Contact Pose Identification for Peg-in-Hole Assembly under Uncertainties

Shiyu Jin, Xinghao Zhu, Changhao Wang, and Masayoshi Tomizuka

Abstract— Peg-in-hole assembly is a challenging contact-rich manipulation task. There is no general solution to identify the relative position and orientation between the peg and the hole. In this paper, we propose a novel method to classify the contact poses based on a sequence of contact measurements. When the peg contacts the hole with pose uncertainties, a tilt-then-rotate strategy is applied, and the contacts are measured as a group of patterns to encode the contact pose. A convolutional neural network (CNN) is trained to classify the contact poses according to the patterns. In the end, an admittance controller guides the peg towards the error direction and finishes the peg-in-hole assembly. Simulations and experiments are provided to show that the proposed method can be applied to the peg-in-hole assembly of different geometries. We also demonstrate the ability to alleviate the sim-to-real gap.
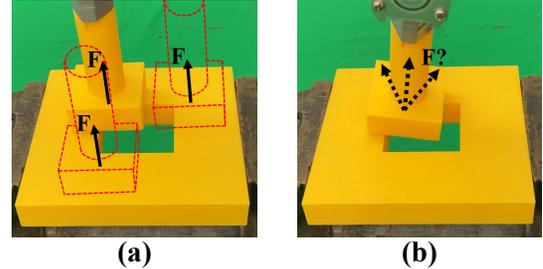
Fig. 1. (a) The same upward contact force could come from many possible contact poses. (b) The same contact pose could generate many possible contact forces within the Coulomb friction cone.

I. INTRODUCTION

Robotic peg-in-hole assembly has been studied for decades. It is challenging because it requires accurate state estimations of the peg and the hole for alignment, and a combination of precise planning and control algorithms for insertion.

Identifying the contact pose, the relative position and orientation between the peg and the hole, is required to align the peg and the hole before insertion. Visual feedback is the most common strategy to identify the pose [1, 2]. However, vision sensors suffer from high precision requirements and occlusions during the assembly task. In order to avoid such problems, search-based algorithms such as random search or spiral search [3] have been proposed to compensate the uncertainties of contact pose. The search strategy generates a search path within the search area for hole localization, which is not efficient especially when the search area is large and the search dimension is high.

For insertion, the clearance between the peg and the hole is usually smaller than the precision of a robot. A tiny position and orientation error could cause workpieces to jam and wedge and may lead to failure or even damage to the workpieces. Compliance, either passive or active, has shown to be effective in handling the small uncertainties of position and orientation. Passive compliance utilizes passive compliance hardwares such as RCC [4, 5] to compensate uncertainties. In constrast, active compliance applies control strategies from software to let the robot mimic the spring-damping behavior [6, 7].

In contact-rich scenarios, force/torque-based method normally conveys more information than vision-based and

Department of Mechanical Engineering, University of California, Berkeley, CA, USA. {jsy, zhuxh, changhaowang, tomizuka}@berkeley.edu

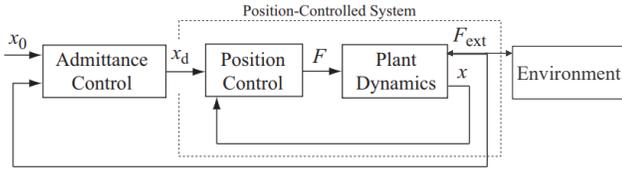search-based methods. Tang [8] analyzed a three-point contact model for round peg and hole. But the method lacked the ability to generalize to complex geometries. Kim proposed a peg shape recognition and hole detection algorithm using the force/torque sensor by inclining the peg in all directions, but their method suffered from the cumulative error [9]. In recent years, many learning-based methods have been proposed to solve the peg-in-hole assembly problem [10]–[14]. They treated the task as a Markov decision process, where the contact feedback at the current time step is used to determine the action of the next step. However, the mapping from the force/torque feedback to the contact pose is not injective as shwon in Fig. 1. On one hand, the same contact forces can be measured at different contact poses. On the other hand, the same contact pose could generate different contact forces, i.e. all the possible forces within the Coulomb friction cone. To deal with the above problem, particle filter was applied to identify the location based on multiple observations in [15, 16]. However, it is time-consuming to generate the force-position mapping in the real world and hard to generalize.

