

An Economic Framework for 6-DoF Grasp Detection

Xiao-Ming Wu^{1,3,†,&}, Jia-Feng Cai^{1,3,&}, Jian-Jian Jiang^{1,3}, Dian Zheng^{1,3}, Yi-Lin Wei^{1,3}, and Wei-Shi Zheng^{1,2,3,4,*}

¹ School of Computer Science and Engineering, Sun Yat-sen University, China;

² Peng Cheng Laboratory, Shenzhen, China; ³ Key Laboratory of Machine Intelligence and Advanced Computing, Ministry of Education, China

⁴ Pazhou Laboratory (Huangpu), Guangzhou, Guangdong, China
{wuxm65, caijf23, jiangjj35, zhengd35, weiylin5}@mail2.sysu.edu.cn,
wszheng@ieee.org

Abstract. Robotic grasping in clutters is a fundamental task in robotic manipulation. In this work, we propose an economic framework for 6-DoF grasp detection, aiming to economize the resource cost in training and meanwhile maintain effective grasp performance. To begin with, we discover that the dense supervision is the bottleneck of current SOTA methods that severely encumbers the entire training overload, meanwhile making the training difficult to converge. To solve the above problem, we first propose an economic supervision paradigm for efficient and effective grasping. This paradigm includes a well-designed supervision selection strategy, selecting key labels basically without ambiguity, and an economic pipeline to enable the training after selection. Furthermore, benefit from the economic supervision, we can focus on a specific grasp, and thus we devise a focal representation module, which comprises an interactive grasp head and a composite score estimation to generate the specific grasp more accurately. Combining all together, the **Economic-Grasp** framework is proposed. Our extensive experiments show that EconomicGrasp surpasses the SOTA grasp method by about **3AP** on average, and with extremely low resource cost, for about **1/4** training time cost, **1/8** memory cost and **1/30** storage cost. Our code is available at <https://github.com/iSEE-Laboratory/EconomicGrasp>.

Keywords: Robotics · 6-DoF Grasp Detection · Economic Framework

1 Introduction

As a fundamental task in robotic manipulations, 6-DoF grasp detection [2, 10, 11, 51] achieves rapid development in recent years, which has broad applications in many real-world scenarios, such as picking [7], assembling [48], fruit harvesting [68] and house serving like cooking and cleaning [13].

† : project lead, &: equal key contributions, *: corresponding author.

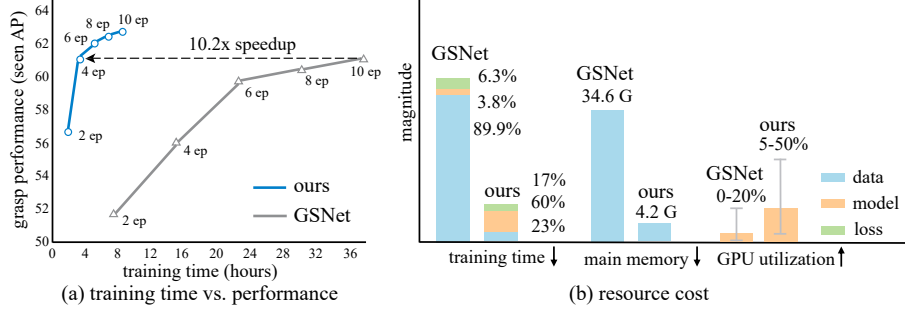


Fig. 1: Economic supervision vs. dense supervision. In our economic framework, the resource cost is minimal and the training is easy to converge. Moreover, with our well design for economic supervision, our framework achieves better performance than the SOTA dense supervision method GSNet [54]. "ep" means epochs. All the costs are tested in an empty machine with one NVIDIA RTX3090 GPU for fair. The results are trained and tested with GraspNet-1Billion [11] on Kinect data.

Recently, with the help of the real-world grasp dataset GraspNet-1Billion [11], which has abundant labels as supervision (we call it **dense supervision**, which is in billion scale), 6-DoF grasp detection methods [10, 11, 15, 34, 54] achieve great grasping performance. However, this type of supervision also occurs some problems. It 1) has huge resource costs in model training, 2) makes the training difficult to learn and converge. As shown in Fig. 1 (a), this supervision is somehow too abundant for the model to train, which causes slow and hard convergence. Furthermore, as shown in Fig. 1 (b), the data processing time of the state-of-the-art method GSNet [54], which is under dense supervision, is very long (about 9x of the sum of model training time and loss calculation time). Moreover, nearly 100 million labels will be loaded in each batch to train the model, resulting in high memory cost (34.6G) and low GPU utilization (20%). Prior to GraspNet-1Billion [11], due to supervision-limited, many grasping methods are designed to learn in **sparse supervision** [2, 29, 38, 42, 51] (in million scale). Although they may have fewer training resources or fast converge speed, they have limited grasp performance due to the inadequate training data. Therefore, how to enable economic grasping with low resource costs and effective grasp performance is a meaningful but rarely explored area in 6-DoF grasp detection, which can promote the extension of the research area and support wide applications in resource-constrained environments.

In this work, we propose an economic framework for 6-DoF grasp detection, which aims to **economize the resource costs in training and meanwhile maintain effective grasp performance**. Firstly, we analyze the reason which causes the gap between sparse and dense supervision methods. We follow the previous sparse supervision methods, gradually "modernizing" it, to see whether it is the module or loss designs that lead the gap. However, the answer is NO and there still exist other problems in sparse supervision. Secondly, we conduct variance analysis in labels, and find that the ambiguity problem is the "chief culprit" causing this gap. Therefore, we propose an economic supervision paradigm for

efficient and effective grasping, including a well-designed supervision selection strategy and a economic pipeline to enable training after selection. Thirdly, economic supervision affords us the opportunity to focus on the learning of a specific grasp. Toward this end, we design a focal representation module to enable it, which includes an interactive grasp head to learn discriminate features for a specific grasp, and a composite score for more accurate score estimation. Combining all together, the **EconomicGrasp** is proposed, which is resource-friendly and performance-effective for robotic grasping.

Extensive experiments show that EconomicGrasp has great grasp performance, surpassing the state-of-the-art grasp method with 3AP on average. More importantly, our framework is economic, which has extremely low resource costs comparing to state-of-the-art grasp method, for about 1/4 training time cost, 1/8 memory cost and 1/30 storage cost. Our code is available at <https://github.com/iSEE-Laboratory/EconomicGrasp>.

