

Automating Kitting Novel 3D Objects into Novel 3D Cavities Using Self-Supervised Learning of SE(3) Transforms

*Shivin Devgon
Ken Goldberg, Ed.*



Electrical Engineering and Computer Sciences
University of California, Berkeley

Technical Report No. UCB/EECS-2021-100

<http://www2.eecs.berkeley.edu/Pubs/TechRpts/2021/EECS-2021-100.html>

May 14, 2021

Copyright © 2021, by the author(s).
All rights reserved.

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. To copy otherwise, to republish, to post on servers or to redistribute to lists, requires prior specific permission.

Acknowledgement

To my parents,

Automating Kitting Novel 3D Objects into Novel 3D Cavities
Using Self-Supervised Learning of $SE(3)$ Transforms

by

Shivin Devgon

A dissertation submitted in partial satisfaction of the

requirements for the degree of

Master of Science

in

Electrical Engineering and Computer Sciences

in the

Graduate Division

of the



University of California, Berkeley

Committee in charge:

Professor Ken Goldberg, Chair
Professor Avidah Zakhori

Spring 2021

The dissertation of Shivin Devgon, titled Automating Kitting Novel 3D Objects into Novel 3D Cavities Using Self-Supervised Learning of $SE(3)$ Transforms, is approved:

Chair	 <small>Ken Goldberg (May 14, 2021 13:36 PDT)</small>	Date	5/14/2021
	 <small>Avideh Zakhor (May 14, 2021 14:15 PDT)</small>	Date	5/14/2021

University of California, Berkeley

Automating Kitting Novel 3D Objects into Novel 3D Cavities
Using Self-Supervised Learning of $SE(3)$ Transforms

Copyright 2021
by
Shivin Devgon

Abstract

Automating Kitting Novel 3D Objects into Novel 3D Cavities Using Self-Supervised Learning of $SE(3)$ Transforms

by

Shivin Devgon

Master of Science in Electrical Engineering and Computer Sciences

University of California, Berkeley

Professor Ken Goldberg, Chair

In industrial part kitting, 3D objects are inserted into cavities for transportation or subsequent assembly. Kitting is a critical step in industrial transportation and assembly, as kitting can decrease downstream processing and handling times and enable lower storage and shipping costs. Incorporating recent advances in deep learning and depth sensing can allow us to better automate this crucial process.

We present Kit-Net, a framework for kitting previously unseen 3D objects into cavities given depth images of both the target cavity and an object held by a gripper in an unknown initial orientation. Kit-Net uses self-supervised deep learning and data-augmentation to train a Convolutional Neural Network (CNN) to robustly estimate 3D rotations between objects and matching concave or convex cavities using a large dataset of simulated depth image pairs. Kit-Net then uses the trained CNN to implement a controller to orient and position novel objects for insertion into novel prismatic and conformal 3D cavities. Experiments in simulation suggest that Kit-Net can orient objects to have a 99% average intersection volume between the object mesh and that of the target cavity. Physical experiments with 3 industrial objects suggest that Kit-Net can successfully insert objects into cavities with a 63% success rate while a baseline which restricts itself to 2D rotations succeeds only 18% of the time.

To my parents,

Contents

Contents	ii
List of Figures	iii
List of Tables	v
1 Introduction	1
2 Related Work	3
2.1 Classical Methods	3
2.2 Relative Pose Estimation	3
2.3 Pose Estimation for Industrial Automation	4
2.4 Kitting Novel Objects	4
3 Self-Supervised Learning of Rotation Transforms	5
3.1 Introduction	5
3.2 Background	5
3.3 Problem Statement	6
3.4 Kit-Net Framework	7
4 Orienting and Kitting Objects	11
4.1 Simulation Experiments	11
4.2 Physical Experiments	14
5 Discussion and Future Work	20
5.1 Discussion	20
5.2 Limitations and Future Work:	20
Bibliography	22

List of Figures

1.1	Physical experiments using the ABB YuMi and a Photoneo depth camera. (Left) A suction gripper holds the handrail bracket, an object unseen during training time, near the kitting cavity. Kit-Net orients the handrail bracket for insertion into the cavity through 5 steps. A) Starting state. B) Rotate the object by 180° to face the 3D camera and minimize occlusion from the gripper. C) Iteratively orient the object into a goal configuration. D) Rotate by 180 degrees and align centroids of the object and cavity to prepare for insertion. E) Insert and release.	2
3.1	Successful Kitting: Visualization of successfully kitting a 3D object into a concave cavity.	6
3.2	Datapoint in the training dataset: The pear object in its starting configuration I^s (left) and goal configuration I^g (right). Here ${}_sR^g$ is a rotation of 23° . These depth images include data augmentation of zeroing out pixels randomly, rectangular cuts, and cropping as described in Subsection 3.4.	8
4.1	Endstop Holder and Target Prismatic Cavity in Simulation: Given I^s , an image of an object in some configuration (top left) and I^g , an image of a target prismatic box to which the object must be aligned (top right), the objective is to find a 3D rotation ${}_s\hat{R}^g$ that would allow the object to fit within the box. In simulation experiments, R^g is a X° rotation from R^s , where $X \in (0, 30)$. The image in the figure shows a 30° rotation, meaning the object must be rotated by 30° to perfectly fit it inside the prism. 3D models corresponding to I^s and I^g are shown in the bottom row for clarity.	13
4.2	Aligning Objects to Prismatic Cavities in Simulation: We evaluate Kit-Nets ability to align objects with prismatic cavities under the percent fit metric introduced in Subsection 4.1 across 512 depth image pairs for each of 174 objects not seen during training. Given (I^s, I^g) , the network predicts ${}_s\hat{R}^g$ that will allow it to fit inside the cavity. We bin results by object eccentricity and observe that the mean percent fit decreases for objects of higher eccentricity. Kit-Net outperforms both the 2D and random baselines by a greater amount as object eccentricity increases.	14
4.3	Examples of Novel Objects for Kit-Net Simulation Experiments: The four test objects are unseen during training and have eccentricity greater than 2, meaning their minimum volume bounding boxes are narrow and long. An outline of the corresponding minimum volume bounding box is shown around each part.	15

