

Zero-Shot Constrained Motion Planning Transformers Using Learned Sampling Dictionaries

Jacob J. Johnson[†], Ahmed H. Qureshi[‡], and Michael C. Yip[†]

Abstract—Constrained robot motion planning is a ubiquitous need for robots interacting with everyday environments, but it is a notoriously difficult problem to solve. Many sampled points in a sample-based planner need to be rejected as they fall outside the constraint manifold, or require significant iterative effort to correct. Given this, few solutions exist that present a constraint-satisfying trajectory for robots, in reasonable time and of low path cost. In this work, we present a transformer-based model for motion planning with task space constraints for manipulation systems. Vector Quantized-Motion Planning Transformer (VQ-MPT) is a recent learning-based model that reduces the search space for unconstrained planning for sampling-based motion planners. We propose to adapt a pre-trained VQ-MPT model to reduce the search space for constraint planning without retraining or finetuning the model. We also propose to update the neural network output to move sampling regions closer to the constraint manifold. Our experiments show how VQ-MPT improves planning times and accuracy compared to traditional planners in simulated and real-world environments. Unlike previous learning methods, which require task-related data, our method uses pre-trained neural network models and requires no additional data for training and finetuning the model making this a *one-shot* process. We also tested our method on a physical Franka Panda robot with real-world sensor data, demonstrating the generalizability of our algorithm. We anticipate this approach to be an accessible and broadly useful for transferring learned neural planners to various robotic-environment interaction scenarios.

I. INTRODUCTION

Constraint motion planning (CMP) is a fundamental challenge in robotics that involves finding a collision-free path between a start and goal configuration while satisfying certain constraints. Constraints may include kinematic constraints on the robot’s joints [1], [2], dynamic constraints like torque limits [3], and task constraints like maintaining end-effector orientation [4]. Trajectories that adhere to these constraints are relevant in fields such as home robotics - to move containers without spilling their content, medical robotics - constrain end-effector torques to interact with humans safely [3], and industrial robotics - articulate objects such as levers and pulleys [5].

The most common algorithms for finding such trajectories are sampling-based motion planners (SMP) [7]. These algorithms build a discrete search graph of collision-free robot states by random sampling in the planning space to connect the start and goal states. They have been particularly successful at finding solutions for robotic systems with high

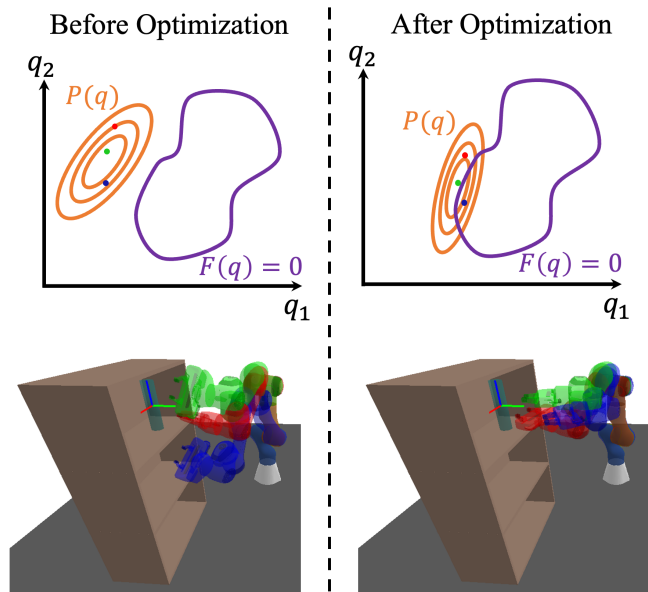


Fig. 1. Concept figure showing joint configurations (bottom) sampled from distributions, $P(q)$, generated by CVQ-MPT and previous transformer-based approach [6]. CVQ-MPT uses the predicted distribution from a pre-trained VQ-MPT model (top left) and uses gradient-based optimization to update the distribution (top right) towards the constraint manifold, $F(q) = 0$. Configurations sampled from this distribution are closer to the given constraint, improving planning times and sampling efficiency.

degrees of freedom, such as manipulators [8]. Previous works have extended SMPs to plan trajectories that satisfy given constraints [9]. However, SMPs can be highly inefficient, especially for higher dimensional spaces. With the added complexity of satisfying constraints, the valid search region shrinks, exacerbating inefficiency.

Recent works have improved the sampling efficiency of SMPs by leveraging prior data [10], [11], [12], [13] for unconstrained motion planning problems. By prudently selecting sampled points, these techniques improve planning times considerably. For CMP, the sampled points not only have to be collision-free but also need to satisfy a constraint function. Oftentimes, fully identifying these regions itself is computationally intractable. Constraint Motion Planning Networks (CoMPNet), an extension of Motion Planning Networks (MPNet), extended the idea of neural samplers for multimodal kinematic constraints [14] to generate sampled points on this constraint manifold, i.e., regions where the constraints are satisfied. However, these methods require the collection of task-specific data to train the models.

