

Safe Self-Supervised Learning in Real of Visuo-Tactile Feedback Policies for Industrial Insertion

Letian Fu^{1,†}, Huang Huang¹, Lars Berscheid³, Hui Li², Ken Goldberg¹, Sachin Chitta²

Abstract—Industrial insertion tasks are often performed repetitively with parts that are subject to tight tolerances and prone to breakage. Learning an industrial insertion policy in real is challenging as the collision between the parts and the environment can cause slippage or breakage of the part. In this paper, we present a safe self-supervised method to learn a visuo-tactile insertion policy that is robust to grasp pose variations. The method reduces human input and collisions between the part and the receptacle. The method divides the insertion task into two phases. In the first *align* phase, a tactile-based grasp pose estimation model is learned to align the insertion part with the receptacle. In the second *insert* phase, a vision-based policy is learned to guide the part into the receptacle. The robot uses force-torque sensing to achieve a safe self-supervised data collection pipeline. Physical experiments on the USB insertion task from the NIST Assembly Taskboard suggest that the resulting policies can achieve 45/45 insertion successes on 45 different initial grasp poses, improving on two baselines: (1) a behavior cloning agent trained on 50 human insertion demonstrations (1/45) and (2) an online RL policy (TD3) trained in real (0/45).

I. INTRODUCTION

Industrial assembly [21] is a precise manipulation task requiring contact between parts. Part feeding, peg insertion and object reorientation (three sub-tasks of industrial assembly) have been extensively studied [7, 18, 19, 21, 24]. Early work considers the mechanical design aspect [18, 24] and motion planning aspect [7, 19, 26]. Through Computer Aided Design (CAD), the order of part assembly can be predetermined in simulation with precise pose information [3], allowing robots to plan the necessary actions to assemble the design [14]. Learning-based approaches recently have shown promise on industrial insertion tasks [34] on the NIST taskboard [12], a standard benchmark that represents common industrial insertion tasks with parts that have complex geometries [17]. However, applying learning-based methods for industrial insertion remains challenging due to the requirement for frequent human inputs during learning [20] or high-precision sensors for collecting training data [34]. There is also a need for a safe training and data collection method for learning insertion tasks since parts are prone to breakage.

Another challenge in an industrial insertion task is that the precise grasp pose is often unknown due to variations in kitting and feeding of parts as they arrive for assembly. As the grasped part is often occluded by the gripper visually, grasp pose estimation is better achieved when using tactile sensing [25]. While recent work has shown improvement

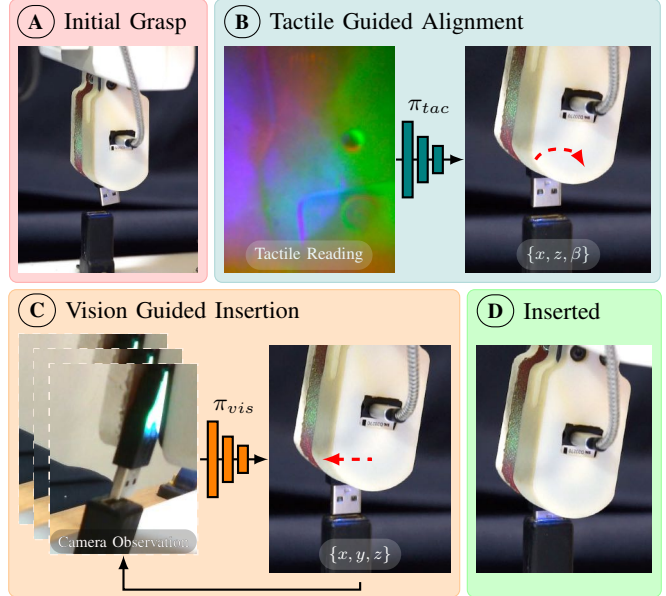


Fig. 1: Overview of the learned two-phase insertion policy: the red arrows indicating the robot actions given by the policies. (A) The robot grasps the part at an initial pose. (B) The tactile guided policy π_{tac} estimates the grasp pose using the tactile image and aligns the z-axis of the part with the insertion axis. (C) A vision guided policy π_{vis} is used to insert the part. (D) The part is inserted successfully into the receptacle.

in simulation accuracy for industrial insertion [23] and successes in Sim2Real transfer for tactile-based insertion tasks [11, 32], the simulation of soft contacts between tactile sensors and objects with complex geometries remains an open problem [33], and often can not transfer to real because the object models are not publicly available. In this work, we present a novel method to safely learn visuo-tactile feedback policies in real for industrial insertion tasks under grasp pose uncertainties, with inexpensive off-the-shelf sensors. Our approach draws on tactile and visual feedback to deal with the grasp pose uncertainty and force-torque sensing for a self-supervised training procedure that is *safe*, minimizing damage during the training phase. We divide the insertion task into two phases (as shown in Fig. 1):

- An initial *Align* phase where a tactile-based policy π_{tac} estimates the grasp pose. The robot reorients and aligns the part with the insertion axis of the receptacle.
- A second *Insert* phase where an RGB image-based policy π_{vis} guides the robot to insert the part.

