

Constrained 6-DoF Grasp Generation on Complex Shapes for Improved Dual-Arm Manipulation

Gaurav Singh^{1*}, Sanket Kalwar^{1*}, Md Faizal Karim¹, Bipasha Sen², Nagamanikandan Govindan¹,
Srinath Sridhar³ and K Madhava Krishna¹

Abstract—Efficiently generating grasp poses tailored to specific regions of an object is vital for various robotic manipulation tasks, especially in a dual-arm setup. This scenario presents a significant challenge due to the complex geometries involved, requiring a deep understanding of the local geometry to generate grasps efficiently on the specified constrained regions. Existing methods only explore settings involving tabletop/small objects and require augmented datasets to train, limiting their performance on complex objects. We propose CGDF: Constrained Grasp Diffusion Fields, a diffusion-based grasp generative model that generalizes to objects with arbitrary geometries, as well as generates dense grasps on the target regions. CGDF uses a part-guided diffusion approach that enables it to get high sample efficiency in constrained grasping without explicitly training on massive constraint-augmented datasets. We provide qualitative and quantitative comparisons using analytical metrics and in simulation, in both unconstrained and constrained settings to show that our method can generalize to generate stable grasps on complex objects, especially useful for dual-arm manipulation settings, while existing methods struggle to do so. More results, code and an extended version of the paper can be found on the project page: <https://constrained-grasp-diffusion.github.io/>

I. INTRODUCTION

Grasping is a critical robot capability serving a wide array of applications spanning industrial automation, household assistance, and beyond. Significant progress has been made in building grasp generation methods [1]–[4]. These works primarily focus on generating stable and collision-free grasps **uniformly distributed** across the object or scene to support pick and place tasks. However, they have certain limitations.

First, uniformly distributed grasps may not be practical since objects may be fragile, unwieldy, or large [6], necessitating grasps that are concentrated more on certain parts of the object. Second, due to suboptimal shape representations and inherent biases, such methods fail to provide dense coverage over objects with complex shapes, making them **sample-inefficient**, i.e., one needs to run the model multiple times and generate a large number of grasps to get some good grasps on the target region. Due to these issues, it is desirable to have a method to generate grasps **constrained** to specific regions of the object. *Given a target region, constrained grasp generation aims to generate dense, sample-efficient grasps only on the desired region.* Constrained grasping for complex objects also facilitates the possibility of generating grasps for more than one arm by constraining the grasps on

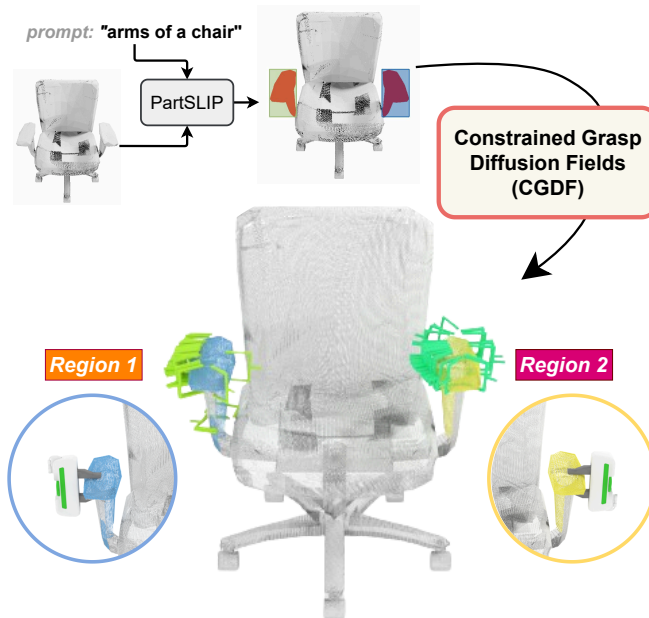


Fig. 1: CGDF: Constrained Grasp Diffusion Fields generates dense grasps on large objects with complex shapes (like a chair), in a dual-arm setting. Given target regions (which can be generated from a text prompt using PartSLIP [5]), CGDF uses a Part-Guided Diffusion strategy to generate sample efficient grasps on the specified regions, enabling grasping for multiple regions for improved multi-arm grasping.

multiple regions of interest corresponding to multiple arms. This thereby enables stable “dual-arm” grasping - an area of growing interest that is yet to be addressed widely in the community.

We introduce **CGDF: Constrained Grasp Diffusion Fields**, a novel method designed for generating constrained 6-DoF grasps tailored to complex shapes. CGDF relies on an improved shape representation allowing the grasp generator to gain detailed descriptors for complex objects. Further, we propose a **part-guided diffusion** strategy to generate such grasps in a sample-efficient manner.

