± 0.05	—	—	—	—	0.58 ± 0.09	0.70 ± 0.09




Fig. 5: Comparison of samples (green) generated by all baseline algorithms on a 2D problem, planning from start (blue) to goal (red).

Fig. 7(b) shows that all learners are worse than HALTON. SHORTESTPATH (Fig. 7(e)) densifies around a particular constrained region while LEGO (Fig. 7(f)) still finds a path due to the DIVERSEPATHSET component sampling in multiple bottleneck regions.

VIII. DISCUSSION

We present a framework for training a generative model to predict roadmaps for sampling-based motion planning. We build upon state-of-the-art methods that train the CVAE using the shortest path as target input. We identify important failure modes such as complex obstacle configurations and train-test mismatch. Our algorithm LEGO directly addresses these issues by training the CVAE using *diverse bottleneck nodes* as target input. We formally define these terms and provide provable algorithms to extract such nodes. Our results indicate that the predicted roadmaps outperform competitive baselines on a range of problems.

Using priors in planning is a double edged sword. While one can get astounding speed ups by focusing search on a tiny portion of C-space [11], any problem not covered

in the dataset can lead to catastrophic failures. This is symptomatic of the fundamental problem of *over-fitting* in machine learning. While one could ensure the training data covers all possible environments [43], an algorithmic solution is to explore regularization techniques for planning. We argue DIVERSEPATHSET can be viewed as a form of regularization.

We can also include a more informed conditioning vector that captures the state of the search, e.g., the length of the current shortest path. This is similar to Informed RRT* [44]. Finally, we wish to scale to problems with varying workspace where a global planner guides the sampler to focus on relevant parts of the workspace [13, 45].

REFERENCES

- [1] L.E. Kavraki, P. Svestka, J.C. Latombe, and M.H. Overmars. Probabilistic roadmaps for path planning in high-dimensional configuration spaces. *IEEE TRO*, 1996.
- [2] S. M. LaValle. *Planning Algorithms*. Cambridge University Press, Cambridge, U.K., 2006. Available at <http://planning.cs.uiuc.edu/>.
- [3] Peter E Hart, Nils J Nilsson, and Bertram Raphael. A formal basis for the heuristic determination of minimum cost paths. *IEEE transactions on Systems Science and Cybernetics*, 1968.
- [4] J. H. Halton. On the efficiency of certain quasi-random sequences




Fig. 6: Comparisons of SHORTESTPATH against LEGO on different problem domains. In each problem domain (column), success rate (top), normalized path lengths (middle) and solutions determined on the roadmaps constructed using samples generated by the two samplers (bottom) are shown.

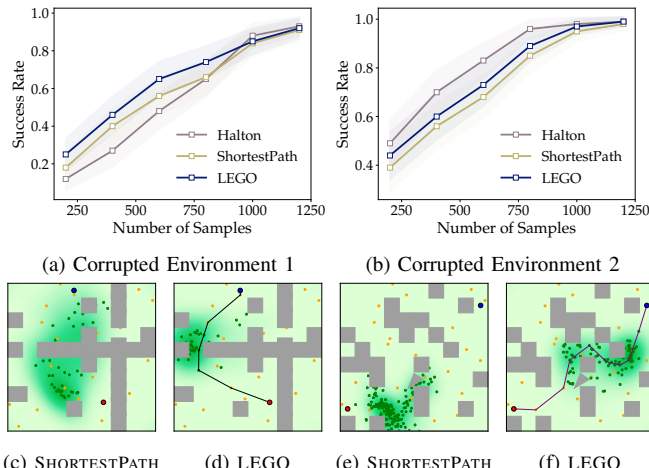


Fig. 7: Comparison of samplers on corrupted environments, i.e., different from training dataset. Success rate on (a) less corrupted environment 1: mixture of walls and random squares and (b) more corrupted environment 2: only squares. Output of SHORTESTPATH and LEGO on environment 1 (c,d) and environment 2 (e,f).

of points in evaluating multi-dimensional integrals. *Numerische Mathematik*, 2:84–90, 1960.

- [5] Lucas Janson, Brian Ichter, and Marco Pavone. Deterministic sampling-based motion planning: Optimality, complexity, and performance. *arXiv preprint arXiv:1505.00023*, 2015.
- [6] Christopher Holleman and Lydia E Kavraki. A framework for using the workspace medial axis in PRM planners. In *ICRA*, 2000.
- [7] D Hsu, G Sánchez-Ante, and Z Sun. Hybrid prm sampling with a cost-sensitive adaptive strategy. In *ICRA*, 2005.
- [8] Brendan Burns and Oliver Brock. Sampling-based motion planning using predictive models. In *ICRA*, 2005.
- [9] D Hsu, J Latombe, and H Kurniawati. On the probabilistic foundations of probabilistic roadmap planning. *IJRR*, 2006.
- [10] H Kurniawati and D Hsu. Workspace-based connectivity oracle: An

adaptive sampling strategy for prm planning. In *Algorithmic Foundation of Robotics VII*, pages 35–51. Springer, 2008.

