# When Transformer Meets Robotic Grasping: Exploits Context for Efficient Grasp Detection

Shaochen Wang<sup>10</sup>, Zhangli Zhou, and Zhen Kan<sup>10</sup>, *Senior Member, IEEE* 

Abstract-In this letter, we present a transformer-based architecture, namely TF-Grasp, for robotic grasp detection. The developed TF-Grasp framework has two elaborate designs making it well suitable for visual grasping tasks. The first key design is that we adopt the local window attention to capture local contextual information and detailed features of graspable objects. Then, we apply the cross window attention to model the long-term dependencies between distant pixels. Object knowledge, environmental configuration, and relationships between different visual entities are aggregated for subsequent grasp detection. The second key design is that we build a hierarchical encoder-decoder architecture with skip-connections, delivering shallow features from the encoder to decoder to enable a multi-scale feature fusion. Due to the powerful attention mechanism, TF-Grasp can simultaneously obtain the local information (i.e., the contours of objects), and model long-term connections such as the relationships between distinct visual concepts in clutter. Extensive computational experiments demonstrate that TF-Grasp achieves competitive results versus state-of-art grasping convolutional models and attains a higher accuracy of 97.99% and 94.6% on Cornell and Jacquard grasping datasets, respectively. Real-world experiments using a 7DoF Franka Emika Panda robot also demonstrate its capability of grasping unseen objects in a variety of scenarios. The code is available at https://github.com/WangShaoSUN/grasp-transformer.

Index Terms—Grasp detection, robotic grasping, vision transformer.

## I. INTRODUCTION

ATA-DRIVEN methodologies such as deep learning have become the mainstream methods for robotic visual sensing tasks such as indoor localization [1], trajectory prediction [2], and robotic manipulation [3], [4], since they require less handcrafted feature engineering and can be extended to many complex tasks. In recent years, as visual sensing is increasingly being used in manufacturing, industry, and medical care, growing research is devoted to developing advanced robot's perception abilities. A typical application of visual sensing is the robotic grasp detection, where the images of objects are used to infer the grasping pose. Considering a grasping task of manipulating a wide diversity of objects, to find the graspable regions, the robots

Manuscript received 23 February 2022; accepted 20 June 2022. Date of publication 29 June 2022; date of current version 7 July 2022. This letter was recommended for publication by Associate Editor S. Jain and Editor M. Vincze upon evaluation of the reviewers' comments. This work was supported by the National Natural Science Foundation of China under Grants U2013601 and 62173314. (*Corresponding author: Zhen Kan.*)

The authors are with the Department of Automation, University of Science and Technology of China, Hefei 230026, China (e-mail: samwang@mail.ustc.edu.cn; zzl1215@mail.ustc.edu.cn; zkan@ustc.edu.cn).

Digital Object Identifier 10.1109/LRA.2022.3187261

have to concentrate on not only partial geometric information but also the entire visual appearance of the object. Particularly in unstructured and cluttered environments, dealing with variations in shape and position (e.g., occlusion) and also the spatial relationship with other objects are critical to the performance of grasp detection. Therefore, this work is particularly motivated to investigate grasp detection that takes into account both local neighbor pixels and long-distance relationships in spatial dimensions.

Most modern grasp detectors [3], [5] are based on convolutional neural networks (CNNs) which emerge as the de facto standard for processing visual robotic grasping. However, current CNNs are composed of individual convolution kernels, which are more inclined to concentrate on local level information. Also, the convolution kernels in a layer of CNN are viewed as independent counterparts without mutual information fusion. Generally, to maintain a large receptive field, CNNs have to repeatedly stack convolutional layers, which reduce the spatial resolution and inevitably results in the loss of global details and degraded performance.

Recently, as a novel approach to handle natural language processing and computer vision, the transformer [6]–[8] demonstrates remarkable success. The widely adopted attention mechanisms [6] of transformers in sequence modeling provide an elegant resolution that can better convey the fusion of information across global sequences. In fact, as robots are deployed in more and more diverse applications such as industrial assembly lines and smart home, the sensing capacity of robotic systems needs to be enriched, not only in local regions, but also in global interaction. Especially when robots frequently interact with objects in the environment, the awareness of global attention is particularly important with respect to safety and reliability. However, most vision transformers are designed for image classification on natural images processing tasks. Few of them are specifically built for robotic tasks.

In this letter, we present a transformer-based visual grasp detection framework, namely TF-Grasp, which leverages the fact that the attention can better aggregate information across the entire input sequences to obtain an improved global representation. More specifically, the information within independent image patches is bridged via self-attention and the encoder in our framework captures these multi-scale low-level features. The decoder incorporates the high-level features through long-range spatial dependencies to construct the final grasping pose. We provide detailed empirical evidence to show that our grasping transformer performs reasonably well on popular

2377-3766 © 2022 IEEE. Personal use is permitted, but republication/redistribution requires IEEE permission.