In this work, we assume that the regions of interest (i.e. constraints) can be easily identified using existing approaches such as today’s VLM-based affordance detection methods [5], [7], [8]. However, while these methods are good at segmenting regions based on text prompts, they lack the ability to propose grasps specific to these regions of interest thereby relying on existing uniform grasp generation methods [9] or simulation-based methods [10] to sample

*Equal Contributions

¹Robotics Research Center, IIIT-Hyderabad

²Massachusetts Institute of Technology

³Brown University

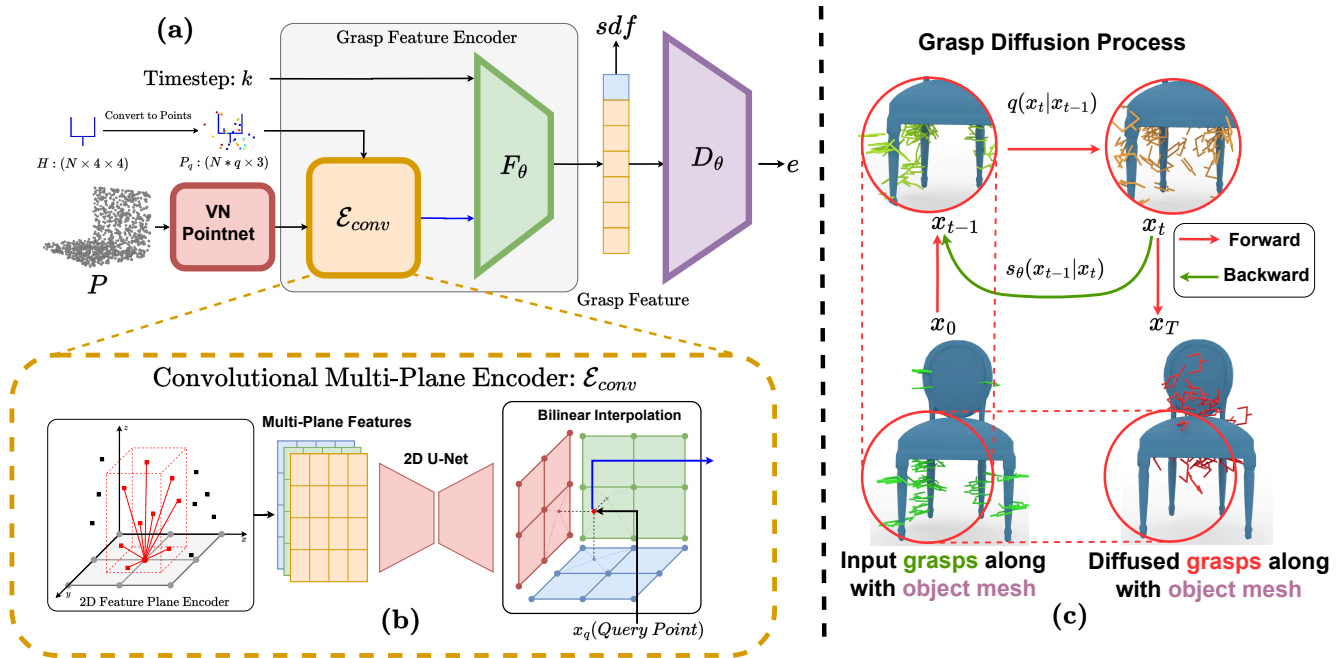


Fig. 2: **Overview:** The figure section (a) shows the architecture of our proposed energy-based model E_θ as explained in III-C and III-A. During this process, the model takes as input a point cloud and the grasp pose, which is subsequently converted into a set of query points. We use a VN-Pointnet-based point cloud encoder, which generates per-point features. (b) shows how these features are then distilled into three 2D feature planes oriented along the XY, XZ, and YZ planes using a convolutional multi-plane encoder. For each grasp pose, feature vectors corresponding to N query points are obtained using bilinear interpolation on the feature planes. Subsequently, the grasp feature vector is derived from F_θ and decoded into an energy value by D_θ . In figure section (c), we show the grasp diffusion process, where grasps are diffused over the object during the forward diffusion process and denoised using the backward diffusion process.

many grasps and prune grasps based on segmented regions. Moreover, for large objects, data-driven uniform grasping methods fail to provide dense coverage, thus becoming a bottleneck to the pipeline. Existing constrained grasping methods like [11] overcome this challenge but do not generalize well to large objects with complex shapes.

We tackle the task of constrained grasp generation on arbitrary regions of large objects with complex shapes. Further, we showcase the efficacy of our approach in handling the complex task of dual-arm grasp generation. The key insight of our method is that convolutional plane features [12] that have shown significant quality improvements in fine-grained implicit 3D reconstruction [13], [14] also encode better grasp pose descriptors that improve grasp generation on complex objects. Existing general constrained grasping methods [11] require expanding existing large-scale grasp datasets like ACRONYM [15] with constrained regions and annotated grasps on them. Due to rich local geometric features, our method, trained in an unconstrained grasping setting, can be directly applied to constrained grasping using a part-guided diffusion strategy *without* losing sample efficiency, thus removing the need for such conditionally labeled datasets.

Our evaluation demonstrates the effectiveness of our method in generating stable grasps within a dual-arm setup, as evidenced by stability metrics, and in simulated environments, highlighting its practical applicability. Furthermore, conceptually our method can be extended to any number

of arms, and also be helpful in multi-arm settings. To summarize, our contributions are:

- 1) We propose CGDF, a method to generate constrained grasps on complex shapes. CGDF is powered by convolutional plane features with the ability to store local geometries efficiently, enabling dense grasp generation on complex shapes.
- 2) We propose a novel part-conditioned generation strategy to generate sample-efficient constrained grasps without explicitly training on conditionally labeled datasets.
- 3) We further show the effectiveness of our approach in the complex setting of dual-arm grasping. We demonstrate that CGDF outperforms existing methods in a dual-arm constrained grasp setting, showcasing the key use case of our method.