A significant challenge in learning tactile feedback policies in real for industrial insertion is the frequent slippage of the part that occurs due to collision with the environment and

¹The AUTOLab at UC Berkeley, ²Autodesk Research, ³Karlsruhe Institute of Technology

[†]Work done as an intern at Autodesk

the smooth surface of the tactile sensor gel pad. This makes RL methods difficult to succeed without human intervention or an automatic reset mechanism to detect and correct slippage. In this work, we develop a self-supervised data collection pipeline that avoids collision between the part and its environment, by recognizing that the insertion operation is reversible only from certain target insertion poses – i.e. starting from such poses, the part can be repeatedly unplugged from and inserted into the receptacle. Prior to data collection, a human free-drives the robot to provide one approximate target pose where the part is inserted. The robot refines this target pose to find such a reversible pose by minimizing the grasping force-torque, which helps minimize collisions during data collection, resulting in a safer training process that is unlikely to damage the insertion part and receptacle.

This paper contributes:

- 1) A safe self-supervised data collection pipeline with force-torque sensing in real for insertion, designed to minimize contact force for data collection.
- 2) A two-phase policy learned from the collected data including a tactile-based alignment policy for orienting the part and an RGB image-based insertion policy;
- 3) Experimental results suggest that the policy achieves 45/45 successes on USB connector insertion, outperforming two baseline methods (1/45 and 0/45).

II. RELATED WORK

Industrial insertion has been central in robotics for 50 years. It is challenging due to occlusions brought by the robot gripper, grasp uncertainty from the process of acquiring the part and its collision with the environment, the fragility of the parts, and the precision required in controlling the robot for insertion. Early work approached this problem using CAD information to infer desired assembly sequences [3] and generating designs of part feeders based on object geometry [24]. Other work approached the problem from an algorithmic design perspective, with a focus on developing motion planning strategies for peg insertion [18, 26].

Recently, learning-based methods have shown success on this task. This includes learning insertion policies with a physical robot via Sim2Real transfer [10], online adaptation with meta-learning [29, 36], reinforcement learning [20, 28], self-supervised data collection with impedance control [30], accurate state estimation [34], or decomposing the insertion algorithm into a residual policy that relies on conventional feedback control [10]. These approaches assume that the parts are grasped with a fixed pose. To overcome this assumption, Wen *et al.* [34] perform accurate pose estimation and motion tracking with a high-precision depth camera and use a behavioral cloning algorithm to insert the part. Spector *et al.* [30, 31] require contact between the part and the environment to occur during data collection, a process that is expensive and often impractical for fragile parts. In comparison, we use inexpensive tactile sensors and a safe self-supervised data collection procedure that does not require such contact.

Grasped parts are often visually occluded by the gripper. Tactile feedback can be an alternative sensing modality

for grasp pose estimation. Recent work uses tactile images from vision-based tactile sensors such as GelSight [35] and DIGIT [15] to estimate the pose of grasped objects. Li *et al.* [16] use Gelsight sensors, BRISK features and RANSAC to estimate grasp pose. Gelsight produces high-quality 3D tactile images and can determine depth imprint, which improves feature detection by isolating the object from the background. DIGIT, a more affordable tactile sensor, provides a 2D RGB image but not the light incident direction (to generate the depth image). Kelestemur *et al.* [11] generates tactile image data in simulation for pose estimation of bottle caps but simulating contact and physical interaction between tactile sensors and objects with more intricate geometry is still challenging [33]. In this work, we collect a dataset of tactile images in real for the USB connector with different grasp poses to train a tactile based policy for grasp pose estimation.