4.4	Kit-Net Simulation Results: We visualize data from 100 runs on each of the 4 objects shown in Figure 4.3. All objects require a 30° rotation to be in alignment with the prismatic target at iteration 0, but their initial percent fits differ due to different eccentricities. Results suggest that Kit-Net is able to successfully align all 4 objects with their respective prismatic cavities while the baseline, which restricts itself to 2D rotations, performs significantly worse on all 4 objects.	15
4.5	Objects for Kit-Net Physical Experiments: We use 3 packaged industrial objects that can be commonly found in a hardware store. These objects were selected for their complex geometries, making precise orientation critical for effective kitting.	16
4.6	Examples of Physical Kitting Cavities: The handrail bracket (bottom) and the ornamental handrail bracket (top), next to the corresponding convex cavity (left) and concave cavity (right).	17
4.7	Generating Negative Goal Images: After taking the original image (left), we segment out all parts that don't belong to the cavity (middle). Then, we project from depth to point cloud, rotate the pointcloud about its centroid, and deproject to depth image to get I^g (right).	18
4.8	Kit-Net Failure Cases: The top-left image shows a configuration of the ornamental handrail bracket where the suction gripper occludes the handle below the base. The bottom-left image shows the sink handle. Although Kit-Net was able to orient the handle correctly for insertion, the centroid matching had a small error in estimating translation and the cavity does not have enough slack to be properly inserted. The center and right images show depth images for the ornamental handrail bracket for the concave conformal cavity and convex conformal cavity, respectively. The inside of the concave cavity is very thin and the angle of the camera makes it hard to perfectly image it, resulting in a poor depth image (center image). This leads to 0 successes for both the baseline and for Kit-Net when the initial rotation is 60° away from the desired rotation for kitting.	19

List of Tables

4.1	Physical Experiments Results for Convex Cavities: We report the number of successful kitting trials for Kit-Net and the 2D baseline over 10 trials for 3 previously unseen objects with initial rotations of 30° and 60° . Results suggest that Kit-Net significantly outperforms the 2D baseline for initial rotations of 60° and outperforms the baseline for two out of three objects for initial rotations of 30°	17
4.2	Physical Experiments Results for Concave Cavities: We report the number of successful kitting trials for Kit-Net and the 2D baseline over 10 trials for 3 previously unseen objects with initial rotations of 30° and 60° . Results suggest that Kit-Net significantly outperforms the baseline in all settings except for the handrail bracket with an initial rotation of 60° , for which neither Kit-Net nor the baseline can successfully kit the object.	18

Acknowledgments

First, I would like to thank my family. Without their sacrifices and support, I would never have been able to have the opportunity to study at Berkeley and pursue my passions.

I also want to thank Professor Ken Goldberg for being a great advisor and mentor for these past two years. Thanks to him and AUTOLAB, I've been able to have invaluable experiences and prove myself with publications.

Huge thanks to Jeffrey Ichnowski, Ashwin Balakrishna, and Michael Danielczuk. You all guided me at every step along the way and I could not have enjoyed these two years of research and published all my papers without your help. More thanks to my peers at AUTOLAB such as Harry Zhang, Priya Sundaresan, Ryan Hoque, and Adi Ganapathi.

Thank you to Daniel Brown, Eduardo Rocha, Shirin Joshi, Eugen Solowjow, for helping me out and guiding me on this research paper. Thank you to Professor Avidesh Zakhori for teaching me the fundamentals of computer vision and for reading my thesis. Thank you to all professors that taught my classes with passion and inspired my love for learning.

Finally, thank you to all my friends that made my time at Berkeley a lifelong memory and gave me such an enjoyable four years.

Chapter 1

Introduction

Kitting is a critical aspect of industrial automation that involves organizing and placing 3D parts into complementary cavities. This process saves time on the manufacturing line and frees up space to reduce shipping and storage cost. Automating kitting requires picking and rotating a part to a desired position and orientation, then inserting it into a cavity that loosely conforms to the object geometry. However, this process is a great challenge, and in industry, most kitting is performed manually.

Given a 3D CAD model of the object to be inserted and the desired object pose, one approach is to directly estimate the object pose and the transformation to the desired pose [2, 23]. However, CAD models may not be available for all objects to be kitted and are time consuming to create for every object, motivating an algorithm that can kit previously unseen objects without requiring such models. Prior work has considered kitting objects without models, but has focused on $SE(2)$ transforms (1D rotation and 2D translation) for extruded 2D polygonal objects [25, 26]. In contrast, we aim to estimate $SE(3)$ rotations for novel objects such that they can be inserted into a novel cavity.

We formalize the problem of rotating and translating a novel 3D object to insert it into a novel kitting cavity and present Kit-Net, a framework for inserting previously unseen 3D objects with unknown geometry into a novel target cavity given depth images of the object in its current orientation and a depth image of either a flipped (convex) or standard (concave) target cavity. Kit-Net extends prior work from Devgon et al. [7], which used simulation and self-supervision to train a deep neural network to directly estimate 3D transformations between the two depth images. Given the trained deep neural network, a depth image of a previously unseen insertion cavity, and a depth image of a previously unseen object, Kit-Net iteratively estimates the $SE(3)$ transform to reorient and insert the object, without requiring detailed knowledge of its geometry. Kit-Net improves on prior work by (a) introducing dataset augmentations that make the controller more robust, (b) using a suction cup gripper to minimize object occlusion during rotation, (c) incorporating 3D translations, and (d) applying the resulting controller to kit novel objects into novel cavities on a physical robot. We evaluate Kit-Net both in simulation and in physical experiments on an ABB YuMi robot with a suction gripper and overhead depth camera. Experiments in simulation suggest that Kit-Net can orient objects to have a 99 % average intersection volume between the object mesh and that of the




Figure 1.1: Physical experiments using the ABB YuMi and a Photoneo depth camera. (Left) A suction gripper holds the handrail bracket, an object unseen during training time, near the kitting cavity. Kit-Net orients the handrail bracket for insertion into the cavity through 5 steps. A) Starting state. B) Rotate the object by 180° to face the 3D camera and minimize occlusion from the gripper. C) Iteratively orient the object into a goal configuration. D) Rotate by 180 degrees and align centroids of the object and cavity to prepare for insertion. E) Insert and release.

target cavity. Physical experiments with 3 industrial objects and cavities suggest that Kit-Net can kit objects at a 63 % success rate from a diverse set of initial orientations.