II. RELATED WORK

A. Grasp Generation

The research for single-arm grasping dealing with small objects has been well explored [1], [2], [16]–[20]. Existing single-arm grasps methods [1], [2] are trained to map the point cloud of an observed object or scene to a diverse set of grasps and incorporate an evaluator network to handle collisions. The grasp generation method described in [21] employs a differentiable sampling procedure, guided by a multi-task optimization objective. On the other hand, [22]

presents an SE(3) equivariant energy-based model, which is able to train with a very small number of demonstrations. [4] is a 6-DoF grasp pose synthesis approach from 2D/2.5D input based on estimated keypoints. [23] proposed a possible dual-arm grasping strategy and a multi-stage learning method for selective dual-arm grasping using Convolutional Neural Networks (CNN) for grasp point prediction and semantic segmentation on the RGB images. However, these methods focus on generating uniform grasps for the entire object or scene, limiting their use case in constrained grasping.

B. Constrained generative grasp sampling

Constrained grasping methods have previously been explored for both single-arm [11], [24] and dual-arm settings [25]. However, [24], [25] focus on task-oriented grasps, and are thus limited by the tasks they train on. We follow the setting of [11], which focuses on constraining arbitrary regions. It proposes a novel constrained 6-DoF generative grasp sampler, VCGS, trained on the CONG dataset that generates dense grasps on arbitrarily specified target regions. The proposed CONG dataset is an augmentation of the ACRONYM [15] dataset, which contains randomly subsampled target grasping regions along with the ground truth grasps on them. These areas can represent, for instance, semantically meaningful locations on the target object, such as the handle of a cup or the bottle cap, but can also cover the entire object. We extend the task setting of [11] to a dual-arm setting, using the DA^2 dataset [6]. However, unlike VCGS, we do not need a conditionally labeled dataset as our model can directly generate constrained grasps without explicitly training on that objective.

C. 3D Representations in Grasping

Convolutional Occupancy Networks [12] represent a 3D surface implicitly as a decision boundary of a neural network classifier and make a popular representation choice for fine-grained implicit 3D reconstruction [13], [14] tasks due to the translation equivariance and their ability to reconstruct finer details. Neural Descriptor Fields [26] represent point and grasp descriptors as the vector of concatenated activations of a conditional occupancy function modeled by a neural network. SE(3) Diffusion Fields [27] formulate grasp pose generation as a gradient-based inverse diffusion process [28] of an SE(3) Diffusion model. This method essentially samples random grasp poses and moves them to “low-cost” regions that represent good grasping poses.

III. CGDF: CONSTRAINED GRASP DIFFUSION FIELDS

Given a object point cloud $P \in \mathbb{R}^{N \times 3}$ along with a target region (i.e. constraint) $P_t \subseteq P$, our goal is to generate M parallel-jaw grasp poses $H_i \in SE(3), i \in [0, M)$ on P , such that H_i is located on P_t . Here M is a number that can be user-defined. Additionally, we make no assumptions regarding the granularity of the target regions, allowing P_t to encompass anything from a small, localized region to the entire point cloud. For an n-arm setting with large objects, n

target regions can be selected $P_T = \{P_t^1, P_t^2 \dots P_t^n\}$. Consequently, our method must learn a robust shape representation to generalize to arbitrary shapes and complex geometries, ensuring its applicability across various constraint regions. Further, the model should generate diverse grasps with full coverage over the target region. Diffusion models have shown excellent performance in generating diverse outputs while learning the data distribution effectively [27]. Thus we choose a diffusion-based architecture as our backbone.

In this section, we **(A)** first dive into constructing a grasp generation diffusion model. **(B)** We then formulate a strategy to generate sample-efficient grasps constrained regions without explicitly training on conditionally labeled data. In **(C)**, we outline the architecture of our model, enabling CGDF to achieve better grasps on complex shapes, with and without constrained regions.

A. SE(3) Diffusion Model

Our diffusion model architecture uses the formulation presented by SE3Diff [27] for grasp generation. To adapt Euclidean diffusion models to the Lie group SE(3) [29], SE3Diff works in the vector space \mathbb{R}^6 , which is isomorphic to the **Lie algebra** $\mathfrak{se}(3)$. This allows one to apply linear algebra to an element H in the Lie group SE(3). H can be moved between the Lie group and the vector space using the logarithmic and exponential maps: $\text{Logmap} : SE(3) \rightarrow \mathbb{R}^6$ and $\text{Expmap} : \mathbb{R}^6 \rightarrow SE(3)$. Please refer to [27] for more details. A **diffusion model in SE(3)** can then be formulated as a vector field s_θ that returns a vector $v \in \mathbb{R}^6$, given a query pose $H \in SE(3)$, pointcloud P and the current noise scale k [30]. Formally: $v_k = s_\theta(H, k, P)$.

Energy Based Model: Following SE3Diff, instead of directly training a model to learn a vector field s_θ , we construct an energy-based model E_θ as shown in Fig 2(a). This learns a scalar field representing the energy of the grasp distribution, as this allows the model to score the generated grasps. The energy e_k of a grasp pose H , given a point cloud P is defined as:

$$e_k = E_\theta(H, k, P) \quad (1)$$

Following the notation in [29], the derivative of a function mapping $SE(3)$ to \mathbb{R} w.r.t to an SE(3) element is a vector $v \in \mathbb{R}^6$. Since $E_\theta : SE(3) \rightarrow \mathbb{R}$, s_θ can then be defined as the derivate of e_k w.r.t H :

$$s_\theta(H, k, P) = \frac{De_k}{DH} \quad (2)$$

This formulation enables a **scoring mechanism** for the grasp poses based on the energy values, where lower energy corresponds to better grasps.