See https://www.ieee.org/publications/rights/index.html for more information.

grasping testbeds, e.g., Cornell and Jacquard grasping datasets. The experimental results demonstrate that the transformer architecture plays an integral role in generating appropriate grasping poses by learning local and global features from different parts of each object. The vision transformer-based grasp detection works well on the real robotic system and shows promising generalization to unseen objects. In addition, our TF-Grasp can generate the required grasping poses for parallel grippers in a single forward pass of the network.

In a nutshell, the contributions of this letter can be summarised in three folds:

- This work presents a novel and neat transformer architecture for visual robotic grasping tasks. To the best of our knowledge, it is one of the first attempts considering vision transformers in grasp detection tasks.
- We consider simultaneous fusion of local and global features and redesign the classical ViT framework for robotic visual sensing tasks.
- Exhaustive experiments are conducted to show the advantages of the transformer-based robotic perception framework. The experimental results demonstrate that our model achieves improved performance on popular grasping datasets compared to the state-of-the-art methods. We further show that our grasping transformer can generate appropriate grasping poses for known or unknown objects in either single or cluttered environments.

### II. RELATED WORK

This section reviews recent advances in the field of robotic grasping and briefly describes the progress of transformers in different areas.

## A. Grasp Detection

The ability to locate the object position and determine the appropriate grasping pose is crucial to stable and robust robotic grasping. Grasp detection, as the name implies, uses the image captured from the camera to infer the grasping pose for the robot manipulator. Using geometry-driven methods, earlier works [9], [10] mainly focus on analyzing the contours of objects to identify grasping points. A common assumption in these methods is that the geometric model of the object is always available. However, preparing the CAD models for graspable objects is time-consuming and impractical for real-time implementation. Recently, deep learning based methods have been successfully applied in visual grasping tasks [3], [5], [11]-[13]. The work of [14] is one of the earliest works that introduces deep neural networks to grasp detection via a two-stage strategy where the first stage finds exhaustive possible grasping candidates and the second stage evaluates the quality of these grasp candidates to identify the best one. However, due to numerous grasping proposals, the method in [14] suffers from relatively slow speed. Many recent works utilize convolutional neural networks to generate bounding box proposals to estimate the grasp pose of objects. Redmon et al. [5] employed an Alexnet-like CNN architecture to regress grasping poses. Kumra et al. [3] explored the use of ResNet-50 as a backbone to incorporate multimodal including depth and RGB information to further improve the grasp performance. Besides, CNN-based grasp quality networks [15], [16] were proposed to evaluate and predict the robustness of grasp candidates. In the same line, GG-CNN [17] developed a fully convolutional neural network to perform grasp detection, which provides a lightweight and real-time solution for visual grasping. Currently, most of the existing grasp detection methods are still heavily inspired by computer vision techniques such as object recognition, object detection, etc. In contrast to classical visual problems where the detected objects are usually well-defined instances in the scene, in grasp detection, the grasp configuration to be generated is continuous, which implies an infinite number of possible grasp options. This places significant challenges in feature extraction to identify a valid grasp configuration from all possible candidates. We argue that the loss of long-term dependencies in feature extraction is a major drawback of current CNNs based grasp detection methods.

#### B. Transformer

Transformer [6] first emerged in machine translation and is rapidly establishing itself as a new paradigm in natural language processing due to its potential to model global information, which learns the high quality features by considering the whole context. Thanks to its excellent global representation and friendly parallel computation, the transformer is competitive in long sequences modeling and gradually replaces RNNs and CNNs.

Motivated by the remarkable success of transformers achieved in natural language processing, more and more researchers are interested in the employment of attention mechanisms in visual tasks. At present, the transformer has been successfully applied to image classification, object detection, and segmentation tasks. However, there still exist many challenges. First, visual signals and word tokens are very different on many scales. Second, the high dimension of pixel-level information may introduce significant computational complexity.

More recently, ViT [7] was presented as a transformer model to tackle natural images recognition, which splits the image into non-overlapping patches. The authors in [8] proposed a hierarchical ViT called Swin-Transformer by calculating the local self-attention with shifted windows. In contrast to the quadratic computation complexity of self-attention in ViT, Swin-Transformer achieves a linear complexity. Inspired by this fashion, many researchers have tried to apply transformer to other fields. For example, TransUNet [18] combines transformer and Unet [19] for medical image diagnosis. Nevertheless, how to exploit the strengths of attention to aggregate information from entire inputs has not been investigated in the task of visual grasp detection. Unlike prior works, we design a transformer based encoder-decoder architecture to predict the grasp posture in an end-to-end manner. It is shown that our method achieves higher grasp success than the state-of-the-art CNNs counterparts.

