

## Language-driven Grasp Detection

An Dinh Vuong<sup>1</sup>, Minh Nhat Vu<sup>2,\*</sup>, Baoru Huang<sup>3,\*</sup>, Nghia Nguyen<sup>1</sup>, Hieu Le<sup>1</sup>, Thieu Vo<sup>4</sup>, Anh Nguyen<sup>5</sup>

<sup>1</sup>FPT Software AI Center, Vietnam <sup>2</sup>Automation & Control Institute, TU Wien, Austria <sup>3</sup>Imperial College London, UK

<sup>4</sup>Ton Duc Thang University, Vietnam <sup>5</sup>University of Liverpool, UK \*Co-Corresponding authors

<https://airvlab.github.io/grasp-anything>



Figure 1. We present a new dataset and method for *language-driven grasp* task.

### Abstract

*Grasp detection is a persistent and intricate challenge with various industrial applications. Recently, many methods and datasets have been proposed to tackle the grasp detection problem. However, most of them do not consider using natural language as a condition to detect the grasp poses. In this paper, we introduce Grasp-Anything++, a new language-driven grasp detection dataset featuring 1M samples, over 3M objects, and upwards of 10M grasping instructions. We utilize foundation models to create a large-scale scene corpus with corresponding images and grasp prompts. We approach the language-driven grasp detection task as a conditional generation problem. Drawing on the success of diffusion models in generative tasks and given that language plays a vital role in this task, we propose a new language-driven grasp detection method based on diffusion models. Our key contribution is the contrastive training objective, which explicitly contributes to the denoising process to detect the grasp pose given the language instructions. We illustrate that our approach is theoretically supportive. The intensive experiments show that our method outperforms state-of-the-art approaches and allows real-world robotic grasping. Finally, we demonstrate our large-scale dataset enables zero-shot grasp detection and is a challenging benchmark for future work.*

### 1. Introduction

Imagine we want an assistant robot to grasp a cup among a clutter of daily objects such as a knife, a fork, a cup, and a pair of scissors. Conventionally, to convey the idea of grasping this specific object, humans use the natural language command, “give me the cup”, for instance. Although humans intuitively know how to grasp the cup given the linguistic command, determining specific grasp actions for objects based on natural language instructions or *language-driven grasp detection* remains challenging for robots [64]. First, natural language is usually overlooked in existing grasp datasets [56] while training vision-and-language neural networks necessitates an excessive number of labeled examples [66]. Second, recent works usually focus on particular manipulation tasks with limited objects [25], imposing a bottleneck for in-the-wild robot execution [63]. Finally, despite recent developments, bridging the gap between language, vision, and control for real-world robotic experiments remains a challenging task [83].

Recently, language-driven robotic frameworks are gaining traction, offering the potential for robots to process natural language, and bridging the gap between robotic manipulations and real-world human-robot interaction [51]. PaLM-E [21], EgoCOT [51], and ConceptFusion [32] are some notable embodied robots with the ability to compre-

hend natural language by harnessing the power of large foundation models such as ChatGPT [54]. However, most works assume the high-level actions of robots and ignore the fundamental grasping actions, restricting the structure for generalization across robotic domains, tasks, and skills [50]. In this paper, we explore training a language-driven agent to implement low-level actions, focusing on the task of object grasping via image observations. Specifically, our hypothesis is centered around the establishment of a robotic system that can execute grasping actions following a given language instruction for any universal object.

We first present **Grasp-Anything++** to serve as a large-scale dataset for language-driven grasp detection. Our dataset is based on the Grasp-Anything [74], and is synthesized from foundation models. Compared to the original Grasp-Anything dataset, we provide more than 10M grasp prompts and 3M associated object masks, 6M ground truth poses at the object part level. Our dataset showcases the ability to facilitate grasp detection using language instructions. We label the ground truth at both the object level and part level, providing a comprehensive understanding of real-world scenarios. For example, our ground truth includes both general instructions “*give me the knife*” and detail ones such as “*grasp the handle of the steak knife*”. We empirically show that our large-scale dataset successfully facilitates zero-shot grasp detection on both vision-based tasks and real-world robotic experiments.

To tackle the challenging language-driven grasp detection task, we propose a new diffusion model-based method. Our selection of diffusion models is motivated by their proven efficacy in conditional generation tasks [28]. These models have shown efficiency beyond image synthesis, including other image-based tasks such as image segmentation [77], and visual grounding [40]. Despite achieving notable success, integrating visual and text features effectively remains a challenge [13] as the majority of existing literature employs latent strategies to combine visual and text features [36]. We address this challenge by employing a new training strategy for learning text and image features, focusing on the use of feature maps as guidance information for grasp pose generation. Our main contribution is a new training objective that incorporates the feature maps and explicitly contributes to the denoising process. In summary, our contributions are three-fold:

- We propose Grasp-Anything++, a large-scale language-driven dataset for grasp detection tasks.
- We propose a diffusion model with a training objective that explicitly contributes to the denoising process to detect the grasp poses.
- We demonstrate that our Grasp-Anything++ dataset and the proposed method outperform other approaches and enable successful robotic applications.

## 2. Related Work

**Grasp Detection.** Grasp detection is a popular task in both computer vision and robotic community [2, 17, 24, 64]. Recently, establishing robotic systems with the ability to follow natural commands has been actively researched [64, 80, 83]. The prevalent solution to the language-driven grasp detection task is to split into two stages: one for grounding the target object, and the other is to synthesize grasp poses from the grounding visual-text correlations [1, 80]. Training in two stages may result in longer inference time [41]. In addition, several works [82, 83, 90] adopt foundation models, such as GroundDINO [42] and GPT-3 [7]. Accessing such commercial foundation models is not always available [71], especially on robotic systems with limited resources or unstable internet connection [37]. In our work, we directly train the model on the large-scale Grasp-Anything++ dataset to inherit the power of a foundation-based dataset, while ensuring a straightforward inference process for the downstream robotic applications.

**Language-driven Grasp Detection Datasets.** While there are many grasp datasets have been introduced [9, 17, 22–24, 33, 38, 45, 48, 49, 55, 78], the majority of them overlook the text modality. Therefore, the grasping of objects out of a clutter typically experiences ambiguities in what object to grasp [88]. DailyGrasp [83] is one of the first grasp datasets employing natural language for scene descriptions; however, the scene description corpus in this dataset is relatively small and does not specify which part of the object should be grasped. In our work, we present Grasp-Anything++, which is a large-scale language-driven grasping dataset. Grasp-Anything++ provides object instance description with a natural word phrase, for example, “*a blue ceramic vase*”. Furthermore, Grasp-Anything++ describes the grasping object at both the part level and object level, providing more information for the robot to execute the grasping [86].