In this thesis, I performed all of the neural network training, simulation experiments, and physical experiments. I received consultation, guidance, and extensive help in writing the thesis from Jeffrey Ichnowski, Michael Danielczuk, Daniel S. Brown, Ashwin Balakrishna, Shirin Joshi, Eduardo M. C. Rocha, Eugen Solowjow, and Ken Goldberg. In particular, Jeffrey Ichnowski generated 3.1 and Michael Danielczuk visualized my results in 4.2 and 4.1

Chapter 2

Related Work

2.1 Classical Methods

Reorienting objects is an important subtask in a variety of domains, such as manufacturing, packing, and assembly. There has been significant prior work on reorienting objects using geometric algorithms. Goldberg [9] proposes a geometric algorithm that orients polygonal parts with known geometry without requiring sensors. Akella and Mason [1] extend the work of Goldberg with sensor-based and sensor-less algorithms for orienting objects with known geometry and shape variation. Kumbla, Marvel, and Gupta [11] propose a method for estimating object pose via computer vision and then reorient the object using active probing. Leveroni and Salisbury [12] optimize robot finger motions to reorient a known convex object while maintaining grasp stability. In contrast to the above works, which require prior knowledge of object geometry, Kit-Net can reorient objects without 3D object models.

2.2 Relative Pose Estimation

Melekhov et al. [16], Suwajanakorn et al. [20], Wen et al. [22], and Devgon et al. [7] use data-driven approaches to estimate the relative pose difference between images of an object in different configurations. Melekhov et al. [16] use a Siamese network to estimate the relative pose between two cameras given an RGB image from each camera. Suwajanakorn et al. [20] propose KeypointNet, a deep learning approach that learns 3D keypoints by estimating the relative pose between two different RGB images of an object of unknown geometry, but known category. Wen et al. [22] considers an object tracking task by estimating a change in pose between an RGBD image of the object at the current timestep and a rendering of the object at the previous timestep, but require a known 3D object model. We use the network architecture from Wen et al. [22] to train Kit-Net, and extend the self-supervised training method and controller from Devgon et al. [7] to kit novel objects into previously unseen cavities. We find that by extending Devgon et al. [7] to be more robust to object translations and using a suction gripper to reduce occlusions, Kit-Net is able to learn more accurate reorientation controllers.

2.3 Pose Estimation for Industrial Automation

There has also been significant interest in leveraging ideas in pose estimation for core tasks in industrial automation. Litvak, Biess, and Bar-Hillel [13] leverage CAD models and assemble gear like mechanisms using depth images taken from a camera on a robotic arm’s end effector. Stevic, Christen, and Hilliges [19] estimate a goal object’s pose to perform a shape assembly task involving inserting objects which conform to a specific shape template into a prismatic cavity. Zachares et al. [24] combines vision and tactile sensorimotor traces for an object fitting task involving known holes and object types. Huang et al. [10] consider the problem of assembling a 3D shape composed of several different parts. This method assumes known part geometry and develops an algorithm to generate the 6-DOF poses that will rearrange the parts to assemble the desired 3D shape. In contrast to the above work, we focus on the problem of designing a controller which can reorient and place a novel object within a previously unseen cavity for industrial kitting tasks.

2.4 Kitting Novel Objects

Novel object kitting has also seen recent interest from the robotics community. Zakka et al. [25] introduce Form2Fit, an algorithm which learns $SE(2)$ transforms to perform pick-and-place for kitting novel planar objects. In contrast, we consider 6DOF transforms of novel 3D objects. Zeng et al. propose a network for selecting suction grasps and grasp-conditioned placement, which can generalize to multiple robotic manipulation tasks, including pick-and-place for novel flat objects. Zeng et al. focuses on $SE(2)$ rotations and translations for pick-and-place tasks involving novel flat, 2D extruded objects. Zeng et al. [26] also presents an algorithm for $SE(3)$ pick-and-place tasks, but only evaluate the algorithm on 2D extruded objects. In contrast, we use Kit-Net to kit novel 3D objects with a wide range of complex geometries.

Chapter 3

Self-Supervised Learning of Rotation Transforms

3.1 Introduction

This chapter formalizes the problem of rotating and translating a novel 3D object to insert it into a novel kitting cavity. To do this, we use a simulation environment to generate depth image pairs of 3D objects in different orientations. Then, we train a Convolutional Neural Network (CNN) to learn to estimate the relative rotation between these orientations given the two depth images. We then describe how this CNN can be used to orient and insert a novel 3D object into a novel cavity.

3.2 Background

We use unit quaternions to represent rotations. A quaternion $q = q_r + q_i i + q_j j + q_k k$ is an extension of complex numbers with a real component q_r and 3 scaled fundamental imaginary units i , j , and k . We represent q using the convention of a vector $[q_r \ q_i \ q_j \ q_k]^T$. A *unit* quaternion has the property that $\|q\|^2 = q_r^2 + q_i^2 + q_j^2 + q_k^2 = 1$, and can represent a rotation with properties we make use of in this work:

Normalization A unnormalized or non-unit quaternion \tilde{q} can be converted to a unit quaternion by dividing by its norm $\tilde{q}/\|\tilde{q}\|_2$

Angle difference The angle of rotation θ between two quaternions q_0 and q_1 is $2 \cos^{-1} |\langle q_0, q_1 \rangle|$

Rotation difference The quaternion rotation between two quaternions is $q_{\text{diff}} = q_0 q_1^{-1}$, where q_1^{-1} is the conjugate (negated imaginary components) of q_1

Slerp The spherical linear interpolation or *slerp* between rotations q_0 and q_1 by a scalar $t \in [0, 1]$ is $\text{Slerp}(q_0, q_1, t) = (q_0 \sin(1-t)\theta + q_1 \sin t\theta) / \sin \theta$, where θ is the angle between the two rotations [18].




Figure 3.1: **Successful Kitting:** Visualization of successfully kitting a 3D object into a concave cavity.

Angle of rotation The angle of rotation of a quaternion is defined by $\text{quat2angle}(q) = 2 \cos^{-1} q_r$

Axis of rotation The axis of rotation of a quaternion is $[q_i \ q_j \ q_k] / \sqrt{1 - q_r^2}$

Double Coverage Quaternions double cover $SO(3)$, in that q and $-q$ represent the same rotation.

3.3 Problem Statement

Let $T^s \in SE(3)$ be the initial 6D pose of a unknown 3D rigid object O in the world coordinate frame, consisting of a rotation $R^s \in SO(3)$ and a translation $t^s \in \mathbb{R}^3$. Given O with starting pose T^s and a kitting cavity K , let $\mathcal{G} \subset SE(3)$ be the set of goal 6-DOF poses of object O that result in successful kitting. The goal is to orient object O to $T^g \in \mathcal{G}$, where T^g consists of rotation R^g and translation t^g . Figure 3.1 shows a simulated example where a 3D object O is successfully kitted into a concave cavity K .

