

STOPNet: Multiview-based 6-DoF Suction Detection for Transparent Objects on Production Lines

Yuxuan Kuang^{1,2*}, Qin Han^{1*}, Danshi Li^{2,3}, Qiyu Dai¹, Lian Ding⁴, Dong Sun⁴, Hanlin Zhao⁴, He Wang^{1†}

Abstract—In this work, we present STOPNet, a framework for 6-DoF object suction detection on production lines, with a focus on but not limited to transparent objects, which is an important and challenging problem in robotic systems and modern industry. Current methods requiring depth input fail on transparent objects due to depth cameras’ deficiency in sensing their geometry, while we proposed a novel framework to reconstruct the scene on the production line depending only on RGB input, based on multiview stereo. Compared to existing works, our method not only reconstructs the whole 3D scene in order to obtain high-quality 6-DoF suction poses in real time but also generalizes to novel environments, novel arrangements and novel objects, including challenging transparent objects, both in simulation and the real world. Extensive experiments in simulation and the real world show that our method significantly surpasses the baselines and has better generalizability, which caters to practical industrial needs.

I. INTRODUCTION

Object picking on production lines is an essential task for robotic systems and is widely used in modern industrial applications, such as logistics sorting and bin picking, *etc.* Building such an autonomous robotic system to pick and place objects on a moving production line can greatly save time, reduce human labor and increase productivity. In addition, as a major end-effector in picking objects, the suction cup has gained much popularity due to its simplicity, efficiency and robustness. Therefore reliable suction detection is also worth studying to solve this task.

Due to the importance of this task, recent years have witnessed great progress in developing such systems to automatically conduct production line pick and place [1]–[7]. These methods rely on low-level image features or 3D CAD models to detect suction poses, which only work on regular or seen objects. There are also learning-based methods on static object suction [8]–[12] that leverage accurate surface information from RGBD/depth input and can generalize to novel diffuse objects.

However, transparent object suction still remains a challenging case for such a system. Depth sensors can hardly sense transparent objects and often generate wrong, missing and fuzzy depth on these objects, which causes the failure of these learning-based methods on transparent objects.

For grasping transparent objects, recent studies have focused on depth restoration [13]–[15], leveraging imperfect

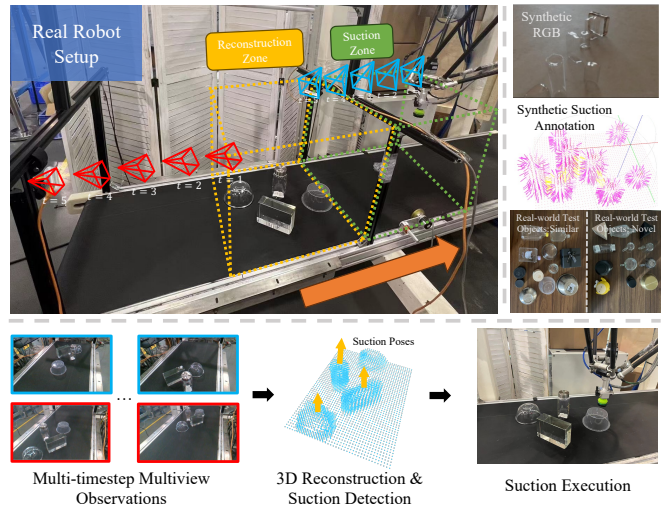


Fig. 1. Overview of our production line real robot setup and our proposed STOPNet. Taking multi-timestep RGB input from two cameras, our framework conducts 3D scene reconstruction and suction detection in real time. Our model is trained on a large-scale synthetic dataset but it can also generalize to real-world objects and environments effectively.

depth [16], [17] and RGB-based grasping [18]–[20]. However, recent works for suction detection [8]–[12] require accurate depth input, since suction relies more on local surface geometry than 6-DoF grasping. Therefore these methods struggle to detect suction on transparent objects without reliable depth input. In that sense, the issue of suction detection for transparent objects is worthy of study.

As to production line pick and place system design, recent works mainly focus on suction detection from a single view [1], [2], [5]. These methods usually suffer from performance degradation in cluttered scenes due to their limited observations. Instead, our method utilizes multi-timestep multiview images to reconstruct the whole scene using volumetric truncated signed distance function (TSDF) representation, making it easier to handle cluttered scenes, which are common in real-world applications.

Our method, STOPNet, addresses the problems mentioned above by “stopping” the production line. Specifically, assuming that the production line is moving at a constant speed, we can transform our multi-timestep multiview RGB observations to static multiview observations, which enable 3D reconstruction. After that, we tackle this problem by combining 2D and 3D, where in 3D space, we conduct scene reconstruction and wrench-collision prediction only given multiview RGB images, and in 2D space, we predict pixel-level seal and normal on original RGB input since these two attributes rely more on local object surface smoothness than

¹Center on Frontiers of Computing Studies, School of Computer Science, Peking University.

²Galbot.

³New York University.

⁴Huawei Cloud Computing Technologies Co., Ltd.

*The first two authors contributed equally.

†Corresponding to hewang@pku.edu.cn.

global scene geometry. By conducting the 2D-3D fusion, we can obtain reliable Top- k suction poses, which will be further executed by robot suction cups.

To achieve generalizability and reduce the sim2real gap while saving training costs, we utilize a domain randomization-based synthetic data generation pipeline to generate a large-scale and diverse synthetic dataset containing over 300K images and 40M suction poses for our training. Combining this with other designs, such as discarding redundant textures and illuminations, extensive experiments demonstrate our method’s generalizability both in simulation and the real world.