Our previous work, Vector Quantized-Motion Planning Transformers (VQ-MPT), has shown promising results in

[†]J.J. Johnson and M.C.Yip are with the Electrical and Computer Engineering Department at University of California San Diego, La Jolla, CA, USA {jjj025, yip}@eng.ucsd.edu

[‡]A.H. Qureshi is with the Department of Computer Science at Purdue University, West Lafayette, IN, USA ahqureshi@purdue.edu

improving the sampling efficiency of SMP planners by narrowing down the search region [6]. VQ-MPT has shown promising generalization capabilities in that trained models can reduce planning times for environments outside the training data. This paper introduces Constraint VQ-MPT (CVQ-MPT), a fast and efficient neural planner for kinematic and task-specific constraints. Unlike previous learning-based approaches, our model requires no additional task-related data for training or finetuning the model and uses a pre-trained VQ-MPT model for planning. The main contributions of this paper are:

- 1) Present a zero-shot planning algorithm that requires no task-specific training data for solving constraint motion planning problems and results in $2\times$ improvement in planning times.
- 2) Formulate a gradient-based update of distributions predicted by VQ-MPT, improving planning performance.
- 3) Provide empirical evidence of how trajectories generated by CVQ-MPT are 35%-40% shorter compared to previous SMPs, resulting in improved task execution times.

II. RELATED WORKS

Numerous methods have been proposed to solve the CMP. Broadly, they can be categorized into optimization-based and sampling-based approaches.

Optimization-based approaches formulate the entire planning problem, including constraints, as an optimization problem. In [15], the authors apply Covariant Hamiltonian Optimization for Motion Planning (CHOMP) [16], a method that uses covariant gradients to solve the unconstrained planning problem to construct unconstrained trajectories and project them onto the constraint manifold. TrajOpt [17] improves optimization-based planning by utilizing Sequential Convex Programming (SCP) and incorporates constraint functions as penalties within the optimization problem to solve CMP. Bonalli et al. [18] extend SCP methods by lifting the manifold constraints to Euclidean spaces. Howell et al. [19] use Augmented Lagrangian- Iterative Linear Quadratic Regulators (AL-iLQR) to build a rough solution and a projection-based method to refine the coarse solution to solve the constraint problem. However, these techniques require hand-tuning penalty functions for different tasks, often resulting in local minima solutions.

On the other hand, sampling-based approaches build discrete search graphs by random sampling in the planning space to connect the start and goal states. SMP planners such as Rapidly Exploring Random Trees (RRT) [20] and its variants [21] have shown to be particularly effective in solving planning problems for robotic systems with higher degrees of freedom, such as manipulators and also easily adapt to a wide range of tasks [22]. Sampling-based methods can solve constraint planning by generating random samples on the constraint manifold mainly through projection-based and continuation-based methods.

The projection-based approach uses a projection operator to map sampled points onto the constraint manifold. The projection operators usually involve first-order gradients to

iteratively move the sampled point toward the constraint manifold using the Jacobian of the constraint function. [23], [24] use projection-based methods for solving CMP. On the other hand, continuation-based approaches approximate constraint manifolds at local regions and use this approximation to sample points and construct local trajectories. Methods such as AtlasRRT [25] and Tangent-Bundle RRT (TB-RRT) [26], [27] simultaneously approximate the constraint manifold using tangent spaces at adhering points and uses a BiRRT to construct a tree between the start and goal points on this approximated manifold. However, the underlying optimization routines and constructing atlas's can make planning computationally expensive.

To improve the efficiency and speed of sampling-based planners, recent methods have utilized learning-based techniques to solve the planning task [28], [10]. Constraint Motion Planning Networks (CoMPNet) [29], [14] was the first of these methods that used a fully connected neural network to generate configurations near constraint surfaces. By selectively sampling points, these methods reduced planning times considerably. Similarly, in [30], a deep neural network models the constraint manifold using prior trajectory data and can compute trajectories quickly. Liu et al. [31] propose using reinforcement learning to find a policy that always satisfies the constraints, which allows the agent to explore the space efficiently and removes the need for a projection operator. Although these methods improve the efficiency and speed of constraint planning, they lack generalizability to new environments and require access to task-specific demonstrations. In this work, we address previous CMP drawbacks without collecting any task-specific data.

III. PROBLEM DEFINITION

Consider an n dimensional planning space defined by $\mathcal{C} \in \mathbb{R}^n$. The space is split into two subspaces $\mathcal{C}_{free} \subset \mathcal{C}$ and $\mathcal{C} - \mathcal{C}_{free}$ such that all states in \mathcal{C}_{free} do not collide with any obstacle in the environment and are considered valid configurations. The objective of the motion planner is to generate a continuous path (or trajectory), σ , such that it connects the given start state (q_s) and a goal region (\mathcal{C}_{goal}), and all states in σ are in \mathcal{C}_{free} . For CMP, all the points in σ must also satisfy a constraint function $F : \mathcal{C} \rightarrow \mathbb{R}^k$, where k is the number of constraints. For a state $q \in \mathcal{C}_{free}$ to satisfy the constraint, $F(q) = \mathbf{0}$, where $\mathbf{0}$ is a vector of zeros. This work focuses on the constraint function defined on joint configurations, \mathcal{C} , and not on the robot kinematics or dynamics.