Assumptions

We assume access to depth images of a rigid object O and a kitting cavity K . The cavity image may be taken with the cavity either in its standard, concave orientation (i.e., open to object insertion), or flipped, convex orientation (i.e., mirroring the shape of the object to be inserted). We also assume that orienting O to a pose in $T^g \in \mathcal{G}$ and releasing the gripper results in a successful kitting action.

Input

Let $I^s \in \mathbb{R}^{H \times W}$ be a depth image observation of the object in initial pose T^s , and $I^k \in \mathbb{R}^{H \times W}$ be the depth image observation of a kitting cavity, K . See Figure 4.6 for physical examples of objects and kitting cavities.

Output

The goal is to successfully kit an unknown 3D object O into a novel 3D cavity K (Fig. 1.1, Fig. 3.1). Thus, we aim to transform the initial pose T^s into a goal pose that fits into the cavity (i.e., $T^g \in \mathcal{G}$). For objects with symmetries, the objective is to estimate and orient objects relative to a (symmetric) orientation that results in successful insertion into the cavity K .

3.4 Kit-Net Framework

We present Kit-Net, a framework that first reorients the object into a pose that can be successful kit in a desired cavity, and then translates and inserts the object into the cavity. We do this by learning to estimate ${}_sT^g \in SO(3)$, a relative transformation consisting of rotation ${}_sR^g$ and translation ${}_s\hat{t}^g$, which transforms the object from T^s to T^g for some $T^g \in \mathcal{G}$. The overall approach is to (1) compute an estimate of ${}_sR^g$, denoted ${}_s\hat{R}^g$, given only image observations I^s and I^k , (2) iteratively reorient the object according to ${}_s\hat{R}^g$ until $\hat{R}^g \in \mathcal{G}$, and (3) translate the object by ${}_s\hat{t}^g$, an estimate of the translation between the start and goal object translations, such that it lies in the kitting cavity.

We first discuss preliminaries (Subsection 3.4) and then describe the key new ideas in training Kit-Net (Subsection 3.4) which make it possible to design a controller to rotate and translate an object to fit it in a cavity (Subsection 3.4).

Preliminaries: Estimating Quaternion Rotations in 3D

Devgon et al. [7] presented a self-supervised deep-learning method to align two 3D objects. The method takes two depth images as input: I^s , an image of the object in its current orientation R^s and I^g , an image of the same object in its desired goal orientation R^g . It trains a deep neural network $f_\theta : \mathbb{R}^{H \times W}, \mathbb{R}^{H \times W} \rightarrow SO(3)$ to estimate the rotation (parametrized by a quaternion) between the pair of images (I^s, I^g) . Then, using a proportional controller, it iteratively rotates the object to minimize the estimated rotational difference. This controller applies $\eta {}_s\hat{R}^g$ until the network predicts that the current object rotation R^s is within $\delta = 0.5^\circ$ of R^g , or until the controller reaches an iteration limit. The tunable constant η is the Spherical-Linear intERPolation (slerp) factor describing the proportion of \hat{R}^g that the controller will apply to O . Devgon et al. use a small slerp value of $\eta = 0.2$ and a maximum iteration limit of 50 rotations.

For training, Devgon et al. generate a dataset consisting 200 pairs of synthetic depth images for each of the 698 training objects with random relative rotations, for a total of 139,600 pairs. To account for parallax effects, each pair of images were generated from a fixed translation relative to the camera. Devgon et al. propose three loss functions to train f_θ : a cosine loss, a symmetry-resilient loss, and a hybrid of the two, with the hybrid loss outperforming the first two. Note that Devgon et al. do not consider cavity insertion tasks, which is complicated by the need to reason about the translations and required alignment with a cavity.

Kit-Net improves on Devgon et al. by (1) introducing dataset augmentations to make the controller more robust, (2) using a suction cup gripper to minimize object occlusion during rotation,




Figure 3.2: **Datapoint in the training dataset:** The pear object in its starting configuration I^s (left) and goal configuration I^g (right). Here ${}_sR^g$ is a rotation of 23° . These depth images include data augmentation of zeroing out pixels randomly, rectangular cuts, and cropping as described in Subsection 3.4.

and (3) incorporating 3D translations into the controller to enable kitting. We discuss these contributions in the following sections.

Kit-Net Dataset Generation, Augmentation, and Training

Kit-Net trains a neural network f_θ with a self-supervised objective by taking as input pairs of depth images (I^s, I^g) and estimating ${}_s\hat{R}^g$ from image pair (I^s, I^g) . As in Devgon et al., f_θ encodes each depth image into a length 1024 embedding, concatenates the embeddings, and passes the result through two fully connected layers to estimate a quaternion representation of the rotational difference between the object poses.

Initial Dataset Generation

In contrast to Devgon et al., we are interested in kitting, rather than just reorienting an object in the robot gripper. Thus, in this thesis we focus on two types of kitting cavities: prismatic cavities (Fig. 4.1) and conformal cavities (Fig. 4.6). We generate a separate dataset for each type of cavity and train a separate network for each dataset. To generate both datasets, we use the set of 698 meshes from Mahler et al. [15]. For each mesh, we generate 512 depth image pairs, for a total of 357,376 pairs. To do this we first generate a pair of rotations (R^s, R^g) , where R^s is generated by applying one rotation sampled uniformly at random from $SO(3)$ to O , and R^g is generated by applying a random rotation with rotation angle less than 30 degrees onto R^s . To generate the conformal cavity dataset, we then obtain a pair of depth images (I^s, I^g) by rendering the object in rotations R^s and R^g from an overhead view. The pair is labeled with the ground truth rotation difference between the images. To generate the prismatic cavity dataset we follow the same process

as above, except we fit and render a prismatic box around the rotated object (Fig. 4.3) and render the depth image pairs (Fig. 4.1). This process results in two datasets, each containing 357,376 total labeled image pairs (I^s, I^g) with ground truth rotation labels ${}_sR^g$.

Data Augmentation

We found that simply training a network directly on the datasets described above results in poor generalization to depth images from the physical system, which contain sensor noise, object occlusions from both the object itself and the arm or gripper, and 3D object translations within the image. To address these three points, we introduce a set of dataset augmentations to ease network transfer from simulated to real depth images. To simulate noise and occlusion in training, we randomly zero out 1 % of the pixels in each depth image and add rectangular cuts of width 30 % of zero pixels to the image, respectively. To simulate translations in training, we translate the object across a range of 10 cm in the x , y , and z axes with respect to the camera in the simulated images. We also crop the images at sizes from 5 % to 25 % greater than the object size with center points offset from the object’s centroid by 5 pixels to simulate I^s, I^g pairs generated from objects and cavities outside the direct overhead view for the kitting task. 3.2 visualizes an example datapoint from the training dataset.