To summarize, our framework has many advantages over current methods on production line object suction: it is generalizable to diverse objects in the real world, achieving success rates up to 90.38% in the real world; it performs well on transparent objects without depth input, which is a challenging case in practice; it utilizes multi-timestep multiview input to conduct 3D reconstruction, thus is able to handle cluttered scenes; it is in real time and can be adapted to practical industrial production; it has a low cost of only 2 RGB cameras; and it can perform 6-DoF suction detection.

II. RELATED WORK

A. Transparent Object Grasping

Grasping transparent objects is a challenging task for robotic manipulation since stereo cameras can hardly capture good-quality depth images of transparent objects. Thus recent studies on object suction [8]–[12] based on RGBD or depth images will inevitably fail on transparent objects due to their imperfect depth.

Recent years have witnessed multiple solutions to this issue. One direct way is to restore depths in advance, such as in [13]–[15]. Some studies focus on using single-view fuzzy RGBD images to learn a grasping policy [16], [17], which cannot handle cluttered scenes due to their single-view input. There are also works focusing on taking multiview RGB as input. For example, GlassLoc [18] constructs a Depth Likelihood Volume (DLV) descriptor from multiview light field observations to represent transparent object clutter scenes. GhostPose [19] conducts transparent object grasping by a model-free pose estimation method based on multiview geometry. And GraspNeRF [20] represents the scene geometry as the generalizable NeRF and jointly trains the NeRF and grasping detection. Similar to [20], our method predicts TSDF and detects suction poses directly from multiview RGB images, thus being able to handle transparent objects.

B. Suction Analytic Models and Suction Detection

Different from grasping by parallel grippers, suction cups have the advantage of high efficiency, flexibility and simplicity, thus are widely used in industrial production lines. Many studies [9], [11], [21] have proposed suction analytic models and their evaluation. As elaborated in [11], a suction pose is defined as a 3D suction point and a direction vector starting from the suction point and pointing outside of the object surface. And a suction pose can be evaluated in dimensions

of seal, wrench and collision. We follow their suction analytic model in our setting.

As to suction detection, database-based or CAD model-based suction detection [1] is one of the most commonly used methods in the industry since it’s simple and efficient. However, it is not generalizable and thus cannot handle a wide variety of objects. Current learning-based methods, such as [8]–[12] are generalizable but only cater to single view static suction and need depth input, which will fail on transparent objects. Our method, instead, not only utilizes multi-timestep multiview RGB images to handle dynamic and cluttered scenes but is also generalizable to a wide variety of objects, including challenging transparent objects.

C. Object Picking on Production Lines

In line with the principles of Industry 4.0 [22] which aims to push for flexibility on target changes and autonomy, object picking on production lines has wide applications in autonomous industrial scenarios, thus was heavily studied in the previous decades [1]–[7]. Recent studies implement object recognition and suction detection mainly by CAD model database matching [1], HOG/SVM [2] or morphological operators [5], which are either non-generalizable or ineffective. By contrast, our method can both generalize to novel objects and achieve high performance, with a low cost of 2 RGB cameras.

III. METHOD

A. Problem Statement and Method Overview

Given a sequence of RGB images of objects on a production line, the goal of the proposed robotic system is to detect the 6-DoF suction poses and then execute the suction to remove objects on the line, as shown in Fig. 1. The production line moves at a known fixed speed from left to right. On the left side is the reconstruction zone, where RGB images are captured by two cameras at regular timesteps. Then 3D reconstruction and suction pose detection are conducted in real time and the final suction pose at this timestep will be sent to the robot arm. When objects move into the robot’s workspace on the right side, *i.e.*, the suction zone, the robot arm will execute the suction based on a position shift. The process is pipelined and repeated until there is no detected valid suction pose.

After the production line is stable, for each timestep, we get the paired images of this timestep, and retrieve paired images of previous $N - 1$ timesteps, obtaining N pairs of multi-timestep images in total. In this way, we formulate the 6-DoF suction detection as a learning problem that maps a set of paired RGB images from N timesteps, *i.e.* $2N$ images, $\{I_i\}_{i=1, \dots, 2N}$, to a set of 6-DoF suction poses $\{S_j | S_j = (p_j, d_j, s_j)\}$, where, for each detected suction S_j , $p_j \in \mathbb{R}^3$ is the suction position, $d_j \in \mathbb{R}^3$ is the suction direction, and $s_j \in [0, 1]$ is the overall suction score, which is composed of 3 parts: seal score, wrench score and collision score, where seal score $s_{seal} \in [0, 1]$ models the seal formation of the suction cup, wrench score $s_{wrench} \in [0, 1]$ models the resistance to the wrench caused by gravity, and the binary