IV. BACKGROUND

A. Task-Space Regions

A task space for manipulators is the space spanned by possible end-effector poses. For example, the task space for the Franka Panda arm is the $SE(3)$ space. For many applications, such as moving a cup without spilling its content, opening doors for shelves, and gasping objects from a table, constraints can be defined with respect to the task space. Task-Space Regions (TSR) [23] defines such

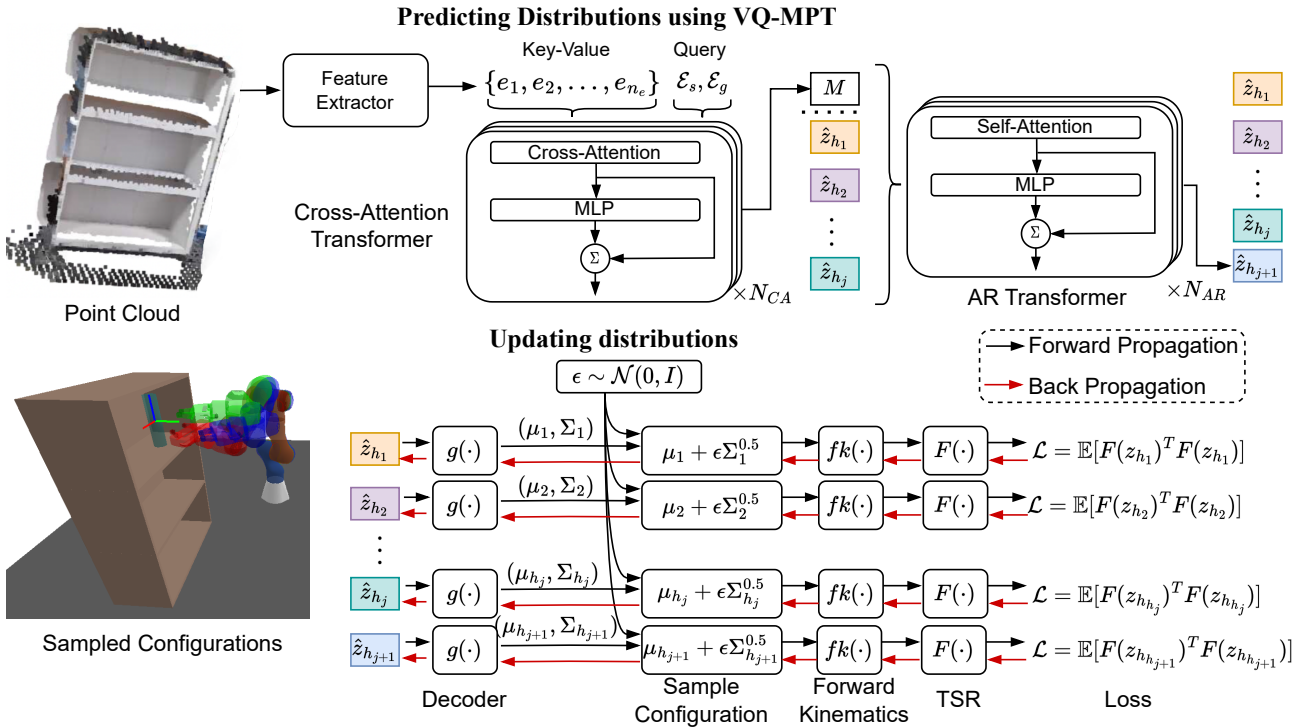


Fig. 2. An outline of the model architecture of CVQ-MPT. Given a point cloud, and start (q_s) and goal (q_g) configurations, a pre-trained transformer model is used to generate a set of distributions parameterized using $\{\hat{z}_{h_1}, \dots, \hat{z}_{h_{j+1}}\}$. The predicted distributions are updated using gradient-based optimization, minimizing the loss term \mathcal{L} , moving it closer to the constraint function, $F(q)$. This allows for sampling configurations closer to the constraint manifold.

constraints as the relative pose between the robot end-effector pose evaluated using a forward kinematics function, $fk(q) : \mathcal{C} \rightarrow SE(3)$, and a given target pose. The relative pose is expressed as a vector in \mathbb{R}^6 where the first three elements represent the relative position (\mathbf{p}^p), and the last three represent the relative orientation (ω^p) in axis-angle form. Thus, the constraint function F is defined using these bounds. Refer to [23] for more details.

B. Vector Quantized-Motion Planning Transformer

The VQ-MPT model consists of two stages - a quantization stage and a prediction stage. The quantization stage segments the planning space as a collection of distributions. VQ-MPT uses a Vector Quantized (VQ) model to generate the collection of distributions of the planning space. VQ models are generative models with an encoder-decoder architecture similar to Variational AutoEncoder (VAE) models but with the latent dimension represented as a collection of learnable vectors - $\mathcal{Z}_Q = \{\hat{z}_1, \hat{z}_2, \dots, \hat{z}_N\}$ referred to as dictionaries. Each dictionary value can be reconstructed back to the planning space using a decoder function $g(\cdot)$.