2 Revisiting 6-DoF Grasp Detection

Task Definition. Robotic grasping is a fundamental skill for robotic manipulation [13, 68, 69], which can be divided into two basic processes, perception and planing [43]. The robot first detects where to grasp (grasp detection), then executes the control algorithm to move to the target position. 6-DoF grasp detection, different from the methods of detecting contact points [26, 46] or rectangles [23, 27, 35, 44, 64], is to generate 6-DoF grasp poses based on the visual input [9, 29, 36, 49, 51, 66], which consists of the 3D positions with its 3D rotations. Formally speaking, giving the point cloud $\mathbf{C} \in \mathbb{R}^{n \times 3}$ as input (n is the number of input points), our goal is to learn how to generate 6-DoF grasps that can successfully grab the object. Following the setting of previous works [5, 10, 11, 33, 54], we represent the 6-DoF grasp as $\mathbf{G} = [\mathbf{c}, v, a, d, w, s]$, where $\mathbf{c} \in \mathbb{R}^3$ is the central point of the grasp, v is an integer ranging from 1 to 300, standing for different approaching directions of the grasp view, a is an integer ranging from 1 to 12, representing different angles of the 2D in-plane rotation, d is also an integer ranging from 1 to 4, indicating the grasp depth, $w \in \mathbb{R}$ is the grasp width and $s \in \mathbb{R}$ is the grasp score describing the grasp quality. For better comprehension, we demonstrate the task definition and the 6-DoF grasp pose in Fig. 2. More details can be found in [11]. To be noted that there are also other types of grippers like dexterous hand [20, 31, 56, 59, 62, 63], human hand [22, 57] and so on, here we mainly focus on the basic two-finger parallel gripper.

Development of 6-DoF Grasp Detection. 6-DoF grasp detection can trace back to many pioneering works [12, 17, 18, 24, 51, 52] that aims to relax the restrict of the rectangle representation from three domains of freedom into the more relaxed ones. Among them, GPD [51] is one of the most popular method that enables high success rate to grasp novel objects, which greatly promotes the research in this field. Inspired by GPD [51], early 6-DoF grasp detection methods apply the sample-then-evaluation pipeline [29, 36, 51] to generate grasp. This pipeline first samples large amounts of grasp candidates from the input scene, then evaluates this candidate to see whether it is a successful grasp. However,




Fig. 2: Task definition and the grasp pose. The input for this task is the single-view point cloud from the depth camera and the model aims to output the successful 6-DoF grasp poses for the input scene.

this pipeline is highly time-consuming in inference, since that we should sample a large amount of grasp candidates and then put them into the evaluator to evaluate. This means that the evaluator (usually a neural network) should run several times to generate successful grasps. To enable efficient grasping, end-to-end methods are proposed [3, 8, 21, 30, 47, 49, 58, 60, 65–67], which directly generates successful 6-DoF grasps based on the visual input.

Recent years, the large-scale real-world dataset GraspNet-1Billion [1] is proposed, which has dense grasp labels (in billion scale) and enables robust grasping in real world [5, 10, 11, 15, 33, 34, 54]. However, the supervision in this pipeline may be too abundant, which may occur 1) huge resource costs in model training, 2) difficulty for the model to train and converge, as shown in Fig. 1. Prior to the GraspNet-1Billion [1], due to the data limit at that time, many methods are designed under sparse supervision [2, 29, 38, 42, 51]. Although they may be resource-friendly or easy to converge, these methods obtain limited grasp performance. Meanwhile, few methods attempt to use partial labels in the large scale dataset GraspNet-1Billion [1] as supervision [19, 28, 32, 41, 55]. But they have neither explored how to conduct high-quality supervision selection for effective performance, nor attempted to pursue low resource costs when using partial supervision. Thus, they have limited grasp performance and not low resource costs. Thus, how to enable economic grasping (with low resource costs and effective grasp performance) is a meaningful but rarely explored area in 6-DoF grasp detection. Toward this end, in this work, we propose an economic framework, EcomonicGrasp, for 6-DoF grasp detection, aiming to economize the resource costs in training and meanwhile maintain effective grasp performance.

3 The Economic Grasp Framework

Our economic grasp framework aims to enable effective grasping with low resource costs. First, we explore whether it is the module and loss designs causing the performance gap between the dense and sparse supervision methods. We find that they are part of the reason, but not the key factor (Sec. 3.1). Second, we conduct variance analysis in labels and discover that the ambiguity problem is the "chief culprit" causing this gap. Thus, we design an economic supervision paradigm to mitigate the ambiguity, comprising a supervision selection strategy and the corresponding training pipeline (Sec. 3.2). Third, under economic supervision, we can focus on the learning of a specific grasp, which inspires us to further design the focal representation module (Sec. 3.3).

3.1 A Vanilla Grasp Framework

Previous to the GraspNet-1Billion dataset [11], there exists many methods training on sparse supervision. Thus, we firstly want to see whether it is the module and loss designs that cause the performance gap between the sparse supervision methods and dense supervision methods. To do so, we reduce the labels in GraspNet-1Billion dataset [11], only keeping the

Table 1: We gradually "modernize" a simple method S4G [42] (line 1) with the cutting-edge approaches, resulting in a vanilla grasp framework (line 4) for sparse supervision. However, its performance is still limited comparing to dense supervision methods. The performance are the seen results on GraspNet-1Billion dataset [11] in Kinect data.

PointNet++	3DConv	region	statistic	seen	similar	novel
✓				18.72	11.82	5.38
✓		✓		37.38	29.50	10.23
✓		✓	✓	39.63	31.72	11.09
	✓	✓	✓	43.59	34.09	13.36

best grasp for each point in each scene, and we implement a simple framework, which follows the PointNet++Grasp [38] or S4G [42] pipeline, using PointNet++ [40] as the backbone and directly predicting the grasp pose for each point, to enable the training under this supervision. However, as shown at the first line in Table 1, this simple baseline only has very limited performance. We further fuse some cutting-edge approaches into the simple baseline, gradually "modernizing" it into a better framework. First, region-based approaches are widely used in recent methods [1, 4, 5, 66], so we apply the region-based method, cylinder grouping, used in [1, 54], into the model. Specifically, for each of the point feature outputted by the backbone, we first predict the grasp view, then group the points near the grasp center and along with the grasp view, forming a more discriminate feature to predict the remaining parts of the grasp pose. As shown in the second line in Table 1, region-based method largely improve the model performance. Furthermore, we apply the global statistic proposed in [54] (the graspness identifier) as an extra task, to identify which points are suitable to grasp, which also improves the model performance. Finally, following many previous works [2, 5, 54], we use the 3D convolution [16] instead of PointNet++ [40], which also enhances the grasp performance. Combining all together, a vanilla grasp framework is proposed, which use sparse supervision to get satisfactory performance. However, there still a huge gap between the vanilla grasp framework and the SOTA method [54] (achieve 61.19, 47.39, 19.01 performance in seen, similar and novel), where the only main difference between them are the supervision modes. To sum up, we have the following conclusion:

Cutting-edge approaches help to bridge the gap between sparse and dense supervision methods. But, a significant performance gap still remains.