#### III. METHOD

*Grasp Representation:* The autonomous visual grasping tasks generally start from collecting visual images of the object by sensory input, which will then be processed to generate an effective grasp configuration to maximise the probability of grasp success. Considering a parallel-plate gripper, the grasp representation g [20] is formulated as a 5-dimensional tuple:

$$g = \{x, y, \theta, w, h\}$$
(1)

where (x, y) are the center coordinates of the grasp rectangle, (w, h) denote the width and height of the grasp rectangle, and  $\theta$  is the orientation of the grasp rectangle with respect to the horizontal axis. Given a gripper with known dimensions, a simplified representation can be expressed as  $g = (p, \phi, w)$ where p = (x, y),  $\phi$  indicates the orientation angle of gripper and w denotes the opening distance of gripper, respectively.

To facilitate grasping, we follow the setting in [17] to represent the grasp in 2-D image space as

$$G = \{Q, W, \Theta\} \in \mathbb{R}^{3 \times W \times H},\tag{2}$$

where the grasp quality Q measures the grasp success of each pixel, and W and  $\Theta$  are the gripper width and orientation angle maps. The value of each pixel in W and  $\Theta$  represents the corresponding width and angle of gripper at that position during the grasping.

Consequently, in the developed TF-Grasp, the grasp detection task boils down to three sub-tasks, namely the problems of predicting grasping position, angle, and width.

Grasp Transformer Overview: A deep motivation of this work is that the treatment of robot perception in complex, dynamic robotic tasks should be global and holistic with information mutual fusion. Specifically, the grasping model can be formulated into an encoder-decoder architecture with a U-shaped structure, as detailed in Fig. 1. The encoder branch aggregates the entire visual input, mutually fuses features by using attention blocks, and then extracts the specific features that are useful for visual robotic grasping. During the decoder process, the model incorporates features delivered via skip-connections and performs a pixel-level grasp prediction by up-sampling. More concretely, the attention modules in the decoder enable more comprehensive processing of local and long-range information, allowing for better multi-scale feature fusion. Each pixel in the prediction heatmap is correlated with the final location and orientation of the end-effector.

To bridge the domain gaps between the transformer and visual robotic grasping tasks, we have carefully designed our grasping transformer in the following aspects for improved grasp detection. (a) Cascade Design. Different from the classic ViT architecture, we adapt a cascaded encoder-decoder structure. The encoder utilizes self-attention to learn a contextual representation that facilitates grasping and the decoder makes use of the extracted features to perform a pixel-level grasp prediction. (b) Local and Global balance. We utilize the swin attention layer to achieve a trade-off between global and local information for better scene perception. Window attention allows cross



Fig. 1. Overview of the TF-grasp model. Our model takes as input the image captured by the camera mounted on the end-effector of the manipulator and generates a pixel-level grasp representation.

window interactions to globally focus on more diverse regions. (c) Feature Fusion. The feature representations at different stages are connected by skip-connections for a multi-scale feature fusion, which acquire both rich semantic and detailed features. (d) Lightweight Design. It is essential for robots to account for efficiency and the real-time performance. We utilize shifted attention blocks and a slimming design for our grasping transformer to reach an ideal trade-off between the performance and speed.

Grasp Transformer Encoder: Before being fed into the encoder, the image is first passed through patch partition layer and is then cut into non-overlapping patches. Each patch is treated as a word token in the text. For example, a 2D image  $I \in \mathbb{R}^{W \times H \times C}$  is split into fixed-size patches  $x \in \mathbb{R}^{N \times (P \times P \times C)}$ , where (H, W) denote the height and width of the original image, C represents the channel of the image, P is the shape size of each image patch, and  $N = H \times W/P^2$  refers to the number of image patches. Then token-based representations can be obtained by passing the images patches into a projection layer.

The encoder is composed by stacking identical transformer blocks. Attentions in the transformer block build long-distance interactions across distant pixels and attend on these positions in the embedding space. At the top of the encoder is a bottleneck block attached to the decoder. The fundamental element in our grasping transformer framework is the multi-head self-attention. The input feature  $\mathbf{X}$  is linearly transformed to derive the query Q, key K, and value V, which are defined as follows:

$$Q = XW_Q, K = XW_K, V = XW_V, \tag{3}$$



Fig. 2. The architecture of our transformer block.

where  $W_Q, W_K, W_V$  are linear projection matrices. Next, we compute the similarity between the query and key by using the dot product to obtain the attention,

Attention
$$(Q, K, V) = \text{SoftMax}\left(\frac{QK^T}{\sqrt{d}} + B\right)V$$
 (4)

where  $\sqrt{d}$  is the scaling factor and B is the learnable relative position encoding.