- [11] Brian Ichter, James Harrison, and Marco Pavone. Learning sampling distributions for robot motion planning. In *ICRA*, 2018.
- [12] Diederik P Kingma and Max Welling. Auto-encoding variational bayes. *arXiv preprint arXiv:1312.6114*, 2013.
- [13] D. Hsu, J.-C. Latombe, and R. Motwani. Path planning in expansive configuration spaces. *International Journal Computational Geometry & Applications*, 4:495–512, 1999.
- [14] Steven A Wilmarth, Nancy M Amato, and Peter F Stiller. MAPRM: A probabilistic roadmap planner with sampling on the medial axis of the free space. In *ICRA*, 1999.
- [15] D Hsu, T Jiang, J Reif, and Z Sun. The bridge test for sampling narrow passages with probabilistic roadmap planners. In *ICRA*, 2003.
- [16] H Kurniawati and D Hsu. Workspace importance sampling for probabilistic roadmap planning. In *IROS*.
- [17] Yuandong Yang and Oliver Brock. Adapting the sampling distribution in prm planners based on an approximated medial axis. In *ICRA*, 2004.
- [18] Jur P Van den Berg and Mark H Overmars. Using workspace information as a guide to non-uniform sampling in probabilistic roadmap planners. *IJRR*, 2005.
- [19] Nancy M Amato, O Burchan Bayazit, and Lucia K Dale. OBPRM: An obstacle-based PRM for 3D workspaces. 1998.
- [20] Valérie Boor, Mark H Overmars, and A Frank Van Der Stappen. The gaussian sampling strategy for probabilistic roadmap planners. In *ICRA*, pages 1018–1023, 1999.
- [21] M Morales, L Tapia, R Pearce, S Rodriguez, and N M Amato. A machine learning approach for feature-sensitive motion planning. In *Algorithmic Foundations of Robotics VI*, pages 361–376. Springer, 2005.
- [22] S Rodriguez, S Thomas, R Pearce, and N M Amato. Resampl: A region-sensitive adaptive motion planner. In *Algorithmic Foundation of Robotics VII*, pages 285–300. Springer, 2008.
- [23] Sébastien Dalibard and Jean-Paul Laumond. Linear dimensionality reduction in random motion planning. *IJRR*, 2011.
- [24] Jia Pan, Sachin Chitta, and Dinesh Manocha. Faster sample-based motion planning using instance-based learning. In *WAFR*, 2012.
- [25] Shushman Choudhury, Christopher M Dellin, and Siddhartha S Srinivasa. Pareto-optimal search over configuration space beliefs for anytime motion planning. In *IROS*, 2016.
- [26] Xinyu Tang, Jyh-Ming Lien, and Nancy Amato. An obstacle-based rapidly-exploring random tree. In *ICRA*, 2006.