3.2 Economic Supervision

Ambiguity Problem. Now we further explore what problems cause the remain performance gap. An interesting idea comes to our mind when we visualize the




Fig. 3: In specific points, there exist many good grasps that with **different poses**. If we reduce the supervision, it may cause ambiguity in the learning process.

good grasps selected in each objects. As shown in Fig. 3, we observe that the good grasps in the specific points in different objects are in different poses (views, angles and depths are all different). In GraspNet-1Billion dataset [11], it labels all the possible views, angles and depths ($300 \times 12 \times 4$) for each point in the object. If we attempt to reduce the labels in each point, it may only keep some of the probabilities within the point. In this case, similar points, especially in neighbor or the same points, may have different labels as supervision, which will confuse the neural networks training. We call this the **ambiguity problem**.

To further quantitatively verify this, we analyze the standard deviation of the view, angle and depth of the good grasps in each point. As shown in the first line in Table 2, the good grasp poses are actually diversified and may cause ambiguity problem when we reduce them. In addition, we can also see that grasp views are the most ambiguous part, which has the largest standard variance due to its wide range. Furthermore, we are curious that if we already determine the grasp view, angle or depth beforehand, how ambiguous will be of the remaining parts of the grasp pose. To this end, we test the standard deviation of the good grasps after determining the view, angle or depths in each point. The results are demonstrated from the two to four lines in Table 2. We can see that if the grasp view is determined beforehand, the angle and depth are also nearly fixed, indicating that the ambiguity problem is basically solved. In contrast, although determining the grasp angle or grasp depth can somehow makes the view less diverse, it is still variant. In summary, we can give the following conclusion:

Table 2: The standard deviation of good grasps within a point.

	view	angle	depth
original	46.81	3.18	0.81
determine the view	-	0.22	0.05
determine the angle	13.45	-	0.27
determine the depth	33.81	2.28	-

The ambiguity problem exists when reducing the supervision, which causes the gap between sparse and dense supervision methods. Moreover, grasp view is the most important part to mitigate this problem.

Economic Supervision Paradigm. Based on the above analysis, we recommend to keep all the view as supervision rather than only keeping one grasp in each point as in the vanilla grasp framework, which can largely mitigate the ambiguity problem. Under this consensus, we propose the economic supervision paradigm to implement this proposal, including the a well-designed supervision selection strategy to reduce supervision, and a pipeline to enable its training under this type of supervision.

First, we detail the procedure of the selection for economic supervision. The first step is grasp pose pruning. We keep the grasp view following the above recommend, and prune the grasp angle and grasp depth for each point. In detail, we only keep the view graspness [54] and the best grasp for each view (300 grasps for each point instead of previous $300 \times 12 \times 4$ grasp poses for each point). Second, we collect the grasp labels of all the objects belonging to each scene to form the scene-level labels, instead of constructing them during training. This process will increase the storage size of the label set. But due to the grasp pose pruning at the first step, this increase is acceptable and will improve data loading speed in training. The final step is point pruning. Since we have already constructed the scene-level labels in the second step, we can further prune the labels in the scene points without graspable grasp labels (with friction coefficient [11,29,37,51] lower than 0.8 or with collision in the scene). This step will reduce the label set to 1/2 of the original. After the above three steps, we can reduce the label set from **55 GB to 1.6 GB**, becoming lower than 1/30 of the original, and meanwhile construct the scene-level labels in advance. Therefore, we can rapidly load the labels when training, which also has less memory and GPU costs, largely economizing the training overload, and meanwhile enabling fast and stable convergence.

To enable the model to train under the above economic supervision, we just need to slightly revise the vanilla grasp framework describe in Sec. 3.1. First, we predict the view scores for each point, and then select the best view based on the predicted scores. Next, we generate the remaining parts of the grasp pose for the selected view. The grasp view training is based on the view graspness [54]. Since we already keep a grasp for each view, we can use the corresponding label to supervise the learning of the remaining parts (angle, depth, width and score). Besides, since we also prune the labels in points without good grasps, the labels are hard to match all the input points (we may not have suitable labels close enough to supervise the training of the grasps of some the input points). Thus, we use a selective match loss, masking the loss for the grasps of input points without suitable labels and preventing back-propagating its loss.

3.3 Focal Representation under Economic Supervision

After economic supervision, we have the chance to give comprehensive consideration to learn a specific grasp. Thus, we design the focal representation module, comprising an interactive grasp head and a composite score estimation module to learn the specific grasp more accurately.

Interactive Grasp Head. The insight of the interactive grasp head is that since we only keep one grasp for each view, we aim to facilitate more interaction among the feature, to be more aligned with the specific grasp. Thus, we design a global and local interactive attentions [53] to achieve this goal.

First, the region-based method described in Sec. 3.1 will output the feature that can represent the region near the point along a specific view, which basically contains the range of all possible grasps to this point along the selected view. However, since we only aim to learn a specific good grasp in this range under economic supervision, the region is still too broad for our task. Therefore, we

design a global interactive attention to learn interaction within the region, which conducts the attention within the regions to construct a more discriminate feature for the learning the specific grasp. After that, following the previous object detection methods [14, 45, 61], we use multiple heads to learn the remaining parts of the grasp r, d, w, s from different features, instead of directly learning them from a common feature space.

Second, the remaining parts of the grasp are dependent, and this dependent relation is not determined beforehand. To be specific, sometimes when we have determined the angle, the depth naturally becomes fixed. However, this is not always the case. In some instances, establishing the depth first is better, which can result in a nearly-fixed angle. To learn this dependent relation, we use a local attention module to enable it, which executes the attention within the four grasp features. For clarity, we demonstrate the interactive grasp head in Fig. 4. To be noted that the global interactive attention is conducted among the region points, following a pooling layer to fuse them into an unified feature, and the local interactive attention is executed among different parts of the grasp pose.

This design is mainly for economic supervision. In previous dense supervision methods, they directly predict the combination of angles and depths, which is hard to learn features focusing on all the grasp poses economically.

Composite Score Estimation. What’s more, we want to predict the score more accurately for the specific grasp. Inspired by the idea from action quality assessment [50], which aims to predict a score to reflect the quality of the actions, we can model the score by classification. To be specific, we model the score by simply $1.1 - \mu$, where μ is the friction coefficient [11, 29, 37, 51]. In this way, the scores will have six value (from 0 to 1 with $\Delta 0.2$) and can be learnt by a six-class classifier. For inference, naive approach is to predict the score using the value corresponding to the highest probability. However, this cannot fully leverage the advantage of learning the score distribution using classification. Instead, we use the composite score estimation, linearly combining the scores in each segment by the probability, which can be formulated as $s = [0, 0.2, 0.4, 0.6, 0.8, 1] \cdot \mathbf{s}_c^T$, where \mathbf{s}_c is the output of the score classifier. The composite score estimation is highly necessary for predicting accurate score, which is evaluated in Sec. 4.3.