The computational complexity of self-attention grows quadratically with respect to the image size. To achieve computational efficiency, we leverage the advantages of CNNs and transformer and adopt the swin-transformer block [8] in our framework. The swin-transformer layer consists of two parts: local attention and global attention. Within the local attention, the calculation of self-attention is restricted to local regions where images patches are divided into non-overlapping local windows. Cross-window attention introduces connections between neighbors by sliding non-overlapping windows. The structure of swin-transformer block is presented in Fig. 2 which is composed of MLP, Layer Norm, window-based MSA and shifted-window MSA. The computation procedure of swin-transformer block is represented as follows:

$$\hat{\mathbf{x}}^{l} = \text{W-MSA}\left(\text{LN}\left(\mathbf{x}^{l-1}\right)\right) + \mathbf{x}^{l-1},$$
$$\mathbf{x}^{l} = \text{MLP}\left(\text{LN}\left(\hat{\mathbf{x}}^{l}\right)\right) + \hat{\mathbf{x}}^{l},$$
$$\hat{\mathbf{x}}^{l+1} = \text{SW-MSA}\left(\text{LN}\left(\mathbf{x}^{l}\right)\right) + \mathbf{x}^{l},$$
$$\mathbf{x}^{l+1} = \text{MLP}\left(\text{LN}\left(\hat{\mathbf{x}}^{l+1}\right)\right) + \hat{\mathbf{x}}^{l+1}$$
(5)

where W-MSA and SW-MSA refer to the local window and global shifted window multi-head self-attention, respectively.  $x^{l-1}$  denotes the feature of output from the previous layer. Then, the features will be sent into the window attention, W-MSA. There is a layer norm before both MLP and attention layer, and residual connections are applied to these modules. Between every two swin transformer blocks, there exists a patch merging operation that reduces the resolution of feature maps. The patch merging layer builds a hierarchical representation by gradually merging consecutive neighboring patches between successive transformer layers.

Grasp Transformer Decoder: The decoder generates an executable grasping configuration that allows the end-effector to move to the corresponding positions. We transform the planar grasp detection problem into a pixel-level prediction. Three grasping heads are attached in parallel to the top of the decoder, including a grasp confidence head Q, a gripper angle head  $\Theta$ , and a gripper width head W. The output of each head is a heat map with the same size as the input visual image. The grasp confidence head outputs a value between 0 and 1, which indicates the probability of the successful grasping at each pixel point. Likewise, the gripper width and angle heads output the width and rotation angle of the gripper when grasping at the corresponding point in the image, respectively. We treat the grasping posture estimation as a regression problem and use our transformer model to learn a mapping  $F: I \to G$  by minimizing the distances between the predicted grasping heatmaps  $G(Q, W, \Theta)$  and the ground truth, where I is the input data. The loss function is defined as follows:

$$\mathcal{L} = \sum_{i}^{N} \sum_{m \in \{Q, W, \Theta\}} \|\tilde{G}_{i}^{m} - L_{i}^{m}\|^{2}$$
(6)

where N is the number of sample size and  $L_i$  is the corresponding label.

The ultimate grasp location is the position with the highest grasp confidence by retrieving the grasp quality heatmap, defined as:

$$\mathcal{G}^*_{pos} = \operatorname{argmax}_{pos} \mathbf{Q},$$
 (7)

where Q is the grasp confidence map. Afterward, we extract the predicted angle  $\theta$  and angle w of the corresponding position from the angle and width heatmaps.

In our grasp detection decoder, we also adopt swin transformer block to reduce the computational complexity. Swin attention aggregates multi-scale features and builds a hierarchical representation. And skip-connections merge the features learned at these different stages for further fusion to produce a better grasp posture. Analogous to U-net [19], skip-connections are implemented by concatenating features from the *i*-th layer of the encoder directly into the layer *i*-th in the decoder. In the decoding phase, following the patch expanding layer, the concatenated features are taken as input to the next attention block stage. Simultaneously, we can learn the relationship between the fused features where the features in the encoder can be used as queries and keys to interact with the counterparts in the decoder for self-attention computing.

A benefit of our pixel-level grasp representation is that only a single forward propagation is required to obtain the best grasp postures within the global visual scene, avoiding the need to generate multiple grasp candidates and saving the computation expense.

#### **IV. EXPERIMENTS**

In this section, extensive experiments are carried out to validate the performance of the proposed TF-Grasp method. We verify the performance of TF-Grasp on two popular grasping



Fig. 3. The visualized attention heatmaps learned by our method, which show that our transformer model can learn the concepts beneficial for grasping.

datasets and then evaluate its effectiveness on a real Franka Panda robotic manipulator.

The goal of this section tends to answer the following questions:

- Is the transformer-based grasp detection model better than CNN-based models?
- If true, what makes the transformer-based grasp detection model outperforming others?

## A. Datasets and Experiment Setup

The Cornell grasping data [14] is a multi-object dataset that contains 885 images. The resolution of each image is  $640 \times 480$ . The whole dataset is relatively small and we use various data augmentation techniques such as rotation, zooms, and random cropping to avoid overfitting. We then validate the performance of TF-Grasp on the Jacquard dataset [21] which is generated in a simulator via CAD models. The Jacquard dataset is fairly large, containing over 50 k images of 11 k object categories, and there are over 1 million annotated grasp labels.