In this paper, we propose a novel method that can identify the contact poses based on a sequence of contact measurements. At initialization, the peg contacts the hole with pose uncertainties. The peg then follows a designed tilt-then-rotate motion to make contact with the hole. The contact measurements are plotted in polar coordinates to generate a group of patterns. An injective mapping between the patterns and contact poses is learned by a convolutional neural network (CNN), which classifies the contact poses based on the error directions. Finally, an admittance controller will guide the peg towards the error direction and finish insertion. There are two main contributions of this paper. 1) We construct the mapping using a sequence of measurements as input instead of feedback at one single time step. This makes the mapping become one-to-one. 2) We classify the contact pose based

Fig. 2. Admittance Control.

$$F + F_{ext} = m\ddot{x} \tag{1}$$

$$F = k_p(x_d - x) - k_d\dot{x} \tag{2}$$

$$F_{ext} = M_d(\ddot{x}_d - \ddot{x}_0) + D_d(\dot{x}_d - \dot{x}_0) + K_d(x_d - x_0) \tag{3}$$

where M_d, D_d, and K_d represent the desired inertia , damping, and stiffness, respectively. k_p and k_d are PD position control gains.

C. Assembly Strategy

Peg-in-hole assembly has been studied for decades. An efficient and widely used assembly strategy divides the task into several stages [9, 17]: initialization, approaching, contact pose estimation, alignment, and insertion. At initialization, the peg and the hole are fixed on the robot manipulator and the table, respectively. A vision system is applied to roughly estimate the pose of the hole. At approaching, the peg approaches to the hole with an admittance controller. With well-tuned controller parameters, the plane contact between the flat surface of the peg and the hole could eliminate the pose uncertainties in 3 dimensions, roll axis, pitch axis, and z-axis. At contact pose estimation, the peg explores along the surface of the hole to estimate the relative position and orientation between the peg and the hole. This stage eliminates the uncertainties in x and y axes. Finally, based on the contact pose estimation, the peg can slide towards the hole and finish insertion with an admittance controller. Small oscillation is added to the yaw axis in this stage, together with admittance control, to compensate small uncertainties of yaw axis. In this paper, we mainly focus on the contact pose estimation stage, which is introduced in section III.

III. PROPOSED METHOD

Peg-in-hole assembly can be accomplished easily by a human even with eyes closed. The human will first use the peg to make contact with the hole. Then he/she will locally move the peg to sense hole's location based on a sequence of contacts instead of just one single contact. If there is a hole in one direction, the tip of the peg could slide into the hole a little bit and the force/torque feedback also have an impulse in that direction. Based on the historical measurements in a sequence of contacts, the human keeps updating the knowledge of the contact pose and eliminating the hole uncertainties.

Inspired by the human strategy, we propose to use a sequence of contact feedback to identify the contact pose under uncertainties (Fig. 3).

A. Tilt-then-Rotate Strategy

The peg contacts the hole after the approaching stage (Fig. 4.1). The peg and the hole have some overlaps but are not aligned well due to the uncertainties of the contact pose. We propose a tilt-then-rotate strategy to identify the contact pose.

We tilt the peg for α degrees in all directions by rotating the peg for 2π (Fig. 4). The tilt-then-rotate trajectory can be

---

on patterns, which improves the generalization ability of the proposed method. It can even tackle the sim-to-real gap.

The remainder of this paper is organized as follows. Section II introduces the background including task description, admittance control, and assembly strategy. Section III describes the proposed contact pose identification method according to contact patterns. Section IV shows the performance of the proposed method by both simulations and real-world experiments. Section V discusses the advantages and disadvantages of the proposed method and proposes future work.

II. BACKGROUND

A. Task Description

We focus on the peg-in-hole assembly task under pose uncertainties. Generally speaking, the pose of the peg and the hole might be noisy due to sensor inaccuracy. To simply the problem, we assume the peg is fixed with the robot end-effector, and the pose can be obtained via forward kinematics. The hole is fixed on the table, and the pose can be estimated by a vision system with uncertainties in 6 degrees of freedom (DOF). The magnitudes of the uncertainties are roughly $\pm 20mm$ and $\pm 3°$ for position and orientation respectively, which are determined by the precision of the visual system. The clearance between the peg and the hole is $1mm$.