Forward diffusion Process: As shown in Fig 2(c), we first sample a perturbed data point by first moving the ground-truth grasp pose H from the Lie group to the vector space using the Logmap function. We then add gaussian noise $\epsilon_k \in \mathbb{R}^6$ to it, with standard deviation σ_k for noise scale k , and map it back to SE(3) using the Expmap function as follows:

$$H_k = \text{Expmap}[\text{Logmap}(H) + \epsilon_k], \epsilon_k \sim \mathcal{N}(0, \sigma_k^2 I) \quad (3)$$




Fig. 3: **Part-Guided Diffusion strategy:** Using 2 instances of CGDF, conditioned on full pointcloud and the target region, we show how our part-guided diffusion works.

Loss Function: The diffusion model can now be trained to predict the noise given a perturbed sample by minimizing the L1 loss objective between the sampled perturbation ϵ_k and the predicted perturbation as follows:

$$\mathcal{L}_{\text{diff}} = ||s_{\theta}(H_k, k, P) - \epsilon_k|| \quad (4)$$

Inverse Diffusion Step: Following Se3diff [27], the inverse diffusion step is formulated as an adapted version of the Euclidean Langevin MCMC [31], as shown in Fig 2(c):

$$H_{k-1} = \text{Expmap} \left(-\frac{\alpha_k^2}{2} s_{\theta}(H_k, k, P) + \alpha_k \epsilon \right) H_k \quad (5)$$

where $\epsilon \in \mathbb{R}^6, \epsilon \sim \mathcal{N}(0, I)$, α_k refers to a non-negative step-dependent coefficient.

B. Constrained Grasp Diffusion

The above diffusion model trains on an unconstrained grasping objective and thus generates grasps covering the entire object pointcloud P . However, if the model is able to learn a rich enough shape representation to model local geometries, we can use the model trained on unconstrained grasping to generate constrained grasps without retraining, by using the energy values predicted by E_{θ} as explained below.

Assume we have an object pointcloud P and a target region pointcloud P_t . Since E_{θ} is an EBM, stable grasp poses have lower energies as compared to invalid or colliding grasps. Therefore, if we treat P_t as a stand-alone pointcloud, then a grasp pose H_t on it is stable if its energy at $k = 0$ ($e'' = E_{\theta}(H_t, 0, P_t)$) is low (using Eqn 1). Now, we can also calculate the energy of the H_t giving P as the context instead of P_t , that is, $e' = E_{\theta}(H_t, 0, P)$. Notice that e' is low only if H_t is also a stable grasp on P . Thus, H_t is a **stable constrained grasp** on P if:

$$e' < \delta \ \& \ e'' < \delta$$

where δ is the user-defined energy threshold for valid grasps. This threshold depends on the training conditions of the

model and can only be defined empirically by looking at the range of energy values returned by E_{θ} .

Based on the above insight, we formulate a method to generate sample efficient constrained grasps using a **part-guided diffusion** strategy, illustrated by Fig. 3. We modify Equation 1 as follows:

$$e_k = \max(e'_k, e''_k) \quad (6)$$

where,

$$e'_k = E_{\theta}(H_k, k, P); e''_k = E_{\theta}(H_k, k, P_t)$$

Equation 2 can now be written as:

$$s_{\theta}(H_k, k, \{P, P_t\}) = \frac{D \max(e'_k, e''_k)}{DH_k}$$

$$\therefore s_{\theta}(H_t, k, \{P, P_t\}) = \begin{cases} s_{\theta}(H_k, k, P) & \text{if } e'_k \geq e''_k \\ s_{\theta}(H_k, k, P_t) & \text{if } e'_k < e''_k \end{cases} \quad (7)$$

As shown in Equations 7 and 6, by taking the maximum value of the energies, we essentially guide the grasp from a random pose to a stable configuration near the target region. During the inverse diffusion step, if the grasp H_k is near the constrained region but collides with the full object or is unstable, the energy e'_k is higher, which makes H_k move to a more stable pose. Conversely, if the grasp is stable but is far from the constrained region, then e''_k is higher, and the grasp moves closer to the constrained region. Eventually, the grasp moves to a pose where the energies for both P and P_t are low, i.e. the grasp is on the constrained region as well as stable.

C. Model Architecture

Our model architecture is inspired by [26] and [27], which show that the joint learning of surface reconstruction and grasping generates rich grasp features. However, these methods only use a global point cloud embedding as the conditional vector. This approach works for smaller objects where smaller details do not significantly affect grasp poses but fails to give good results when working with larger objects with complex shapes. Since part-guided diffusion is heavily dependent on the energy values, it requires that the EBM E_{θ} learns a shape representation that is expressive enough to distinguish between arbitrary shapes at a fine-grained level. To achieve this required fidelity, we build on top of the architecture of [27] and incorporate local geometric features inspired from [12], as shown in Fig 2. Given an input pointcloud $P \in \mathbb{R}^{N \times 3}$, we first use a VN-Pointnet [32] combined with a convolutional plane encoder to generate three canonical feature planes. Given a query point x_q , we obtain the point-feature vector z_{x_q} at that point using by sampling from the feature planes using bilinear interpolation. The feature encoder F_{θ} is an MLP that outputs the sdf along with the feature vector at that point as follows:

$$f_{x_q}^k = \{\text{sdf}_{x_q}, \psi_{x_q}^k\} = F_{\theta}(z_{x_q} \oplus x_q, k)$$

where \oplus means concatenation. The grasp pose $H \in SE(3)$ is represented as $P_H = H \cdot P_q$, where $P_q \in \mathbb{R}^{N \times 3}$ is a