**Diffusion Models for Robotic Applications.** Diffusion models [28] have emerged as the new state-of-the-art method of generative tasks [84]. Recent have witnessed growing attention for utilizing diffusion models in robotic applications [61]. Liu *et al.* [43] propose a diffusion model to handle the language-guided object rearrangement task. Diffusion models are also applied to other robotic tasks such as motion planning [10], and trajectory optimization [31]. The authors in [70] present a diffusion model to determine grasp poses by minimizing a SDF loss. Overall, the diffusion models employed in previous works often combine visual and text features in a latent mechanism [6], which may cause interpretability problems [83] for robotic systems that require low-level controls [20]. To tackle this challenge, we propose a training objective that explicitly contributes to the denoising process. We demonstrate that our proposed strategy is more effective than the latent strategy.





| Step             |                 | Description  |  |
|------------------|-----------------|--|--|
| Scene Generation | User            | Please help me generate scene descriptions for natural arrangements of daily objects. Each description has the following form: <Object.1><Object.2>...<Verb><Container_Object>. Please also ensure the incorporation of a rich and varied lexicon in the scene descriptions.   |   |
|                  | Sample          | A steel knife, a polished fork and a pristine ceramic plate on a wooden table.   |  |
|                  | Text-to-Image   | We use Stable Diffusion [60] to proceed text-to-image generation.  |  |
| Object Masking   | User            | For object part-level description, given an input list {<Object.1>, <Object.2>, ...}, the output will be a list that describes the parts of objects as: {<Object.1>: [<Part.1.1>, <Part.1.2>, ...], <Object.2>: [<Part.2.1>, <Part.2.2>, ...]}.  |   |
|                  | Sample          | {knife: [handle, blade], fork: [handle, neck, stem, tines], plate: [rim, base]}  |  |
|                  | Post Process    | We use OFA [75] and SAM [34] to locate the region describing the objects.  |  |
| Part Masking     | User            | Given the object list and part lists of each scene description, you will generate for me all prompts with the following format: {<Manipulation_Action><Object.ID><Part.ID>}. The part that is more suitable for human grasping is positioned at the start of the list to represent the grasping actions.                       |   |
|                  | Sample          | Give me the steel knife; Grasp the knife at its handle.  |  |
|                  | Post Process    | We leverage VLPART [67] to locate the region describing the parts of objects.  |  |
| Grasp Generation | User            | Generate for me a scene description with grasp instructions following the templates.   |  |
|                  | Sample          | <b>Scene description:</b> A steel knife, a polished fork and a pristine ceramic plate on a wooden table. <b>Object list:</b> {knife, fork, plate}. <b>Part lists:</b> {knife: [handle, blade], fork: [handle, neck, stem, tines], plate: [rim, base]}. <b>Prompts:</b> Give me the steel knife; Grasp the knife at its handle. |  |
|                  | Grasp Labelling | We utilize a pretrained RAGT-3/3 [9] to generate grasp poses corresponding to the located region.  |  |

Table 1. **Grasp-Anything++ creation pipeline.** We utilize ChatGPT to generate scene descriptions and grasp instructions from the user input. We generate images given scene descriptions and automatically synthesize the grasp poses.

### 3. The Grasp-Anything++ Dataset

We utilize large-scale foundation models to create the Grasp-Anything++. Our dataset offers open-vocabulary grasping commands and images with associated groundtruth. There are three key steps in establishing our dataset: *i*) prompting procedure, *ii*) image synthesis and grasp poses annotation, and *iii*) post-processing.

#### 3.1. Prompting Procedure

We first establish prompt-based procedures to generate a large-scale scene description corpus as well as grasp prompt instructions. In particular, we utilize ChatGPT to generate the prompts for two tasks: *i*) Scene descriptions: Sentences capturing the scene arrangement, including the extracted object and part lists, and *ii*) Grasp instructions: Prompts directing the robot to grasp specific objects or parts.

We follow a procedure in Table 1 to implement ChatGPT’s output templates. The reference target in the grasp instruction may be either an object or an object’s part. When the reference is an object’s part, that part is directly selected as the reference in the grasp instruction sentence. If the reference is an object, we determine the grasping region on the part of the object that is likely to be grasped in everyday scenarios as described in affordance theory [53].

#### 3.2. Image Synthesis and Grasp Annotation

**Image Synthesis.** Given the scene description corpus, we first utilize a large-scale pretrained text-to-image model, namely, Stable Diffusion [60] to generate images from scene descriptions. Next, we perform a series of visual grounding and image segmentation using OFA [75], Segment-Anything [34], and VLPART [67] to locate the referenced object or part to the grasp instruction.

**Grasp Annotation.** Grasp poses are represented as 2D rectangles, consistent with prior research and practical compatibility with real-world parallel grippers [17, 33]. Utilizing a pretrained network [9], we annotate grasp poses based on part segmentation masks. Since potential inaccuracies in these candidate poses could occur, we follow the procedure as defined in [74] to evaluate the quality of generated grasp poses to discard unreasonable grasp poses.

Specifically, grasp quality is evaluated through net torque  $\mathcal{T} = (\tau_1 + \tau_2) - RMg$ , where resistance at contact points is  $\tau_i = K\mu_s F \cos \alpha_i$ . With constants such as  $M$  (mass),  $g$  (gravitational acceleration),  $K$  (geometrical characteristics),  $\mu_s$  (static friction coefficient), and  $F$  (applied force), accurately determining  $\mathcal{T}$  directly is challenging due to the physical difficulties in precisely measuring  $M$ ,  $K$ , and  $\mu_s$ . Thus, we employ a surrogate measure,




Figure 2. **Data statistics.** We analyze object categories and the use of natural language in the Grasp-Anything++ dataset.

$$\tilde{T} = \frac{\cos \alpha_1 + \cos \alpha_2}{R},$$

as an alternative. As a result in [12], antipodal grasps score higher on  $\tilde{T}$ , indicating better quality. Consequently, grasps are evaluated based on  $\tilde{T}$ , with positive values indicating positive grasps and others considered as negative.

**Post Processing.** Despite training on extensive datasets, Stable Diffusion [60] may produce subpar content, commonly termed as hallucination [30] when generating images from the text prompts. To address this, we perform manual reviews to filter out such images, with qualitative examples in our figures. Our process includes checks at every stage to prevent duplicate or hallucinated content. However, manual inspection introduces biases, which we counter with guidelines focusing on abnormal structures or implausible gravity (Fig. 3), aligns with approach in the literature [62].



Figure 3. **Failure image generation cases.** Images generated by Stable Diffusion [60] may exhibit hallucinatory artifacts such as sunglasses lacking a lens, scissors with an anomalous structure, and a spoon not resting properly on a table.

Additionally, ChatGPT-generated scene prompts often duplicate [44]. To address this, we use duplication checking, filtering out identical prompts with BERTScore [89], which assesses sentence similarity through cosine similarities of token embeddings. We remove sentences with a BERTScore above 0.85 as in prior study [76].

### 3.3. Data Statistics

**Number of Categories.** To evaluate the object category diversity, we apply a methodology akin to that in [16]. Utilizing 300 categories from LVIS dataset [26], we employ a pretrained model [91] to identify 300 candidate objects from our dataset for each category. We then curate a subset comprising 90,000 objects, refining it by excluding items that do not align semantically with their designated categories. A category is considered significant if it has more

than 40 objects. Fig. 2a shows the results. Overall, our dataset spans over 236 categories from LVIS dataset, indicating a notable degree of object diversity in our dataset.

**Scene Descriptions.** Fig. 2b shows the distribution of scene descriptions based on sentence length. The analysis reveals a wide range of sentence lengths, spanning from 10 words to 100 words per sentence. On average, each scene description consists of approximately 54 words, indicative of detailed and descriptive sentences. These scene descriptions correspond to sets of grasp instructions. Fig. 2c further shows the objects and object parts in scene descriptions.

**Diversity Analysis.** We assess the diversity of occlusion and lighting conditions in the dataset. Regarding occlusion, we use a pretrained YOLOv5 model to identify objects within images. The results indicate that 93.8% of images have a substantial overlap of five or more bounding boxes, which suggests a diverse range of occlusion within the Grasp-Anything++ dataset. Regarding lighting conditions, we convert images to YCbCr to analyze Y channel (luminance) and find that GraspLIM has the most diverse lighting conditions, identifying by the lowest Gini coefficient (a metric to measure the inequality of a distribution) of 0.26, compared to VMRD [87] (0.31), OCID-grasp [2] (0.32), Cornell [33] (0.62), Jacquard [17] (0.91).