To be noted that this idea is also tailored for economic supervision. If we introduce this idea into the dense framework, the composite score estimation should transform the regression head into a classification head, this will make the angle-depth combination in dense supervision methods six times larger, which is highly uneconomical.

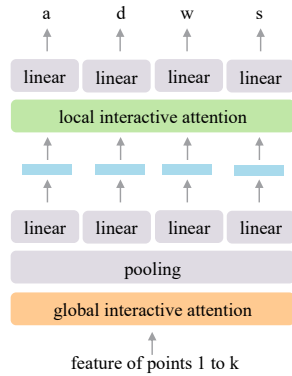


Fig. 4: Interactive grasp head.

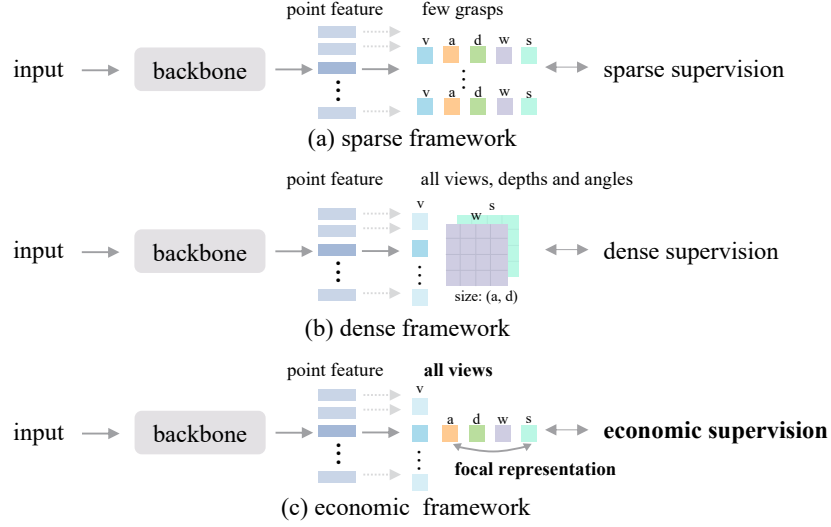


Fig. 5: Frameworks overview. Benefit from the economic supervision and focal representation, our economic framework achieves effective performance with low costs. We highlight the main contributions of our paper in bold face.

3.4 Framework Overall

The overall economic framework is described in Fig. 5. Previous sparse framework predicts one or few grasps for each point. However, this learning is somehow ambiguous since that the output just contains partial probabilities of all the good grasp, which results in limited performance, as analyzed in Sec.3.2. The dense framework will output and learn the combination of grasp angles and grasp depths, which has high resource costs and slow convergence speed, as shown in Fig. 1. Our economic framework innovatively keeps all views to mitigate the ambiguity problem, and design a focal representation module to generate the specific grasp more accurately.

3.5 Dataset and Details

Dataset. We use GraspNet-1Billion [11] for training and evaluation. GraspNet-1Billion [11] is a large-scale real-world dataset with 88 objects and 190 scenes captured by two cameras (Kniect and RealScenes), which can reflect the real-world grasp performance accurately [10, 54]. There are 100 scenes as the training set and 90 scenes to be the testing set. In addition, the testing set are further equally divided into three splits for seen, similar and novel objects respectively. To be noted that this is the only dense real-world dataset to fully analyze the performance and costs of our framework with other sparse or dense frameworks.

Metric. As for the metric, we adopt AP as the metric following GraspNet-1Billion [11]. In detail, we first generate 6-DoF grasps by our model, then evaluate the grasps with high scores by reconstructing the test scenes and calculating the

Table 3: Performance comparison in Kinect camera. Best results are in bold face. We highlight the important comparison metrics by gray cell.

supervision	methods	seen			similar			novel		
		AP	AP _{0.8}	AP _{0.4}	AP	AP _{0.8}	AP _{0.4}	AP	AP _{0.8}	AP _{0.4}
sample	GPD [51]	24.38	30.16	13.46	23.18	28.64	11.32	9.58	10.14	3.16
	PointNetGPD [29]	27.59	34.21	17.83	24.38	30.84	12.83	10.66	11.24	3.21
sparse	S4G [42]	18.72	22.90	14.58	11.82	14.74	7.67	5.38	6.69	2.19
	TransGrasp [32]	35.97	41.69	31.86	29.71	35.67	24.19	11.41	14.42	5.87
	GraNet [55]	41.38	49.84	33.86	35.29	43.15	26.89	11.57	14.31	5.24
	GraspNet Baseline [11]	29.88	36.19	19.31	27.84	33.19	16.62	11.51	12.92	3.56
dense	TSB [34]	49.42	58.38	43.66	41.49	50.29	34.60	15.35	19.18	7.98
	GSNet [54]	61.19	71.46	56.04	47.39	56.78	40.43	19.01	23.73	10.60
economic	EconomicGrasp	62.59	73.89	55.99	51.73	62.70	43.45	19.54	24.24	11.12

Table 4: Performance comparison in RealSense camera. Best results are in bold face. We highlight the important comparison metrics by gray cell.

supervision	methods	seen			similar			novel		
		AP	AP _{0.8}	AP _{0.4}	AP	AP _{0.8}	AP _{0.4}	AP	AP _{0.8}	AP _{0.4}
sample	GPD [51]	22.87	28.53	12.84	21.33	27.83	9.64	8.24	8.89	2.67
	PointNetGPD [29]	25.96	33.01	15.37	22.68	29.15	10.76	9.23	9.89	2.74
sparse	S4G [42]	25.71	31.19	20.84	18.45	23.14	12.86	9.04	11.23	3.61
	TransGrasp [32]	39.81	47.54	36.42	29.32	34.80	25.19	13.83	17.11	7.67
	GraNet [55]	43.33	52.56	34.04	39.98	48.66	32.00	14.90	18.66	7.76
	GraspNet Baseline [11]	27.56	33.43	16.95	26.11	34.19	14.23	10.55	11.25	3.98
dense	TSB [34]	58.95	68.18	54.88	52.97	63.24	46.99	22.63	28.53	12.00
	GSNet [54]	65.70	76.25	61.08	53.75	65.04	45.97	23.98	29.93	14.05
economic	EconomicGrasp	68.21	79.60	63.54	61.19	73.60	53.77	25.48	31.46	13.85

force-closure metric [11, 29, 37, 51] by giving different friction coefficients μ . By doing so, we will get a binary metric to indicate whether each grasp can successfully grab the object under the friction coefficient μ . AP_μ means the average success rate of the predicted top-k grasps with k ranging from 1 to 50 under μ , and AP is the average of AP_μ for μ ranging from 0.2 to 1.0 with $\Delta\mu = 0.2$.