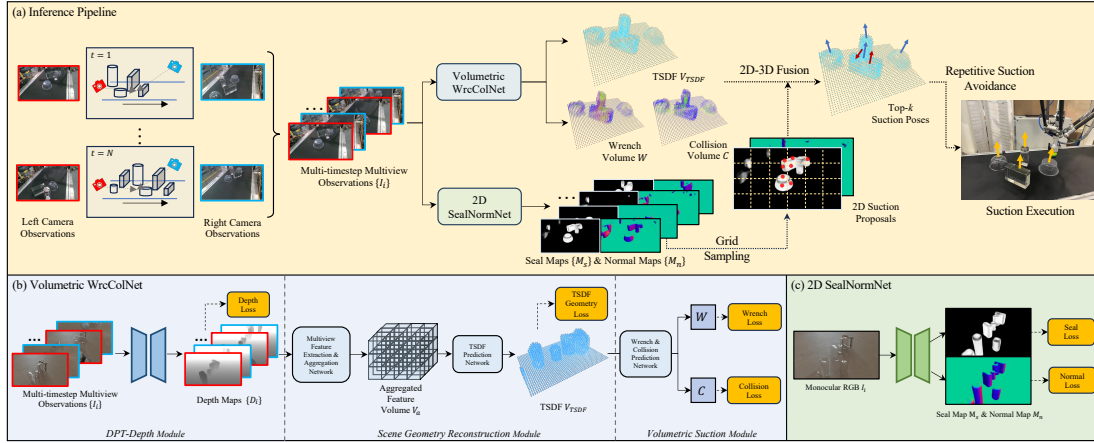


Fig. 2. The framework of our proposed STOPNet: (a) **Inference Pipeline**: multi-timestep multiview RGB images are captured by two cameras as the input to our framework. Volumetric TSDF, wrench and collision volumes are predicted from our *Volumetric WrcColNet*, while seal and normal maps are predicted from our *2D SealNormNet*. We then use the grid sampling method to sample proposal suctions from predicted 2D seal maps and conduct 2D-3D fusion to obtain Top- k suction poses for robot execution. (b) **Volumetric WrcColNet**: from RGB input we predict depth maps, from which we extract and aggregate features to construct the aggregated feature volume V_a . Then we predict the TSDF volume V_{TSDF} and conduct volumetric wrench-collision predictions; (c) **2D SealNormNet**: from monocular RGB images, we jointly predict seal and normal maps.

collision score $s_{collision} \in \{0, 1\}$ denotes whether a suction pose is in collision with other objects in the scene. The overall score for a suction pose is then computed by:

$$s_j = s_{seal} \times s_{wrench} \times s_{collision} \quad (1)$$

Given that the production line is moving at a fixed and known speed, we can transform the observed multi-timestep multiview images into static multiview images. This is because we can consider the objects to be stationary and the cameras to be moving, as illustrated in Fig. 1. It allows us to calculate the transformed camera extrinsics of each timestep based on the production line’s speed and original camera calibration. We can then use these images, along with the transformed camera extrinsics, to perform 3D reconstruction and suction pose detection.

Our proposed framework is composed of two branches, *Volumetric WrcColNet* and *2D SealNormNet*, as illustrated in Fig. 2. For each timestep, in *Volumetric WrcColNet*, we leverage our *DPT-Depth Module* to conduct monocular depth estimation on separate RGB inputs and feed depth predictions into a multiview *Scene Geometry Reconstruction Module* to output volumetric TSDF of the reconstruction zone. Taking TSDF as input, the following *Volumetric Suction Module* learns to predict grid-level wrench scores and collision scores of the whole scene. In *2D SealNormNet*, we take RGB images as input to jointly predict pixel-level seal scores and surface normals. The two branches are separately trained and fused together at inference time to output the Top- k suction poses for each timestep.

B. Volumetric WrcColNet: Scene Reconstruction and Wrench-Collision Prediction

To better detect suction poses on a moving production line, it is crucial to wholly and accurately reconstruct the scene geometry in a generalizable way. Therefore our proposed *Volumetric WrcColNet*, which is composed of *DPT-Depth Module*, *Scene Geometry Reconstruction Module* and *Volumetric Suction Module*, utilizes multi-timestep multiview

images to reconstruct the whole 3D scene to learn better suction detection. All the modules above are jointly trained in a fully end-to-end manner.

1) *DPT-Depth Module*: Since real-world data are expensive and labor intensive, it is more common to get training data from simulation environments, thus the sim2real gap is an inevitable issue to tackle. Under these circumstances, we propose to first learn a module to unify synthetic data and real data into the depth domain.

Dense Prediction Transformers, *i.e.* DPT [23] is a large-scale vision transformer that can be applied to many dense image-to-image prediction tasks, such as semantic segmentation, monocular depth estimation, *etc.* DPT’s network architecture contains an encoder that is composed of a series of embedding layers, transformer blocks, reassemble layers and fusion layers to encode the RGB input into dense image feature maps, and a convolutional decoder network to output dense prediction results of the same resolution as input.

Specifically, we leverage a DPT model pretrained on *MIX 6* [23], a large-scale real-world monocular depth estimation dataset containing about 1.4 million images, and finetune it in our setting. Therefore the network inherits the abundant real-world knowledge in advance and thus better generalizes to the real world. Taking multi-timestep RGB observations $\{I_i\}$ as input, the module outputs depth map predictions $\{D_i\}$ for 3D reconstruction. In this way we discard redundant textures and illuminations from input and unify input domains into the depth domain, which helps to bridge the sim2real gap. To supervise the *DPT-Depth Module*, an L1 pixel-level **Depth Loss** L_d is utilized to directly supervise depth estimations.

2) *Scene Geometry Reconstruction Module*: Inspired by [20], first we propose a *Multiview Feature Extraction and Aggregation Network* to encode the depth predictions into a volumetric representation. To be specific, we divide our 3D reconstruction zone into grids of size X, Y and Z and utilize a CNN encoder to extract input depth prediction feature maps. Each grid point $x \in \mathbb{R}^{XYZ}$ is then projected to the $2N$ depth

feature maps and queries the corresponding features, which are then aggregated among views using MLPs to form an aggregated feature volume $V_a \in \mathbb{R}^{XYZ \times C}$.

After obtaining aggregated feature volume V_a , similar to [24] and [20], we propose a *TSDF Prediction Network* to predict volumetric TSDF $V_{TSDF} \in [-1, 1]^{XYZ}$ from V_a , supervised by a L1 **TSDF Geometry Loss** L_g .

3) *Volumetric Suction Module*: Motivated by [25], taking predicted TSDF as input, we propose a 3D CNN encoder-decoder *Wrench and Collision Prediction Network* to predict volumetric wrench scores $W \in [0, 1]^{XYZ}$ and collision scores $C \in \{0, 1\}^{XYZ}$. We explicitly supervise the L2 **Wrench Loss** L_w and binary cross entropy **Collision Loss** L_c on grids.