## 4. Language-driven Grasp Detection

**Motivation.** The use of diffusion model for language-driven grasp detection is motivated by its efficiency in various generative tasks [15, 28, 43, 84]. Conditional generation, such as our language-driven grasp detection task, aligns seamlessly with diffusion models’ capabilities [29]. Moreover, language-driven grasp detection represents a fine-grained problem in which the outputs strongly depend on the text input [4]. For example, “grasp the steak knife” and “grasp the kraft knife” refer to two different objects on the image. To this end, we propose using *contrastive loss with diffusion model* to tackle this task, as contrastive learning is a popular solution for fine-grained tasks [8, 14, 85].

### 4.1. Contrastive Loss for Diffusion Model

We represent the target grasp pose as  $\mathbf{x}_0$  in the diffusion model. The objective of our diffusion process of language-driven grasp detection involves denoising from a noisy state

$\mathbf{x}_T$  to the original grasp pose  $\mathbf{x}_0$ , conditioned on the input image and grasp instruction represented by  $y$ .

In a diffusion process [28], assume that  $q(\mathbf{x}_{1:T}|\mathbf{x}_0)$  is the forward process and we parameterize the reverse process by  $p_\theta(\mathbf{x}_{0:T})$ . The conditional diffusion process [19] assumes  $\hat{q}$  is the forward process but with the inclusion of a condition  $y$ . The goal of the reverse process is to optimize the variational bound on negative log likelihood [28]

$$\mathcal{L} = \mathbb{E} \left[ -\log p_\theta(\mathbf{x}_T) - \sum_{t \geq 1} \log \frac{p_\theta(\mathbf{x}_{t-1}|\mathbf{x}_t)}{q(\mathbf{x}_t|\mathbf{x}_{t-1})} \right]. \quad (1)$$

We prove in the Appendix that

$$\mathcal{L} = \mathbb{E} \left[ \underbrace{C - \log p_\theta(\mathbf{x}_T)}_{\text{Constant}} + \underbrace{\log \hat{q}(\mathbf{x}_0|\mathbf{x}_T, y)}_{\text{Contrastive}} + \underbrace{\sum_{t > 1} D_{\text{KL}}(\hat{q}(\mathbf{x}_{t-1}|\mathbf{x}_t, \mathbf{x}_0, y) \| p_\theta(\mathbf{x}_{t-1}|\mathbf{x}_t)) - \log p_\theta(\mathbf{x}_0|\mathbf{x}_1, y)}_{\text{Denoising score}} \right]. \quad (2)$$

The terms  $D_{\text{KL}}(\hat{q}(\mathbf{x}_{t-1}|\mathbf{x}_t, \mathbf{x}_0, y) \| p_\theta(\mathbf{x}_{t-1}|\mathbf{x}_t))$  and  $\log p_\theta(\mathbf{x}_0|\mathbf{x}_1, y)$  in Equation 2 are similar to the concept of denoising score used in [28]. Thus, we can represent the quantity  $\sum_{t > 1} D_{\text{KL}}(\hat{q}(\mathbf{x}_{t-1}|\mathbf{x}_t, \mathbf{x}_0, y) \| p_\theta(\mathbf{x}_{t-1}|\mathbf{x}_t)) - \log p_\theta(\mathbf{x}_0|\mathbf{x}_1, y)$  by the loss utilized in [68, 69]

$$\mathcal{L}_{\text{diffusion}} = \mathbb{E}_{\mathbf{x}_0 \sim q(\mathbf{x}_0|y), t \sim [1, T]} [\mathbf{x}_0 - f(\mathbf{x}_{t+1}, t+1, \tilde{\mathbf{x}}_0)]^2. \quad (3)$$

Since  $q(\mathbf{x}_T|\cdot)$  is equivalent to an isotropic Gaussian distribution as  $T \rightarrow +\infty$ , the estimation quantity  $\log p_\theta(\mathbf{x}_T)$  converges to a constant when  $\theta \rightarrow \theta^*$ . Therefore, we can ignore the first term of Equation 2.

Finally, the term  $\hat{q}(\mathbf{x}_0|\mathbf{x}_T, y)$  of Equation 2 provides more information about the relation between  $\mathbf{x}_T$  and  $\mathbf{x}_0$ . As this term is intractable [65], we parameterize  $\hat{q}(\mathbf{x}_0|\mathbf{x}_T, y)$  as  $p_{\psi, y}(\mathbf{x}_0, \mathbf{x}_T)$ . This estimation resembles the noise-contrastive estimation [27], where the  $\mathbf{x}_T, \mathbf{x}_0$  can be considered as a pair of contrastive estimation and  $\psi$  can be estimated by a contrastive loss.

In [19, 68, 73], the authors indicate that predicting  $\mathbf{x}_0$  is often infeasible but predicting an estimation  $\tilde{\mathbf{x}}_0$  is tractable and can be used as a ‘pseudo’ estimation of  $\mathbf{x}_0$ . We denote  $\tilde{\mathbf{x}}_0$  as an estimation of  $\mathbf{x}_0$ . The loss term  $\hat{q}(\mathbf{x}_0|\mathbf{x}_T, y)$  can be approximated by using the following contrastive loss

$$\mathcal{L}_{\text{contrastive}} = \max \left( 0, \left\| \frac{\sqrt{\alpha_T} \tilde{\mathbf{x}}_0 - \mathbf{x}_T}{\sqrt{1 - \alpha_T}} \right\|_2^2 - M \right), \quad (4)$$

where  $M$  is the number of dimension of  $\mathbf{x}_0$ , and  $\alpha_t$  is the variance schedule at timestep  $t$  ( $t = \bar{1}; \bar{T}$ ).

**Proposition 1.** Suppose that  $\tilde{\mathbf{x}}_0, \mathbf{x}_0$  and  $\epsilon$  are independent, and that

$$\left\| \frac{\sqrt{\alpha_T} \tilde{\mathbf{x}}_0 - \mathbf{x}_T}{\sqrt{1 - \alpha_T}} \right\|_2^2 \geq M.$$

Then there exists  $C > 0$  such that: for arbitrary  $\delta > 0$ , if  $\mathcal{L}_{\text{contrastive}} < \delta$ , then

$$\mathbb{E} [\|\tilde{\mathbf{x}}_0 - \mathbf{x}_0\|_2^2] < C\delta.$$

*Proof.* See Supplementary Material.  $\square$

**Remark 1.1.** Proposition 1 suggests that if the contrastive loss  $\mathcal{L}_{\text{contrastive}}$  tends to zero, then the prediction  $\tilde{\mathbf{x}}_0$  will approach the ground truth  $\mathbf{x}_0$ .

**Remark 1.2.** The condition  $\left\| \frac{\sqrt{\alpha_T} \tilde{\mathbf{x}}_0 - \mathbf{x}_T}{\sqrt{1 - \alpha_T}} \right\|_2^2 \geq M$  is suitable for our language-driven grasp detection task as  $\tilde{\mathbf{x}}_0$  and  $\mathbf{x}_T$  are two contrastive quantities, therefore, we can assume there is a minimum distance between  $\tilde{\mathbf{x}}_0$  and  $\mathbf{x}_T$ . In addition, in the proof of Proposition 1, we see that  $\mathbb{E} \left[ \left\| \frac{\sqrt{\alpha_T} \tilde{\mathbf{x}}_0 - \mathbf{x}_T}{\sqrt{1 - \alpha_T}} \right\|_2^2 - M \right] = \beta^2 \mathbb{E} [\|\tilde{\mathbf{x}}_0 - \mathbf{x}_0\|_2^2]$ , which is always nonnegative. Therefore, it is both theoretically and experimentally reasonable to add this assumption.