- [27] Brendan Burns and Oliver Brock. Toward optimal configuration space sampling. In *RSS*, 2005.
- [28] Rosen Diankov and James Kuffner. Randomized statistical path planning. In *IROS*, 2007.
- [29] Matt Zucker, James Kuffner, and J Andrew Bagnell. Adaptive workspace biasing for sampling-based planners. In *ICRA*, 2008.
- [30] Yen-Ling Kuo, Andrei Barbu, and Boris Katz. Deep sequential models for sampling-based planning. *arXiv preprint arXiv:1810.00804*, 2018.
- [31] Paul Vernaza and Daniel D Lee. Learning dimensional descent for optimal motion planning in high-dimensional spaces. In *AAAI*, 2011.
- [32] Carl Doersch. Tutorial on variational autoencoders. *arXiv preprint arXiv:1606.05908*, 2016.
- [33] N Chen, M Karl, and P van der Smagt. Dynamic movement primitives in latent space of time-dependent variational autoencoders. In *Humanoids*, 2016.
- [34] Jung-Su Ha, Hyeok-Joo Chae, and Han-Lim Choi. Approximate inference-based motion planning by learning and exploiting low-dimensional latent variable models. *IEEE RAL*, 2018.
- [35] Brian Ichter and Marco Pavone. Robot motion planning in learned latent spaces. *arXiv preprint arXiv:1807.10366*, 2018.
- [36] C Zhang, J Huh, and D D Lee. Learning implicit sampling distributions for motion planning. *arXiv preprint arXiv:1806.01968*, 2018.
- [37] A H Qureshi and M C Yip. Deeply informed neural sampling for robot motion planning. In *2018 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, pages 6582–6588. IEEE, 2018.
- [38] K Sohn, H Lee, and X Yan. Learning structured output representation using deep conditional generative models. In *NIPS*, 2015.
- [39] Uriel Feige. A threshold of $\ln n$ for approximating set cover. *Journal of the ACM (JACM)*, 1998.
- [40] S. Srinivasa, D. Berenson, M. Cakmak, A. C. Romea, M. Dogar, A. Dragan, R. Knepper, T. Niemueller, K. Strabala, J. M. Vandeweghe, and J. Ziegler. Herb 2.0: Lessons learned from developing a mobile manipulator for the home. *Proceedings of the IEEE*, 100(8):1–19, July 2012.
- [41] Martín Abadi et al. TensorFlow: Large-scale machine learning on heterogeneous systems. URL <https://www.tensorflow.org/>.
- [42] R Vernwal, A Mandalika, S Choudhury, and S Srinivasa. Supplementary Material. 2018. URL http://adityavk.com/LEGO_Supplementary.pdf.
- [43] J Tobin, R Fong, A Ray, J Schneider, W Zaremba, and P Abbeel. Domain randomization for transferring deep neural networks from simulation to the real world. In *IROS*, 2017.
- [44] J D Gammell, S S Srinivasa, and T D Barfoot. Informed RRT*: Optimal incremental path planning focused through an admissible ellipsoidal heuristic. In *IROS*, 2014.
- [45] Stephan Zheng, Yisong Yue, and Jennifer Hobbs. Generating long-term trajectories using deep hierarchical networks. In *NIPS*, 2016.
- [46] Serena Yeung, Anitha Kannan, Yann Dauphin, and Li Fei-Fei. Tackling over-pruning in variational autoencoders. *arXiv preprint arXiv:1706.03643*, 2017.

APPENDIX A CVAE FRAMEWORK

We refer the reader to [32] for technical details and a comprehensive tutorial on CVAE. In Section A-A we describe the CVAE architecture implemented to train LEGO and SHORTESTPATH algorithms. In Sections A-B and A-C, we study two parameters that determine the performance of the CVAE generative model.

A. Architecture

The entire CVAE module (Fig. 8) takes as input the training samples X , which in case of LEGO are the samples in bottleneck regions and along diverse paths. Additionally, the CVAE also takes as input a vector of external features y , upon which the generative model is also conditioned upon. In the problems we consider, these features include information regarding the environment such as the poses of the obstacles and the start-goal pair. A standard CVAE model consists of an encoder and a decoder, often represented by neural networks. The input samples and the external features are used to train the encoder and the decoder networks end-to-end.

During training, The encoder network takes as input the high-dimensional vector of features including the training sample and the other external features and encodes it into a low-dimensional latent variable vector. The latent variable is then fed into the decoder network along with the vector of external features as an input which outputs a sample in the configuration space. This sample output by the decoder is used to minimize an objective function which aims to fundamentally reduce the divergence between the probability distribution of the training samples and the learned generative model to be able to closely reconstruct the training samples set. During testing, only the decoder network is used to generate the required samples. The decoder takes as input a latent variable sampled from standard normal distribution as well as the vector of external features to generate useful samples.

In our implementation of the CVAE, both encoder and decoder networks have two fully connected hidden layers with 512 unites each. The specifics of the external features used in each of the planning problems considered in Section VII are discussed in Section B-A. The behavior of the generative model, in addition to the features used, also depends on certain parameters. We study the effect of these parameters and their design choices in our implementation in the following subsections.

B. Dimensionality of the Latent Variable

The latent variable captures the information available to the model through the training examples in a lower dimensional latent space. The dimensionality of the latent variable denotes how efficiently the model can capture the sources of variability required to regenerate data similar to the training examples. Theoretically, a model with larger latent dimension is at

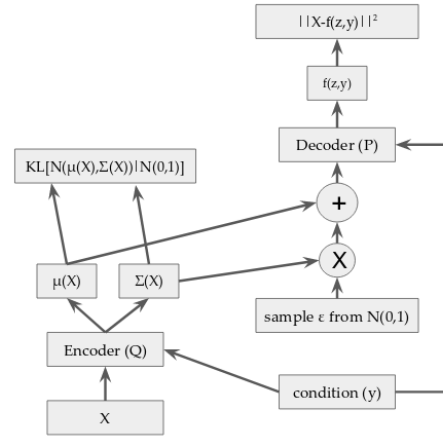


Fig. 8: A simple illustration of the CVAE framework setup for training with X and y together denoting the input to the CVAE.