The goal of the task is to compensate the uncertainties of contact pose and achieve the peg-in-hole assembly. The contact surfaces of both the peg and the hole are assumed to be flat.

B. Admittance Control

Admittance control [6, 7] is widely used in robotic manipulation tasks to handle contact dynamics. By adding a virtual spring-damping system, the contact between the robot and the environment becomes soft, which improves the manipulation performance and prevents from damaging either the robot or the environment. We apply admittance control to the following assembly strategy to track the desired peg trajectory and compensate small uncertainties in assembly.

In admittance control, the desired pose x_0 and measured external force/torque F_{ext} are inputs to the admittance control block (Fig. 2), which generates the reference pose x_d for the PD position control.
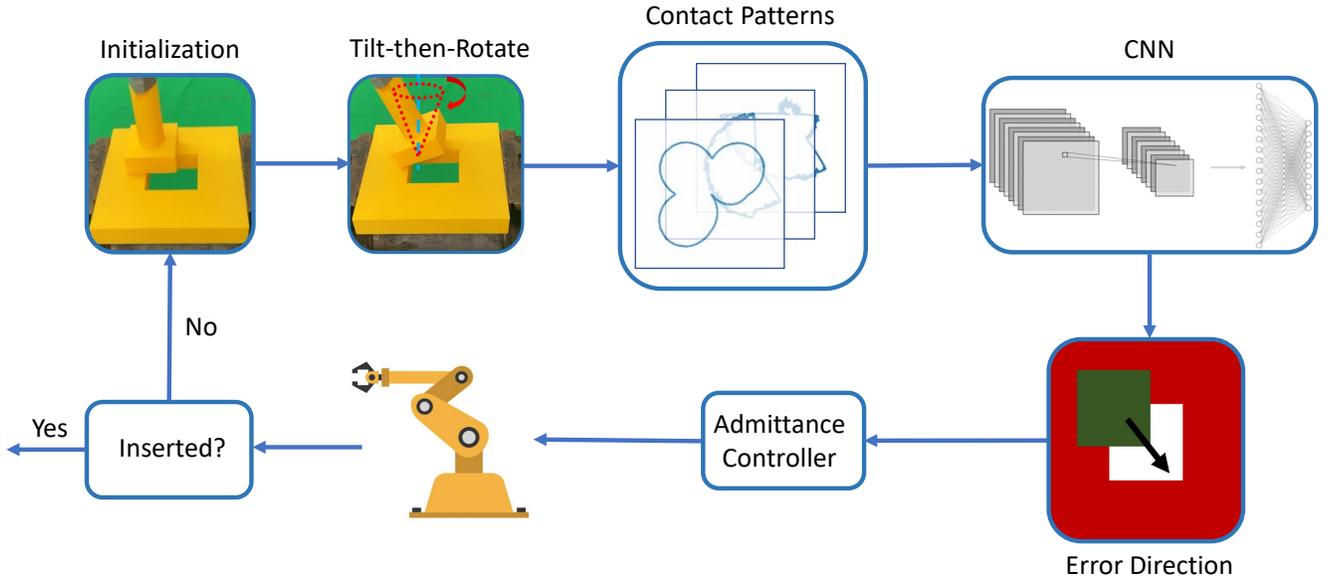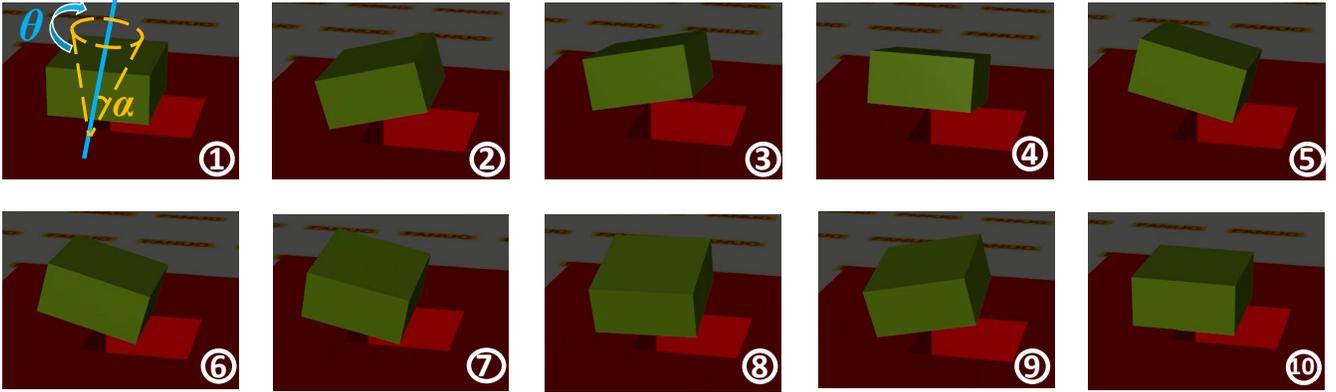
Fig. 3. Framework of the proposed method.