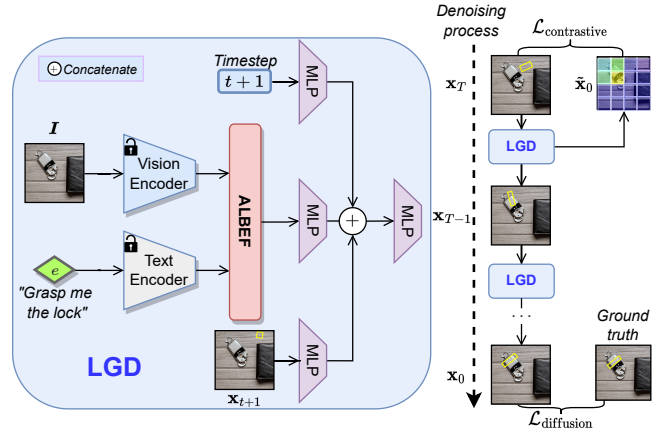


Figure 4. **Language-drive Grasp Detection (LGD) network.** We present the network architecture (left) and the proposed training objectives of the denoising process (right).

## 4.2. Language-driven Grasp Detection Network

**Network.** Our network operates on two conditions: an image denoted as  $I$  and a corresponding text prompt represented as  $e$ . To process these conditions, we employ a vision encoder to extract visual features from  $I$  and a text encoder to derive textual embeddings from  $e$ . The resulting feature vectors, denoted as  $I'$  and  $e'$ , are subsequently subjected to a fusion module, ALBEF [39]. We leverage the attention mask generated by the ALBEF fusion module




Figure 5. Language-driven grasp detection results visualization.

as the estimation  $\tilde{x}_0$  of  $x_0$ . Next, we aggregate three elements: the estimation region  $\tilde{x}_0$ , the grasp pose at the current timestep  $x_{t+1}$ , and the timestep  $t + 1$ . These inputs are combined using MLP layers, similar to the approach outlined in [68]. Specifically, the output operation can be expressed as:  $x_t = f(x_{t+1}, t + 1, \tilde{x}_0)$ , where function  $f$  encompasses a composition of multiple MLP layers. Additional specifics regarding these universal MLP layers are provided in the Supplementary Material.

**Training Objective.** In our context, conditioned grasp detection models the distribution  $p(x_0|y)$  as the reversed diffusion process of gradually cleaning  $x_{t+1}$ . Instead of predicting  $x_t$  as formulated by [28], we follow Ramesh *et al.* [58] and predict the signal itself, i.e.,  $x_t = f(x_{t+1}, t + 1, \tilde{x}_0)$  with the simple objective [28]. To this end, we utilize the contrastive loss as in Equation 4 to explicitly improve the learning objective of the denoising process:

$$\mathcal{L}_{\text{total}} = \mathcal{L}_{\text{contrastive}} + \mathcal{L}_{\text{diffusion}}. \quad (5)$$

## 5. Experiments

We conduct experiments to evaluate our proposed method and Grasp-Anything++ dataset using both the vision-based metrics and real robot experiments. We then demonstrate zero-shot grasp results and discuss the challenges and open questions for future works.

### 5.1. Language-driven Grasp Detection Results

**Baselines.** We compare our language-driven grasp detection method (LGD) with the linguistically supported versions of GR-CNN [35], Det-Seg-Refine [2], GG-CNN [47], CLIPORT [64] and CLIP-Fusion [80]. In all cases, we employ a pretrained CLIP [57] or BERT [18] as the text embedding. The implementation details of all baselines can be found in our Supplementary Material.

|                                | Baseline | Seen        | Unseen      | H           |
|--------------------------------|----------|-------------|-------------|-------------|
| GR-ConvNet [35] + CLIP [57]    |          | 0.37        | 0.18        | 0.24        |
| Det-Seg-Refine [2] + CLIP [57] |          | 0.30        | 0.15        | 0.20        |
| GG-CNN [47] + CLIP [57]        |          | 0.12        | 0.08        | 0.10        |
| CLIPORT [64]                   |          | 0.36        | 0.26        | 0.29        |
| CLIP-Fusion [80]               |          | 0.40        | 0.29        | 0.33        |
| LGD (ours) + BERT [18]         |          | 0.44        | 0.38        | 0.41        |
| LGD (ours) + CLIP [57]         |          | <b>0.48</b> | <b>0.42</b> | <b>0.45</b> |

Table 2. Language-driven grasp detection results.

**Setup.** To assess the generalization of all methods trained on Grasp-Anything++, we utilize the concept of base and new labels [92] in zero-shot learning. We categorize LVIS labels from Section 3.3 to form labels for our experiment. In particular, we select 70% of these labels by frequency for ‘Base’ and assign the remaining 30% to ‘New’. We also use the harmonic mean (‘H’) to measure the overall success rates [92]. Our primary evaluation metric is the success rate, defined similarly to [35], necessitating an IoU score of the predicted grasp exceeding 25% with the ground truth grasp and an offset angle less than 30°.

**Main Results.** Table 2 shows the results of language-driven grasp detection on the Grasp-Anything++ dataset. The findings indicate a notable performance advantage of our LGD over other baseline approaches, with LGD outperforming the subsequent best-performing baselines (CLIP-Fusion) by margins of 0.14 on Grasp-Anything++ dataset.

|                                  | Baseline | Seen        | Unseen      | H           |
|----------------------------------|----------|-------------|-------------|-------------|
| LGD w/o predicting $\tilde{x}_0$ |          | 0.15        | 0.08        | 0.10        |
| LGD w/o contrastive loss         |          | 0.45        | 0.40        | 0.42        |
| LGD w contrastive loss           |          | <b>0.48</b> | <b>0.42</b> | <b>0.45</b> |

Table 3. Contrastive loss analysis.




Figure 6. Loss visualization.




Figure 7. t-SNE visualization. We apply t-SNE to cluster the vision-and-language features with and without contrastive loss.

**Contrastive Loss Analysis.** Table 3 presents the performance of LGD under varied configurations. The outcomes emphasize the substantial influence of the training objective (contrastive loss) and the importance of language instructions in enhancing LGD performance on both seen and unseen classes in the Grasp-Anything++ dataset.

Fig. 6 shows the contrastive loss  $\mathcal{L}_{\text{contrastive}}$  approaching towards 0 during training, indicating the grasp pose estimation  $\hat{x}_0$  aligns with the ground truth  $x_0$ , as anticipated by Proposition 1. The subsequent attention maps visualization in Fig. 8 shows the attention region is meaningful and improves the results when employing our proposed contrastive loss compared to its absence. Moreover, we employ t-SNE for vision-and-language embedding visualization, as in [46], by processing 2,000 samples from the Grasp-Anything++ dataset through the ALBEF module. The outcomes reveal that our contrastive loss facilitates better object classification, as evidenced in Fig. 7 by clearer segregation of pixel embeddings across various semantic classes, underscoring contrastive loss’s role in refining embeddings’ differentiation for improved class distinctions.

**Qualitative Results.** Fig. 5 presents qualitative results of the language-driven grasp detection task, suggesting that our LGD method generates more semantically plausible than other baselines. Despite satisfactory performance, LGD occasionally predicts incorrect results, with a detailed analysis of these cases available in our Appendix.