The second stage generates sampling regions by predicting indexes (\mathcal{H}) from the dictionary set for a given planning problem and sensor data. It comprises two models - a cross-attention model with a feature extractor to embed start and goal pair, and the environment embedding into latent vectors (M) and a transformer-based Auto-Regressive (AR) model to predict the dictionary indexes. Transformer networks are ideal for cross-attention and AR models because they are agnostic to the environment size and can generate latent

embeddings for point clouds of varying sizes and output sequences of varying lengths. The architecture of the second stage is outlined in Fig. 2. The environment representation (i.e., point cloud data) is passed through a feature extractor to construct the environment encodings $\mathcal{E} = \{e_1, e_2, \dots, e_{n_e}\}$ where $e_i \in \mathbb{R}^d$ and d is the size of the latent dimension. The feature extractor reduces the dimensionality of the environment representation and captures local environment structures as latent variables using set-abstraction proposed in PointNet++ [32]. The start and goal states (q_s and q_g) are projected to the start and goal embedding ($\mathcal{E}_s \in \mathbb{R}^d$ and $\mathcal{E}_g \in \mathbb{R}^d$) using a Multi-Layer Perceptron (MLP) network. A cross-attention model uses the environment embedding, \mathcal{E} , and the start and goal embedding, $\{\mathcal{E}_s, \mathcal{E}_g\}$ to generate latent vectors M . The cross-attention model learns a feature embedding that fuses the given start and goal pair with the given planning environment. It uses the vector in \mathcal{E} as key-value pairs, and \mathcal{E}_s and \mathcal{E}_g as query vectors to generate M .

The second model of stage 2 uses an AR transformer model, $\pi(\cdot)$, to predict the dictionary indexes \mathcal{H} . The transformer architecture allows the model to make long-horizon connections. For each index h_j , the model outputs a probability distribution over $\mathcal{Z}_Q \cup \{z_g\}$ given dictionary values of previous predictions $\{\hat{z}_{h_1}, \hat{z}_{h_2}, \dots, \hat{z}_{h_{j-1}}\}$ and the planning context M :

$$\pi(h_j = i | \hat{z}_{h_1}, \dots, \hat{z}_{h_{j-1}}, M) = p_i \quad \text{where} \quad \sum_{i=1}^{N+1} p_i = 1 \quad (1)$$

The model stops when it predicts the goal key (z_g). Using the learned \hat{z} decoder from Stage 1, we can convert each \hat{z}_{h_j}

Algorithm 1: Project(q)

```
1  $x \leftarrow F(q)$ ;
2 while  $\|x\| > \epsilon$  do
3    $J \leftarrow \nabla_q F(q)$ ;
4    $\Delta q \leftarrow -J^T(JJ^T)^{-1}x$ ;
5    $q \leftarrow q + \Delta q$ ;
6    $x \leftarrow F(q)$ ;
7 end
8 return  $q$ 
```

value into a Gaussian distribution in the planning space.

$$g(\hat{z}_{h_j}) = \mathcal{N}(\mu_{h_j}, \Sigma_{h_j}) \quad (2)$$

This work uses a pre-trained VQ-MPT model trained on unconstrained shortest trajectory data collected in simulation.

V. CONSTRAINT VQ-MPT

To improve the efficiency of CMP, we sample from the distributions generated by VQ-MPT and project them onto the constraint manifold. The following section details our sampler for constraint planning, optimization for updating predicted distributions, and planner for solving CMP.

A. Generating samples

Given a start (q_s) and goal (q_g) state, we use the trained Stage 2 model of VQ-MPT to generate a sequence of dictionary indexes $\mathcal{H} = \{h_1, \dots, h_{n_h}\}$, where h_{n_h} is the goal index. A beam-search algorithm similar to those used in language model tasks [33] is used to optimize the following:

$$P(h_1, \dots, h_{n_h} | M) = \prod_{i=1}^{n_h} \pi(h_i | h_1, \dots, h_{i-1}, M) \quad (3)$$

to generate the most probable sequence, \mathcal{H} . Here π is the probability from Eqn. 1. We use a Gaussian Mixture Model (GMM) with uniform mixing coefficients reconstructed from predicted sequence \mathcal{H} for sampling points. Individual Gaussian distribution is reconstructed from \mathcal{H} using the decoder model from Eqn. 2.

$$\mathcal{P}(q) = \sum_{i=1}^{n_h-1} \frac{1}{n_h-1} g(\hat{z}_{h_i}) \quad (4)$$

To ensure that the sampled points from $\mathcal{P}(q)$ satisfy the given constraints $F(q)$, we project them on the constraint manifold using the first-order gradient projection operator given in Algorithm 1.

B. Improving sampling efficiency

The samples generated from Eqn. 4 are not guaranteed to lie on the constraint manifold and might require multiple optimization iterations to project them onto it. We propose a novel optimization-based update of the latent vectors in \mathcal{H} such that the new distribution lies closer to the constraint, reducing projection times and consequently planning performance (Fig. 2). If q^* adheres to the constraint surface, it must also satisfy the following:

$$q^* = \min_q F(q)^T F(q) f \quad (5)$$

Algorithm 2: UpdateDistribution(z)

```
1 repeat
2    $z_{cur} \leftarrow z$ ;
3    $\mu, \Sigma \leftarrow g(z)$ ;
4    $\mathcal{L} \leftarrow 0$ ;
5   for  $i \leftarrow 1$  to  $N$  do
6      $\epsilon \sim \mathcal{N}(0, I)$ ;
7      $q \leftarrow \mu + \epsilon \Sigma^{0.5}$ ;
8      $\mathcal{L} \leftarrow \mathcal{L} + (F(q)^T F(q)/N)$ ;
9   end
10   $\Delta z \leftarrow \nabla_z \mathcal{L}$ ;
11   $z \leftarrow z_{cur} - \eta \Delta z$ ;
12 until  $\|\Delta z\| < \delta$ ;
13 return  $z$ 
```