Fig. 4. Snapshots of the tilt-then-rotate strategy. The blue line is z-axis. The yellow cone represents the designed trajectory for rotation. Tilt (2) then rotate (2-9) the peg for 2π. While the peg is being rotated, a constant downward force is applied to maintain a single point contact (3,5,9), line contact (2,4,6,8), or two points contact (7) between the peg and the hole.

described as continuously changing θ from 0 to 2π in order to change the roll and the pitch angle:

$$\{roll, pitch\} = \{\alpha sin(\theta), \alpha cos(\theta)\}, \quad \theta \in [0, 2\pi) \quad (4)$$

The desired tilt-then-rotate trajectory is tracked by an admittance controller. At the same time, a constant downward force is applied to the peg in order to maintain contact with the hole. During the procedure, contact force and torque are measured by a force/torque sensor. As the peg is tilted in all directions, the contact keeps switching between one point contact, two points contact, and line contact (Fig. 4.2-4.9). The tip of the peg could go into the hole when the peg tilts towards the hole and the force/torque measurements would also have an impulse. Different contact poses will result in different sequences of measurements along the designed tilt-then-rotate trajectory. Comparing with one measurement at a single time step, the mapping from a sequence of measurements to contact poses becomes an injective mapping.

B. Contact Pattern Generation

The tilt-then-rotate strategy generates a sequence of measurements in 12 dimensions including force ($\mathbb{R}^3$), torque ($\mathbb{R}^3$), and peg pose ($\mathbb{R}^6$). For different control forces or different sizes of the parts, those measurements can be different in the order of magnitude. Human can sense the contact pose in different scenarios by the same exploring strategy. There must be some high-level features we can extract from the measurements.

We propose to plot the measurements of each dimension in polar coordinate as one channel. The data in each channel is normalized, then smoothed by moving average. The normalization makes the data invariant to control forces and sizes of the parts. The moving average reduces the sensor noises. We utilize the plotted image with 12 channels as one contact pattern, which encodes high-level features about the contact pose. Fig. 5 shows z-axis channel of the contact pattern for different contact poses.
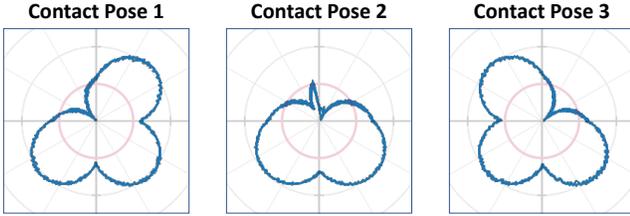
Fig. 5. Contact patterns in polar coordinates for 3 different contact poses. Only z-axis channel is shown.
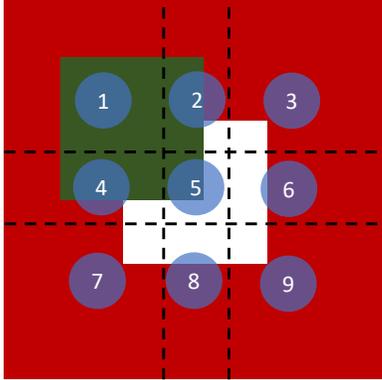

Fig. 6. The contact poses are classified into 9 classes according to which edge of the peg contacts the hole.