*Evaluation Metric:* A predicted grasp is regarded as correct if the following conditions are satisfied.

i) The discrepancy between the predicted grasping angle and the ground truth is within  $30^{\circ}$ .

ii) The Jaccard index defined in (8) is greater than 0.25.

$$J\left(\mathcal{R}^*, \mathcal{R}\right) = \frac{|\mathcal{R}^* \cap \mathcal{R}|}{|\mathcal{R}^* \cup \mathcal{R}|} \tag{8}$$

TF-Grasp takes a  $224 \times 224$  image as input and outputs three pixel-wise maps with the same resolution as the input. The input is normalized by subtracting its mean and dividing the standard deviation. We follow the common strategy to train the grasp transformer. Both the encoder and decoder contain four swinattention blocks and each consists of 1, 2, 4, 8 attention heads. The window size is 7. At each training step, a batch of samples is randomly sampled from the training set and we use the ground truth as the target values to train our neural network. Concretely, we utilize the mean squared error as the loss function and apply AdamW [28] as the optimizer. The default size of batch size is set to 64. The patch partition layer is implemented by convolutions with kernels of  $p \times p$  and a stride p. In our implementation, p is set to 4. In order to preserve a one-to-one mapping of the angle

TABLE I THE ACCURACY ON CORNELL GRASPING DATASET

Method	Input A	Accuracy(%)		Time (ms)
Wiethod		IW	OW	Time (IIIs)
Fast Search [20]	RGB-D	60.5	58.3	5000
GG-CNN [17]	D	73.0	69.0	19
SAE [14]	RGB-D	73.9	75.6	1350
Two-stage closed-loop [22]	RGB-D	85.3	-	140
AlexNet, MultiGrasp [5]	RGB-D	88.0	87.1	76
STEM-CaRFs [23]	RGB-D	88.2	87.5	-
GRPN [24]	RGB	88.7	-	200
ResNet-50x2 [3]	RGB-D	89.2	88.9	103
GraspNet [12]	RGB-D	90.2	90.6	24
ZF-net [25]	RGB-D	93.2	89.1	-
E2E-net [26]	RGB	98.2	-	63
GR-ConvNet [27]	D	93.2	94.3	19
GR-ConvNet [27]	RGB	96.6	95.5	19
GR-ConvNet [27]	RGB-D	97.7	96.6	20
	D	95.2	94.9	41.1
TF-Grasp	RGB	96.78	95.0	41.3
	RGB-D	97.99	96.7	41.6

 $\Theta$  between  $\left[-\frac{\pi}{2}, \frac{\pi}{2}\right]$ , we decode the learning of angle into two components,  $\sin(2\Theta)$  and  $\cos(2\Theta)$ . In this way, the final angle is obtained by  $\arctan(\frac{\sin 2\Theta}{\cos 2\Theta})/2$ . TF-Grasp is implemented by PyTorch, and the entire grasp detection system is running on the Ubuntu 18.04 desktop with Intel Core i9 CPU and NVIDIA 3090 GPU.

### B. Experimental Results and Analysis

To show its effectiveness, our approach is compared with a number of baselines under the same experimental conditions, i.e., evaluation metric. The results of image-wise (IW) and object-wise (OW) settings in the public Cornell grasping dataset are present in Table I. Since the Cornell dataset is relatively small, we follow the setting of previous works [3], [5], [14] by adopting a five-fold cross-validation. Also, to make the comparison fair and comprehensive, the input modalities and running time are considered. For all compared baselines, we use the data reported in their original papers. Taking as input only the depth information, our TF-Grasp achieves an accuracy of 95.2% which is competitive to the state-of-the-arts. When using both depth and RGB data, our model obtains 97.99% accuracy. For

TABLE II THE ACCURACY ON JACQUARD GRASPING DATASET

Authors	Method	Input	Accuracy (%)
Depierre [21]	Jacquard	RGB-D	74.2
Morrison [17]	GG-CNN2	D	84
Zhou [29]	FCGN, ResNet-101	RGB	91.8
Alexandre [16]	GQ-STN	D	70.8
Zhang [11]	ROI-GD	RGB	90.4
Stefan [26]	Det Seg	RGB	92.59
Stefan [26]	Det Seg Refine	RGB	92.95
Kumra [27]	GR-ConvNet	D	93.7
Kumra [27]	GR-ConvNet	RGB	91.8
Kumra [27]	GR-ConvNet	RGB-D	94.6
	TF-Grasp	D	93.1
Our	TF-Grasp	RGB	93.57
	TF-Grasp	RGB-D	94.6

Table II, we use 90% data of the Jacquard dataset as the training set and the remaining 10% as the validation set. In addition, our model takes about 41 ms to process a single image using the Intel Core i9-10900X CPU processor, which is competitive with the state-of-art approaches and basically meets the real-time requirements. The transformer grasping model exhibits a better accuracy on both datasets compared to conventional CNN models. Our proposed approach achieves a higher accuracy of 94.6% which is on-par or superior to previous methods. The results on the Cornell and Jacquard datasets all indicate that the model with the attention mechanism is more suitable for visual grasping tasks.