Algorithm 3: CVQMPTBiPlanner($q_s, q_g, \mathcal{P}, K, b$)

```
1  $\tau_s \leftarrow \{q_s\}, \tau_g \leftarrow \{q_g\}$ ;
2 for  $k \leftarrow 0$  to  $K$  do
3    $q_{rand} \leftarrow \text{Sample}(\mathcal{P})$ ;
4    $q_{near}^s \leftarrow \text{NearestNode}(q_{rand}, T_s)$ ;
5    $q_{reach}^s \leftarrow \text{ConstrainedExtent}(q_{rand}, q_{near}^s, T_s)$ ;
6    $q_{near}^g \leftarrow \text{NearestNode}(q_{rand}, T_g)$ ;
7    $q_{reach}^g \leftarrow \text{ConstrainedExtent}(q_{rand}, q_{near}^g, T_g)$ ;
8   if  $\text{Connect}(q_{reach}^s, q_{reach}^g)$  then
9      $\tau \leftarrow \text{ExtractPath}(\tau_s, \tau_g)$ ;
10    return  $\text{Simplify}(\tau)$ ;
11  else
12     $\text{Swap}(\tau_s, \tau_g)$ ;
13  end
14 end
15 return  $\Phi$ 
```

Thus, we can use the function $G(q) = F(q)^T F(q)$ to evaluate constraint adherence and use the same to update our latent vectors, \hat{z}_{h_i} . Since $G : \mathcal{C} \rightarrow \mathbb{R}^+$ is always positive, using Markov Inequality, we can upper-bound the probability of G lying outside a threshold δ for the distribution $g(\hat{z}_{h_i})$.

$$P(G(q) > \delta) \leq \frac{\mathbb{E}_{g(\hat{z}_{h_i})}[G(q)]}{\delta} \quad (6)$$

By minimizing the upper bound in Eqn. 6, we can implicitly reduce the samples that are further away from the constraint manifold, improving the sampling efficiency of the planner. We can use Monte Carlo estimates of $G(\cdot)$ by sampling points on the distribution $g(\hat{z}_{h_i})$ to evaluate the upper bound in Eqn. 6.

$$\mathcal{L} = \mathbb{E}_{g(\hat{z}_h)}[G(q)] \approx \frac{1}{n} \sum_{k=1}^n G(q_k) \quad q_k \sim g(\hat{z}_h) \quad (7)$$

To optimize Eqn. 7, we want to differentiate the objective function with respect to latent variable \hat{z}_h . We use the reparameterization trick [34] to express the random variable q_k as a function of μ_h, Σ_h , and a normal distribution ϵ .

$$q_k = \mu_h + LD\epsilon \quad \epsilon \sim \mathcal{N}(0, I), \Sigma_h = LD^2L^T \quad (8)$$

where L and D are a lower triangular and diagonal matrix, respectively. To generate points on $g(\hat{z}_h)$, we can sample points on $\mathcal{N}(0, I)$ and transform it using Eqn. 8. By substituting Eqn. 8 in Eqn. 7, the gradient of \mathcal{L} becomes a function of deterministic parameters.




Fig. 3. The histogram of the objective function, $G(q)$, before and after optimizing for two different dictionary values predicted by VQ-MPT for the place task. We can update the manifolds to generate samples closer to the constraint manifold and not push away distributions that are already closer.

$$\nabla_{\hat{z}_h} \mathcal{L} \approx \frac{1}{n} \sum_{k=1}^n \nabla_q F(q_k) \frac{\delta(\mu_{h_j} + LD\epsilon_k)}{\delta \hat{z}_h} \quad (9)$$

Using the derivatives from Eqn. 9, any optimization methods, such as Gradient Descent (GD), can be used to minimize the upper bound. Algorithm 2 provides a general outline for updating the latent variable. Fig. 3 shows the histogram of the objective function before and after optimization.

C. Planning

Any SMP can generate the trajectory using the sampling strategy from Section V-A and V-B. We provide Algorithm 3, a bidirectional planning algorithm, to generate a path using samples from the distribution in Eqn. 4. The *CVQMPTBi-Planner* function takes the start and goal state (q_s and q_g), the GMM model (\mathcal{P}), the number of samples to generate (K), and a threshold value (b) to sample the goal state and returns a valid trajectory. The function *NearestNode* finds the closest node on the tree to the sampled node, while the function *ConstrainedExtend* extends the tree from the nearest node toward the sampled configuration until a collision or constraint violation occurs. For start and goal regions defined by TSR, which occur during grasping or object-placement tasks, we sample a pose from a TSR and use an Inverse Kinematics (IK) solver [35] to generate a valid collision-free configuration. This configuration is used to identify the sampling region using the VQ-MPT planner. Since the VQ-MPT planner can generate paths rapidly, multiple configurations can be evaluated successively to find a valid path.

In [23], [36], the authors have proved the probabilistic completeness of projection-based planners. They show that any RRT-based algorithm with a non-zero probability of sampling in an n -dimensional ball centered at a node of the existing tree together with the projection operator in Algorithm 1 is probabilistically complete. Since we use a set of Gaussian Distributions to sample, which spans the entire planning space, our planner also satisfies these conditions, making it probabilistically complete as well.




Fig. 4. An example of the trajectory planned using CVQ-MPT for the place task for a given start (green) and goal (red) configurations. The constraint is to hold the can upright during the motion.