Fig. 4: Qualitative results of PartSLIP to generate the constrained region based on given text prompts followed by CGDF to generate grasps on the proposed regions.

fixed query pointcloud. The pose descriptors of H can then be obtained as $f_H^k = \bigoplus_{i=0}^N [f_{x_i}^k], x_i \in P_H$. The decoder is an MLP that takes the grasp pose descriptors and outputs the energy of the grasp, $e_k = D_\theta(f_H^k)$. The model trains on the main objective as defined in Equation 4, along with an auxiliary loss to predict the SDF as mentioned in [27].

IV. EXPERIMENTS

In this section, we evaluate CGDF on Constrained and Unconstrained Generation for complex and large objects in a dual-arm setting. Since grasping large objects with a single arm is impractical, an evaluation from a dual-arm perspective gives us better performance insights.

We sample pointclouds (#points=1000, irrespective of object size) for all objects and use them for constrained and unconstrained grasping. For constrained grasping, two target regions are generated by sampling query points and obtaining $k = 100$ nearest neighbors. These regions vary based on geometry and query point location. Constrained grasps are generated as per Section III, while unconstrained grasps consider the entire object point cloud. For dual-arm grasp generation, we prune out grasps where the 2 grippers are close to each other and calculate dual-arm stability metrics. We also show ablation results showcasing the effects of how each design decision helps our method.

Datasets: We train and evaluate models on grasps and objects from the DA² Dataset [6], consisting of dual-arm grasps on large object meshes from [33]. Unlike the ACRONYM dataset [15], which scales objects to tabletop size, DA² retains original object scales, challenging shape prior learning. To train [11], we augment the DA² dataset with part annotations, as mentioned in [11]Section V. Qualitative results show our model’s generalization using PartSLIP [5] on the PartNetE Dataset [34], [35], where a part is segmented using a text prompt. We then use our model with the part-guided diffusion strategy (trained on DA²) to generate grasps on the region.

Metrics: We evaluate the performance of these models using three metrics: Force Closure, Grasp Success Rate and

Metric	Methods	Unconstrained	Constrained
FC (%)↑	VCGS	3.28	3.96
	SE3Diff	14.22	-
	CGDF (Ours)	43.51	44.8
GSR (%)↑	VCGS	41.01	43.36
	SE3Diff	46.2	-
	CGDF (Ours)	60.3	60.88
TG (%)↑	VCGS	100	76.2
	CGDF (Ours)	100	91.86

TABLE I: Quantitative comparison CGDF with SE3Diff [27] and VCGS [11] in unconstrained and constrained setting. Force Closure (FC) and Grasp Success Rate(GSR) are explained in Sec IV. Target Grasp (TG)% refers to the ratio of grasps on the target regions as explained in Sec IV.

Target Grasps. *Target Grasps (TG)*: Target Grasps measures the percentage of grasps that are within 6cm, which is the assumed gripper finger length of the target region.

Force Closure (FC): We evaluate the stability of the generated dual-arm grasps using the force closure metric as defined in [6], [36], which defines a set of contact forces (4 in our case, 2 for the gripper of each hand) to be force closure if $GG' \succcurlyeq \epsilon I_{6 \times 6}$; and $\|Gc\| < \delta'$, where G is the grasp matrix,, c is the axis of the friction cone, S is the object surface and $x_i \in S$. Please refer to [36] for more details. We transform the gripper mesh to the predicted pose, prune grasps that intersect with the object mesh, and sample rays from the gripper fingers to find intersections (contact points) with the object’s ground truth mesh. The FC metric is the percentage of predicted grasps that satisfy force closure and do not collide. *Grasp Success Rate (GSR)*: We evaluate the grasp success rate in Isaac Gym [37]. Similar to the setting in [11], we use two free-floating *Franka Emika Panda* grippers with extended fingers (6cm), to grasp a free-floating object. We place the two grippers at the specified grasp poses, close them until the object is secured, activate gravity, and lift the grippers. If the object remains between the grippers after lifting, the grasp is considered successful.

Results As shown in Table I, CGDF outperforms the existing methods on all metrics as well as sample efficiency. The existing methods rely only on the global shape representation. This makes them prone to errors where the object includes finer details like multiple thin structures. As seen from the qualitative outputs shown in Fig 5, VCGS-Sampler is able to generate grasps close to the constrained region. However, it fails to model the fine geometry of the region and generates multiple grasps that slightly touch the objects causing collisions and lowering the FC metric drastically. VCGS-Evaluator, which is also trained with a global shape representation, fails to model the accurate shapes well and thus classifies many of the slightly colliding candidate grasps as positive. However, these slight collisions do not have a drastic effect on the grasp success as such grasps cause the object to shift slightly, while still being successful. Even though SE3Diff generates relatively fewer colliding grasps, it still achieves almost the same level of grasp success as VCGS, showcasing the limitations of using global embeddings.