Training

We adopt the network architecture from Wen et al. [22], as it is designed to be trained in simulation and demonstrates state-of-the-art performance on object tracking [6] by regressing the relative pose between two images. We use the hybrid quaternion loss proposed by Devgon et al. [7]. The network is trained with the Adam optimizer with learning rate 0.002, decaying by a factor of 0.9 every 5 epochs with an L2 regularization penalty of 10^{-9} .

Kit-Net Suction Gripper

Kit-Net uses an industrial uncontact suction gripper from Mahler et al. [14] to grasp the object for kitting. In contrast, Devgon et al. used a parallel jaw gripper. We find the suction gripper to be better suited for kitting because it reduces gripper occlusions and enables the robot to position the object directly inside the kitting cavity.

Kit-Net Controller

The Kit-Net controller consists of two stages: rotation and translation.

Rotation

Kit-Net first re-orientes an object using the depth image of the current object pose I^s and a depth image of the goal cavity pose I^g . In preliminary experiments, we found the rotation parameters used by Devgon et al. [7] (Subsection 3.4) to be overly conservative. Thus, to speed up the alignment

process, we use a larger slerp factor of $\eta = 0.8$. If the network predicts a rotation difference of less than $\delta = 5^\circ$, then we assume that the object is close enough to the required pose for kitting into the cavity and terminate the rotation controller. Because of the larger slerp value, Kit-Net is able to quickly reorient the object, thus we terminate the rotation controller after a maximum of 8 sequential rotations (8 iterations).

Translation

Once the rotational alignment is computed, Kit-Net computes a 2D translation to move the object directly over the target cavity, and then lowers the object and releases it into the cavity. To calculate the 2D translation, we perform centroid matching between the final depth image I^s of the object after rotation and the depth image of the cavity I^s to find ${}_s\hat{t}^s$.

Chapter 4

Orienting and Kitting Objects

4.1 Simulation Experiments

We first discuss metrics to evaluate performance in Subsection 4.1. We then introduce a baseline algorithm (Subsection 4.1) with which to compare Kit-Net and present experimental results in Subsection 4.1. In experiments, we first evaluate Kit-Net on re-orienting novel objects with unknown geometry into a target prismatic box in simulation (Subsection 4.1).

Baselines

Random Baseline

We also compare Kit-Net with a baseline that applies a randomly sampled rotation but with the correct rotation angle to evaluate how important precise reorientation is for successful kitting.

2D Baseline

To evaluate the importance of estimating 3D rotations for successful kitting, we compare Kit-Net to a baseline inspired by Form2Fit [25], which only considers 2D rotations when orienting objects for kitting. The baseline (1) aligns the centroids of the point clouds of the object and the cavity and (2) searches over all possible z -axis rotations at a 1° discretization to find the rotation that minimizes Chamfer distance between the centroid-aligned point clouds.

Metrics

Object Eccentricity

We categorize test objects by their *eccentricity*, which provides a measure of kitting difficulty. This categorization follows the intuition that objects that are more elongated along certain dimensions than others have a smaller set of acceptable orientations in which they can be successfully kit into a cavity. Let the eccentricity ε of a 3D object be $\varepsilon = A - 1$, where A is the aspect ratio (ratio of

longest side to shortest side) of the minimum volume bounding box of the object. This definition generalizes the 2D definition of eccentricity from [8] to 3D. Under this definition, a sphere has $\varepsilon = 0$, and if one axis is elongated by a factor p , then the resulting ellipsoid has eccentricity $p - 1$. This definition is also consistent with the intuition provided earlier, as a sphere is entirely rotationally symmetric, and thus does not require any reorientation for kitting. By contrast, the ellipsoid will require reorientation to ensure that its longer side is aligned to a region with sufficient space in the cavity. Thus, we use objects with high eccentricity in evaluating both Kit-Net and the baselines, as these objects pose the greatest challenge for kitting in practice.

Results

When evaluating Kit-Net in simulation, we have access to ground-truth object and cavity geometry. Thus, we evaluate kitting performance using the following percent fit metric:

$$\hat{f}(I^s, I^g, {}_s\hat{R}^g) = \frac{1}{N} \sum_{i=1}^N \mathbf{1}_{(p_i \in K)}, \quad (4.1.1)$$

where we sample N points within the object volume at configuration ${}_s\hat{R}^g R^s$ (after the target object has been rotated for insertion) and count the proportion of sampled points that also lie within cavity. This metric can be efficiently computed using ray tracing and effectively estimates how much of the object fits inside the target mesh after the predicted rotation. In experiments, we use $N = 10000$ sampled points to evaluate \hat{f} . Assuming the true percent fit metric is f , a 95 % confidence interval for f is $\hat{f} \pm 1.96 \sqrt{\frac{\hat{f}(1-\hat{f})}{10000}}$. For example, if $\hat{f} = 0.99$, then f lies between (0.988, 0.992) with 95 % confidence.

Simulated Kitting into a Prismatic Target

We first study whether Kit-Net can orient objects into alignment with a prismatic cavity that loosely conforms to their 3D geometry in simulation. Precisely, we first generate the prismatic cavity for the target by creating a mesh with faces corresponding to its minimum volume bounding box. We then rotate both the prismatic cavity and target to random orientations within 30 degrees of each other. The objective is to apply a rotation ${}_s\hat{R}^g$ that will allow the object to fit into the cavity. An example image pair of an object and an associated prismatic cavity is shown in Figure 4.1.