**Robotic Validation.** We provide quantitative results by integrating our language-driven grasp detection pipeline for a robotic grasping application with a KUKA LBR iiwa

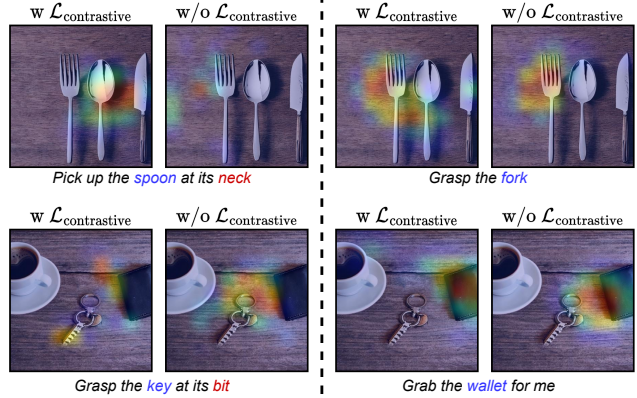


Figure 8. Attention map visualization. We compare the attention map when utilizing our proposed contrastive loss and when not.

|                                | Baseline | Single      | Cluttered   |
|--------------------------------|----------|-------------|-------------|
| GR-ConvNet [35] + CLIP [57]    |          | 0.33        | 0.30        |
| Det-Seg-Refine [2] + CLIP [57] |          | 0.30        | 0.23        |
| GG-CNN [47] + CLIP [57]        |          | 0.10        | 0.07        |
| CLIPORT [64]                   |          | 0.27        | 0.30        |
| CLIP-Fusion [80]               |          | 0.40        | 0.40        |
| LGD (ours)                     |          | <b>0.43</b> | <b>0.42</b> |

Table 4. Robotic language-driven grasp detection results.

R820 robot. Using the RealSense D435i camera, the grasp pose inferred from approaches in Table 4 is transformed into the 6DoF grasp pose, similar to [35]. The optimization-based trajectory planner in [5,72] is employed to execute the grasps. Experiments are conducted for two scenarios, i.e., the single object scenario and the cluttered scene scenario, of a set of 20 real-world daily objects. In each scenario, we run 30 experiments using baselines listed in Table 4 and a predefined grasping prompt corpus. The results exhibit that our LGD outperforms other baselines. Furthermore, although LGD is trained on our Grasp-Anything++ which is a solely synthesis dataset created by foundation models, it still shows reasonable results on real-world objects.

## 5.2. Zero-shot Grasp Detection

Our proposed Grasp-Anything++ is a large-scale dataset. Apart from the language-driven grasp detection task, we believe it can be used for other purposes. In this experiment, we seek to answer the question: Can Grasp-Anything++ be useful in the traditional grasp detection task without text? Consequently, we verify our Grasp-Anything++ and LGD (no text) with other existing datasets and grasping methods.

**Setup.** We setup an LGD (no text) version, and other state-of-the-art grasp detection methods GR-ConvNet [35], Det-Seg-Refine [2], GG-CNN [47]. We use five datasets: our Grasp-Anything++, Jacquard [17], Cornell [33], VMRD [87], and OCID-grasp [2] in this experiment.

| Baseline             | Grasp-Anything++ (ours) |             |             | Jacquard [17] |             |             | Cornell [33] |             |             | VMRD [87]   |             |             | OCID-grasp [2] |             |             |
|----------------------|-------------------------|-------------|-------------|---------------|-------------|-------------|--------------|-------------|-------------|-------------|-------------|-------------|----------------|-------------|-------------|
|                      | Base                    | New         | H           | Base          | New         | H           | Base         | New         | H           | Base        | New         | H           | Base           | New         | H           |
| GR-ConvNet [35]      | 0.71                    | 0.59        | 0.64        | 0.88          | 0.66        | 0.75        | 0.98         | 0.74        | 0.84        | 0.77        | 0.64        | 0.70        | 0.86           | 0.67        | 0.75        |
| Det-Seg-Refine [2]   | 0.62                    | 0.57        | 0.59        | 0.86          | 0.60        | 0.71        | <b>0.99</b>  | <b>0.76</b> | <b>0.86</b> | 0.75        | 0.60        | 0.66        | 0.80           | 0.62        | 0.70        |
| GG-CNN [47]          | 0.68                    | 0.57        | 0.62        | 0.78          | 0.56        | 0.65        | 0.96         | 0.75        | 0.84        | 0.69        | 0.53        | 0.59        | 0.71           | 0.63        | 0.67        |
| LGD (no text) (ours) | <b>0.74</b>             | <b>0.63</b> | <b>0.68</b> | <b>0.89</b>   | <b>0.69</b> | <b>0.77</b> | 0.97         | <b>0.76</b> | 0.85        | <b>0.79</b> | <b>0.66</b> | <b>0.72</b> | <b>0.88</b>    | <b>0.68</b> | <b>0.76</b> |

Table 5. Base-to-new zero-shot grasp detection results.

**Zero-shot Results.** Table 5 summarizes the base-to-new grasp detection results on five datasets. Overall, the performance of LGD even without the language branch is better than other baselines across all datasets. Furthermore, this table also shows that our Grasp-Anything++ dataset is more challenging to train as the detection results are lower than related datasets using the same approaches due to the greater coverage of unseen objects in the testing phase.

| Train \ Test            | Test        |             |             |             |                  |
|-------------------------|-------------|-------------|-------------|-------------|------------------|
|                         | Jacquard    | Cornell     | VMRD        | OCID-grasp  | Grasp-Anything++ |
| Jacquard [17]           | 0.87        | 0.51        | 0.13        | 0.21        | <b>0.17</b>      |
| Cornell [33]            | 0.07        | 0.98        | 0.20        | 0.12        | 0.13             |
| VMRD [87]               | 0.06        | 0.21        | <u>0.79</u> | 0.11        | 0.10             |
| OCID-grasp [2]          | 0.09        | 0.12        | 0.20        | <u>0.74</u> | 0.11             |
| Grasp-Anything++ (ours) | <b>0.41</b> | <b>0.63</b> | <b>0.30</b> | <b>0.39</b> | <u>0.65</u>      |

Table 6. Cross-dataset grasp detection results.

**Cross-dataset Evaluation.** To further verify the usefulness of our Grasp-Anything++ dataset, we conduct the cross-dataset validation in Table 6. We use the GR-ConvNet [35] to reuse its results on existing grasp datasets. GR-ConvNet is trained on a dataset (row) and evaluated on another dataset (column). For example, training on Jacquard and testing on Cornell yields an accuracy of 0.51. Notably, training with our dataset improves performance by approximately 10 – 33% compared to other datasets.

**In the wild grasp detection.** Fig. 9 shows visualization results using LGD (no text) trained on our Grasp-Anything++ dataset on random internet images and other datasets images. We can see that the detected grasp poses are adequate in quality and quantity. This demonstrates that although our Grasp-Anything++ is fully created by foundation models without having any real images, models trained on our Grasp-Anything++ dataset still generalize well on real-world images.

### 5.3. Discussion

Our experiments indicate that Grasp-Anything++ can serve as a foundation dataset for both language-driven and traditional grasp detection tasks. However, there are certain limitations. First, our dataset lacks depth images for directly being applied to robotic applications [52]. Second, we remark that the creation of our dataset is time-consuming and relies on access to the ChatGPT API. Fortunately, future research can reuse our provided assets (images, prompts,




Figure 9. In the wild grasp detection. Top row images are from GraspNet [24], YCB-Video [78], NBMOD [9] datasets; bottom row shows internet images.

etc.) without starting from scratch. Furthermore, our experiments show that adding language to the grasp detection task (Table 2) poses a more challenging problem compared to standard grasp detection task (Table 5).

We see several interesting future research directions. First, future work could investigate the use of text or image-to-3D models [79] or image-to-depth [59] and reuse our dataset’s prompts and images to construct 3D language-driven grasp datasets. Additionally, beyond linguistic grasp instruction adherence, our dataset holds potential for varied applications, including scene understanding [81] and scene generation [3], hallucination analysis [30], and human-robot interaction [11].