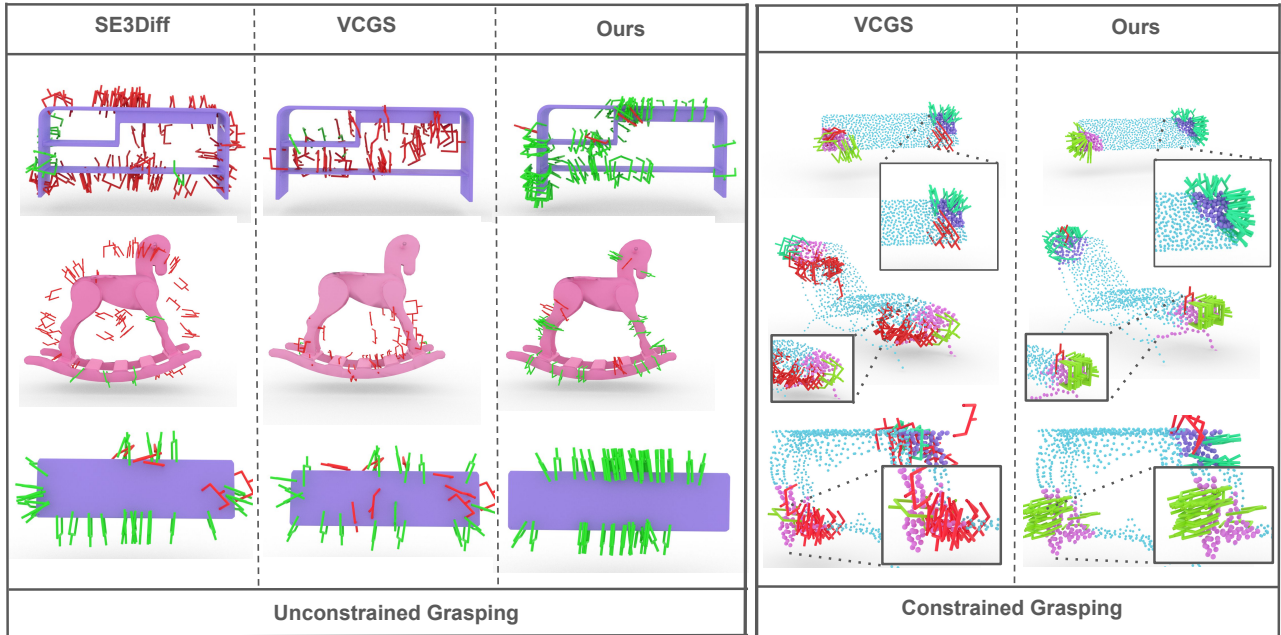


Fig. 5: Qualitative comparison of *unconstrained* and *constrained* grasp generation in a dual-arm setting on CGDF and the baselines: VCGS [11] and SE3Diff [27]. All the baselines perform well for objects with simple shapes (like planar or elongated shapes). However, for relatively complex geometries like chairs, tables, etc. CGDF generates dense, constrained and unconstrained grasps, whereas the baselines struggle to do so. The green grasps are valid, and the red grasps are invalid. **Zoom in for a better experience.**

Ablations: CGDF has 2 key design decisions: (1) Convolutional Plane features and (2) Part-Guided Diffusion. We evaluate the contribution of both of these decisions, which enable our network to improve grasp performance.

(A) *Convolutional Plane features:* Convolutional plane features [12] represent local geometries efficiently by distilling features from existing pointcloud encoders like VN-PointNet, as shown in Figure 2. Since we learn a joint representation for shapes and grasps (by converting them into points) [26], the features of a grasp pose H are similar to the target region of the grasp. Thus, we analyze how well this 3D representation encodes the object geometry. To do this, we compare vanilla SE3Diff [27] with our model. Since both these models generate SDF as an auxiliary task, we analyze the role of convolutional features using Chamfer’s Distance (CD) ¹ between pointclouds sampled from the predicted SDF and the ground truth pointcloud. As shown in Table II, CGDF significantly improves the reconstruction quality compared to SE3Diff, showcasing the effectiveness of Convolutional plane features.

(B) *Part Conditioned Diffusion:* It is possible to generate grasps on constrained regions with CGDF without using the part-guided strategy, by employing a naive albeit sample-inefficient approach mentioned in Section III-B, which is to simply generate grasps on the target region pointcloud P_t and threshold grasps based on energy values. In this experiment, we try to study the effect of part-guided diffusion over this naive method in terms of sample efficiency. For this, we measure the percentage of grasps left after the

Modification	Metric	Ours	Modified
w/o Conv	CD ↓	14.04	60.21
w/o PD	SE (%)↑	93.6	37.4

TABLE II: Ablation study: We show the contribution of different components of our model. In w/o Conv, we evaluate the quality shape representations of SE3Diff [27] and CGDF by comparing the reconstructions of both models. Chamfer’s Distance (CD) is scaled by 10^4 . SE (sample efficiency) refers to the percentage of grasps remaining after the energy thresholding step.

thresholding step for grasps generated with and without part-guided diffusion (w/o PD). Since the bounds of energy are not fixed [28], we set the energy threshold δ such that the FC metric of the grasps from both the models is in a similar range as in Table I. As shown in Table II, part-guided diffusion generates **2.5 times** more valid grasps than the baseline (w/o PD).