Therefore the whole *Volumetric WrcColNet*'s objective L_{3D} can be formulated as:

$$L_{3D} = \alpha_1 L_d + \alpha_2 L_g + \alpha_3 L_w + \alpha_4 L_c \quad (2)$$

C. 2D SealNormNet: Suction Pose Detection

As SuctionNet-1Billion [11] indicates, each suction pose contains 6 degrees of freedom and can be evaluated through three dimensions: seal, wrench and collision. Wrench and collision mainly focus on the global characteristics of objects and scenes, while seal heavily relies on the local characteristics of object surfaces. Therefore a prediction of seal scores from unreliable 3D representations can do great harm to the performance, and so do surface normals, which are the default choices of suction directions. To this end, we propose to predict seal scores and surface normal in a 2D manner.

Specifically, taking RGB as input, we utilize a DPT model that shares the transformer encoder layers to extract image features and leverage different convolutional decoder layers to jointly predict the seal map M_s and the normal map M_n . And we directly supervise the dense pixel loss, which is:

$$L_{2D} = \beta_1 \|M_s - M_s^*\|_2 + \beta_2 \|M_n - M_n^*\|_1 \quad (3)$$

where M_s^* and M_n^* are the GT seal and normal map.

D. Inference Pipeline

At inference time, we need to conduct 2D-3D fusion to acquire Top- k suction poses to execute. First we do grid sampling on 2D seal maps to get pixel positions of high seal scores and query corresponding normals on predicted normal maps to obtain our 2D suction proposals. Based on the predicted TSDF volume V_{TSDF} , we can render the corresponding de-noised depth maps using the Marching Cube Algorithm [26]. This allows us to obtain 3D positions $\{p_j\}$ and normals $\{d_j\}$ from 2D pixels. Given 3D points of seal scores $s_{seal} \in [0, 1]$, both wrench scores $s_{wrench} \in [0, 1]$ and collision scores $s_{collision} \in \{0, 1\}$ of the nearest grids from the predicted wrench and collision volumes will be assigned to these points. Finally, we calculate the overall scores by Eq. 1 and rank them to get $\{s_j\}$. In this way we obtain the 3D Top- k suction poses $\{S_j | S_j = (p_j, d_j, s_j)\}$, which can be further executed by robot suction cups.

As Top- k suction poses are unaware of object identities, multiple suction can be predicted on the same object. To

avoid this, we propose *Repetitive Suction Avoidance*, where we predict instance masks with MobileSAM [27] and apply heuristic filters to ensure each object is only sucked once. At each timestep, if there exist valid and non-repetitive suction poses, one final suction pose will be selected and sent to the robot arm. It will be executed once the object appears in the end-effector's workspace.

E. Synthetic Data Generation and Sim2real Gap Reduction

Since our case is specific, no existing dataset can satisfy our needs. And as real-world data with annotations can be expensive and labor intensive, training on synthetic datasets is more costless and scalable. Hence, motivated by [20], we propose an automatic data generation pipeline and a large-scale production line suction dataset, containing over 1000 production line layouts, 10000 multiview scenes, 300K images and 40M suction poses of diverse objects.

To generate the data, we first collect a wide range of CAD models and textures. Then we randomly place these objects on a moving production line in PyBullet simulator [28]. Based on simulation output, we render photo-realistic multi-view RGB images, depth and normal maps in Blender [29]. For suction annotation, we use APIs provided by [11] to generate GT seal, wrench and collision scores. We then render the corresponding 2D seal maps using PyTorch3D [30]. Example data can be found in Fig. 1 upper right.

To better bridge the sim2real gap, we leverage domain randomization to randomize object arrangements, illuminations, camera poses and backgrounds. Combining this with our finetuned *DPT-Depth Module*, we largely reduce sim2real gaps, as models trained on our synthetic dataset can easily generalize to the real world and achieve higher performance than other methods, which will be covered in Section IV.

IV. EXPERIMENTS

In this section, we evaluate the performance of our proposed method for production line suction tasks in both simulation and real-world environments. We also perform ablation studies in the real world to analyze the effectiveness of different modules of our framework.

A. Implementation Details

As indicated in Section III, we set the size of the reconstruction zone grid as $X = 50, Y = 40, Z = 30$ in *cm*, and the number of timesteps used in one reconstruction is $N = 5$. We train our model for a maximum number of 300K iterations utilizing the Adam optimizer [31] with a learning rate of 1×10^{-4} and a learning rate decay of 0.5 per 100K iterations.

B. Experiment Setup

1) *Object Sets*: In simulation, we utilize 160 hand-scale objects from [32] and assign them with transparent textures. We then replenish the object set with 80 package objects of similar scales and assign them with diverse real-world textures in order to learn suction detection for practical industrial usage. Following [33] and [11], objects are split into seen, similar, and novel sets. We train our model only