## 6. Conclusion

We introduce Grasp-Anything++, a large-scale dataset with 1M images and 10M grasp prompts for language-driven grasp detection tasks. We propose LGD, a diffusion-based method to tackle the language-driven grasp detection task. Our diffusion model employs a contrastive training objective, which explicitly contributes to the denoising process. Empirically, we have shown that Grasp-Anything++ serves as a foundation grasp detection dataset. Finally, our LGD improves the performance of other baselines, and the real-world robotic experiments further validate the effectiveness of our dataset and approach.



## References

- [1] Michael Ahn, Anthony Brohan, Noah Brown, Yevgen Chebotar, Omar Cortes, Byron David, Chelsea Finn, Chuyuan Fu, Keerthana Gopalakrishnan, Karol Hausman, et al. Do as i can, not as i say: Grounding language in robotic affordances. *arXiv preprint arXiv:2204.01691*, 2022. 2
- [2] Stefan Ainetter and Friedrich Fraundorfer. End-to-end trainable deep neural network for robotic grasp detection and semantic segmentation from rgb. In *ICRA*, 2021. 2, 4, 6, 7, 8
- [3] Dor Arad Hudson and Larry Zitnick. Compositional transformers for scene generation. *NeurIPS*, 2021. 8
- [4] Umar Asif, Jianbin Tang, and Stefan Herrer. Graspnet: An efficient convolutional neural network for real-time grasp detection for low-powered devices. In *IJCAI*, 2018. 4
- [5] Florian Beck, Minh Nhat Vu, Christian Hartl-Nesic, and Andreas Kugi. Singularity avoidance with application to on-line trajectory optimization for serial manipulators. *IFAC-PapersOnLine*, 56(2):284–291, 2023. 7
- [6] Andreas Blattmann, Robin Rombach, Huan Ling, Tim Dockhorn, Seung Wook Kim, Sanja Fidler, and Karsten Kreis. Align your latents: High-resolution video synthesis with latent diffusion models. In *CVPR*, 2023. 2
- [7] Tom Brown, Benjamin Mann, Nick Ryder, Melanie Subbiah, Jared D Kaplan, Prafulla Dhariwal, Arvind Neelakantan, Pranav Shyam, Girish Sastry, Amanda Askell, et al. Language models are few-shot learners. *NeurIPS*, 33, 2020. 2
- [8] Guy Bukchin, Eli Schwartz, Kate Saenko, Ori Shahar, Rogério Feris, Raja Giryes, and Leonid Karlinsky. Fine-grained angular contrastive learning with coarse labels. In *CVPR*, 2021. 4
- [9] Boyuan Cao, Xinyu Zhou, Congmin Guo, Baohua Zhang, Yuchen Liu, and Qianqiu Tan. Nbmmod: Find it and grasp it in noisy background. *arXiv preprint arXiv:2306.10265*, 2023. 2, 3, 8
- [10] Joao Carvalho, An T Le, Mark Baierl, Dorothea Koert, and Jan Peters. Motion planning diffusion: Learning and planning of robot motions with diffusion models. *arXiv preprint arXiv:2308.01557*, 2023. 2
- [11] Paola Cascante-Bonilla, Hui Wu, Letao Wang, Rogério S Feris, and Vicente Ordonez. Simvqa: Exploring simulated environments for visual question answering. In *CVPR*, 2022. 8
- [12] I-Ming Chen and Joel W Burdick. Finding antipodal point grasps on irregularly shaped objects. *T-RA*, 1993. 4
- [13] Sijia Chen and Baochun Li. Language-guided diffusion model for visual grounding. *arXiv preprint arXiv:2308.09599*, 2023. 2
- [14] Hongjun Choi, Anirudh Som, and Pavan Turaga. Amc-loss: Angular margin contrastive loss for improved explainability in image classification. In *CVPR workshops*, 2020. 4
- [15] Florinel-Alin Croitoru, Vlad Hondru, Radu Tudor Ionescu, and Mubarak Shah. Diffusion models in vision: A survey. *TPAMI*, 2023. 4
- [16] Matt Deitke, Dustin Schwenk, Jordi Salvador, Luca Weihs, Oscar Michel, Eli VanderBilt, Ludwig Schmidt, Kiana Ehsani, Aniruddha Kembhavi, and Ali Farhadi. Objaverse: A universe of annotated 3d objects. In *CVPR*, 2023. 4
- [17] Amaury Depierre, Emmanuel Dellandréa, and Liming Chen. Jacquard: A large scale dataset for robotic grasp detection. In *IROS*, 2018. 2, 3, 4, 7, 8
- [18] Jacob Devlin, Ming-Wei Chang, Kenton Lee, and Kristina Toutanova. Bert: Pre-training of deep bidirectional transformers for language understanding. *arXiv preprint arXiv:1810.04805*, 2018. 6
- [19] Prafulla Dhariwal and Alexander Nichol. Diffusion models beat gans on image synthesis. *NeurIPS*, 34, 2021. 5
- [20] Finale Doshi-Velez and Been Kim. Towards a rigorous science of interpretable machine learning. *arXiv preprint arXiv:1702.08608*, 2017. 2
- [21] Danny Driess, Fei Xia, Mehdi SM Sajjadi, Corey Lynch, Aakanksha Chowdhery, Brian Ichter, Ayzaan Wahid, Jonathan Tompson, Quan Vuong, Tianhe Yu, et al. Palm-e: An embodied multimodal language model. *arXiv preprint arXiv:2303.03378*, 2023. 1
- [22] Clemens Eppner, Arsalan Mousavian, and Dieter Fox. A billion ways to grasp: An evaluation of grasp sampling schemes on a dense, physics-based grasp data set. In *ISRR*, 2019. 2
- [23] Clemens Eppner, Arsalan Mousavian, and Dieter Fox. Acronym: A large-scale grasp dataset based on simulation. In *ICRA*, 2021. 2
- [24] Hao-Shu Fang, Chenxi Wang, Minghao Gou, and Cewu Lu. Graspnet-1billion: A large-scale benchmark for general object grasping. In *CVPR*, 2020. 2, 8
- [25] Maximilian Gilles, Yuhao Chen, Tim Robin Winter, E Zhixuan Zeng, and Alexander Wong. Metagraspnet: A large-scale benchmark dataset for scene-aware ambidextrous bin picking via physics-based metaverse synthesis. In *CASE*, 2022. 1
- [26] Agrim Gupta, Piotr Dollar, and Ross Girshick. Lvis: A dataset for large vocabulary instance segmentation. In *CVPR*, 2019. 4
- [27] Michael Gutmann and Aapo Hyvärinen. Noise-contrastive estimation: A new estimation principle for unnormalized statistical models. In *AISTAT*, 2010. 5
- [28] Jonathan Ho, Ajay Jain, and Pieter Abbeel. Denoising diffusion probabilistic models. *NeurIPS*, 2020. 2, 4, 5, 6
- [29] Jonathan Ho, Tim Salimans, Alexey Gritsenko, William Chan, Mohammad Norouzi, and David J Fleet. Video diffusion models. *arXiv preprint arXiv:2204.03458*, 2022. 4
- [30] Lei Huang, Weijiang Yu, Weitao Ma, Weihong Zhong, Zhangyin Feng, Haotian Wang, Qianglong Chen, Weihua Peng, Xiaocheng Feng, Bing Qin, et al. A survey on hallucination in large language models: Principles, taxonomy, challenges, and open questions. *arXiv preprint arXiv:2311.05232*, 2023. 4, 8
- [31] Michael Janner, Yilun Du, Joshua Tenenbaum, and Sergey Levine. Planning with diffusion for flexible behavior synthesis. In *ICML*, 2022. 2
- [32] Krishna Murthy Jatavallabhula, Alihusein Kuwajerwala, Qiao Gu, Mohd Omama, Tao Chen, Shuang Li, Ganesh Iyer, Soroush Saryazdi, Nikhil Keetha, Ayush Tewari, et al. Conceptfusion: Open-set multimodal 3d mapping. *arXiv preprint arXiv:2302.07241*, 2023. 1