Fig. 4.2 shows the percent fit across 174 unseen test objects. We use the eccentricity ε of the objects to sort them into 5 bins of increasing difficulty (increasing ε). We find that Kit-Net is able to reliably kit novel objects, significantly outperforming the 2D rotation baseline. When averaged across all eccentricities, Kit-Net achieves an average fit of 99 % compared to an average fit of 94 % for the 2D baseline and 83 % when applying a random 30° quaternion. These results demonstrate the need for 3D rotations to solve complex kitting problems. Figure 4.2 demonstrates that Kit-Net is robust to highly eccentric objects which require the most precision for kitting. Kit-Net achieves an average fit of 90. % for objects with eccentricity greater than 8. The 2D rotation baseline performs




Figure 4.1: **Endstop Holder and Target Prismatic Cavity in Simulation:** Given I^s , an image of an object in some configuration (top left) and I^g , an image of a target prismatic box to which the object must be aligned (top right), the objective is to find a 3D rotation ${}_{s}\hat{R}^g$ that would allow the object to fit within the box. In simulation experiments, R^g is a X° rotation from R^s , where $X \in (0, 30)$. The image in the figure shows a 30° rotation, meaning the object must be rotated by 30° to perfectly fit it inside the prism. 3D models corresponding to I^s and I^g are shown in the bottom row for clarity.

especially poorly for these difficult objects and achieves an average fit of only 73 % while applying a random 30° quaternion results in an average fit of just 37 %.

As described in Subsection 3.4, Kit-Net iteratively orients each object using the controller until ${}_{s}\hat{R}^g \leq 5^\circ$ or until we hit the stopping condition of 8 rotations. Our previous results in Fig. 4.2 suggests that Kit-Net can consistently align objects within 5 controller steps. To better visualize the ability of Kit-Net to rapidly reorient an object for kitting, we plotted the per-iteration performance of Kit-Net for 4 test objects unseen during training time with high eccentricity ($\epsilon \geq 2$). Fig. 4.3 shows renderings of these objects along with outlines of the corresponding prismatic kitting cavities. Fig. 4.4 shows the average per-iteration percent fit across 100 controller rollouts of randomly sampled (I^s, I^g) pairs for each object. We find that Kit-Net is able to consistently align objects with their target prismatic cavities, and achieves a median fit percentage of 99.4 % after only 3 successive reorientations. By contrast, the 2D baseline is not able to surpass an average fit of 90 % for any of the objects. The results in Fig. 4.4 validate the importance of iteratively reorienting parts and




Figure 4.2: **Aligning Objects to Prismatic Cavities in Simulation:** We evaluate Kit-Nets ability to align objects with prismatic cavities under the percent fit metric introduced in Subsection 4.1 across 512 depth image pairs for each of 174 objects not seen during training. Given (I^s, I^g) , the network predicts ${}_s\hat{R}^g$ that will allow it to fit inside the cavity. We bin results by object eccentricity and observe that the mean percent fit decreases for objects of higher eccentricity. Kit-Net outperforms both the 2D and random baselines by a greater amount as object eccentricity increases.

demonstrates that applying multiple iterations of the rotation output by the trained network can greatly help to reduce the error between ${}_s\hat{R}^g R^s$ and R^g as compared to a single iteration.

4.2 Physical Experiments

Our previous experiments studied the effectiveness of Kit-Net for insertion tasks involving prismatic cavities. However, as shown in Fig. 4.5 and Fig. 4.6, many physical kitting tasks involve non-prismatic cavities. In this section, we study how Kit-Net can be used to kit objects in physical trials using depth images of the types of cavities shown in Fig. 4.6. We call these *conformal cavities*, as they “conform” to some degree to the object shape.

In these experiments we use a quaternion prediction network trained to predict the quaternion that will rotate a simulated depth image of an object to another simulated depth image of the same object in a different pose. We propose two possible methods for applying this trained network to kitting. Our first method is designed to work well with the clamshell cavities shown in Fig. 4.6. Rather than image the hole of the cavity, we define a *convex conformal cavity* to be the depth image of the inverted cavity. To obtain these depth images, we flip the cavity so the hole is pointing down and take a depth image of the positive mass of the cavity. The left image in Fig. 4.6 shows examples




Figure 4.3: **Examples of Novel Objects for Kit-Net Simulation Experiments:** The four test objects are unseen during training and have eccentricity greater than 2, meaning their minimum volume bounding boxes are narrow and long. An outline of the corresponding minimum volume bounding box is shown around each part.




Figure 4.4: **Kit-Net Simulation Results:** We visualize data from 100 runs on each of the 4 objects shown in Figure 4.3. All objects require a 30° rotation to be in alignment with the prismatic target at iteration 0, but their initial percent fits differ due to different eccentricities. Results suggest that Kit-Net is able to successfully align all 4 objects with their respective prismatic cavities while the baseline, which restricts itself to 2D rotations, performs significantly worse on all 4 objects.

of these convex conformal cavities. Our second method works with a *concave conformal cavity*, like that shown in Fig. 3.1 and the right image in Fig. 4.6, that are formed as impressions into a surface. These types of cavities cannot simply be flipped upside down to obtain a depth image of their shape. Instead, we take a depth image of the actual cavity (where the cavity has negative mass) and rotate it 180° about its principal axis.

We discuss the results for applying Kit-Net to novel convex conformal cavities in Subsection 4.2 and to novel concave conformal cavities Subsection 4.2. For the physical kitting experiments we



Figure 4.5: **Objects for Kit-Net Physical Experiments:** We use 3 packaged industrial objects that can be commonly found in a hardware store. These objects were selected for their complex geometries, making precise orientation critical for effective kitting.

measure success using a binary success metric for insertion by visually inspecting whether or not the object is completely contained in the target cavity.

Physical Kitting into Convex Conformal Cavities

We also evaluate Kit-Net in physical kitting trials on an ABB-YuMi robot with a Photoneo depth camera, shown in Fig. 1.1, using 4 packaged objects widely available in hardware stores and which are unseen during training (Fig. 4.5). To prepare objects for kitting, we carefully extract each tool and kitting shell from its packaging and spray paint the shell to facilitate depth sensing as shown in Fig. 4.6. We then place the kitting cavity open end down and image the cavity to generate I^g before flipping it to expose its opening for the insertion task. For each trial, we insert the object into the cavity by hand, grasp it using the robot’s suction gripper, translate it to be directly under the camera, and apply a random rotation of either 30° or 60° , uniformly sampled from $SO(3)$ to simulate grasping the object from a bin in a non-uniform pose. Then, we flip the object such that the object faces the overhead depth camera and the suction cup grasp is occluded from the camera by the object. This process is illustrated in Fig. 1.1.

Kit-Net then orients the object using the learned controller, and matches centroids between the object and cavity for insertion before flipping it again and attempting to kit it. Table 4.1 shows the number of successful kitting trials (out of 10 per object) of Kit-Net and the 2D baseline across 3 objects. We report a kitting trial as successful if the object is fully contained within the cavity from visual inspection. We observe that Kit-Net outperforms the baseline for 30° initial rotations on 2 of the 3 objects, performing similarly to the baseline on the sink handle. We find that Kit-Net



Figure 4.6: **Examples of Physical Kitting Cavities:** The handrail bracket (bottom) and the ornamental handrail bracket (top), next to the corresponding convex cavity (left) and concave cavity (right).

significantly outperforms the baseline on all objects for 60° initial rotations.