TABLE I
EVALUATION FOR DIFFERENT METHODS IN THE SIMULATION ENVIRONMENT
EACH CELL: TRANSPARENT / MIXED

| Metrics | | TSDF MAE ↓ | Surf. TSDF MAE ↓ | Seal MAE ↓ | Collision Acc. (%) ↑ | AP@Top-1 ↑ | AP@Top-5 ↑ |
|---------|------------------|--------------------|--------------------|----------------------|----------------------|--------------------|--------------------|
| Seen | DPT-Fusion [23] | 4.73 / 2.88 | 20.03 / 12.25 | 0.018 / 0.020 | 93.93 / 94.44 | 0.55 / 0.44 | 0.55 / 0.42 |
| | SuctionNeRF [20] | 0.88 / 0.71 | 3.09 / 2.03 | 0.280 / 0.269 | 95.89 / 95.25 | 0.05 / 0.04 | 0.05 / 0.04 |
| | Ours | 0.55 / 0.47 | 2.08 / 1.43 | 0.018 / 0.020 | 96.32 / 95.68 | 0.65 / 0.80 | 0.58 / 0.71 |
| Similar | DPT-Fusion [23] | 6.12 / 3.78 | 23.32 / 12.43 | 0.025 / 0.024 | 93.36 / 94.60 | 0.45 / 0.55 | 0.30 / 0.56 |
| | SuctionNeRF [20] | 1.41 / 1.01 | 4.02 / 2.52 | 0.237 / 0.222 | 95.29 / 95.69 | 0.00 / 0.10 | 0.10 / 0.07 |
| | Ours | 0.97 / 0.65 | 2.82 / 1.59 | 0.025 / 0.024 | 95.57 / 96.51 | 0.55 / 0.85 | 0.59 / 0.81 |
| Novel | DPT-Fusion [23] | 3.30 / 3.23 | 18.77 / 16.62 | 0.032 / 0.034 | 93.08 / 93.06 | 0.42 / 0.51 | 0.22 / 0.43 |
| | SuctionNeRF [20] | 0.67 / 0.85 | 2.43 / 3.14 | 0.394 / 0.371 | 95.56 / 93.20 | 0.00 / 0.10 | 0.06 / 0.07 |
| | Ours | 0.47 / 0.59 | 1.84 / 2.19 | 0.032 / 0.034 | 95.87 / 93.34 | 0.55 / 0.80 | 0.44 / 0.65 |

on the seen set and test them on all three sets in novel arrangements to check the generalizability of our framework.

In the real world, we gather around 30 different types of diverse objects, including transparent ones and other common items, to put on the production line. All the objects are unseen to our model so we split them into similar and novel sets based on their geometry, as shown in Fig. 1.

2) *Real Robot Setup*: As in Fig. 1, we set two RealSense D415 cameras diagonally to capture RGB images (depth is not used in our method) and set the speed of the production line to 100mm/s. For motion planning, since the objects are moving at a fixed speed, we utilize a heuristic algorithm that moves the robot arm above to track the suction position in the workspace. When the two trajectories coincide, the suction cup will push down to execute the suction.

For the final suction pose selection from Top- k proposals at each timestep, we take heights and directions of suction poses into account to adapt 6-DoF suction poses to our 4-DoF AtomRobot D3P-1100-P3 parallel robot arm. We set the timestep length to 1s, and the inference time of our whole pipeline is around 0.5s using a single GPU, leading to continuous and pipelined suction on the production line in 1Hz, which is compatible with our robot arm’s cycle time.

C. Baseline Methods

We compare our method with the following baselines in simulation or real-world environments.

- **DPT-Fusion** [23]. As multiview depths can be directly used to fuse a TSDF in a non-learning way, we utilize the toolbox provided by [34] to fuse a TSDF volume from predicted depths by finetuned DPT [23] in place of our *Scene Geometry Reconstruction Module*. We utilize our *2D SealNormNet* for seal and normal prediction.
- **SuctionNeRF** [20]. Based on GraspNeRF [20], a multi-view 3D reconstruction and grasp detection framework using generalizable NeRF, we modify its VGN [25] grasping detection head layer to predict seal, wrench and collision scores to adapt to our setting.
- **SuctionNet** [11]. An object suction detection framework that takes RGBD images as input. It gets a partial point cloud from depth input and a 2D heatmap from RGBD input, and then executes suction by the heatmap scores and projected point coordinates.
- **2D Object Detector** [35]. In industry, a common practice is to use a 2D object detector to detect

objects’ centroids. To ensure generalizability, we utilize YOLOv8 [35] for object detection and finetuned DPT [23] for depth estimation.

D. Simulation Experiments

1) *Metrics*: In simulation environments, we mainly focus on different methods’ performance on 3D object reconstruction and 6-DoF suction prediction precision, so we compare our method with DPT-Fusion [23] and SuctionNeRF [20] using multiple metrics, including the **TSDF MAE** and **Surface TSDF MAE** (converted to millimeters) to evaluate reconstruction quality of the whole scene and object surfaces, **Seal MAE** to evaluate seal score prediction and **Collision Accuracy** to evaluate collision score prediction. Following [11], we also evaluate the overall suction score prediction using **AP@Top- k** , the average of AP_s with threshold s ranging from 0.2 to 0.8, with an interval of $\delta_s = 0.2$. We report the AP value under Top-1 and Top-5.

2) *Results and Analysis*: As presented in Table I, in the simulation environment, our method outperforms all the baselines for all combinations of object sets (seen/similar/novel) and object types (transparent/mixed), demonstrating the superiority of our method.

Compared to non-learning DPT-Fusion, we perform significantly better due to our method’s great improvement in reconstruction quality, which verifies that our *Scene Geometry Reconstruction Module* can significantly reduce the inconsistency of multiple monocular depth estimations.

Also, compared to SuctionNeRF, our method improves reconstruction quality because of the prior knowledge on depth estimation embodied in the pretrained *DPT-Depth Module*. Moreover, better reconstruction quality benefits suction collision detection, as our method has better performance on collision accuracy.

For seal prediction and AP@Top- k , our method greatly surpasses SuctionNeRF because instead of taking flawed TSDF as input, we predict seal scores in a 2D manner, which takes RGB as input so that the model can better learn local characteristics of object surfaces, which is crucial for seal scores. Our method also beats DPT-Fusion since our reconstruction quality is better.