Implementation Details. We now describe the implementation details of our EconomicGrasp framework. We use a 14-layers 3D UNet as our backbone, which is implemented by the Minkowski Engine [6]. The output feature dimension is 512 and the hidden dimensions are [32, 64, 128, 256, 192, 192, 192, 192]. For the design of the loss function, we use the smooth L1 loss for view prediction. Regarding the angle, depths, widths and score, we use cross entropy for angle, depth and score learning and smooth L1 loss for width learning. In terms of the optimization, we use Adam [25] as our optimizer with start learning rate $1e^{-3}$, and use the cosine learning rate decay schedule for training. The training batch size is 4 and the training epoch is 10. All the models are implemented with PyTorch [39] on NVIDIA RTX3090 GPUs.

4 Experiments

4.1 Grasp Performance

To evaluate the performance of our economic framework, we conduct performance study comparing with other 6-DoF grasping methods, including GPD [51],

Table 5: Training resource costs comparison. Best results are in bold face.

supervision	methods	training			mAP \uparrow	
		time (h) \downarrow	memory (G) \downarrow	GPUs (G) \downarrow	Kinect	RealSense
sparse	TransGrasp [32]	41.2	4.8	7.61	25.69	29.48
	GraNet [55]	75.6	48.2	8.02	29.41	35.63
dense	TSB [34]	63.6	51.4	16.21	35.42	48.27
	GSNet [54]	37.8	35.4	9.15	42.53	47.81
economic	EconomicGrasp	8.3	4.2	5.81	44.62	51.63

PointNetGPD [29], S4G [42], GraspNet-Baseline [11], TransGrasp [32], GraNet [55], TSB [34], GSNet [54]. Among them, GPD [51], PointNetGPD [29] are sample-based methods that trains by sampling different grasps and evaluating its score. S4G [42], TransGrasp [32] and GraNet [55] are sparse supervision methods that trains on a partial of the GraspNet-1billion dataset, which output one or few grasps for each point. GraspNet-Baseline [11], TSB [34] and GSNet [54] are dense supervision methods that trains on dense labels, which learns the grasp from all the (300 view \times 12 rotation \times 4 depth) grasp labels for each points.

The testing results of Kinect and RealSense cameras are shown in Fig. 3 and Fig. 4 respectively. From the table we can see that our EconomicGrasp outperforms all the SOTA methods. For example, we outperform GSNet [54] by 1.40, 4.32, 0.53 AP increases on seen, similar and novel scenes in Kinect camera, and by 2.51, 7.42, 1.50 AP increases on seen, similar and novel scenes in RealSense camera. On average, there is about 3 AP enhancement, showing great potential for the economic framework. The improvement could be attributed to the following reasons: 1) our economic framework innovatively maintains grasp views for training, which mitigates the ambiguity problem; 2) our focal representation module is effective and specifically designed for our economic supervision. The result shows great potential for the economic framework.

4.2 Training Resource Cost

To evaluate the economy of our economic framework, we also conduct the training costs experiments comparing with some effective 6-DoF grasping methods, including TransGrasp [32], GraNet [55], TSB [34] and the current SOTA dense supervision method GSNet [54]. Specifically, we test the resources of total training time, memory costs and GPU costs for training the GraspNet-1Billion dataset on Kinect data. All the experiments are conducted in an empty machine without other running processes for fair comparison.

As shown in Table 5, we can see that our method has the lowest training resource costs in training time, memory utilization and GPU utilization comparing to other methods, meanwhile surpasses all methods in grasp performance. For example, comparing to the SOTA method GSNet [54], our method has 3 AP improvement on average in performance, and only uses 1/4 training time, 1/8 memory costs. This is due to our economic framework that using less but important supervision to train the model, and design a tailored focal representation module for the economic training.

Table 6: Real world experiments. Best is in bold face.

IDs	objects	GSNet [54]		EconomicGrasp	
		attempts	success	attempts	success
1, 4, 9, 11, 15	5	6	83.3%	5	100%
2, 3, 6, 7, 12	5	5	100%	6	83.3%
5, 8, 10, 11, 13, 14	6	7	85.2%	7	85.2%
1, 2, 3, 11, 12, 15	6	7	85.2%	6	100%
5, 6, 7, 9, 13, 14, 15	7	8	87.5%	8	87.5%
1, 3, 6, 8, 10, 11, 12	7	8	87.5%	7	100%
total	36	41	87.8%	39	92.3%

Table 7: Analysis of bad cases. It shows the number of testing scenes without any successful grasps in the generated top50 grasps. There are 7680 scenes in total. Best results are in bold face.

GSNet [54]	EconomicGrasp
109/7680	36 /7680

Table 8: Main ablation study. Important comparison metric is in gray cell.

economic supervision	interactive head	composite score	seen	similar	novel	mean
			43.59	34.09	13.36	30.34
✓			60.07	48.16	18.70	42.31
✓	✓		63.08	50.61	18.74	44.14
✓		✓	59.81	48.45	19.01	42.42
✓	✓	✓	62.59	51.73	19.54	44.62

Table 9: Ablation study of the economic supervision paradigm. Important comparison metric is in gray cell.

selective loss	seen	similar	novel	mean
	59.67	48.68	19.27	42.54
✓	62.59	51.73	19.54	44.62

4.3 Ablation Study

Moreover, we conduct a series of ablation studies to test the effectiveness of each design of our framework. All the experiments are trained on GraspNet-1Billion dataset [11] on the Kinect data.

Main Ablation Study. We first conduct ablation study to test the main components of our framework. As shown in the first line in Table 6, the economic supervision paradigm gains the major improvements for our method, since it mitigates the key ambiguity problem when reducing the supervision. Moreover, from the second line in Table 6, we can see that interactive head also helps to enhance the performance by 1.8 mean AP, showing that interactively learning focal and discriminate features for a specific grasp is useful. However, as shown in the third line in Table 6, the composite score estimation cannot amplify the performance along. But with the help of the interactive head to learn focal representation for the specific grasp, the composite score can be used together with the interactive grasp head and improve the performance, as shown in the forth line in Table 6. Due to the space limit, the ablation study for resource costs of the main components is put in the supplementary material.