One contact pose corresponds to one contact pattern with 12 channels. In order to construct an informative mapping, we need to perform hundreds of tilt-then-rotate motions for all contact poses. This is not only time-consuming but also inaccurate due to the limitation of pose sensing in the real world. We propose to generate the contact pattern in the MuJoCo physics engine. The simulated environment can perform hundreds of trails in a short time. In addition, ground truth contact pose can be obtained easily in simulation (Fig. 4).

C. Contact Pose Classification Neural Network

Contact poses of a square peg-hole can be classified into 9 classes according to which edge of the peg contacts the hole (Fig. 6). Each class of contact pose has a different error direction. Classifying the contact poses from the contact patterns is an image recognition problem. CNN has shown great success in image recognition in terms of efficiency and accuracy [18]. We train one simple CNN to classify the contact patterns. The CNN has two convolutional layers, two pooling layers, and one fully-connected layer. The input data are the 3 most informative channels out of the 12-channel pattern. The output is the class of contact pose, which has 9 error directions for a square peg-hole and 11 error directions for a pentagonal peg-hole. Once the contact pose is identified, the peg will be guided towards the error direction with admittance control and inserted into the hole.

D. Failure Recovery

From the experiments, we observe failure cases even with the method described above. The reason is either the contact pose classification model predicts a wrong error direction or the admittance control fails to compensate small uncertainties. To increase the robustness of the proposed method, we add a failure recovery module. If we fail to insert the peg into the hole, the peg will be initialized to a slightly different pose than the original one, and redo the tilt-then-rotate strategy again.

IV. SIMULATIONS AND EXPERIMENTS

A. Simulations

1) Simulation Setup: The simulated environment in MuJoCo is shown in Fig. 4. The environment includes a peg and a hole, where the hole is fixed on the ground, and the peg is controlled by a well-tuned admittance controller. The side length of the hole is 50mm and the side length of the peg is 49mm (clearance = 1mm). Contact force/torque is measured at the peg's center of mass.

2) Data Collection: A self-supervised scheme is applied to collect the data and build the contact pose mapping. As mentioned in II-C, once the peg contacts the hole on the flat surface, the uncertainties in roll, pitch, and z-axis are eliminated. We only consider the remaining uncertainties in the x, y, and yaw axis. The contact poses are uniformly sampled from $x \in [-20, +20]mm$, $y \in [-20, +20]mm$, and $yaw \in [-3°, +3°]$. After the approaching stage in II-C, the tilt-then-rotate strategy is applied, and α in equation (4) is set to 15°. The tilt-then-rotate motion is executed by the admittance controller in N time steps, where $N = 2000$. The 12-dimension peg pose and contact force/torque are recorded in a matrix $A \in \mathbb{R}^{N \times 12}$. The data of each dimension is normalized then smoothed by moving average with a window length $n = 20$, and the contact pattern is recorded in polar coordinates as a $12 \times 200 \times 200$ binary image. We label the contact patterns of a square peg-hole with 9 classes according to the initial contact poses (Fig. 6). The uncertainty in the yaw axis is compensated by the admittance controller and small oscillations in the yaw axis. We also add 5% noise to the parameters of the admittance controller in order to introduce variance to the collected data. We perform the self-supervised data collection for 5000 trails. The computation time is around 10 minutes. We split 80% data as the training set and 20% data as the test set.

3) Model Training: From the 12 channels contact patterns, we select 3 channels $A' \in \mathbb{R}^{N \times 3}$ including the position in z axis X_z, the torque in roll axis M_x, and the torque in pitch axis M_y as the input to the CNN. The reason that we select these 3 channels is that we experimentally find that these channels contain more features than other channels. We downsample the contact patterns into $3 \times 20 \times 20$ images. We use an NVIDIA GeForce GTX 1080 Ti GPU for training. The training time is around 1 minute.