In addition, our method’s improvement on similar and novel object sets demonstrates that our method has better generalizability due to our pretrained *DPT-Depth Module* and using predicted depth as input of 3D reconstruction.

TABLE II
EVALUATION FOR DIFFERENT METHODS IN THE REAL WORLD
EACH CELL: TRANSPARENT / MIXED

| Metrics | | SR (%) \uparrow | DR (%) \uparrow |
|---------|-------------------------|----------------------|----------------------|
| Similar | DPT-Fusion [23] | 13.76 / 27.64 | 12.50 / 26.25 |
| | SuctionNeRF [20] | 48.78 / 42.56 | 46.87 / 39.33 |
| | SuctionNet [11] | 27.27 / 40.00 | 22.72 / 33.33 |
| | 2D Object Detector [35] | 27.77 / 28.57 | 14.28 / 15.38 |
| | Ours | 87.64 / 90.38 | 70.90 / 78.77 |
| Novel | DPT-Fusion [23] | 6.81 / 14.01 | 6.38 / 12.71 |
| | SuctionNeRF [20] | 19.56 / 13.15 | 19.14 / 11.36 |
| | SuctionNet [11] | 9.66 / 15.38 | 8.10 / 12.50 |
| | 2D Object Detector [35] | 10.00 / 11.76 | 10.00 / 9.75 |
| | Ours | 69.23 / 85.52 | 61.36 / 73.86 |




Fig. 3. Visualization of TSDF reconstructed by different methods. The geometry from DPT-Fusion, which is directly fused from depth estimations has severe flaws, indicating that our method can greatly reduce the inconsistency of multiple monocular depth estimations. In comparison to SuctionNeRF, which uses direct RGB input and produces floaters in 3D reconstruction, our method results in more accurate reconstruction and better suction detection.

E. Real Robot Experiments

We evaluate our method and all baselines in the real world with a production line and objects configured as in Fig. 1. Each experiment consists of 3 rounds with each round featuring 120 objects which are sequentially placed in the same arrangement. A suction is considered successful if it removes the object from the production line.

1) *Metrics*: As the GT labels are not available in the real world, following [20], we measure the performance by **Success Rate (SR)**: the ratio of the successful suction number to the attempt number; and **Declutter Rate (DR)**: the average percentage of removed objects to all objects.

2) *Results and Analysis*: Table II presents SR & DR of real-world object suction on the production line with the same objects and arrangements. Results show that our method beats all the baselines for all combinations of object sets (similar/novel) and object types (transparent/mixed). The great improvement in the novel object set also indicates our method’s generalizability. In addition, in Fig. 3 we provide real-world reconstruction results using different methods.

Based on the results, it is clear that our approach performs much better than DPT-Fusion and SuctionNeRF, as explained in Section IV-D.2. Additionally, the visualization in Fig. 3 demonstrates that our method can create more precise and superior 3D scenes in real-world settings. This is in contrast to other methods that generate artifacts or floaters that hinder wrench prediction and collision detection.

Compared to SuctionNet, which needs depth for input, our method performs significantly better in SR and DR, especially in transparent objects, indicating that errors in transparent object depths greatly harm the performance of RGBD-based methods.

Finally, our method outperforms the basic 2D object detector method because we leverage multi-timestep images

to wholly reconstruct the scene, making it more capable of handling cluttered scenes. Also, rather than just selecting the centroids of objects, our method can better predict seal scores on objects, contributing to better performance.

However, we also notice some failure cases. For example, occasionally there will be a small offset in reconstruction and suction detection because of the unstable condition of production lines.

F. Ablation Studies

To analyze our method design, we conduct ablation studies on different configurations in the real world. As the effectiveness of our *2D SealNormNet* has been analyzed at Section IV-D.2, this section focuses mainly on our *Volumetric WrcColNet*. Results can be found in Table III.

TABLE III
ABLATION STUDIES IN THE REAL WORLD
EACH CELL: TRANSPARENT / MIXED

| Metrics | | SR (%) \uparrow | DR (%) \uparrow |
|---------|-----------------|----------------------|----------------------|
| Similar | RGB | 80.37 / 74.71 | 65.64 / 65.19 |
| | w/o pretraining | 68.33 / 79.60 | 58.15 / 72.62 |
| | 1/5 Data | 70.12 / 76.43 | 49.09 / 62.82 |
| | Ours | 87.64 / 90.38 | 70.90 / 78.77 |
| Novel | RGB | 50.94 / 52.08 | 48.21 / 42.01 |
| | w/o pretraining | 50.00 / 51.85 | 43.63 / 41.17 |
| | 1/5 Data | 60.00 / 70.00 | 50.00 / 57.14 |
| | Ours | 69.23 / 85.52 | 61.36 / 73.86 |

First, we examine the effect of taking depth maps as input to our *Scene Geometry Reconstruction Module*. We change the module’s input from predicted depth maps to raw RGB images, and the performance degradation indicates that removing redundant elements, such as textures and illuminations from input helps to bridge the sim2real gap.

Then we analyze the effect of the pretrained *DPT-Depth Module* on *MIX 6* [23] dataset. Results show that compared with training *DPT-Depth Module* from scratch, our method can better generalize to the real world due to the prior knowledge from pretraining.

To show the necessity of the large scale of our synthetic dataset, we cut down training data to 1/5 of the original data. The resulting decrease in performance highlights the significance of a large training data scale for model performance.

V. CONCLUSIONS

In this work, we propose STOPNet, a multiview RGB-based 3D reconstruction and 6-DoF suction detection framework for transparent objects on production lines. We also present a large-scale synthetic dataset and corresponding data generation pipeline for object suction on production lines. We also leverage multiple practices to reduce the sim2real gap. Experiments in simulation and the real world demonstrate our method’s superiority over current methods and great generalizability to the real world. We hope our work can benefit related research and industrial practices.