Kit-Net’s main failure modes are due to errors in the centroid matching procedure, as illustrated in Fig. 4.8. On the 30 degree sink handle task, Kit-Net aligned it correctly every time, but the centroid matching had it about 0.5 cm off, and there is no slack at the top of the cavity.

Object	Angle	2D Baseline	Kit-Net
Handrail bracket	30°	3/10	10/10
Ornamental handrail bracket	30°	8/10	10/10
Sink handle	30°	4/10	3/10
Handrail bracket	60°	1/10	9/10
Ornamental handrail bracket	60°	2/10	7/10
Sink handle	60°	0/10	7/10

Table 4.1: **Physical Experiments Results for Convex Cavities:** We report the number of successful kitting trials for Kit-Net and the 2D baseline over 10 trials for 3 previously unseen objects with initial rotations of 30° and 60° . Results suggest that Kit-Net significantly outperforms the 2D baseline for initial rotations of 60° and outperforms the baseline for two out of three objects for initial rotations of 30° .




Figure 4.7: **Generating Negative Goal Images:** After taking the original image (left), we segment out all parts that don't belong to the cavity (middle). Then, we project from depth to point cloud, rotate the pointcloud about its centroid, and deproject to depth image to get I^g (right).

Physical Kitting into Concave Conformal Cavities

The positive cavity task provides a detailed depth image, but requires the cavity to first be presented at a flipped, 180° rotation for imaging, which may not be feasible in an industrial environment. Here, we perform the same experiment as in Section 4.2, but instead generate I^g directly from an image of the cavity without flipping. Precisely, we segment out the cavity from an overhead depth image, deproject the depth image into its point cloud representation, and rotate the point cloud 180° around its center of mass. Then, we project the rotated point cloud to the depth image I^g . This process is illustrated in Fig. 4.7.

Object	Angle	2D Baseline	Kit-Net
Handrail bracket	30°	0/10	9/10
Ornamental handrail bracket	30°	0/10	7/10
Sink handle	30°	1/10	3/10
Handrail bracket	60°	0/10	7/10
Ornamental handrail bracket	60°	0/10	0/10
Sink handle	60°	1/10	4/10

Table 4.2: **Physical Experiments Results for Concave Cavities:** We report the number of successful kitting trials for Kit-Net and the 2D baseline over 10 trials for 3 previously unseen objects with initial rotations of 30° and 60° . Results suggest that Kit-Net significantly outperforms the baseline in all settings except for the handrail bracket with an initial rotation of 60° , for which neither Kit-Net nor the baseline can successfully kit the object.




Figure 4.8: **Kit-Net Failure Cases:** The top-left image shows a configuration of the ornamental handrail bracket where the suction gripper occludes the handle below the base. The bottom-left image shows the sink handle. Although Kit-Net was able to orient the handle correctly for insertion, the centroid matching had a small error in estimating translation and the cavity does not have enough slack to be properly inserted. The center and right images show depth images for the ornamental handrail bracket for the concave conformal cavity and convex conformal cavity, respectively. The inside of the concave cavity is very thin and the angle of the camera makes it hard to perfectly image it, resulting in a poor depth image (center image). This leads to 0 successes for both the baseline and for Kit-Net when the initial rotation is 60° away from the desired rotation for kitting.

Table 4.2 shows results from experiments with 3 novel objects from Fig. 4.5 across 10 controller rollouts. We observe that Kit-Net outperforms the baseline for initial rotations of both 30 and 60 degrees on the handrail bracket and sink handle, and for an initial rotation of 30° for the ornamental handrail bracket. For the ornamental handrail bracket, the depth image from the concave cavity is low quality as shown in Fig. 4.8 (center image), causing Kit-Net to fail when the object is 60 degrees away the correct insertion orientation. We examined this failure and found that it occurs because the cavity for the neck of the bracket is very thin, making it difficult to obtain a good depth image. Kit-Net also has low performance on the sink handle due to small errors in centroid matching, as discussed in the prior section. Fig. 4.8 (bottom left) shows an example failure case where the sink handle is correctly oriented but the translation is slightly off. There were also occasional cases (Fig. 4.8 (top left)) where the suction gripper occludes the handle of the ornamental handrail bracket. In these cases, the robot can only see the base, resulting in failure.

Chapter 5

Discussion and Future Work

5.1 Discussion

We present Kit-Net, a framework that uses self-supervised deep learning in simulation to kit novel 3D objects into novel 3D cavities. Results in simulation experiments suggest that Kit-Net can kit unseen objects with unknown geometries into a prismatic target in less than 5 controller steps with a median percent fit of 99 %. In physical experiments kitting novel 3D objects into novel 3D cavities, Kit-Net is able to successfully kit novel objects 63 % of the time while a 2D baseline which only considers $SE(2)$ transforms only succeeds 18 % of the time.

During this thesis we have learned how to incorporate modern advances in deep learning to facilitate automation of kitting novel 3D objects into novel cavities. We also have discussed critical assumptions about R^g , cavity geometry, and input depth images in order for Kit-Net to work in practice.

5.2 Limitations and Future Work:

In this thesis we do not consider suction grasps that would require the object to be regrasped (grasps that prevent both rotational alignment and insertion); however, this functionality may be necessary when grasping objects from a heap when R^g exceeds 60° . In future work, we will work to improve performance by using the predicted error from Kit-Net to regrasp the object in a new stable pose [21, 3, 4] before reattempting the kitting task, study Kit-Net’s performance with other depth sensors, and apply Kit-Net to kit objects that are initially grasped from a heap [17, 5].

We would also like to accurately quantify R^g for objects dropped into a cavity. We plan to do this in a dynamics simulation environment where an object in varying poses is dropped into a cavity. Furthermore, we want to train Kit-Net on a wider subset of $SE(3)$. While large rotations out of the plane result in depth image pairs of disparate parts of the object, in plane rotations don’t have this issue. We will also explore techniques to allow us to reduce the rotation prediction error for rotations with larger angles. We also observe that while Kit-Net was trained on depth image pairs where $R^g \leq 30^\circ$, it was occasionally able to kit objects where $R^g = 60^\circ$. We would like to quantify

this generalization and estimate which subsets of $SE(3)$ are feasible for the kitting task. We also conjecture that other methods for computing ${}_s\hat{t}^g$ such as least squares on a small set of corresponding points could lead to higher accuracy than doing the current method of centroid matching.