4) Results: The test accuracy of the contact pose classification neural network is 97.4%. Most of the failure cases are the contact pose at the boundary between two classes. We test on a second data set by collecting 1000 data from a smaller square peg-hole, where the side length of the hole is 32mm (clearance = 1mm). The test accuracy is 96.8%. This shows

# of attempts	1	2	3	> 3	total	success rate
square (50*mm*)	96	3	1	0	100	100%
pentagon (37*mm*)	82	11	3	4	100	96%

# of attempts	clearance	1	2	3	> 3	total	success rate
square (50*mm*)	1*mm*	16	3	1	0	20	100%
square (32*mm*)	0.5*mm*	10	5	2	3	20	85%
pentagon (37*mm*)	1*mm*	15	3	1	1	20	95%

the generalization ability of the proposed method. Although the sizes of the parts, the contact measurements such as force, torque are different, the model still works very well. The reason is that we predict the contact pose according to the contact pattern, which is invariant to the size of the parts.

We perform another simulation experiment on a pentagonal peg-hole. The side length of the hole is 37*mm* (clearance = 1*mm*). Because the contact pattern highly depends on the geometry of the peg-hole, we cannot apply the model learned from square peg-hole to pentagonal peg-hole. We redo the data collection and model training on the pentagonal pen-hole. Everything is the same as square peg-hole, except the number of contact pose classes becomes 11. The test accuracy is 91.0%.

We test the entire peg-in-hole assembly framework using the proposed method. We perform 100 trials on both square and pentagonal peg-hole. If the peg fails to be inserted into the hole, the failure recovery module will initialize the peg to a slightly different pose than the original one, and redo this trial again. If it requires more than 3 attempts to finish the task, we claim it fails. Table I shows the number of attempts needed to finish assembly in simulation. The high success rate shows that the proposed framework works well.

B. Experiments

1) Experimental Setup: The experiment environment (Fig. 7) includes a 6 DOF FANUC LR-Mate 200iD, an ATI Mini45 F/T sensor, and 3D printed peg-holes. The F/T sensor is embedded in the robot end-effector to measure the force and torque during assembly. The force/torque measured at the robot wrist can be transfer to the force/torque at the peg's center of mass. The peg is fixed on the robot end-effector and the hole is fixed on a vise. The peg's pose can be controlled with an admittance controller at 125*Hz*. The hole is randomly initialized with position and orientation uncertainties $\pm 20mm$ and $\pm 3°$, respectively. Three pairs of 3D printed peg-holes are tested, including a 50*mm* square hole (clearance = 1*mm*), a 32*mm* square hole (clearance = 0.5*mm*), and a 37*mm* pentagonal hole (clearance = 1*mm*).

2) Results: Fig. 8 shows the comparison of the contact patterns generated from tilt-then-rotate strategy in simulation and real-world experiments. They are generated from the same class of contact pose. The data collected from the real-world has much noise than from simulation. Although there is a huge sim-to-real gap [19] between the simulated environment and the real world in terms of friction coefficient, inertia, stiffness, damping ratio, etc., we observe that the contact patterns do share similar features.

The contact pattern classification model learned in the simulation are applied to real-world experiments. Fig. 7 shows

the snapshots of the assembly experiments. We perform 20 experiments on 3 different pairs of peg-hole respectively. Table II shows the number of attempts needed to finish assembly in real-world experiments. The model learned in simulation (50*mm* square, clearance = 1*mm*) can be successfully applied to real-world peg-hole of different sizes (32*mm*) and smaller clearance (0.5*mm*). This shows that the proposed method is able to tackle the sim-to-real gap. Supplementary videos can be found in [20].

V. DISCUSSION

In this paper, we propose a novel framework to identify contact pose for peg-in-hole assembly under uncertainties. The proposed method utilizes a tilt-then-rotate strategy to generate contact patterns. A CNN is utilized to classify the contact poses and guide the robot to achieve the assembly task with admittance control. Simulation and experiment results are provided to demonstrate the effectiveness of the proposed method. The main advantages of the proposed method include:

- The injective mapping from the contact pattern to the contact pose.
- Robustness to sensor noise.
- The contact pose classification model is easy to obtain. All the training data can be quickly generated in simulation with a self-supervised scheme.
- Good generalization ability and small sim-to-real gap. Since the contact data is normalized and recorded in a polar coordinate, the pattern is sensitive neither to the size of the object nor the parameters of the admittance controller. A model learned from a larger peg-hole can be successfully applied to smaller ones as long as the geometries are the same. Furthermore, the model learned in simulation can be adapted to the real world, despite the huge sim-to-real gap.