VI. EXPERIMENTS

This section covers our experiment setup and compares CVQ-MPT with three other planners - CBiRRT [23], AtlasRRT [25], and TB-RRT [26] for a 7D Franka Panda arm on simulated and physical scenes. We also compared the performance of a pre-trained MPNet [10], which used the projection operator for steering.

A. Setup

We used pre-trained VQ-MPT and MPNet models from our prior work [6] for a 7D Franka Panda robot arm. The model was trained on simulated data using unconstrained trajectories. No constraint trajectories were used to finetune these models. All planners used in this work were implemented using the Open Motion Planning Library (OMPL) [37] on a system with an AMD Ryzen Threadripper 1950X CPU with an Nvidia RTX 3090 GPU.

B. Place Task

First, we compared our planner's performance against traditional and learning-based SMP algorithms in solving a placement task (Fig 4). A can is randomly placed on the kitchen counter, and the objective is to plan a path to place it on the shelf without tilting it. To quantify planning performance, we measured four metrics: planning time - the time it takes for the planner to generate a valid trajectory; vertices - the number of collision-free vertices required to find the trajectory, accuracy - the percentage of planning problems solved before a given cutoff time, and path length - the sum of all the Euclidean distance between adjacent joint states. The number of vertices will help us to determine the efficiency of different constraint planning techniques since projecting points on the constraint manifold can be computationally expensive [22]. We tested 120 different planning problems; Table I summarizes the results.

CVQ-MPT and Opt-CVQ-MPT achieve similar or better accuracy than previous constraint planners while planning shorter trajectories. CVQ-MPT produces shorter paths because the underlying VQ-MPT model was trained to identify



Fig. 5. Snapshots of the trajectory planned using CVQ-MPT for the Panda Arm completing the multi-sequence task (from left to right). CVQ-MPT can plan trajectories for various types of planning constraints in complex environments.

TABLE I. PLACE TASK AND PHYSICAL ROBOT EXPERIMENTS: ACCURACY, PLANNING TIME, VERTICES, & PATH LENGTH

Environments		CBiRRT	Atlas-RRT	TB-RRT	MPNet (w/ proj)	CVQ-MPT	Opt-CVQ-MPT
Place Task	Accuracy	93.7%	100%	89.76%	33.85%	97.63%	98.42%
	Time (s)	22.95	21.80	25.31	20.61	11.25	11.03
	Vertices	40	37	32	35.98	26	24
	Path Length	8.715	5.85	7.28	8.934	5.080	5.507
Physical robot	Accuracy	10/10	10/10	10/10	4/10	10/10	10/10
	Time (s)	8.37	11.01	13.98	7.35	9.89	8.37
	Vertices	38	28	34	5	56	41
	Path Length	12.643	10.613	12.659	18.747	7.530	7.313

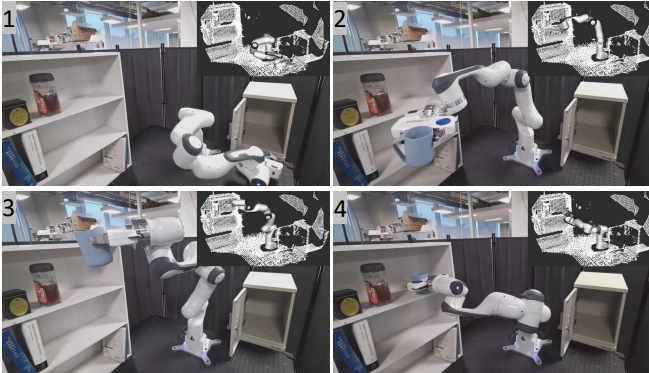


Fig. 6. Sequences of the trajectory planned using CVQ-MPT for the Panda Arm on the physical robot and the corresponding point cloud used to represent the environment (1→2→3→4). CVQ-MPT can plan trajectories using physical sensor data and achieve the same level of performance observed in simulated environments.

TABLE II. MULTI-SEQUENCE TASK: PLANNING, EXECUTING TIMES

	CBiRRT	Atlas-RRT	TB-RRT	CVQ-MPT	Opt-CVQ-MPT
Planning (s)	19.71	23.46	19.24	13.10	10.98
Execution (s)	22.11	20.48	22.49	20.86	18.66
Total Time (s)	41.82	43.94	41.73	33.96	29.64

sampling regions in \mathcal{C} where the shortest unconstrained path exists. Thus, our planner projects the unconstrained shortest path on the constraint manifold. We can also observe that by reducing the search space, CVQ-MPT can plan constraint trajectories almost $2\times$ faster than previous planners. The MPNet model on the other hand achieves poor accuracy since the model is not generalizable to unseen environments [6].

C. Multi-Sequence Task

Our next experiment compared the planning and execution of CVQ-MPT for a sequence of tasks with varying constraints. The task involves opening the kitchen cabinet, grasping a can from the shelf, and placing it on the kitchen counter (Fig. 5). We compared 10 different tasks, where the robot's start position and the can's final goal position were random. Due to the poor performance of MPNet on the place task, we did not compare against it for this experiment. The

average total planning and executing time results for the tasks are reported in Table II. CVQ-MPT reduces total planning time by 45%, while planning a shorter path can reduce task execution times by 18%. This also shows how our planner can adapt to diverse types of task constraints.