Despite the fact that our model is trained on a single object dataset, it can be well adapted to multi-object environments with the help of attention mechanisms. In addition, to evaluate the advantages of the transformer versus CNNs for visual grasping tasks, we use the original convolution layers, residual layers, and our transformer as feature extractors to test detection accuracy on different objects on the Cornell dataset. We apply an object-wise split to the Cornell dataset and Fig. 5 shows the detection accuracy of objects not seen during the training phase. All objects are subsets of the Cornell dataset and are evaluated 5 times. All models shown in Fig. 5 employ an encoder-decoder architecture with 4 stages in order to guarantee a fair comparison, where the original-conv is a fully convolutional neural network and resnet-conv is to replace the original convolution layer with the residual block. The result of different models is shown in Fig. 5. Note that the transformer outperforms original convolutions on all selected objects and is marginally better or on-par with the residual network.

These results demonstrate that the transformer improves robotic grasp detection. We conjecture that prior methods that rely on local operations of the convolution layers might ignore the dependencies between long-range pixels. Instead, our approach leverages the attention mechanism to exploit both local and global information and integrates features that are useful for grasping. To better demonstrate whether the transformer-based grasping model can model the relationships between objects and across the scene, we present the multi-object grasping results and grasping quality heatmaps of the transformer and CNN in Fig. 4. Our aim is to verify that the transformer is preferred over CNN for visual grasping tasks and is better at capturing global and local information. From Fig. 4, we can see that the grasp rectangles predicted by CNN have the right grasp position in most cases, but the predicted gripper angle and width are often not appropriate. In some cases, CNN even generates grasping rectangles in the background. With the attention mechanism, our transformer-based model is able to clearly identify the objects from the background. In the second row of Fig. 4, the grasping quality images show that the CNN-based approach can not identify the graspable area and consider the entire region of objects as a graspable zone with high success probabilities. Instead, as shown in the fourth row of Fig. 4, the transformer-based model is prone to capture the area that is easy to grasp due to its larger receptive field. For each attention block, the attention operation establishes the inter-element relationships through self-attention, and the subsequent multi-layer-perceptron (MLP) module further models the inherent relation between each element. The layer normalization and residual connections that interleave these two operations keep the training stable and efficient. In contrast, in CNN, the receptive field of each convolutional kernel is limited. To build a larger receptive field, the model often needs to repeatedly stack convolutional layers to gain global and semantically rich features. However, such a method in general results in the loss of detailed feature information such as the position and shape information of objects that are essential for grasping tasks. Therefore, we exploit a transformer-based model which can better capture not only the global information but also detailed features (e.g., the position and shape information).

## C. Visualization Analysis

To clarify why the transformer architecture is helpful for grasp detection tasks, we visualize the heatmaps of attention maps, detailed in Fig. 3. From these heat maps, we can discover that the self-attention modules can readily learn the area that is easy for grasping, such as the edges of objects, ignore irrelevant details, and pay more attention on the contour and shape of the objects. Meanwhile, the model focuses on more general characteristics rather than individual features. For example, for the chairs shown in Fig. 3, our method evaluates the edge of the chairs with a higher grasp quality. We further provide more concrete examples of real-world grasping, and the experimental results show that the attention mechanism is more likely to achieve a better understanding of the grasping scenario, generate more accurate grasping rectangles, and work well on both household and novel objects. In Fig. 6, we illustrate a pick-and-place task based on our TF-Grasp on the Franka manipulator. Our grasp detection system works well for novel objects that have not been seen during training procedure and also locates graspable objects in cluttered environments.

In conclusion, the visualization results indicate that our TF-Grasp can produce a more general and robust prediction, which contributes to improving the detection accuracy.

## D. Ablation Studies

To understand the role of skip-connections in our transformer model on the visual grasping problems, we conduct experiments



(d) Predicted grasp quality heatmaps by Transformer

Fig. 4. Visualization comparison of the CNN and transformer-based grasping models.



Fig. 5. The accuracy of different models as feature extractors on selected objects.

TABLE III COMPARISON BETWEEN USING AND NOT USING SKIP-CONNECTIONS

The accuracy on Cornell Grasping Results				
	With Skip-connections	Without Skip-connections		
RGB	96.78%	95.7%		
Depth	95.2%	94.3%		
RGB+Depth	97.99%	96.1%		
The accuracy on Jacquard Grasping Results				
	÷ .	1 0		
	With Skip-connections	Without Skip-connections		
RGB	With Skip-connections 93.57%	Without Skip-connections 92.4%		
RGB Depth	With Skip-connections   93.57%   93.1%	Without Skip-connections   92.4%   91.8%		

on the Cornell and Jacquard grasping datasets with and without skip-connections using our transformer, respectively. The detailed experimental results are shown in Table III. The use of skip-connections is better than not using skip-connections in all input modes. The attention mechanism in the transformer builds inter-relationships in each layer, incorporates global features, and achieves promising results. Through skip-connections, the multi-scale representations at different stages are further fused globally. The empirical evidence shows that these further refinement and contextual features contribute to the quality of final grasp prediction.