least as good as a model with lower latent dimension. However, in practice, when the latent variable dimension is high, it becomes computationally expensive for methods like stochastic gradient descent to reduce the KL divergence between the true and the approximated distributions over the latent variables conditioned on the training examples. Fig. 9 shows the behaviors exhibited by the trained generative model for different latent variable dimensions. We choose latent variable dimension of 3 for \mathbb{R}^2 , \mathbb{R}^5 problems and 5 for \mathbb{R}^7 , \mathbb{R}^8 and \mathbb{R}^9 problems.




Fig. 9: Samples generated by CVAE trained with different latent variable dimensions (a) 1 (b) 3 and (c) 7.

C. Regularization Parameter

Although VAEs are generally devoid of regularization parameters, one could introduce the parameter in modifying the objective function the CVAE aims to minimize when learning the generative model. The objective function in a CVAE is given by:

$$\text{Reconstruction Loss} + \lambda \times \text{KL Divergence} \quad (8)$$

The reconstruction loss ensures that the training data can be explained with the data generated by the model and therefore minimizing it ensures proper reconstruction of the training examples. On the other hand, the second term captures the divergence between the prior distribution over latent variable and the posterior given the training examples. Minimizing it ensures that the two distributions are similar. When the value of λ is zero, the behavior of the corresponding VAE is similar to a traditional autoencoder in its capability to reconstruct the training examples. When the value of λ is equal to 1, the

objective function is as in a VAE. However this often leads to *over-pruning* [46] where many of the dimensions of the latent variable are ignored in an attempt to reduce the KL divergence. By tuning the value of λ between 0 and 1, one could weigh the two objectives appropriately to obtain the desired generative model behavior (Fig. 10).




Fig. 10: Example learned distributions for the narrow passage problem for different values of regularization parameter (λ), (a) 2×10^{-8} (b) 2×10^{-2} (c) 2×10^{-4} (chosen value of λ).

APPENDIX B EXPERIMENTS

In this section, we discuss the offline computation involved in training the CVAE for different planning environments considered in Section VII.

A. Training Procedure

a) 2D Point Robot Planning: The training data consisted of 20 randomly generated environments as shown in Fig. 11 with 20 planning problems (start-goal pairs) in each of the environments. The environments were randomized in positions of the vertical and horizontal walls and the narrow passages through them. The CVAE was conditioned upon a vector of 102 features which included the start-goal pair (4 features) as well the 10×10 occupancy grid (100 features). The dataset generation took 4-5 hours while the training time was around 25 minutes. The CVAE was trained using samples from G_{dense} with 3000 samples. The CVAE was trained to sample configurations (in \mathbb{R}^2) of the point robot.




Fig. 11: Environments sampled in \mathbb{R}^2 to train the CVAE.

b) N-Link Arm Planning: The training procedure for the robot in \mathbb{R}^3 , \mathbb{R}^5 consisted of a G_{dense} with 6000 samples which was used to plan for 20 planning problems in each of 20 randomly generated 2D environments. Fig. 12 visualizes some of the environments sampled to train the CVAE. The red and blue positions show the start and goal states respectively. The environment has randomly placed obstacles. The CVAE was

conditioned on a vector of features which included the start-goal pair as well the 10×10 occupancy grid (100 features). The dataset generation took 6-7 hours while the training time was close to 30 minutes.




Fig. 12: Environments sampled in \mathbb{R}^3 to train the CVAE.

c) Snake Robot Planning: For \mathbb{R}^5 , the training procedure was similar to that in the \mathbb{R}^2 problems. The training procedure for the robot in \mathbb{R}^9 consisted of a G_{dense} with 6000 samples which was used to plan for 20 planning problems in each of 20 randomly generated 2D environments. Fig. 13 visualizes some of the environments sampled to train the CVAE. The red and blue positions show the start and goal states respectively. The environments were modified in the wall being horizontal or vertical, the offset in its position, and the position of the narrow passage through it. The CVAE was conditioned on a vector of 118 features which included the start-goal pair (18 features) as well the 10×10 occupancy grid (100 features). The dataset generation took 6-7 hours while the training time was close to 30 minutes. The CVAE was trained to sample configurations of the snake robot that included the base location as well as the revolute joint angles between each of the links.