D. Real-world environment

To evaluate the performance of our planner on physical sensor data, we tested it in a real-world environment (Fig. 6). The environment was represented using point cloud data from Azure Kinect sensors, and collision checking was done using the Octomap collision checker from Moveit [38]. Camera to robot base transform was estimated using markerless pose estimation technique [39]. We tested on 10 random start and goal configurations, and the results are summarized in Table I. We observe that CVQ-MPT and Opt-CV-MPT outperform traditional planners, similar to the Place Task experiment. This experiment shows that CVQ-MPT models can also generalize well to physical sensor data without further training or fine-tuning. Such generalization will benefit the larger robotics community since other researchers can use trained models in diverse settings without collecting new data or fine-tuning the model.

VII. CONCLUSION

We introduced CVQ-MPT, a zero-shot planning framework for solving constraint motion planning problems in this work. By reducing the search space for sampling-based planners, we improve constraint planning efficiency and speed, simultaneously generating shorter trajectories than previous planners. We further refine our search space by optimizing the predicted distributions to be closer to the constraint manifold, further improving planning performance. As we have shown, CVQ-MPT also improves task execution times and success rate which will enable future robotic systems to handle more complex tasks with intricate planning sequences. Future works can explore the application of CVQ-MPT to such task and motion planning problems.