ACKNOWLEDGMENTS

This project is supported by the National Natural Science Foundation of China (No. 62306016).

REFERENCES

- [1] L. Duc Hanh, N. Luat, and L. Bich, "Combining 3d matching and image moment based visual servoing for bin picking application," *International Journal on Interactive Design and Manufacturing (IJI-DeM)*, vol. 16, 03 2022.
- [2] S. D'Avella, C. A. Avizzano, and P. Tripicchio, "Ros-industrial based robotic cell for industry 4.0: Eye-in-hand stereo camera and visual servoing for flexible, fast, and accurate picking and hooking in the production line," *Robotics and Computer-Integrated Manufacturing*, vol. 80, p. 102453, 2023. [Online]. Available: <https://www.sciencedirect.com/science/article/pii/S0736584522001351>
- [3] M. Pfeffer, C. Goth, D. Craiovan, and J. Franke, "3d-assembly of molded interconnect devices with standard smd pick & place machines using an active multi axis workpiece carrier," in *2011 IEEE International Symposium on Assembly and Manufacturing (ISAM)*, 2011, pp. 1–6.
- [4] R. Reif and W. A. Günthner, "Pick-by-vision: augmented reality supported order picking," *Vis. Comput.*, vol. 25, no. 5-7, pp. 461–467, 2009. [Online]. Available: <https://doi.org/10.1007/s00371-009-0348-y>
- [5] M. Naeem, S. Aslam, M. Suhaib, S. Gul, Z. Murtaza, and M. J. Khan, "Design and implementation of pick and place manipulation system for industrial automation," in *2021 International Conference on Artificial Intelligence and Mechatronics Systems (AIMS)*, 2021, pp. 1–6.
- [6] K. Castelli, A. M. A. Zaki, and H. Giberti, "Development of a practical tool for designing multi-robot systems in pick-and-place applications," *Robotics*, vol. 8, no. 3, 2019. [Online]. Available: <https://www.mdpi.com/2218-6581/8/3/71>
- [7] I. Akinola, J. Xu, S. Song, and P. K. Allen, "Dynamic grasping with reachability and motion awareness," in *2021 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, 2021, pp. 9422–9429.
- [8] A. Zeng, S. Song, K.-T. Yu, E. Donlon, F. R. Hogan, M. Bauza, D. Ma, O. Taylor, M. Liu, E. Romo, N. Fazeli, F. Alet, N. C. Daffe, R. Holladay, I. Morona, P. Q. Nair, D. Green, I. Taylor, W. Liu, T. Funkhouser, and A. Rodriguez, "Robotic pick-and-place of novel objects in clutter with multi-affordance grasping and cross-domain image matching," *The International Journal of Robotics Research*, vol. 41, no. 7, pp. 690–705, 2022. [Online]. Available: <https://doi.org/10.1177/0278364919868017>
- [9] J. Mahler, M. Matl, X. Liu, A. Li, D. Gealy, and K. Goldberg, "Dexnet 3.0: Computing robust vacuum suction grasp targets in point clouds using a new analytic model and deep learning," in *2018 IEEE International Conference on Robotics and Automation (ICRA)*, 2018, pp. 5620–5627.
- [10] J. Mahler, M. Matl, V. Satish, M. Danielczuk, B. DeRose, S. McKinley, and K. Goldberg, "Learning ambidextrous robot grasping policies," *Science Robotics*, vol. 4, no. 26, p. eaau4984, 2019. [Online]. Available: <https://www.science.org/doi/abs/10.1126/scirobotics.aau4984>
- [11] H. Cao, H.-S. Fang, W. Liu, and C. Lu, "Suctionnet-1billion: A large-scale benchmark for suction grasping," *IEEE Robotics and Automation Letters*, vol. 6, no. 4, pp. 8718–8725, 2021.
- [12] Y. Fu, T. Sun, L. Wang, S. Yu, L. Deng, B. Chen, L. Li, Y. Xie, S. Deng, and H. Yin, "Rgb-d instance segmentation-based suction point detection for grasping," in *2022 IEEE International Conference on Robotics and Biomimetics (ROBIO)*, 2022, pp. 1643–1650.
- [13] S. Sajjan, M. Moore, M. Pan, G. Nagaraja, J. Lee, A. Zeng, and S. Song, "Clear grasp: 3d shape estimation of transparent objects for manipulation," in *2020 IEEE International Conference on Robotics and Automation (ICRA)*, 2020, pp. 3634–3642.
- [14] H. Fang, H.-S. Fang, S. Xu, and C. Lu, "Transcg: A large-scale real-world dataset for transparent object depth completion and a grasping baseline," *IEEE Robotics and Automation Letters*, vol. 7, no. 3, pp. 7383–7390, 2022.
- [15] Q. Dai, J. Zhang, Q. Li, T. Wu, H. Dong, Z. Liu, P. Tan, and H. Wang, "Domain randomization-enhanced depth simulation and restoration for perceiving and grasping specular and transparent objects," in *Computer Vision – ECCV 2022*, S. Avidan, G. Brostow, M. Cissé, G. M. Farinella, and T. Hassner, Eds. Cham: Springer Nature Switzerland, 2022, pp. 374–391.
- [16] T. Weng, A. Pallankize, Y. Tang, O. Kroemer, and D. Held, "Multi-modal transfer learning for grasping transparent and specular objects," *IEEE Robotics and Automation Letters*, vol. 5, no. 3, pp. 3791–3798, 2020.