Other Ablation Studies. We also conduct other ablation studies for each component. As shown in Table 9, the selective loss can improve the model performance with about 2.08 AP, indicating that masking the unmatched points for loss calculation is important for training under economic supervision. In addition, the ablation study for the interactive grasp head is shown in Table 10. It demonstrates that by gradually adding the global interactive attention

Table 10: Ablation study of the interactive **Table 11:** Ablation study of the compos-grasp head. Important comparison metric is its score estimation. Important comparison metric is in gray cell.

global interactive	local interactive	seen	similar	novel	mean
		59.81	48.45	19.01	42.42
✓		61.96	49.09	19.44	43.49
✓	✓	62.59	51.73	19.54	44.62

classification	composite estimation	seen	similar	novel	mean
		63.08	50.61	18.74	44.14
✓		51.66	38.71	14.58	34.98
✓	✓	62.59	51.73	19.54	44.62




Fig. 6: Qualitative analysis.

and the local interactive attention, the model performance can be elevated step by step, showing that learning interaction to focus on a specific grasp is effective for economic supervision. Furthermore, the ablation study of the composite score estimation is demonstrated in Table 11. From the table, we can observe that composite score estimation at inference is necessary for predicting accurate scores. Predicting the score without it will cause nearly 10 AP decrease. We also conduct several ablation studies about keeping different numbers of views and comparing different types of sparse supervision, which can be seen in the supplementary material.

4.4 Real-world Testing

For a deeper assessment of our performance, we further conduct the real-world experiments comparing with the SOTA method GSNet [54]. In detail, we arrange six scenes with 5,5,6,6,7,7 objects and test the number of grasps required to empty the scenes (a pick and place task). The experiments are conducted by Franka Emika robot hand with a two-finger gripper. The experiment setup and the testing objects are shown in Fig. 7, which contains 4 trained objects (NO. 3, 4, 6, 10) and 11 untrained objects (NO. 1, 2, 5, 7, 8, 9, 11, 12, 13, 14, 15). The testing results are shown in Table 6. It can be seen that our method can achieve robust real-world grasping ability, slightly surpassing the SOTA method GSNet [54] in the real world experiments, showing great poten-




Fig. 7: Real-world experiments setup.

tial for our economic framework. We also make a video demo, showing the real world grasping practice with our economic framework, which can be found in the supplementary material.

4.5 Analysis

Qualitative Analysis. We conduct qualitative analysis to further evaluate the robustness and effectiveness of our method. We randomly sample some scenes from the test set of GraspNet-1Billion dataset [1] and visualize the top30 grasps generated by our method. The results are shown in Fig. 6. From the figure, we can see that the grasps generated by our method are robust and basically reasonable, demonstrating the effectiveness of our economic framework. In addition, the generated grasp poses are diverse and flexible, showing the effectiveness of 6-DoF grasping. We also test the diversity of our generated grasp in the supplementary material to show that we maintain the diversity although we reduce the dataset.

Failure Cases Analysis. Additionally, we also calculate the number of failure cases of our model comparing with the SOTA method GSNet [54]. To be specific, we analyze the number of scenes in the seen test set that has no successful grasps in the predicted top50 grasps with low friction coefficients $\mu = 0.2$, which can evaluate the high-quality grasps generating ability of our method.

As demonstrated in Table 7, we can see that in the total of 7680 scenes, there are 109 failure cases for GSNet [54], and only 36 failure cases for our model, which indicates that selecting high-quality labels to supervise the learning is to some extent robust to generate high-quality grasps, which is also showing the great potential for the exploration of economic grasping. We also conduct plug-and-play experiments in the supplementary material, showing great potential for our economic procedure.

5 Conclusion

In this work, we propose an economic framework for 6-DoF grasp detection, to economize resource costs in training and maintain effective grasp performance. We find that the dense supervision is the bottleneck to encumber the entire training overload, and meanwhile make the training hard to converge. Toward this end, we propose the economic supervision paradigm to reduce the supervision and meanwhile maintain comparable performance. Furthermore, economic supervision gives us the chance to focus on the learning of the specific grasp, thus we devise a focal representation module to enable it. Combining all together, we propose the EconomicGrasp framework. Our extensive experiments show that our framework surpasses the SOTA grasping method with low resources cost (1/4 training time cost, 1/8 memory cost and 1/25 storage cost), showing great potential for the economic framework.

In the future, exploring how to build economic grasping supervision from natural, instead of simplifying from a dense dataset, will be interesting and meaningful, which can be the future work to do. And we believe this work can also give insights to the further exploration about it.

Acknowledgements

This work was supported partially by the National Key Research and Development Program of China (2023YFA1008503), NSFC(U21A20471), Guangdong NSF Project (No. 2023B1515040025, 2020B1515120085). Additionally, I sincerely thank Jia-Run Du's help for the valuable suggestions for the paper.

References

1. Asif, U., Tang, J., Harrer, S.: Graspnet: An efficient convolutional neural network for real-time grasp detection for low-powered devices. In: International Joint Conferences on Artificial Intelligence (2018) 4, 5, 14
2. Breyer, M., Chung, J.J., Ott, L., Siegwart, R., Nieto, J.: Volumetric grasping network: Real-time 6 dof grasp detection in clutter. In: Conference on Robot Learning (2021) 1, 2, 4, 5
3. Cai, J., Cen, J., Wang, H., Wang, M.Y.: Real-time collision-free grasp pose detection with geometry-aware refinement using high-resolution volume. IEEE Robotics and Automation Letters (2022) 4
4. Chen, S., Tang, W., Xie, P., Yang, W., Wang, G.: Efficient heatmap-guided 6-dof grasp detection in cluttered scenes. IEEE Robotics and Automation Letters (2023) 5
5. Chen, Z., Liu, Z., Xie, S., Zheng, W.S.: Grasp region exploration for 7-dof robotic grasping in cluttered scenes. In: IEEE/RSJ International Conference on Intelligent Robots and Systems (2023) 3, 4, 5
6. Choy, C., Gwak, J., Savarese, S.: 4d spatio-temporal convnets: Minkowski convolutional neural networks. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (2019) 10
7. Correll, N., Bekris, K.E., Berenson, D., Brock, O., Causo, A., Hauser, K., Okada, K., Rodriguez, A., Romano, J.M., Wurman, P.R.: Analysis and observations from the first amazon picking challenge. IEEE Transactions on Automation Science and Engineering (2016) 1
8. Dai, Q., Zhu, Y., Geng, Y., Ruan, C., Zhang, J., Wang, H.: Graspnerf: Multiview-based 6-dof grasp detection for transparent and specular objects using generalizable nerf. In: IEEE International Conference on Robotics and Automation (2023) 4
9. Eppner, C., Mousavian, A., Fox, D.: Acronym: A large-scale grasp dataset based on simulation. In: IEEE International Conference on Robotics and Automation (2021) 3
10. Fang, H.S., Wang, C., Fang, H., Gou, M., Liu, J., Yan, H., Liu, W., Xie, Y., Lu, C.: Anygrasp: Robust and efficient grasp perception in spatial and temporal domains. IEEE Transactions on Robotics (2023) 1, 2, 3, 4, 9
11. Fang, H.S., Wang, C., Gou, M., Lu, C.: Graspnet-1billion: A large-scale benchmark for general object grasping. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (2020) 1, 2, 3, 4, 5, 6, 7, 8, 9, 10, 11, 12
12. Fischinger, D., Vincze, M.: Empty the basket-a shape based learning approach for grasping piles of unknown objects. In: IEEE/RSJ International Conference on Intelligent Robots and Systems (2012) 3
13. Fu, Z., Zhao, T.Z., Finn, C.: Mobile aloha: Learning bimanual mobile manipulation with low-cost whole-body teleoperation. arXiv preprint arXiv:2401.02117 (2024) 1, 3