Most prior work on tight tolerance insertion tasks [4, 5, 16, 34] leverages a single modality, such as vision, tactile, or force-torque, limiting the accuracy of the system due to occlusion, perspective effect, and sensory inaccuracy. Multi-modal systems have been explored to improve the robustness of automated insertion. Spector *et al.* [30, 31] use RGB cameras and a force-torque sensor for learning contact and impedance control. Chaudhury *et al.* [2] couple vision and tactile data to perform localization and pose estimation, and demonstrate that vision helps with disambiguating tactile signals for objects without distinctive features. Ichiwara *et al.* [9] leverage tactile and vision for deformable bag manipulation by performing auto-regressive prediction. Hansen *et al.* [8] use a contact-gated tactile, vision and proprioceptive observation to train reinforcement learning policies. Okumura *et al.* [25] also tackle the problem of grasp pose uncertainty for insertion by using Newtonian Variational Autoencoders to combine camera observations and tactile images. They demonstrate results for USB insertion accounting for grasp pose uncertainty in one translation direction. In this work, we separate the insertion problem into an alignment phase and an insertion phase, decoupling vision and tactile inputs and also present a novel safe self-supervised approach to data collection. We are able to handle both grasp pose rotation and translation uncertainty for the USB insertion task.

III. PROBLEM STATEMENT

Overview: We consider a part insertion task using a 7-DoF robot, equipped with a parallel-jaw gripper with a tactile sensor mounted on one jaw. The end-effector has a wrist-mounted RGB camera, and the robot provides reliable force-torque readings at the end-effector. The objective is to learn a policy that can robustly insert the part into the receptacle with an unknown part’s pose within the gripper, while minimizing human inputs and part-receptacle collisions during training. Fig. 2 shows the experiment setup and the coordinate frames.

Assumptions: We make the following assumptions:

- 1) A human provides one top-down grasp pose of the part inserted in the receptacle.
- 2) The robot can accurately measure force and torque either with an external sensor or internal current sensing;




Fig. 2: Experiment setup and coordinate system. The x , y , z axes are labeled by red, green, blue respectively. We label the gripper frame, part frame, human-provided target pose frame, and robot frame as F_G , F_p , F_h , F_R respectively. The insertion direction is defined as the z -axis of F_h . When the part is inserted, $F_h = F_p$.

- 3) An experiment begins with the part pre-grasped by the robot gripper with the part grasp pose within a range relative to the human-provided pose.
- 4) During data collection, the robot operates in a rectangular collision-free configuration space above the receptacle.

Problem Setup: Given a tight-fitting receptacle for the part to be inserted into (Fig. 2), we find the target insertion pose $T_{R,h}$ of the part (grasped at an unknown pose) from one human-provided imperfect demonstration. At any time step t , we have access to the RGB observation $o_\rho(t)$ from the wrist mounted camera, the RGB tactile image $o_\psi(t)$ from the DIGIT tactile sensor, and reliable force $\vec{f}(t)$ and torque $\tau(t)$ readings from the robot. Since we know the insertion axis, we parameterize the action space with 4 degrees of freedom: gripper translation in robot frame and gripper y -axis rotation.


IV. METHOD

A. Hardware

We design a novel parallel jaw gripper mount to accommodate the DIGIT tactile sensor [15] and camera mount (Fig. 2, 3). The elastomer gel on the DIGIT sensor deforms, causing torque applied to the part. This torque and the force applied from the receptacle to the part during insertion often produce undesired slippage and rotation. Therefore, we develop an asymmetric mounting setup where we mount the DIGIT sensor on one jaw with reinforcement to prevent outward bending, while keeping the other jaw flat. We apply sandpaper on the surface of the non-tactile gripper, increasing the friction to reduce slippage. We find that we can predict the part's grasp pose of a USB connector using a single DIGIT sensor.

B. Self-Supervised Data Collection

1) *One Human Provided Imperfect Target Pose:* The target pose $T_{R,h}$ is provided by a human free-driving the robot with a pre-grasped part to insert it in the receptacle from top down. Since this target pose may not have a perfect axis alignment with the receptacle, the system performs a z -axis alignment of the target pose. To account for the change in grasp center after axis alignment, we refine $T_{R,h}$ by finding a target pose that minimizes gripper force-torque using a



(a) Parallel Jaw Gripper

(b) Camera Mount

Fig. 3: CAD models for the parallel jaw grippers and camera mount. grid search through a set of translations and rotations $\{T_\Delta\}$. Formally, we find

$$\tilde{T}_{R,h} = \arg \min_{\tilde{T}_{R,h} \in \{T_\Delta \cdot T_{R,h}\}} \left\| \vec{f}(\tilde{T}_{R,h}) \right\| + \left\| \tau(\tilde{T}_{R,h}) \right\|. \quad (\text{IV.1})$$


Here $\vec{f}(\tilde{T}_{R,h})$ and $\tau(\tilde{T}_{R,h})$ denote the 3-DoF force and torque vectors respectively when the gripper is at $\tilde{T}_{R,h}$. Intuitively, this objective minimizes the external force applied on the part when being unplugged, increasing the likelihood of the insertion process being reversible. Practically, we perform grid sampling over 5 values of $x \in [-1, 1]$ mm, 5 values of $y \in [-1, 1]$ mm and 4 values of $\gamma \in [-\frac{\pi}{180}, \frac{\pi}{180}]$ rad (x, y, γ are in F_h , refer to Fig. 2). The pose with minimum gripper external force-torque is recorded as the refined demonstration pose $\tilde{T}_{R,h}$, and we have $F_h = F_G$ at $\tilde{T}_{R,h}$.