Here are the limitations of the proposed framework:

- The proposed method can only find the directions of the error, while it is unable to obtain the magnitude. In order to compensate for the error, the admittance controller needs to be well-tuned.
- The contact pose classification model can handle only position uncertainties, but it cannot classify the orientation uncertainties in the yaw axis.

For future works, we plan to improve the algorithm so that it can handle orientation uncertainties and test it in more challenging scenarios. We also intend to incorporate active and adaptive sensing strategies to our framework.

(a) Square peg-hole
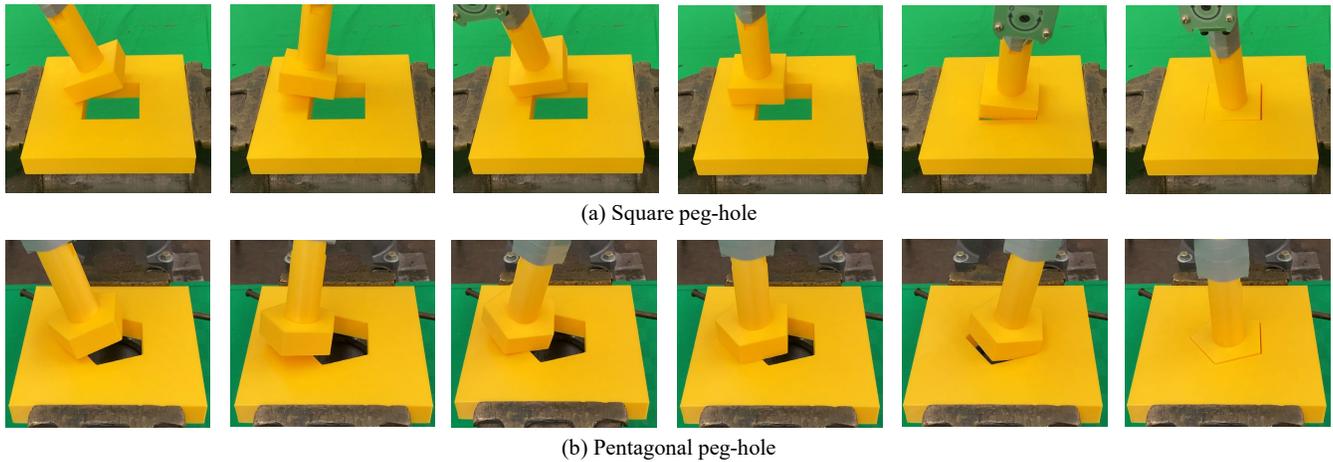

(b) Pentagonal peg-hole

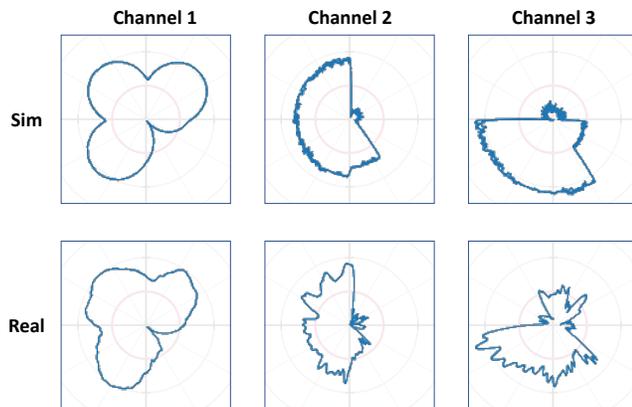Fig. 7.   Snapshots of the experiments.

Fig. 8.   Comparison of the contact patterns in simulations and real-world experiments. They are generated from the same class of contact pose

REFERENCES

[1] Y. Xiang, T. Schmidt, V. Narayanan, and D. Fox, "Posecnn: A convolutional neural network for 6d object pose estimation in cluttered scenes," in *Robotics: Science and Systems (RSS)*, 2018.