14. Girshick, R.: Fast r-cnn. In: Proceedings of the IEEE International Conference on Computer Vision (2015) [8](#)
15. Gou, M., Fang, H.S., Zhu, Z., Xu, S., Wang, C., Lu, C.: Rgb matters: Learning 7-dof grasp poses on monocular rgbd images. In: IEEE International Conference on Robotics and Automation (2021) [2](#), [4](#)
16. Graham, B., Engelcke, M., Van Der Maaten, L.: 3d semantic segmentation with submanifold sparse convolutional networks. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (2018) [5](#)
17. Gualtieri, M., Ten Pas, A., Saenko, K., Platt, R.: High precision grasp pose detection in dense clutter. In: IEEE/RSJ International Conference on Intelligent Robots and Systems (2016) [3](#)
18. Herzog, A., Pastor, P., Kalakrishnan, M., Righetti, L., Asfour, T., Schaal, S.: Template-based learning of grasp selection. In: IEEE International Conference on Robotics and Automation (2012) [3](#)
19. Hoang, D.C., Stork, J.A., Stoyanov, T.: Context-aware grasp generation in cluttered scenes. In: International Conference on Robotics and Automation (2022) [4](#)
20. Huang, S., Wang, Z., Li, P., Jia, B., Liu, T., Zhu, Y., Liang, W., Zhu, S.C.: Diffusion-based generation, optimization, and planning in 3d scenes. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (2023) [3](#)
21. Jauhri, S., Lunawat, I., Chalvatzaki, G.: Learning any-view 6dof robotic grasping in cluttered scenes via neural surface rendering. arXiv preprint arXiv:2306.07392 (2023) [4](#)
22. Jiang, H., Liu, S., Wang, J., Wang, X.: Hand-object contact consistency reasoning for human grasps generation. In: Proceedings of the IEEE/CVF International Conference on Computer Vision (2021) [3](#)
23. Jiang, Y., Moseson, S., Saxena, A.: Efficient grasping from rgbd images: Learning using a new rectangle representation. In: IEEE International Conference on Robotics and Automation (2011) [3](#)
24. Kappler, D., Bohg, J., Schaal, S.: Leveraging big data for grasp planning. In: IEEE International Conference on Robotics and Automation (2015) [3](#)
25. Kingma, D.P., Ba, J.: Adam: A method for stochastic optimization. arXiv preprint arXiv:1412.6980 (2014) [10](#)
26. Le, Q.V., Kamm, D., Kara, A.F., Ng, A.Y.: Learning to grasp objects with multiple contact points. In: IEEE International Conference on Robotics and Automation (2010) [3](#)
27. Lenz, I., Lee, H., Saxena, A.: Deep learning for detecting robotic grasps. The International Journal of Robotics Research (2015) [3](#)
28. Li, Y., Kong, T., Chu, R., Li, Y., Wang, P., Li, L.: Simultaneous semantic and collision learning for 6-dof grasp pose estimation. In: IEEE/RSJ International Conference on Intelligent Robots and Systems (2021) [4](#)
29. Liang, H., Ma, X., Li, S., Görner, M., Tang, S., Fang, B., Sun, F., Zhang, J.: Pointnetgpd: Detecting grasp configurations from point sets. In: International Conference on Robotics and Automation (2019) [2](#), [3](#), [4](#), [7](#), [8](#), [10](#), [11](#)
30. Liu, C., Shi, K., Zhou, K., Wang, H., Zhang, J., Dong, H.: Rgbgrasp: Image-based object grasping by capturing multiple views during robot arm movement with neural radiance fields. arXiv preprint arXiv:2311.16592 (2023) [4](#)
31. Liu, M., Pan, Z., Xu, K., Ganguly, K., Manocha, D.: Deep differentiable grasp planner for high-dof grippers. arXiv preprint arXiv:2002.01530 (2020) [3](#)