A cascaded impedance controller, implemented within the robot's real-time control loop, allows fine-grained force control. In case of a force violation, our system calculates a trajectory to a safe state within a single control cycle. After refinement of the target pose, we search for the minimum offset z_{\min} for the part to be unplugged from the receptacle. Finding the minimum height helps to determine the boundary for data collection and allows the pipeline to collect more data closer to the receptacle while reducing collisions. Iteratively, the robot moves the gripper by $-\Delta_z$ in F_G . We then move the gripper by Δ_x (in the F_G frame) and measure \vec{f}_x (x -component of the gripper force in the F_G frame). If $\vec{f}_x \leq \eta$, we register the total upward distance traveled as the minimum height z_{\min} for removal of the part. Empirically, we find setting $\Delta_x = \Delta_z = 1$ mm, and $\eta = 3.5$ N works well.

2) *Data Collection for Alignment:* The part remains inserted throughout data collection. We explore grasp pose variations in 3-DoF (x , z translation and y -axis rotation β) in F_G (Fig. 2). We perform uniform random sampling over the range $[-3, 3]$ mm, $[-8, -2]$ mm, $[-\frac{\pi}{15}, \frac{\pi}{15}]$ rad for x, z, β , with 5, 10 and 20 samples respectively. The robot closes the gripper with a force of 70N at each of the sampled poses and records the pair of tactile image readings and x, z, β . To account for the noise in the DIGIT tactile sensor, we take a median filter over 5 consecutively captured tactile images. We collect 2000 data points in 120 minutes.

3) *Data Collection for Insertion:* Upon completing the tactile image collection phase, we collect robot poses and RGB images for training the insertion policy for different grasps. We perform grid sampling with 5 samples each of x, z, β in F_G within the same range as in the previous stage,

Tactile Guided Alignment Policy Data Collection



Vision Guided Insertion Policy Data Collection




Fig. 4: Data collection for alignment (Top) and insertion policies (Bottom). Data collection for the **Alignment policy** starts with the part inserted into the receptacle (Top (A)). The robot then samples and records different grasp poses and the corresponding tactile images (Top (B)). Data collection for the **Insertion policy** starts with a sampled refined grasp (Bottom (A)) and unplugs the part to apply sampled transformations (Bottom (B)). Then the robot inserts the part (Bottom (C)) and starts the next round of data collection with a different grasp pose.

resulting in a total of 125 grasps. To account for the difference between the sampled grasp g and $\tilde{T}_{R,h}$, we perform minimum force-torque refinement on the sampled grasp to calculate the grasp-specific target pose $\tilde{T}_{R,h}(g)$. $\tilde{T}_{R,h}(g)$ translated by an offset of z_{\min} gives us the unplugged part pose $T_{\text{unplug}}(g)$.

For each grasp g , we collect image data for the visuoservo policy by moving the gripper to sampled points on a grid above the target pose. In particular, we uniformly sample 5 values each of $x \in [-5, 5]\text{mm}$, $y \in [-5, 5]\text{mm}$ and $z \in [-5, 0]\text{mm}$ in F_R with respect to $T_{\text{unplug}}(g)$, resulting in 125 different translations for the gripper. For each translation, we collect *one* data point that does not contain additional rotation and sample *two* gripper y -axis rotation conditioned on z to avoid collision. Specifically, given a height z , the two rotations are sampled from the uniform distribution $\frac{z}{5} \cdot \mathcal{U}[-\frac{\pi}{15}, 0]\text{rad}$ and $\frac{z}{5} \cdot \mathcal{U}[0, \frac{\pi}{15}]\text{rad}$. These rotated data points provide the system with additional data for camera pose variation with respect to the target pose, which leads to a balanced dataset with 375 distinct gripper poses. Each data point is composed of the gripper pose (translated and rotated away from the target pose) and the corresponding RGB image observation at that pose. Upon visiting all 375 gripper poses for a given grasp, the robot moves to $T_{\text{unplug}}(g)$, performs a vertical movement to $\tilde{T}_{R,h}(g)$, and opens the gripper jaws, thereby resetting the part in the receptacle. The system repeats this data collection process for all 125 grasps without human supervision.