- [33] Yun Jiang, Stephen Moseson, and Ashutosh Saxena. Efficient grasping from rgb-d images: Learning using a new rect-angle representation. In *ICRA*, 2011. 2, 3, 4, 7, 8
- [34] Alexander Kirillov, Eric Mintun, Nikhila Ravi, Hanzi Mao, Chloe Rolland, Laura Gustafson, Tete Xiao, Spencer Whitehead, Alexander C Berg, Wan-Yen Lo, et al. Segment anything. *arXiv preprint arXiv:2304.02643*, 2023. 3
- [35] Sulabh Kumra, Shirin Joshi, and Ferat Sahin. Antipodal robotic grasping using generative residual convolutional neural network. In *IROS*, 2020. 6, 7, 8
- [36] Jaewoong Lee, Sangwon Jang, Jaehyeong Jo, Jaehong Yoon, Yunji Kim, Jin-Hwa Kim, Jung-Woo Ha, and Sung Ju Hwang. Text-conditioned sampling framework for text-to-image generation with masked generative models. *arXiv preprint arXiv:2304.01515*, 2023. 2
- [37] Meng-Lun Lee, Sara Behdad, Xiao Liang, and Minghui Zheng. Task allocation and planning for product disassembly with human-robot collaboration. *Robotics and Computer-Integrated Manufacturing*, 76, 2022. 2
- [38] Sergey Levine, Peter Pastor, Alex Krizhevsky, Julian Ibarz, and Deirdre Quillen. Learning hand-eye coordination for robotic grasping with deep learning and large-scale data collection. *IJRR*, 2018. 2
- [39] Junnan Li, Ramprasaath Selvaraju, Akhilesh Gotmare, Shafiq Joty, Caimeing Xiong, and Steven Chu Hong Hoi. Align before fuse: Vision and language representation learning with momentum distillation. *NeurIPS*, 2021. 5
- [40] Yuheng Li, Haotian Liu, Qingyang Wu, Fangzhou Mu, Jianwei Yang, Jianfeng Gao, Chunyuan Li, and Yong Jae Lee. Gligen: Open-set grounded text-to-image generation. In *CVPR*, 2023. 2
- [41] Jirong Liu, Ruo Zhang, Hao-Shu Fang, Minghao Gou, Hongjie Fang, Chenxi Wang, Sheng Xu, Hengxu Yan, and Cewu Lu. Target-referenced reactive grasping for dynamic objects. In *CVPR*, 2023. 2
- [42] Shilong Liu, Zhaoyang Zeng, Tianhe Ren, Feng Li, Hao Zhang, Jie Yang, Chunyuan Li, Jianwei Yang, Hang Su, Jun Zhu, et al. Grounding dino: Marrying dino with grounded pre-training for open-set object detection. *arXiv preprint arXiv:2303.05499*, 2023. 2
- [43] Weiyu Liu, Tucker Hermans, Sonia Chernova, and Chris Paxton. Structdiffusion: Object-centric diffusion for semantic rearrangement of novel objects. *arXiv preprint arXiv:2211.04604*, 2022. 2, 4
- [44] Yiheng Liu, Tianle Han, Siyuan Ma, Jiayue Zhang, Yuanyuan Yang, Jiaming Tian, Hao He, Antong Li, Mengshen He, Zhengliang Liu, et al. Summary of chatgpt-related research and perspective towards the future of large language models. *Meta-Radiology*, page 100017, 2023. 4
- [45] Jeffrey Mahler, Jacky Liang, Sherdil Niyaz, Michael Laskey, Richard Doan, Xinyu Liu, Juan Aparicio Ojea, and Ken 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. 2
- [46] Jiaxu Miao, Zongxin Yang, Leilei Fan, and Yi Yang. Fedseg: Class-heterogeneous federated learning for semantic segmentation. In *CVPR*, 2023. 7
- [47] Douglas Morrison, Peter Corke, and Jürgen Leitner. Closing the loop for robotic grasping: A real-time, generative grasp synthesis approach. *arXiv preprint arXiv:1804.05172*, 2018. 6, 7, 8
- [48] Douglas Morrison, Peter Corke, and Jürgen Leitner. Egad! an evolved grasping analysis dataset for diversity and reproducibility in robotic manipulation. *RA-L*, 2020. 2
- [49] Arsalan Mousavian, Clemens Eppner, and Dieter Fox. 6-dof graspnet: Variational grasp generation for object manipulation. In *ICCV*, 2019. 2
- [50] Yao Mu, Shunyu Yao, Mingyu Ding, Ping Luo, and Chuang Gan. Ec2: Emergent communication for embodied control. In *CVPR*, 2023. 2
- [51] Yao Mu, Qinglong Zhang, Mengkang Hu, Wenhai Wang, Mingyu Ding, Jun Jin, Bin Wang, Jifeng Dai, Yu Qiao, and Ping Luo. Embodiedgpt: Vision-language pre-training via embodied chain of thought. *arXiv preprint arXiv:2305.15021*, 2023. 1
- [52] Rhys Newbury, Morris Gu, Lachlan Chumbley, Arsalan Mousavian, Clemens Eppner, Jürgen Leitner, Jeannette Bohg, Antonio Morales, Tamim Asfour, Danica Kragic, et al. Deep learning approaches to grasp synthesis: A review. *T-RO*, 2023. 8
- [53] Toan Nguyen, Minh Nhat Vu, An Vuong, Dzung Nguyen, Thieu Vo, Ngan Le, and Anh Nguyen. Open-vocabulary affordance detection in 3d point clouds. In *IROS*, 2023. 3
- [54] OpenAI. Introducing ChatGPT. Software. Accessed: July 6th 2023. 2
- [55] Lerrel Pinto and Abhinav Gupta. Supersizing self-supervision: Learning to grasp from 50k tries and 700 robot hours. In *ICRA*, 2016. 2
- [56] Robert Platt. Grasp learning: Models, methods, and performance. *Annual Review of Control, Robotics, and Autonomous Systems*, 2023. 1
- [57] Alec Radford, Jong Wook Kim, Chris Hallacy, Aditya Ramesh, Gabriel Goh, Sandhini Agarwal, Girish Sastry, Amanda Askell, Pamela Mishkin, Jack Clark, et al. Learning transferable visual models from natural language supervision. In *ICML*, 2021. 6, 7
- [58] Aditya Ramesh, Prafulla Dhariwal, Alex Nichol, Casey Chu, and Mark Chen. Hierarchical text-conditional image generation with clip latents. *arXiv preprint arXiv:2204.06125*, 2022. 6
- [59] René Ranftl, Alexey Bochkovskiy, and Vladlen Koltun. Vision transformers for dense prediction. In *CVP4*, 2021. 8
- [60] Robin Rombach, Andreas Blattmann, Dominik Lorenz, Patrick Esser, and Björn Ommer. High-resolution image synthesis with latent diffusion models. In *CVPR*, 2022. 3, 4
- [61] Kallol Saha, Vishal Mandadi, Jayaram Reddy, Ajit Srikanth, Aditya Agarwal, Bipasha Sen, Arun Singh, and Madhava Krishna. Edmp: Ensemble-of-costs-guided diffusion for motion planning. *arXiv preprint arXiv:2309.11414*, 2023. 2
- [62] Schuhmann et al. Laion-5b: An open large-scale dataset for training next generation image-text models. *NeurIPS*, 2022. 4
- [63] Dhruv Shah, Błażej Osiński, brian ichter, and Sergey Levine. Lm-nav: Robotic navigation with large pre-trained models of