REFERENCES

- [1] J. J. Johnson, L. Li, F. Liu, A. H. Qureshi, and M. C. Yip, "Dynamically constrained motion planning networks for non-holonomic robots," in *2020 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, 2020, pp. 6937–6943.
- [2] J. J. Johnson and M. C. Yip, "Chance-constrained motion planning using modeled distance- to-collision functions," in *2021 IEEE 17th International Conference on Automation Science and Engineering (CASE)*, 2021, pp. 1582–1589.
- [3] E. Peiros, Z.-Y. Chiu, Y. Zhi, N. Shinde, and M. C. Yip, "Finding biomechanically safe trajectories for robot manipulation of the human body in a search and rescue scenario," in *IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, 2023.
- [4] J. Hu, A. Adeleye, and H. I. Christensen, "Place-and-pick-based re-grasping using unstable placement," in *Robotics Research*, 2023.
- [5] A. Pettinger, F. Alambeigi, and M. Pryor, "A versatile affordance modeling framework using screw primitives to increase autonomy during manipulation contact tasks," *IEEE Robotics and Automation Letters*, vol. 7, no. 3, pp. 7224–7231, 2022.
- [6] J. J. Johnson, A. H. Qureshi, and M. Yip, "Learning sampling dictionaries for efficient and generalizable robot motion planning with transformers," 2023.
- [7] S. M. LaValle, *Planning Algorithms*. USA: Cambridge University Press, 2006.
- [8] J. D. Gammell, S. S. Srinivasa, and T. D. Barfoot, "Batch informed trees (BIT*): Sampling-based optimal planning via the heuristically guided search of implicit random geometric graphs," in *2015 IEEE Int. Conf. Robot. Autom.*, 2015.
- [9] L. Jaillet and J. Porta, "Efficient asymptotically-optimal path planning on manifolds," *Robotics and Autonomous Systems*, 2013.
- [10] A. H. Qureshi, Y. Miao, A. Simeonov, and M. C. Yip, "Motion planning networks: Bridging the gap between learning-based and classical motion planners," *IEEE Trans. on Robotics*, 2020.
- [11] C. Yu and S. Gao, "Reducing collision checking for sampling-based motion planning using graph neural networks," in *Advances in Neural Information Processing Systems*, 2021.
- [12] P. Lehner and A. Albu-Schäffer, "The repetition roadmap for repetitive constrained motion planning," *IEEE Robot. and Autom. Letters*, 2018.
- [13] C. Chamzas, Z. Kingston, C. Quintero-Peña, A. Shrivastava, and L. E. Kavraki, "Learning sampling distributions using local 3d workspace decompositions for motion planning in high dimensions," in *IEEE Int. Conf. on Robot. and Autom.*, 2021.
- [14] A. H. Qureshi, J. Dong, A. Choe, and M. C. Yip, "Neural manipulation planning on constraint manifolds," *IEEE Robotics and Automation Letters*, vol. 5, no. 4, pp. 6089–6096, 2020.
- [15] A. D. Dragan, N. D. Ratliff, and S. S. Srinivasa, "Manipulation planning with goal sets using constrained trajectory optimization," in *2011 IEEE International Conference on Robotics and Automation*, 2011, pp. 4582–4588.
- [16] N. Ratliff, M. Zucker, J. A. Bagnell, and S. Srinivasa, "Chomp: Gradient optimization techniques for efficient motion planning," in *2009 IEEE International Conference on Robotics and Automation*, 2009, pp. 489–494.
- [17] J. Schulman, Y. Duan, J. Ho, A. Lee, I. Awwal, H. Bradlow, J. Pan, S. Patil, K. Goldberg, and P. Abbeel, "Motion planning with sequential convex optimization and convex collision checking," *Int. J. Rob. Res.*, vol. 33, no. 9, p. 1251–1270, aug 2014. [Online]. Available: <https://doi.org/10.1177/0278364914528132>
- [18] R. Bonalli, A. Cauligi, A. Bylard, T. Lew, and M. Pavone, "Trajectory optimization on manifolds: A theoretically-guaranteed embedded sequential convex programming approach," in *Robotics: Science and Systems XV, University of Freiburg, Freiburg im Breisgau, Germany, June 22-26, 2019*, A. Bicchi, H. Kress-Gazit, and S. Hutchinson, Eds., 2019. [Online]. Available: <https://doi.org/10.15607/RSS.2019.XV.078>
- [19] T. A. Howell, B. E. Jackson, and Z. Manchester, "Altro: A fast solver for constrained trajectory optimization," in *2019 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, 2019, pp. 7674–7679.
- [20] S. M. LaValle and J. James J. Kuffner, "Randomized kinodynamic planning," *The International Journal of Robotics Research*, 2001.
- [21] A. H. Qureshi and Y. Ayaz, "Intelligent bidirectional rapidly-exploring random trees for optimal motion planning in complex cluttered environments," *Robotics and Autonomous Systems*, vol. 68, pp. 1–11, 2015. [Online]. Available: <https://www.sciencedirect.com/science/article/pii/S0921889015000317>
- [22] Z. Kingston, M. Moll, and L. E. Kavraki, "Sampling-based methods for motion planning with constraints," *Annual Review of Control, Robotics, and Autonomous Systems*, vol. 1, no. 1, pp. 159–185, 2018.
- [23] D. Berenson, S. Srinivasa, and J. Kuffner, "Task space regions: A framework for pose-constrained manipulation planning," *The International Journal of Robotics Research*, vol. 30, no. 12, pp. 1435–1460, 2011.
- [24] P. Englert, I. R. Fernández, R. Ramachandran, and G. Sukhatme, "Sampling-based motion planning on sequenced manifolds," in *Robotics: Science and Systems XVII*. Robotics: Science and Systems Foundation, jul 2021. [Online]. Available: <https://doi.org/10.15607/2Frss.2021.xvii.039>
- [25] L. Jaillet and J. M. Porta, "Path planning under kinematic constraints by rapidly exploring manifolds," *IEEE Transactions on Robotics*, vol. 29, no. 1, pp. 105–117, 2013.
- [26] C. Suh, T. T. Um, B. Kim, H. Noh, M. Kim, and F. C. Park, "Tangent space rrt: A randomized planning algorithm on constraint manifolds," in *2011 IEEE International Conference on Robotics and Automation*, 2011, pp. 4968–4973.
- [27] B. Kim, T. T. Um, C. Suh, and F. C. Park, "Tangent bundle rrt: A randomized algorithm for constrained motion planning," *Robotica*, vol. 34, no. 1, p. 202–225, 2016.
- [28] C.-M. Hung, S. Zhong, W. Goodwin, O. P. Jones, M. Engelcke, I. Havoutis, and I. Posner, "Reaching through latent space: From joint statistics to path planning in manipulation," *IEEE Robotics and Automation Letters*, vol. 7, no. 2, pp. 5334–5341, 2022.
- [29] A. H. Qureshi, J. Dong, A. Baig, and M. C. Yip, "Constrained motion planning networks x," *IEEE Transactions on Robotics*, vol. 38, no. 2, pp. 868–886, 2022.
- [30] P. Kicki, P. Liu, D. Tateo, H. Bou-Ammar, K. Walas, P. Skrzypczyński, and J. Peters, "Fast kinodynamic planning on the constraint manifold with deep neural networks," 2023.
- [31] P. Liu, D. Tateo, H. B. Ammar, and J. Peters, "Robot reinforcement learning on the constraint manifold," in *Proceedings of the 5th Conference on Robot Learning*, ser. Proceedings of Machine Learning Research, A. Faust, D. Hsu, and G. Neumann, Eds., vol. 164. PMLR, 08–11 Nov 2022, pp. 1357–1366. [Online]. Available: <https://proceedings.mlr.press/v164/liu22c.html>
- [32] C. R. Qi, L. Yi, H. Su, and L. J. Guibas, "Pointnet++: Deep hierarchical feature learning on point sets in a metric space," in *Advances in Neural Information Processing Systems*, 2017.
- [33] J. Devlin, M. Chang, K. Lee, and K. Toutanova, "BERT: pre-training of deep bidirectional transformers for language understanding," in *Proceedings of the Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*, 2019.
- [34] D. P. Kingma and M. Welling, "Auto-encoding variational bayes," 2022.
- [35] P. Beeson and B. Ames, "Trac-ik: An open-source library for improved solving of generic inverse kinematics," in *2015 IEEE-RAS 15th International Conference on Humanoid Robots (Humanoids)*, 2015, pp. 928–935.
- [36] Z. Kingston, M. Moll, and L. E. Kavraki, "Exploring implicit spaces for constrained sampling-based planning," *The International Journal of Robotics Research*, vol. 38, no. 10-11, pp. 1151–1178, 2019. [Online]. Available: <https://doi.org/10.1177/0278364919868530>
- [37] I. A. Şucan, M. Moll, and L. E. Kavraki, "The Open Motion Planning Library," *IEEE Robotics & Auto. Magazine*, 2012.
- [38] D. Coleman, I. A. Sucan, S. Chitta, and N. Correll, "Reducing the barrier to entry of complex robotic software: a moveit! case study," *ArXiv*, vol. abs/1404.3785, 2014. [Online]. Available: <https://api.semanticscholar.org/CorpusID:13939653>
- [39] J. Lu, F. Richter, and M. C. Yip, "Markerless camera-to-robot pose estimation via self-supervised sim-to-real transfer," 2023.