## E. Grasping in Real World Scenarios

*Physical Setting:* The Franka Panda robot manipulation and the RealSense D435 RGB-D camera are used in our physical experiment. The camera is attached to the end-effector to keep a good visual coverage of graspable objects. In each grasp attempt, our TF-Grasp receives the visual signals from the depth camera mounted on the robot end-effector and outputs an optimal grasping posture. Next, the end-effector approaches the optimal target grasping posture based on the trajectory planned by a motion planning method, and then closes the gripper. Such a transformer-based grasp detection system can be easily adapted to other hardware platforms. During the grasp process, the raw depth sensor is filled with a portion of missing pixels that have NaN values. We generate the mask of NaN values, normalize the depth image, and apply *cv2.inpaint* [30] for further depth completion.

We perform a total of 165 grasping attempts, of which the robot performs successful grasp 152 times, achieving a success rate of 92.1%. Table IV lists the results of learning-based methods on real robot grasping. These results indicate that the transformer-based grasp detection system also behaves well on real robots.



Fig. 6. Screenshots of physical grasping in clutter.

TABLE IV THE RESULTS FOR PHYSICAL SETUP

Authors	Physical grasp	Success rate (%)
Lenz [14]	89/100	89%
Pinto [31]	109/150	73%
Morrison [17]	110/120	92%
Chu [32]	89/100	89%
TF-Grasp(Ours)	152/165	92.1%

#### V. DISCUSSION AND CONCLUSION

In this work, we develop a novel architecture for visual grasping. Although CNN and its variants are still the dominant models in visual robotic grasping, we show the powerful potential of transformers in grasp detection. Compared with CNN-based counterparts, the transformer-based grasp detection models are better at capturing global dependencies and learning powerful feature representation. The results show that our proposed approach outperforms original CNN-based models. The contexts can be better represented by attention propagation. Nevertheless, the current approach is limited to the parallel gripper. Future research will focus on developing a universal transformer-based grasp detection method for other types of grippers, such as the five finger dexterous hand.

#### REFERENCES

- J. Song, M. Patel, and M. Ghaffari, "Fusing convolutional neural network and geometric constraint for image-based indoor localization," *IEEE Robot. Automat. Lett.*, vol. 7, no. 2, pp. 1674–1681, Apr. 2022.
- [2] D. Zhao and J. Oh, "Noticing motion patterns: A temporal CNN with a novel convolution operator for human trajectory prediction," *IEEE Robot. Automat. Lett.*, vol. 6, no. 2, pp. 628–634, Apr. 2021.
- [3] S. Kumra and C. Kanan, "Robotic grasp detection using deep convolutional neural networks," in *Proc. IEEE/RSJ Int. Conf. Intell. Robots Syst. IROS*, 2017, pp. 769–776.
- [4] X. Zhu, Y. Zhou, Y. Fan, and M. Tomizuka, "Learn to grasp with less supervision: A data-efficient maximum likelihood grasp sampling loss," 2021, arXiv:2110.01379.
- [5] J. Redmon and A. Angelova, "Real-time grasp detection using convolutional neural networks," in *Proc. IEEE Int. Conf. Robot. Automat. ICRA*, 2015, pp. 1316–1322.
- [6] A. Vaswani et al., "Attention is all you need," in Proc. Annu. Conf. Neural Inf. Process. Syst., Long Beach, CA, USA, 2017, pp. 5998–6008.
- [7] A. Dosovitskiy et al., "An image is worth 16x16 words: Transformers for image recognition at scale," in Proc. Int. Conf. Learn. Representations, 2021.
- [8] Z. Liu *et al.*, "Swin transformer: Hierarchical vision transformer using shifted windows," in *Proc. IEEE/CVF Int. Conf. Comput. Vis.*, 2021, pp. 10012–10022.
- [9] R. M. Murray, Z. Li, and S. S. Sastry, A Mathematical Introduction to Robotic Manipulation. Boca Raton, FL, USA: CRC Press, 2017.
- [10] A. Bicchi and V. Kumar, "Robotic grasping and contact: A review," in *Proc. IEEE Int. Conf. Robot. Automat.*, San Francisco, CA, USA, 2000, pp. 348–353.