V. CONCLUSIONS

Existing data-driven grasp generation methods primarily focus on uniformly generating stable and collision-free grasps across target objects, limiting their effectiveness in scenarios requiring constrained grasps on complex geometries. To address this limitation, we introduce CGDF, a diffusion-based grasp generative model. CGDF utilizes a part-guided diffusion approach, enabling high sample efficiency in constrained grasping without the need for extensive constraint-augmented datasets. Through qualitative and quantitative evaluations in both unconstrained and constrained settings, we demonstrate CGDF’s ability to generate stable grasps on complex objects for dual-arm manipulation, surpassing existing methods. Future work could involve possible extensions to scenarios involving more than two

¹<https://pdal.io/en/stable/apps/chamfer.html>

arms, further improving the representation, and potentially expanding CGDF's application to a broader range of robotic manipulation tasks.

REFERENCES

- [1] A. Mousavian, C. Eppner, and D. Fox, "6-dof graspnet: Variational grasp generation for object manipulation," in *Proceedings of the IEEE/CVF International Conference on Computer Vision (ICCV)*, October 2019.
- [2] M. Sundermeyer, A. Mousavian, R. Triebel, and D. Fox, "Contact-graspnet: Efficient 6-dof grasp generation in cluttered scenes," in *2021 IEEE International Conference on Robotics and Automation (ICRA)*, 2021.
- [3] J. Mahler, J. Liang, S. Niyaz, M. Laskey, R. Doan, X. Liu, J. Aparicio, and K. Goldberg, "Dex-Net 2.0: Deep Learning to Plan Robust Grasps with Synthetic Point Clouds and Analytic Grasp Metrics," in *Robotics: Science and Systems XIII*. Robotics: Science and Systems Foundation, Jul. 2017.
- [4] Y. Chen, R. Xu, Y. Lin, H. Chen, and P. A. Vela, "Kgnv2: Separating scale and pose prediction for keypoint-based 6-dof grasp synthesis on rgb-d input," *IEEE International Conference on Intelligent Robots and Systems (IROS)*, 2023.
- [5] M. Liu, Y. Zhu, H. Cai, S. Han, Z. Ling, F. Porikli, and H. Su, "Partslip: Low-shot part segmentation for 3d point clouds via pretrained image-language models," in *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 2023, pp. 21 736–21 746.
- [6] G. Zhai, Y. Zheng, Z. Xu, X. Kong, Y. Liu, B. Busam, Y. Ren, N. Navab, and Z. Zhang, "Da² dataset: Toward dexterity-aware dual-arm grasping," *IEEE Robotics and Automation Letters*, vol. 7, no. 4, pp. 8941–8948, 2022.
- [7] R. Mirjalili, M. Krawez, S. Silenzi, Y. Blei, and W. Burgard, "Langrasp: Using large language models for semantic object grasping," *arXiv preprint arXiv:2310.05239*, 2023.
- [8] A. Rashid, S. Sharma, C. M. Kim, J. Kerr, L. Y. Chen, A. Kanazawa, and K. Goldberg, "Language embedded radiance fields for zero-shot task-oriented grasping," in *7th Annual Conference on Robot Learning*, 2023.
- [9] H.-S. Fang, C. Wang, M. Gou, and C. Lu, "Graspnet-1billion: A large-scale benchmark for general object grasping," in *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, June 2020.
- [10] A. Miller and P. Allen, "Graspit! a versatile simulator for robotic grasping," *IEEE Robotics and Automation Magazine*, vol. 11, no. 4, pp. 110–122, 2004.
- [11] J. Lundell, F. Verdoja, T. N. Le, A. Mousavian, D. Fox, and V. Kyrki, "Constrained generative sampling of 6-dof grasps," in *2023 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, 2023, pp. 2940–2946.
- [12] S. Peng, M. Niemeyer, L. Mescheder, M. Pollefeys, and A. Geiger, "Convolutional occupancy networks," in *Computer Vision—ECCV 2020: 16th European Conference, Glasgow, UK, August 23–28, 2020, Proceedings, Part III 16*. Springer, 2020, pp. 523–540.
- [13] J. J. Park, P. Florence, J. Straub, R. Newcombe, and S. Lovegrove, "DeepSDF: Learning continuous signed distance functions for shape representation," in *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, June 2019.
- [14] L. Mescheder, M. Oechsle, M. Niemeyer, S. Nowozin, and A. Geiger, "Occupancy networks: Learning 3d reconstruction in function space," in *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, 2019, pp. 4460–4470.
- [15] C. Eppner, A. Mousavian, and D. Fox, "Acronym: A large-scale grasp dataset based on simulation," in *2021 IEEE International Conference on Robotics and Automation (ICRA)*. IEEE, 2021, pp. 6222–6227.
- [16] I. Lenz, H. Lee, and A. Saxena, "Deep learning for detecting robotic grasps," *The International Journal of Robotics Research*, vol. 34, no. 4-5, pp. 705–724, 2015.
- [17] S. Levine, P. Pastor, A. Krizhevsky, J. Ibarz, and D. Quillen, "Learning hand-eye coordination for robotic grasping with deep learning and large-scale data collection," *The International journal of robotics research*, vol. 37, no. 4-5, pp. 421–436, 2018.