C. Learning to Insert

While human demonstrations usually serve as “expert policies” for industrial insertion tasks, the self-supervised data collection pipeline allows us to collect ground truth

Algorithm 1: Policy Execution Procedure

```

1 Input: Tactile Image  $o_\rho(t)$ , Camera Image  $o_\psi(t)$ ,
   tactile based grasp pose estimation network  $\pi_{tac}$ , and
   visuoservo insertion policy  $\pi_{vis}$ , Target Pose  $\tilde{T}_{R,h}$ ,
   Minimum Wrench Height  $z_{\min}$ , action norm
   threshold  $\epsilon$ ,  $z$  direction step size  $d_z$ 
2  $T_{p,G} = \pi_{tac}(o_\rho(t))$ 
3 Move gripper to  $\tilde{T}_{R,h}T_{p,G} + 2z_{\min}$ 
4 attempts = 0
5  $(\Delta_x, \Delta_y, \Delta_\beta) = \pi_{vis}(o_\psi(t), T_{R,G}(t))$ 
6 while True do
7   while  $\|[\Delta_x, \Delta_y]\| > \epsilon$  and attempts <  $H$  do
8     Move gripper by  $(\Delta_x, \Delta_y)$ 
9     attempts = attempts + 1
10     $(\Delta_x, \Delta_y, \Delta_\beta) = \pi_{vis}(o_\psi(t), T_{R,G}(t))$ 
11    Translate gripper in insertion direction by  $d_z$ 
12    Measure gripper force-torque in  $z$ -axis as  $F_z$ 
13    if  $F_z > 15N$  or attempts  $\geq H$  then
14      Terminate

```

actions at scale. This allows us to formulate two supervised-learning problems based on Sec. IV-B.2 and Sec. IV-B.3.

1) *Alignment Policy:* We use the data collected from Sec. IV-B.2 to train an alignment policy π_{tac} that, given the tactile image, outputs the desired displacement of the gripper $T_{p,G}$ to align the part with the receptacle (Fig. 1.B). We augment the tactile images by randomly jittering the brightness and contrast over the range $\mathcal{U}[0.7, 1.3]$.

2) *Insertion Policy:* We use data collected from Sec. IV-B.3 to train a visuoservo insertion policy π_{vis} taking normalized camera observation, gripper y -axis rotation β and x, y translations in F_h as inputs, and predict the action: desired translation Δ_x, Δ_y and rotation Δ_β in F_R (Fig. 1.C). We augment camera observations by randomly jittering the brightness and contrast over the range $\mathcal{U}[0.7, 1.3]$.

We use RegNet 3.2GF [27] as the backbone for both policies. For the alignment policy, we replace the last layer of RegNet with a linear layer with 3 outputs. For the insertion policy, we concatenate the robot’s pose with the latent vector of the image and replace the last layer with a linear layer with 3 outputs. For both networks, we use a batch size of 64, a learning rate of 1e-3, and a learning rate decay of 0.99 for every 100 gradient steps. We pick the mean squared error as the loss function and use the Adam optimizer [13].

D. Execution of Insertion

To avoid catastrophic failure (i.e. collision between the part and the surrounding environment or moving out of the training set distribution), we deliberately formulate the visuoservo policy to only control x, y -axis translation but not z -axis translation since the insertion direction is already aligned with the z -axis by the *Align* policy. The formulation is detailed in Algorithm 1. The execution procedure starts by inferring the grasp pose from the tactile image via π_{tac} (line 2). The system then performs the insertion axis alignment of the part with the receptacle (line 3). We measure the z -direction force based on the gripper force-torque sensor (line 12). We then calculate

Algorithms	IL	TD3	Proposed Approach
Success/Total	1/45	0/45	45/45

TABLE I: Results suggest that (1) the IL trained on 50 human demonstrations is insufficient for training an accurate part pose estimation model, and (2) frequent slippage and rotations of the USB caused by collisions with the receptacle lead to failure in training TD3. Our approach outperforms both baseline policies.

rotation and translation based on the camera observation via π_{vis} (line 5, 10) and execute the corresponding actions (line 8) until the action has a norm smaller than ϵ (line 7). Note that the rotation prediction from the *Insertion Policy* is not used, because the part is aligned with the insertion axis by the *Alignment Policy*. When the action converges, we lower the gripper in the z direction by a step size of d_z (line 11) and continue to query π_{vis} until the force constraint is satisfied or the number of attempts exceeds the horizon H (line 13-14). We empirically set $d_z = 1.5\text{mm}$. Empirically, we find setting the action norm $\epsilon = 0.0005$, $d_z = 1.5\text{mm}$ and $H = 200$ works well, as the force-based termination condition is usually triggered first for a successful or unsuccessful insertion.