- [11] H. Zhang, X. Lan, S. Bai, X. Zhou, Z. Tian, and N. Zheng, "ROI-based robotic grasp detection for object overlapping scenes," in *Proc. IEEE/RSJ Int. Conf. Intell. Robots Syst. IROS*, 2019, pp. 4768–4775.
- [12] U. Asif, J. Tang, and S. Harrer, "GraspNet: An efficient convolutional neural network for real-time grasp detection for low-powered devices," in *Proc. 27th Int. Joint Conf. Artif. Intell.*, 2018, pp. 4875–4882.
- [13] X. Zhu, L. Sun, Y. Fan, and M. Tomizuka, "6-DoF contrastive grasp proposal network," in *Proc. IEEE Int. Conf. Robot. Automat. ICRA*, 2021, pp. 6371–6377.
- [14] I. Lenz, H. Lee, and A. Saxena, "Deep learning for detecting robotic grasps," *Int. J. Robot. Res.*, vol. 34, no. 4-5, pp. 705–724, 2015.
- [15] J. Mahler *et al.*, "Dex-net 2.0: Deep learning to plan robust grasps with synthetic point clouds and analytic grasp metrics," in *Proc. Robot.: Sci. Syst. XIII*, Cambridge, MA, USA, 2017.
- [16] A. Gariépy, J.-C. Ruel, B. Chaib-Draa, and P. Giguere, "GQ-STN: Optimizing one-shot grasp detection based on robustness classifier," in *Proc. IEEE/RSJ Int. Conf. Intell. Robots Syst. IROS*, 2019, pp. 3996–4003.
- [17] D. Morrison, P. Corke, and J. Leitner, "Learning robust, real-time, reactive robotic grasping," *Int. J. Robot. Res.*, vol. 39, no. 2-3, pp. 183–201, 2020.
- [18] J. Chen et al., "TransUNet: Transformers make strong encoders for medical image segmentation," 2021, arXiv:2102.04306.
- [19] O. Ronneberger, P. Fischer, and T. Brox, "U-Net: Convolutional networks for biomedical image segmentation," in *Proc. Int. Conf. Med. Image Comput. Comput.-Assisted Intervention*, 2015, pp. 234–241.
- [20] Y. Jiang, S. Moseson, and A. Saxena, "Efficient grasping from RGBD images: Learning using a new rectangle representation," in *Proc. IEEE Int. Conf. Robot. Automat.*, 2011, pp. 3304–3311.
- [21] A. Depierre, E. Dellandréa, and L. Chen, "Jacquard: A large scale dataset for robotic grasp detection," in *Proc. IEEE/RSJ Int. Conf. Intell. Robots Syst. IROS*, 2018, pp. 3511–3516.
- [22] Z. Wang, Z. Li, B. Wang, and H. Liu, "Robot grasp detection using multimodal deep convolutional neural networks," *Adv. Mech. Eng.*, vol. 8, no. 9, 2016, Art. no. 1687814016668077.
- [23] U. Asif, M. Bennamoun, and F. A. Sohel, "RGB-D object recognition and grasp detection using hierarchical cascaded forests," *IEEE Trans. Robot.*, vol. 33, no. 3, pp. 547–564, Jun. 2017.
- [24] H. Karaoguz and P. Jensfelt, "Object detection approach for robot grasp detection," in *Proc. IEEE Int. Conf. Robot. Automat. ICRA*, 2019, pp. 4953– 4959.
- [25] D. Guo, F. Sun, H. Liu, T. Kong, B. Fang, and N. Xi, "A hybrid deep architecture for robotic grasp detection," in *Proc. IEEE Int. Conf. Robot. Automat. ICRA*, 2017, pp. 1609–1614.
- [26] S. Ainetter and F. Fraundorfer, "End-to-end trainable deep neural network for robotic grasp detection and semantic segmentation from RGB," in *Proc. IEEE Int. Conf. Robot. Automat. ICRA*, 2021, pp. 13452–13458.
- [27] S. Kumra, S. Joshi, and F. Sahin, "Antipodal robotic grasping using generative residual convolutional neural network," in *Proc. IEEE/RSJ Int. Conf. Intell. Robots Syst. IROS*, 2020, pp. 9626–9633.
- [28] I. Loshchilov and F. Hutter, "Decoupled weight decay regularization," in Proc. Int. Conf. Learn. Representations, 2018.
- [29] X. Zhou, X. Lan, H. Zhang, Z. Tian, Y. Zhang, and N. Zheng, "Fully convolutional grasp detection network with oriented anchor box," in *Proc. IEEE/RSJ Int. Conf. Intell. Robots Syst. IROS*, 2018, pp. 7223–7230.
- [30] G. Bradski, "The OpenCV library," Dr Dobb's J.: Softw. Tools Professional Programmer, vol. 25, no. 11, pp. 120–123, 2000.
- [31] L. Pinto and A. Gupta, "Supersizing self-supervision: Learning to grasp from 50k tries and 700 robot hours," in *Proc. IEEE Int. Conf. Robot. Automat. ICRA*, 2016, pp. 3406–3413.
- [32] F.-J. Chu, R. Xu, and P. A. Vela, "Real-world multiobject, multigrasp detection," *IEEE Robot. Automat. Lett.*, vol. 3, no. 4, pp. 3355–3362, Oct. 2018.