- [18] J. Mahler, J. Liang, S. Niyaz, M. Laskey, R. Doan, X. Liu, J. A. Ojea, and K. Goldberg, "Dex-net 2.0: Deep learning to plan robust grasps with synthetic point clouds and analytic grasp metrics," *arXiv preprint arXiv:1703.09312*, 2017.
- [19] D. Morrison, P. Corke, and J. Leitner, "Egad! an evolved grasping analysis dataset for diversity and reproducibility in robotic manipulation," *IEEE Robotics and Automation Letters*, vol. 5, no. 3, pp. 4368–4375, 2020.
- [20] H. Zhang, X. Lan, S. Bai, X. Zhou, Z. Tian, and N. Zheng, "Roi-based robotic grasp detection for object overlapping scenes," in *2019 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*. IEEE, 2019, pp. 4768–4775.
- [21] A. Alliegro, M. Rudorfer, F. Frattin, A. Leonardis, and T. Tommasi, "End-to-end learning to grasp via sampling from object point clouds," *IEEE Robotics and Automation Letters*, vol. 7, no. 4, pp. 9865–9872, 2022.
- [22] H. Ryu, H. in Lee, J.-H. Lee, and J. Choi, "Equivariant descriptor fields: Se(3)-equivariant energy-based models for end-to-end visual robotic manipulation learning," in *The Eleventh International Conference on Learning Representations*, 2023. [Online]. Available: <https://openreview.net/forum?id=dnjZSPGmY5O>
- [23] S. Kitagawa, K. Wada, S. Hasegawa, K. Okada, and M. Inaba, "Multi-stage learning of selective dual-arm grasping based on obtaining and pruning grasping points through the robot experience in the real world," in *2018 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, 2018, pp. 7123–7130.
- [24] A. Murali, W. Liu, K. Marino, S. Chernova, and A. Gupta, "Same object, different grasps: Data and semantic knowledge for task-oriented grasping," in *Conference on robot learning*. PMLR, 2021, pp. 1540–1557.
- [25] Y. Zhao, R. Wu, Z. Chen, Y. Zhang, Q. Fan, K. Mo, and H. Dong, "Dualafford: Learning collaborative visual affordance for dual-gripper manipulation," in *International Conference on Learning Representations*, 2023.
- [26] A. Simeonov, Y. Du, A. Tagliasacchi, J. Tenenbaum, A. Rodriguez, P. Agrawal, and V. Sitzmann, "Neural descriptor fields: Se(3)-equivariant object representations for manipulation," *2022 International Conference on Robotics and Automation (ICRA)*, vol. null, pp. 6394–6400, 2021.
- [27] J. Urain, N. Funk, J. Peters, and G. Chalvatzaki, "Se(3)-diffusionfields: Learning smooth cost functions for joint grasp and motion optimization through diffusion," *2023 IEEE International Conference on Robotics and Automation (ICRA)*, vol. null, pp. 5923–5930, 2022.
- [28] Y. Song and D. P. Kingma, "How to train your energy-based models," *arXiv preprint arXiv:2101.03288*, 2021.
- [29] J. Sola, J. Deray, and D. Atchuthan, "A micro lie theory for state estimation in robotics," *arXiv preprint arXiv:1812.01537*, 2018.
- [30] Y. Song and S. Ermon, "Generative modeling by estimating gradients of the data distribution," *Advances in neural information processing systems*, vol. 32, 2019.
- [31] R. M. Neal *et al.*, "Mcmc using hamiltonian dynamics," *Handbook of markov chain monte carlo*, vol. 2, no. 11, p. 2, 2011.
- [32] C. Deng, O. Litany, Y. Duan, A. Poulencard, A. Tagliasacchi, and L. J. Guibas, "Vector neurons: A general framework for so (3)-equivariant networks," in *Proceedings of the IEEE/CVF International Conference on Computer Vision*, 2021, pp. 12 200–12 209.
- [33] A. X. Chang, T. Funkhouser, L. Guibas, P. Hanrahan, Q. Huang, Z. Li, S. Savarese, M. Savva, S. Song, H. Su *et al.*, "Shapenet: An information-rich 3d model repository," *arXiv preprint arXiv:1512.03012*, 2015.
- [34] K. Mo, S. Zhu, A. X. Chang, L. Yi, S. Tripathi, L. J. Guibas, and H. Su, "PartNet: A large-scale benchmark for fine-grained and hierarchical part-level 3D object understanding," in *The IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, June 2019.
- [35] F. Xiang, Y. Qin, K. Mo, Y. Xia, H. Zhu, F. Liu, M. Liu, H. Jiang, Y. Yuan, H. Wang, L. Yi, A. X. Chang, L. J. Guibas, and H. Su, "SAPIEN: A simulated part-based interactive environment," in *The IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, June 2020.
- [36] T. Liu, Z. Liu, Z. Jiao, Y. Zhu, and S.-C. Zhu, "Synthesizing diverse and physically stable grasps with arbitrary hand structures using differentiable force closure estimator," *IEEE Robotics and Automation Letters*, vol. 7, no. 1, pp. 470–477, 2021.
- [37] V. Makoviychuk, L. Wawrzyniak, Y. Guo, M. Lu, K. Storey, M. Macklin, D. Hoeller, N. Rudin, A. Allshire, A. Handa *et al.*, "Isaac gym: High performance gpu-based physics simulation for robot learning," *arXiv preprint arXiv:2108.10470*, 2021.