V. EXPERIMENTS

A. Experiment Setup

We focus on the USB insertion task on the NIST taskboard. We use a 7-DoF Franka Emika robot with a parallel gripper, where one DIGIT tactile sensor is mounted on the inside of one of the fingers. An Intel RealSense camera is mounted offset from the gripper (Fig. 3). To control the Cartesian position of the Franka robot, a time-optimal trajectory respecting velocity, acceleration, and jerk constraints is applied to the policy’s positional output [1]. We use grid sampling to obtain 5 values of β ranging from $[-\frac{\pi}{20}, \frac{\pi}{20}]$ rad, 3 translations in x and z from the range $[-3, 3]\text{mm}$ and $[-6, -2]\text{mm}$ in F_h , resulting in a test set of 45 different grasp configurations that lie in the training distribution of the algorithms.

B. Experimental Procedure

At the beginning of each test experiment, the USB connector (part) is pre-grasped by the robot with a grasp pose selected from the test set and located at a position with a z -axis translation of $2z_{\min}$ relative to F_h . At this starting pose, the gripper is aligned vertically down while the USB connector is misaligned with the receptacle in both translation and rotation as in Fig. 1(A). The robot first executes the *Align* policy to estimate the part grasp pose and aligns it with the insertion direction of the receptacle as in Fig. 1(B). It then uses the *Insert* policy to visuoservo and inserts the part into the receptacle as in Fig. 1(C-D). The robot then resets to the next grasp by releasing the part, re-grasping it, raising it to a start pose as outlined above, and executing the *Align* and *Insert* policies for this new grasp. It steps through all the test grasp poses using the same procedure.

An experiment terminates if the gripper frame (F_G) force in the z direction $\vec{f}_z(t)$ exceeds 15 N. An experiment trial is considered successful if the gripper is within 5 mm of the target pose in $H = 200$ iterations, upon which we also visually inspect whether the insertion is successful. In this set of experiments, the refined human-provided target pose

	x (mm)	z (mm)	β (rad)
Mean Error	8.97e-2	1.46e-1	5.59e-3
Standard Deviation	4.89e-3	6.62e-2	4.89e-3

TABLE II: Mean and standard deviation of the error in predicting part pose (x, z, β) by the tactile-based alignment policy on the test set of 45 grasps.

$\tilde{T}_{R,h}$ is provided as an input to the policies. In Sec. V-E.2, we perform ablation studies on noisy target poses.

C. Comparison

The method is designed with two objectives: (1) minimizing human intervention so that ideally no human needs to be involved in data collection or training of the policy, and (2) minimizing collision among the robot, the part, and the environment. We compare our approach with two baseline learning methods described below with the same environment setup (using the same grippers and camera mount).

1) *Twin Delayed Deep Deterministic Policy Gradient(TD3)*: An off-policy, online reinforcement learning policy [6] that learns the end-to-end part insertion. This baseline satisfies objective (1) but violates objective (2) — i.e. it is incapable of avoiding collisions among the robot, the part, and the environment. We simplify the problem to a fixed, axis-aligned grasp pose and restrict the action space to translations only, so that the policy only have to learn insertion instead of both alignment and insertion. The learned policy runs with a frequency of 10 Hz. Since the policy outputs Cartesian position changes, we controls the Cartesian velocity of the robot, which is equivalent to Cartesian position changes per time step. Due to the part’s axis-alignment with the receptacle, the policy’s input can be restricted to the (low-dimensional) robot pose, velocity, and the force-torque at frame F_G . We use the default TD3 implementation of Ray RLLib [22]. If the force applied on the gripper exceeds 15 N, the episode terminates and the robot resets to a safe starting pose.

2) *Imitation learning (IL)*: Imitation learning from 50 human demonstrations of insertion trained with behavior cloning. Each human demonstration starts with a randomly sampled grasp pose as in Sec. IV-B.2. This baseline algorithm violates objective (1) but satisfies objective (2), where the human demonstrator selects actions that minimize collision between the part and the environment. It takes about 30 minutes to provide all these human demonstrations.