- [17] H. Cao, J. Huang, Y. Li, J. Zhou, and Y. Liu, "Fuzzy-depth objects grasping based on fsg algorithm and a soft robotic hand," in *2021 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, 2021, pp. 3948–3954.
- [18] Z. Zhou, T. Pan, S. Wu, H. Chang, and O. C. Jenkins, "Glass-loc: Plenoptic grasp pose detection in transparent clutter," in *2019 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, 2019, pp. 4776–4783.
- [19] J. Chang, M. Kim, S. Kang, H. Han, S. Hong, K. Jang, and S. Kang, "Ghostpose: Multi-view pose estimation of transparent objects for robot hand grasping," in *2021 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, 2021, pp. 5749–5755.
- [20] Q. Dai, Y. Zhu, Y. Geng, C. Ruan, J. Zhang, and H. Wang, "Grasprnerf: Multiview-based 6-dof grasp detection for transparent and specular objects using generalizable nerf," in *2023 IEEE International Conference on Robotics and Automation (ICRA)*, 2023, pp. 1757–1763.
- [21] J. Hudoklin, S. Seo, M. Kang, H. Seong, A. T. Luong, and H. Moon, "Vacuum suction cup modeling for evaluation of sealing and real-time simulation," *IEEE Robotics and Automation Letters*, vol. 7, no. 2, pp. 3616–3623, 2022.
- [22] H. Lasi, P. Fettke, H.-G. Kemper, T. Feld, and M. Hoffmann, "Industry 4.0," *Business & Information Systems Engineering*, vol. 6, no. 4, pp. 239–242, Aug 2014. [Online]. Available: <https://doi.org/10.1007/s12599-014-0334-4>
- [23] R. Ranftl, A. Bochkovskiy, and V. Koltun, "Vision transformers for dense prediction," in *Proceedings of the IEEE/CVF International Conference on Computer Vision (ICCV)*, October 2021, pp. 12 179–12 188.
- [24] P. Wang, L. Liu, Y. Liu, C. Theobalt, T. Komura, and W. Wang, "Neus: Learning neural implicit surfaces by volume rendering for multi-view reconstruction," in *Advances in Neural Information Processing Systems*, M. Ranzato, A. Beygelzimer, Y. Dauphin, P. Liang, and J. W. Vaughan, Eds., vol. 34. Curran Associates, Inc., 2021, pp. 27 171–27 183. [Online]. Available: https://proceedings.neurips.cc/paper_files/paper/2021/file/e41e164f7485ec4a28741a2d0ea41c74-Paper.pdf
- [25] M. Breyer, J. J. Chung, L. Ott, R. Siegwart, and J. Nieto, "Volumetric grasping network: Real-time 6 dof grasp detection in clutter," in *Proceedings of the 2020 Conference on Robot Learning*, ser. Proceedings of Machine Learning Research, J. Kober, F. Ramos, and C. Tomlin, Eds., vol. 155. PMLR, 16–18 Nov 2021, pp. 1602–1611. [Online]. Available: <https://proceedings.mlr.press/v155/breyer21a.html>
- [26] W. E. Lorensen and H. E. Cline, "Marching cubes: A high resolution 3d surface construction algorithm," in *SIGGRAPH*, M. C. Stone, Ed. ACM, 1987, pp. 163–169. [Online]. Available: <http://dblp.uni-trier.de/db/conf/siggraph/siggraph1987.html#LorensenC87>
- [27] C. Zhang, D. Han, Y. Qiao, J. U. Kim, S.-H. Bae, S. Lee, and C. S. Hong, "Faster segment anything: Towards lightweight sam for mobile applications," *arXiv preprint arXiv:2306.14289*, 2023.
- [28] E. Coumans and Y. Bai, "Pybullet, a python module for physics simulation for games, robotics and machine learning," <http://pybullet.org>, 2016–2021.
- [29] B. O. Community, *Blender - a 3D modelling and rendering package*, Blender Foundation, Stichting Blender Foundation, Amsterdam, 2018. [Online]. Available: <http://www.blender.org>
- [30] N. Ravi, J. Reizenstein, D. Novotny, T. Gordon, W.-Y. Lo, J. Johnson, and G. Gkioxari, "Accelerating 3d deep learning with pytorch3d," *arXiv:2007.08501*, 2020.
- [31] D. P. Kingma and J. Ba, "Adam: A method for stochastic optimization," in *3rd International Conference on Learning Representations, ICLR 2015, San Diego, CA, USA, May 7-9, 2015. Conference Track Proceedings*, Y. Bengio and Y. LeCun, Eds., 2015. [Online]. Available: <http://arxiv.org/abs/1412.6980>
- [32] Z. Jiang, Y. Zhu, M. Svetlik, K. Fang, and Y. Zhu, "Synergies between affordance and geometry: 6-dof grasp detection via implicit representations," in *Robotics: Science and Systems XVII, Virtual Event, July 12-16, 2021*, D. A. Shell, M. Toussaint, and M. A. Hsieh, Eds., 2021. [Online]. Available: <https://doi.org/10.15607/RSS.2021.XVII.024>
- [33] H.-S. Fang, C. Wang, M. Gou, and C. Lu, "Graspnet-1billion: A large-scale benchmark for general object grasping," in *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, June 2020.
- [34] A. Zeng, S. Song, M. Niessner, M. Fisher, J. Xiao, and T. Funkhouser, "3dmatch: Learning local geometric descriptors from rgb-d reconstruc-

tions,” in *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, July 2017.

- [35] G. Jocher, A. Chaurasia, and J. Qiu, “YOLO by Ultralytics,” Jan. 2023. [Online]. Available: <https://github.com/ultralytics/ultralytics>