Fig. 13: Environments sampled in \mathbb{R}^9 to train the CVAE.

d) Manipulator Arm Planning: The training data consisted of 20 random environments where the obstacles in the environment were arbitrarily repositioned. In each of the randomly generated environment, 50 planning problems were considered as an input to the train the CVAE model. Fig. 14 visualized three such environments, where the positions of the table and that of the obstacle on the table are modified along with start and goal configurations. The CVAE in the constrained problem was conditioned on a vector of 46⁶ features which included the start and goal configurations (14 features) and the poses of the table and the obstacle represented as 4×4 homogeneous matrices (32 features). The dataset was generated in 7-8 hours while the training

⁶48 in the unconstrained problem since the configuration of the robot includes an additional degree of freedom.

took around an hour. Samples from a G_{dense} with 30,000 configurations were used to train the CVAE. The CVAE learned to sample the robot configurations which included the joint angles at the seven revolute joints of the arm in the constrained example. The unconstrained \mathbb{R}^8 example consisted of an additional prismatic joint value denoting where the stick is held in the hand.




Fig. 14: Manipulator arm environments sampled to train the CVAE. The solutions obtained using samples generated by LEGO are also visualized.

B. Additional Experiment Results

a) **BOTTLENECKNODE and DIVERSEPATHSET**: In addition to the qualitative observations presented in Section VII (O1 and O2) and Fig. 4, we present here the analysis of the performance of the foundational algorithms of LEGO, namely **BOTTLENECKNODE** and **DIVERSEPATHSET** when compared to **SHORTESTPATH**. Fig. 15a shows that on a \mathbb{R}^2 world, **BOTTLENECKNODE** has a significantly higher success rate than **SHORTESTPATH**, almost converging to 1.0 by 400 samples. Fig. 15b shows that in terms of path length, **SHORTESTPATH** is initially better but both are eventually comparable. This is expected because of the near-optimality objective of **BOTTLENECKNODE** (4). Fig. 15c shows that **DIVERSEPATHSET** has a better success rate. Fig. 15d shows that while both algorithms are comparable in terms of path length, **DIVERSEPATHSET** has a smaller variance.

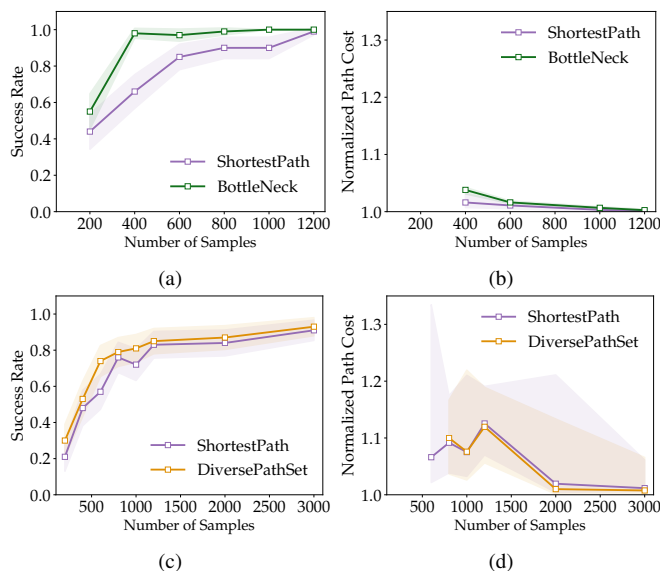


Fig. 15: Comparison of **SHORTESTPATH** against **BOTTLENECKNODE** (top row) and **DIVERSEPATHSET** (bottom row) on success rate (left column) and normalized path length (right column).

C. Roadmap Construction




Fig. 16: Performance of LEGO on different roadmaps. The parameter p denotes the ratio of Halton samples to learned samples in the roadmap.

To evaluate the performance of LEGO, we construct sparse roadmaps, G_{sparse} . The sparse graph consisted of 200 samples in \mathbb{R}^2 , \mathbb{R}^5 problems and 300 samples in case of \mathbb{R}^7 , \mathbb{R}^8 and \mathbb{R}^9 problems. Not however, that this sparse roadmap contains both the learned samples as well as samples generated from Halton sequence. While the learned samples are concentrated near the bottleneck regions and along diverse paths, Halton samples ensure the coverage over the free regions of the configuration space as well. We analyze different proportions of Halton samples and learned samples. Fig. 16 shows the performance characteristics of LEGO on roadmaps constructed with different proportions of Halton and learned samples for the 2D point robot example. We observe that LEGO over a roadmap of 200 samples with just 30% learned samples significantly outperforms LEGO over a Halton graph ($p = 1$). Fig. 17 visualizes the samples generated by LEGO represented by the end-effector positions (blue) in the workspace.




Fig. 17: Samples generated by LEGO for manipulator arm planning. The blue dots represent the end effector positions corresponding to the samples.