D. Results

The results are summarized in Table. I. The imitation learning agent (IL) is only able to perform a single successful insertion out of the 45 grasp poses. Intuitively, 50 different grasp configurations from human demonstrations are not sufficient for training an accurate part pose estimation model; additionally, for different grasp poses, at the same gripper pose, two different human demonstrations may exist. The multi-modality in the distribution of target insertion poses contributes to the failure of the IL policy. Training the TD3 policy in the physical environment led to divergent results in all 5 training trials we attempted. In all cases, the part collides with the receptacle, leading to a drastic change in the grasp pose. This cannot be corrected directly since there is no

Setup	Trial 1	Trial 2	Trial 3	Mean±Standard Error
ZA	0/125	0/125	0/125	0.0±0.0%
ZAWF	125/125	125/125	34/125	75.7±19.8%
ZAWFG	125/125	125/125	125/125	100.0±0.0%

TABLE III: Comparing data collection success rate. We measure the number of successful insertions until failure for 125 different grasps configurations. We compare Human Demonstration with axis alignment (**ZA**), Single Minimum Force-Torque Refinement (**ZAWF**), and Minimum Force-Torque Refinement for all grasps (**ZAWFG**). We report the mean success rate and the standard error for three distinct human-provided target poses.

reset procedure that can systematically recover the gripper to the original state without human supervision. Our approach succeeds for every single grasp pose tested. Empirically, we find that the part rotation and translation predicted from tactile images are fairly accurate (refer to Table. II).

E. Ablation Studies

1) *Effects of Leveraging Force-Torque Sensing in Data Collection:* We compare the completion rate of the data collection process for insertion with or without grasp pose refinement. We consider the following three different methods for refining the target pose gained from the human demonstration $T_{R,h}$: 1) **ZA**: apply z -axis alignment on $T_{R,h}$, 2) **ZAWF**: perform minimum force-torque refinement only once for $T_{R,h}$ after z -axis alignment (this step is only performed for the first grasp and the results reused for all grasps) and 3) **ZAWFG**: perform z -axis alignment for $T_{R,h}$ and apply minimum force-torque refinement separately for each grasp.

At the beginning of each experiment, a human provides $T_{R,h}$ by free-driving the robot with one pre-grasped part to insert the part. A total of 125 different grasp poses are sampled. For each grasp pose $T_{h,G}$, we calculate the grasp pose in robot frame by $T_{R,G} = \tilde{T}_{R,h}T_{h,G}$ with $\tilde{T}_{R,h}$ determined by one of the three methods **ZA**, **ZAWF** or **ZAWFG**. The robot grasps the part with the pose $T_{R,G}$, lifts the part, and tries to re-insert the part. If insertion is successful, the robot executes the next grasp otherwise the experiment terminates. We report total number of successful insertions before termination (Table. III). We repeat the experiment three times for each method with different human demonstrations.

After applying z -axis alignment for the human-provided target pose (**ZA**), the insertion fails as the center of the grasp is not aligned with the center of the receptacle. **ZAWF** addresses this issue by using minimum pose refinement, and can perform successful insertions. However, since the pose refinement is specific to the human demonstration, the refinement is not sufficiently granular, leading to failures when the new grasp configuration has a large y -axis rotation. **ZAWFG** performs pose refinement for each of the grasps, resulting in consistent insertion performance. **ZAWFG** needs to wait for the force measurements to settle and thus takes longer to execute.

2) *Exploring Utility of Tactile and Vision Information:* We perform study the relative benefits of using tactile and vision for insertion tasks. We test 3 different approaches: (1) A Tactile Only approach (2) A Vision Only approach trained using a limited amount of the camera observation data and (3) a Combined Approach. This ablation study differs from our earlier experiments by injecting a uniformly sampled noise

Algorithm	Tactile Only	Vision Only (No Rot)	Tactile + Vision (No Rot)
Success/Total	21/45	8/45	40/45

TABLE IV: Ablation study with noisy target poses comparing single-phase Tactile Only, modified Vision Only, and a Combined two-phase approach leveraging tactile and visual information.

in the range $\pm 1\text{mm}$ into the target pose’s x, y translation to imitate imprecise knowledge of the target pose.

The Tactile Only approach attempts the entire insertion task in a single phase using the tactile information to align the USB connector with the receptacle and then move straight down to insert it. For the purpose of this study, we train a new Vision Only model using a third of the collected camera observation data set that have no additional gripper rotation. This mimics a mono-view visuoservo model for receptacle localization and top-down insertion. We modify the insertion motion from Sec. IV-C.2 so the model only uses camera observation and the y -axis rotation of the gripper, and it only outputs translation in x and y based on the same regression objective as in Alg. IV-C.2. The combined approach sequences a Tactile Only approach in an *Align* phase with the modified Vision Only approach in an *Insert* phase. We perform experiments with the three different approaches with the same test set as in Sec. V and report results in Table IV.