[2] B. Tekin, S. N. Sinha, and P. Fua, "Real-Time Seamless Single Shot 6D Object Pose Prediction," in *CVPR*, 2018.

[3] S. Chhatpar and M. Branicky, "Search strategies for peg-in-hole assemblies with position uncertainty," in *2001 IEEE/RSJ International Conference on Intelligent Robots and Systems*, Maui, Hawaii, USA, 2001.

[4] S. Drake, "Using compliance in lieu of sensory feedback for automatic assembly." 1978.

[5] D. E. Whitney, "Quasi-Static Assembly of Compliantly Supported Rigid Parts," *Journal of Dynamic Systems, Measurement, and Control*, vol. 104, no. 1, pp. 65–77, 03 1982.

[6] S. Bruno and O. Khatib, "Springer handbook of robotics," 2008.

[7] C. Ott, R. Mukherjee, and Y. Nakamura, "Unified impedance and admittance control," in *2010 IEEE International Conference on Robotics and Automation*, 2010, pp. 554–561.

[8] T. Tang, H. Lin, Yu Zhao, Wenjie Chen, and M. Tomizuka, "Autonomous alignment of peg and hole by force/torque measurement for robotic assembly," in *2016 IEEE International Conference on Automation Science and Engineering (CASE)*, 2016, pp. 162–167.

[9] Y. Kim, H. Song, and J. Song, "Hole detection algorithm for chamferless square peg-in-hole based on shape recognition using f/t sensor," *International Journal of Precision Engineering and Manufacturing*, vol. 15, no. 3, pp. 425–432, Mar. 2014.

[10] T. Tang, H. Lin, Y. Zhao, Y. Fan, W. Chen, and M. Tomizuka, "Teach industrial robots peg-hole-insertion by human demonstration," in *2016 IEEE International Conference on Advanced Intelligent Mechatronics (AIM)*, 2016, pp. 488–494.

[11] S. Levine, C. Finn, T. Darrell, and P. Abbeel, "End-to-end training of deep visuomotor policies," in *The Journal of Machine Learning Research*, 2016.

[12] Y. Fan, J. Luo, and M. Tomizuka, "A learning framework for high precision industrial assembly," in *2019 International Conference on Robotics and Automation (ICRA)*, 2019, pp. 811–817.

[13] J. C. Triyonoputro, W. Wan, and K. Harada, "Quickly inserting pegs into uncertain holes using multi-view images and deep network trained on synthetic data," *CoRR*, vol. abs/1902.09157, 2019.

[14] M. A. Lee, Y. Zhu, K. Srinivasan, P. Shah, S. Savarese, L. Fei-Fei, A. Garg, and J. Bohg, "Making sense of vision and touch: Self-supervised learning of multimodal representations for contact-rich tasks," in *2019 International Conference on Robotics and Automation (ICRA)*, 2019, pp. 8943–8950.

[15] S. R. Chhatpar and M. S. Branicky, "Particle filtering for localization in robotic assemblies with position uncertainty," in *2005 IEEE/RSJ International Conference on Intelligent Robots and Systems*, 2005, pp. 3610–3617.

[16] U. Thomas, S. Molkenstruck, R. Iser, and F. M. Wahl, "Multi sensor fusion in robot assembly using particle filters," in *Proceedings 2007 IEEE International Conference on Robotics and Automation*, 2007, pp. 3837–3843.

[17] L. Johannsmeier, M. Gerchow, and S. Haddadin, "A framework for robot manipulation: Skill formalism, meta learning and adaptive control," in *2019 International Conference on Robotics and Automation (ICRA)*, 2019, pp. 5844–5850.

[18] A. Krizhevsky, I. Sutskever, and G. E. Hinton, "Imagenet classification with deep convolutional neural networks," in *Advances in Neural Information Processing Systems 25*.   Curran Associates, Inc., 2012, pp. 1097–1105.

[19] J. Tobin, R. Fong, A. Ray, J. Schneider, W. Zaremba, and P. Abbeel, "Domain randomization for transferring deep neural networks from simulation to the real world," in *2017 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, 2017, pp. 23–30.

[20] Supplementary videos of the tilt-then-rotate strategy., https://shiyujin0.github.io/TiltThenRotate/ACC2021.html.