32. Liu, Z., Chen, Z., Xie, S., Zheng, W.S.: Transgrasp: A multi-scale hierarchical point transformer for 7-dof grasp detection. In: International Conference on Robotics and Automation (2022) [4](#), [10](#), [11](#)
33. Liu, Z., Chen, Z., Zheng, W.S.: Simulating complete points representations for single-view 6-dof grasp detection. IEEE Robotics and Automation Letters (2024) [3](#), [4](#)
34. Ma, H., Huang, D.: Towards scale balanced 6-dof grasp detection in cluttered scenes. In: Conference on Robot Learning (2023) [2](#), [4](#), [10](#), [11](#)
35. Morrison, D., Corke, P., Leitner, J.: Closing the loop for robotic grasping: A real-time, generative grasp synthesis approach. arXiv preprint arXiv:1804.05172 (2018) [3](#)
36. Mousavian, A., Eppner, C., Fox, D.: 6-dof graspnet: Variational grasp generation for object manipulation. In: Proceedings of the IEEE/CVF International Conference on Computer Vision (2019) [3](#)
37. Nguyen, V.D.: Constructing force-closure grasps. The International Journal of Robotics Research (1988) [7](#), [8](#), [10](#)
38. Ni, P., Zhang, W., Zhu, X., Cao, Q.: Pointnet++ grasping: Learning an end-to-end spatial grasp generation algorithm from sparse point clouds. In: IEEE International Conference on Robotics and Automation (2020) [2](#), [4](#), [5](#)
39. Paszke, A., Gross, S., Massa, F., Lerer, A., Bradbury, J., Chanan, G., Killeen, T., Lin, Z., Gimelshein, N., Antiga, L., et al.: Pytorch: An imperative style, high-performance deep learning library. Advances in Neural Information Processing Systems (2019) [10](#)
40. Qi, C.R., Yi, L., Su, H., Guibas, L.J.: Pointnet++: Deep hierarchical feature learning on point sets in a metric space (2017) [5](#)
41. Qin, R., Ma, H., Gao, B., Huang, D.: Rgb-d grasp detection via depth guided learning with cross-modal attention. arXiv preprint arXiv:2302.14264 (2023) [4](#)
42. Qin, Y., Chen, R., Zhu, H., Song, M., Xu, J., Su, H.: S4g: Amodal single-view single-shot se (3) grasp detection in cluttered scenes. In: Conference on robot learning (2020) [2](#), [4](#), [5](#), [10](#), [11](#)
43. Quigley, M., Conley, K., Gerkey, B., Faust, J., Foote, T., Leibs, J., Wheeler, R., Ng, A.Y., et al.: Ros: an open-source robot operating system. In: ICRA workshop on open source software. vol. 3, p. 5. Kobe, Japan (2009) [3](#)
44. Redmon, J., Angelova, A.: Real-time grasp detection using convolutional neural networks. In: IEEE International Conference on Robotics and Automation (2015) [3](#)
45. Ren, S., He, K., Girshick, R., Sun, J.: Faster r-cnn: Towards real-time object detection with region proposal networks (2015) [8](#)
46. Saxena, A., Driemeyer, J., Ng, A.Y.: Robotic grasping of novel objects using vision. The International Journal of Robotics Research (2008) [3](#)
47. Shi, Y., Tang, Z., Cai, X., Zhang, H., Hu, D., Xu, X.: Symmetrygrasp: Symmetry-aware antipodal grasp detection from single-view rgb-d images. IEEE Robotics and Automation Letters (2022) [4](#)
48. Suárez-Ruiz, F., Zhou, X., Pham, Q.C.: Can robots assemble an ikea chair? Science Robotics (2018) [1](#)
49. Sundermeyer, M., Mousavian, A., Triebel, R., Fox, D.: Contact-graspnet: Efficient 6-dof grasp generation in cluttered scenes. In: IEEE International Conference on Robotics and Automation (2021) [3](#), [4](#)
50. Tang, Y., Ni, Z., Zhou, J., Zhang, D., Lu, J., Wu, Y., Zhou, J.: Uncertainty-aware score distribution learning for action quality assessment. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (2020) [8](#)

51. Ten Pas, A., Gualtieri, M., Saenko, K., Platt, R.: Grasp pose detection in point clouds. *The International Journal of Robotics Research* (2017) [1](#), [2](#), [3](#), [4](#), [7](#), [8](#), [10](#), [11](#)
52. Ten Pas, A., Platt, R.: Using geometry to detect grasp poses in 3d point clouds. *Robotics Research* (2018) [3](#)
53. Vaswani, A., Shazeer, N., Parmar, N., Uszkoreit, J., Jones, L., Gomez, A.N., Kaiser, Ł., Polosukhin, I.: Attention is all you need (2017) [7](#)
54. Wang, C., Fang, H.S., Gou, M., Fang, H., Gao, J., Lu, C.: Graspness discovery in clutters for fast and accurate grasp detection. In: *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition* (2021) [2](#), [3](#), [4](#), [5](#), [7](#), [9](#), [10](#), [11](#), [12](#), [13](#), [14](#)
55. Wang, H., Niu, W., Zhuang, C.: Granet: A multi-level graph network for 6-dof grasp pose generation in cluttered scenes. In: *IEEE/RSJ International Conference on Intelligent Robots and Systems* (2023) [4](#), [10](#), [11](#)
56. Wang, R., Zhang, J., Chen, J., Xu, Y., Li, P., Liu, T., Wang, H.: Dexgraspnet: A large-scale robotic dexterous grasp dataset for general objects based on simulation. In: *IEEE International Conference on Robotics and Automation* (2023) [3](#)
57. Wang, Y.K., Xing, C., Wei, Y.L., Wu, X.M., Zheng, W.S.: Single-view scene point cloud human grasp generation. In: *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition* (2024) [3](#)
58. Wei, W., Luo, Y., Li, F., Xu, G., Zhong, J., Li, W., Wang, P.: Gpr: Grasp pose refinement network for cluttered scenes. In: *IEEE International Conference on Robotics and Automation* (2021) [4](#)
59. Wei, Y.L., Jiang, J.J., Xing, C., Tan, X., Wu, X.M., Li, H., Cutkosky, M., Zheng, W.S.: Grasp as you say: Language-guided dexterous grasp generation. *arXiv preprint arXiv:2405.19291* (2024) [3](#)
60. Weng, T., Held, D., Meier, F., Mukadam, M.: Neural grasp distance fields for robot manipulation. In: *IEEE International Conference on Robotics and Automation* (2023) [4](#)
61. Wu, Y., Chen, Y., Yuan, L., Liu, Z., Wang, L., Li, H., Fu, Y.: Rethinking classification and localization for object detection. In: *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition* (2020) [8](#)
62. Xu, G.H., Wei, Y.L., Zheng, D., Wu, X.M., Zheng, W.S.: Dexterous grasp transformer. In: *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition* (2024) [3](#)
63. Xu, Y., Wan, W., Zhang, J., Liu, H., Shan, Z., Shen, H., Wang, R., Geng, H., Weng, Y., Chen, J., et al.: Unidexgrasp: Universal robotic dexterous grasping via learning diverse proposal generation and goal-conditioned policy. In: *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition* (2023) [3](#)
64. Yu, S., Zhai, D.H., Xia, Y.: Egnet: Efficient robotic grasp detection network. *IEEE Transactions on Industrial Electronics* (2022) [3](#)
65. Zhai, G., Huang, D., Wu, S.C., Jung, H., Di, Y., Manhardt, F., Tombari, F., Navab, N., Busam, B.: Monograspnet: 6-dof grasping with a single rgb image. In: *IEEE International Conference on Robotics and Automation* (2023) [4](#)
66. Zhao, B., Zhang, H., Lan, X., Wang, H., Tian, Z., Zheng, N.: Regnet: Region-based grasp network for end-to-end grasp detection in point clouds. In: *IEEE International Conference on Robotics and Automation* (2021) [3](#), [4](#), [5](#)
67. Zheng, L., Cai, Y., Lu, T., Wang, S.: Vgpn: 6-dof grasp pose detection network based on hough voting. In: *IEEE/RSJ International Conference on Intelligent Robots and Systems* (2022) [4](#)

- 68. Zhou, H., Wang, X., Au, W., Kang, H., Chen, C.: Intelligent robots for fruit harvesting: Recent developments and future challenges. *Precision Agriculture* (2022) 1, 3
- 69. Zhou, J., Ma, T., Lin, K.Y., Qiu, R., Wang, Z., Liang, J.: Mitigating the human-robot domain discrepancy in visual pre-training for robotic manipulation. *arXiv preprint arXiv:2406.14235* (2024) 3