- language, vision, and action. In Karen Liu, Dana Kulic, and Jeff Ichnowski, editors, *Proceedings of The 6th Conference on Robot Learning*, volume 205 of *PMLR*. PMLR, 2023. 1
- [64] Mohit Shridhar, Lucas Manuelli, and Dieter Fox. Cliport: What and where pathways for robotic manipulation. In *CoRL*, 2022. 1, 2, 6, 7
- [65] Jascha Sohl-Dickstein, Eric Weiss, Niru Maheswaranathan, and Surya Ganguli. Deep unsupervised learning using nonequilibrium thermodynamics. In *ICML*, 2015. 5
- [66] Chan Hee Song, Jiaman Wu, Clayton Washington, Brian M Sadler, Wei-Lun Chao, and Yu Su. Llm-planner: Few-shot grounded planning for embodied agents with large language models. In *ICCV*, 2023. 1
- [67] Peize Sun, Shoufa Chen, Chenchen Zhu, Fanyi Xiao, Ping Luo, Saining Xie, and Zhicheng Yan. Going denser with open-vocabulary part segmentation. *arXiv preprint arXiv:2305.11173*, 2023. 3
- [68] Guy Tevet, Sigal Raab, Brian Gordon, Yoni Shafir, Daniel Cohen-or, and Amit Haim Bermano. Human motion diffusion model. In *ICLR*, 2022. 5, 6
- [69] Jonathan Tseng, Rodrigo Castellon, and Karen Liu. Edge: Editable dance generation from music. In *CVPR*, 2023. 5
- [70] Julen Urain, Niklas Funk, Jan Peters, and Georgia Chalvatzaki. Se (3)-diffusionfields: Learning smooth cost functions for joint grasp and motion optimization through diffusion. In *ICRA*, 2023. 2
- [71] Sai Vemprala, Rogerio Bonatti, Arthur Buckner, and Ashish Kapoor. Chatgpt for robotics: Design principles and model abilities. *Microsoft Auton. Syst. Robot. Res*, 2023. 2
- [72] Minh Nhat Vu, Florian Beck, Michael Schwegel, Christian Hartl-Nesic, Anh Nguyen, and Andreas Kugi. Machine learning-based framework for optimally solving the analytical inverse kinematics for redundant manipulators. *Mechatronics*, 2023. 7
- [73] An Vuong, Minh Nhat Vu, Toan Tien Nguyen, Baoru Huang, Dzung Nguyen, Thieu Vo, and Anh Nguyen. Language-driven scene synthesis using multi-conditional diffusion model. *arXiv preprint arXiv:2310.15948*, 2023. 5
- [74] An Dinh Vuong, Minh Nhat Vu, Hieu Le, Baoru Huang, Binh Huynh, Thieu Vo, Andreas Kugi, and Anh Nguyen. Grasp-anything: Large-scale grasp dataset from foundation models. *arXiv preprint arXiv:2309.09818*, 2023. 2, 3
- [75] Peng Wang, An Yang, Rui Men, Junyang Lin, Shuai Bai, Zhikang Li, Jianxin Ma, Chang Zhou, Jingren Zhou, and Hongxia Yang. Ofa: Unifying architectures, tasks, and modalities through a simple sequence-to-sequence learning framework. In *ICML*, 2022. 3
- [76] Sabine Wehnert, Viju Sudhi, Shipra Dureja, Libin Kuttu, Saijal Shahania, and Ernesto W De Luca. Legal norm retrieval with variations of the bert model combined with tf-idf vectorization. In *Proceedings of the eighteenth international conference on artificial intelligence and law*, pages 285–294, 2021. 4
- [77] Julia Wolleb, Robin Sandkühler, Florentin Bieder, Philippe Valmaggia, and Philippe C Cattin. Diffusion models for implicit image segmentation ensembles. In *MIDL*, 2022. 2
- [78] Yu Xiang, Tanner Schmidt, Venkatraman Narayanan, and Dieter Fox. Posecnn: A convolutional neural network for 6d object pose estimation in cluttered scenes. *arXiv preprint arXiv:1711.00199*, 2017. 2, 8
- [79] Dejia Xu, Yifan Jiang, Peihao Wang, Zhiwen Fan, Yi Wang, and Zhangyang Wang. Neurlift-360: Lifting an in-the-wild 2d photo to a 3d object with 360deg views. In *CVPR*, 2023. 8
- [80] Kechun Xu, Shuqi Zhao, Zhongxiang Zhou, Zizhang Li, Huaijin Pi, Yifeng Zhu, Yue Wang, and Rong Xiong. A joint modeling of vision-language-action for target-oriented grasping in clutter. *arXiv preprint arXiv:2302.12610*, 2023. 2, 6, 7
- [81] Yiteng Xu, Peishan Cong, Yichen Yao, Runnan Chen, Yue-nan Hou, Xinge Zhu, Xuming He, Jingyi Yu, and Yuexin Ma. Human-centric scene understanding for 3d large-scale scenarios. In *ICCV*, 2023. 8
- [82] Zhixuan Xu, Kechun Xu, Rong Xiong, and Yue Wang. Object-centric inference for language conditioned placement: A foundation model based approach. In *ICARM*, 2023. 2
- [83] Jiange Yang, Wenhui Tan, Chuhao Jin, Bei Liu, Jianlong Fu, Ruihua Song, and Limin Wang. Pave the way to grasp anything: Transferring foundation models for universal pick-place robots. *arXiv preprint arXiv:2306.05716*, 2023. 1, 2
- [84] Ling Yang, Zhilong Zhang, Yang Song, Shenda Hong, Runsheng Xu, Yue Zhao, Wentao Zhang, Bin Cui, and Ming-Hsuan Yang. Diffusion models: A comprehensive survey of methods and applications. *ACM Computing Surveys*, 2022. 2, 4
- [85] Serin Yang, Hyunmin Hwang, and Jong Chul Ye. Zero-shot contrastive loss for text-guided diffusion image style transfer. *ICCV*, 2023. 4
- [86] Yang Yang, Xibai Lou, and Changhyun Choi. Interactive robotic grasping with attribute-guided disambiguation. In *ICRA*, 2022. 2
- [87] Hanbo Zhang, Xuguang Lan, Site Bai, Xinwen Zhou, Zhiqiang Tian, and Nanning Zheng. Roi-based robotic grasp detection for object overlapping scenes. In *IROS*, 2019. 4, 7, 8
- [88] Hanbo Zhang, Yunfan Lu, Cunjun Yu, David Hsu, Xuguang La, and Nanning Zheng. Invigorate: Interactive visual grounding and grasping in clutter. *arXiv preprint arXiv:2108.11092*, 2021. 2
- [89] Tianyi Zhang, Varsha Kishore, Felix Wu, Kilian Q Weinberger, and Yoav Artzi. Bertscore: Evaluating text generation with bert. In *ICLR*, 2019. 4
- [90] Kaizhi Zheng, Xiaotong Chen, Odest Chadwicke Jenkins, and Xin Wang. Vlmbench: A compositional benchmark for vision-and-language manipulation. *NeurIPS*, 35, 2022. 2
- [91] Yiwu Zhong, Jianwei Yang, Pengchuan Zhang, Chunyuan Li, Noel Codella, Liunian Harold Li, Luwei Zhou, Xiyang Dai, Lu Yuan, Yin Li, et al. Regionclip: Region-based language-image pretraining. In *CVPR*, 2022. 4
- [92] Kaiyang Zhou, Jingkang Yang, Chen Change Loy, and Ziwei Liu. Conditional prompt learning for vision-language models. In *CVPR*, 2022. 6