Additionally, we want to consider cavities which may differ significantly in geometry from the object being inserted. In order for Kit-Net to apply in practice, the cavity could be strategically designed to conform to the object geometry and to be effectively imaged by depth cameras.

Bibliography

- [1] Srinivas Akella and Matthew T. Mason. “Parts Orienting with Shape Uncertainty”. In: *International Conference on Robotics and Automation (ICRA)*. 1998.
- [2] C. Choy, W. Dong, and V. Koltun. “Deep Global Registration”. In: *2020 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*. 2020, pp. 2511–2520. DOI: 10.1109/CVPR42600.2020.00259.
- [3] Nikhil Chavan Dafle et al. “Extrinsic dexterity: In-hand manipulation with external forces”. In: *2014 IEEE International Conference on Robotics and Automation (ICRA)*. IEEE. 2014, pp. 1578–1585.
- [4] Michael Danielczuk et al. “Exploratory Grasping: Asymptotically Optimal Algorithms for Grasping Challenging Polyhedral Objects”. In: *Conference on Robot Learning (CoRL)*. 2020.
- [5] Michael Danielczuk et al. “X-ray: Mechanical search for an occluded object by minimizing support of learned occupancy distributions”. In: *arXiv preprint arXiv:2004.09039* (2020).
- [6] X. Deng et al. “PoseRBPF: A Rao–Blackwellized Particle Filter for 6-D Object Pose Tracking”. In: *IEEE Transactions on Robotics* (2021), pp. 1–15. DOI: 10.1109/TRO.2021.3056043.
- [7] S. Devgon et al. “Orienting Novel 3D Objects Using Self-Supervised Learning of Rotation Transforms”. In: *2020 IEEE 16th International Conference on Automation Science and Engineering (CASE)*. 2020, pp. 1453–1460. DOI: 10.1109/CASE48305.2020.9217018.
- [8] K. Goldberg, M.H. Overmars, and A. F. van der Stappen. “Geometric Eccentricity and the Complexity of Manipulation Plans”. In: *Algorithmica* 26.3-4 (Mar. 2000), pp. 494–514. DOI: 10.1007/s004539910022. URL: <https://doi.org/10.1007/s004539910022>.
- [9] Kenneth Y Goldberg. “Orienting polygonal parts without sensors”. In: *Algorithmica* 10.2-4 (1993), pp. 201–225.
- [10] J. Huang et al. “Generative 3D Part Assembly via Dynamic Graph Learning”. In: *ArXiv abs/2006.07793* (2020).
- [11] Nithyananda B Kumbla, Jeremy A Marvel, and Satyandra K Gupta. “Enabling Fixtureless Assemblies in Human-Robot Collaborative Workcells by Reducing Uncertainty in the Part Pose Estimate”. In: *2018 IEEE 14th International Conference on Automation Science and Engineering (CASE)*. IEEE. 2018, pp. 1336–1342.

- [12] Susanna Leveroni and Kenneth Salisbury. “Reorienting Objects with a Robot Hand Using Grasp Gaits”. In: *Int. S. Robotics Research (ISRR)*. 1996.
- [13] Yuval Litvak, A. Biess, and Aharon Bar-Hillel. “Learning Pose Estimation for High-Precision Robotic Assembly Using Simulated Depth Images”. In: *2019 International Conference on Robotics and Automation (ICRA)* (2019), pp. 3521–3527.
- [14] Jeffrey Mahler et al. “Dex-Net 3.0: Computing Robust Robot Suction Grasp Targets in Point Clouds using a New Analytic Model and Deep Learning”. In: *arXiv preprint arXiv:1709.06670* (2017).
- [15] Jeffrey Mahler et al. “Learning ambidextrous robot grasping policies”. In: *Science Robotics* 4.26 (2019), eaau4984.
- [16] Iaroslav Melekhov et al. “Relative camera pose estimation using convolutional neural networks”. In: *Advanced Concepts for Intelligent Vision Systems - 18th International Conference, ACIVS 2017, Proceedings*. Springer Verlag, 2017, pp. 675–687.
- [17] Adithyavairavan Murali et al. “6-dof grasping for target-driven object manipulation in clutter”. In: *2020 IEEE International Conference on Robotics and Automation (ICRA)*. IEEE. 2020, pp. 6232–6238.
- [18] Ken Shoemake. “Animating Rotation with Quaternion Curves”. In: *SIGGRAPH Comput. Graph.* 19.3 (1985), pp. 245–254.
- [19] Stefan Stevic, Sammy Christen, and Otmar Hilliges. “Learning to Assemble: Estimating 6D Poses for Robotic Object-Object Manipulation”. In: *IEEE Robotics and Automation Letters* 5 (2020), pp. 1159–1166.
- [20] Supasorn Suwajanakorn et al. “Discovery of Latent 3D Keypoints via End-to-end Geometric Reasoning”. In: *Proc. Advances in Neural Information Processing Systems*. Ed. by S. Bengio et al. Curran Associates, Inc., 2018, pp. 2059–2070.
- [21] Pierre Tournassoud, Tomás Lozano-Pérez, and Emmanuel Mazer. “Regrasping”. In: *Proceedings. 1987 IEEE international conference on robotics and automation*. Vol. 4. IEEE. 1987, pp. 1924–1928.
- [22] Bowen Wen et al. “SE(3)-TrackNet: Data-driven 6D Pose Tracking by Calibrating Image Residuals in Synthetic Domains”. In: *2020 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)* (2020), pp. 10367–10373.
- [23] Yu Xiang et al. “PoseCNN: A convolutional neural network for 6D object pose estimation in cluttered scenes”. In: *arXiv preprint arXiv:1711.00199* (2017).
- [24] Peter A. Zachares et al. “Interpreting Contact Interactions to Overcome Failure in Robot Assembly Tasks”. In: *ArXiv abs/2101.02725* (2021).
- [25] Kevin Zakka et al. “Form2Fit: Learning Shape Priors for Generalizable Assembly from Disassembly”. In: *2020 IEEE International Conference on Robotics and Automation (ICRA)* (2020), pp. 9404–9410.

- [26] Andy Zeng et al. “Transporter Networks: Rearranging the Visual World for Robotic Manipulation”. In: *Conference on Robot Learning (CoRL)* (2020).