With noisy target pose, the Tactile Only model succeeds only 21/45 times since the model does not estimate the target receptacle state. Since the Vision Only model is constrained to translation actions and is trained on a limited set of data that does not include additional gripper rotation, it inserts the part when the grasps do not have any rotation (8/9 successes) but fails otherwise (for a total of 8/45 successes). A separate Vision Only model trained with all the data is similarly unsuccessful (11/45 successes), indicating the importance of the ability to correct for grasp pose variation using tactile data. The combination of the two models outperforms either model by leveraging the part rotation prediction (using tactile) and implicitly estimating the environment state (using visual information), suggesting that tactile and vision observation jointly reduce the uncertainties in the insertion problem.

VI. LIMITATIONS AND CONCLUSION

Despite promising results, this method has not been tested for generalization to other types of assembly tasks, objects with more complex geometries, or objects that are larger than the tactile sensor. Parts made of different materials may require distinct maximal forces; the grid search for finetuning the insertion pose lengthens the data collection process. The robot must also unplug the part, which can pose a challenge as some parts are designed to be difficult to remove (i.e. an Ethernet connector). This work did not measure the time required for data collection. Future research can improve on the time required for collecting data. In summary, we present a safe, self-supervised method for learning a visuo-tactile insertion policy in real industrial settings with unknown grasp poses. We achieve this by using force-torque sensing to refine human-demonstrated target poses and constructing a two-phase approach to insertion that separates the task into alignment and insertion based on tactile and visual feedback.

REFERENCES

- [1] L. Berscheid and T. Kröger, “Jerk-limited real-time trajectory generation with arbitrary target states,” *Robotics: Science and Systems XVII*, 2021.
- [2] A. N. Chaudhury, T. Man, W. Yuan, and C. G. Atkeson, “Using collocated vision and tactile sensors for visual servoing and localization,” *IEEE Robotics and Automation Letters*, vol. 7, no. 2, pp. 3427–3434, 2022.
- [3] L. H. De Mello and A. C. Sanderson, “A correct and complete algorithm for the generation of mechanical assembly sequences,” in *1989 IEEE International Conference on Robotics and Automation*, IEEE Computer Society, 1989, pp. 56–57.
- [4] Y. Fan, J. Luo, and M. Tomizuka, “A learning framework for high precision industrial assembly,” *2019 International Conference on Robotics and Automation (ICRA)*, pp. 811–817, 2019.
- [5] P. Florence, C. Lynch, A. Zeng, O. Ramirez, A. Wahid, L. Downs, A. Wong, J. Lee, I. Mordatch, and J. Tompson, “Implicit behavioral cloning,” *Conf. on Robot Learning (CoRL)*, 2021.
- [6] S. Fujimoto, H. Hoof, and D. Meger, “Addressing function approximation error in actor-critic methods,” in *International conference on machine learning*, PMLR, 2018, pp. 1587–1596.
- [7] K. Y. Goldberg, “Orienting polygonal parts without sensors,” *Algorithmica*, vol. 10, no. 2, pp. 201–225, 1993.
- [8] J. Hansen, F. Hogan, D. Rivkin, D. Meger, M. Jenkin, and G. Dudek, “Visuotactile-rl: Learning multimodal manipulation policies with deep reinforcement learning,” in *2022 International Conference on Robotics and Automation (ICRA)*, IEEE, 2022, pp. 8298–8304.
- [9] H. Ichiwara, H. Ito, K. Yamamoto, H. Mori, and T. Ogata, “Contact-rich manipulation of a flexible object based on deep predictive learning using vision and tactility,” in *2022 International Conference on Robotics and Automation (ICRA)*, IEEE, 2022, pp. 5375–5381.
- [10] T. Johannink, S. Bahl, A. Nair, J. Luo, A. Kumar, M. Loskyll, J. A. Ojea, E. Solowjow, and S. Levine, “Residual reinforcement learning for robot control,” in *2019 International Conference on Robotics and Automation (ICRA)*, IEEE, 2019, pp. 6023–6029.
- [11] T. Kelestemur, R. Platt, and T. Padir, “Tactile pose estimation and policy learning for unknown object manipulation,” *Int. Conf. on Autonomous Agents and Multiagent Systems (AAMAS)*, 2022.
- [12] K. Kimble, K. V. Wyk, J. Falco, E. Messina, Y. Sun, M. Shibata, W. Uemura, and Y. Yokokohji, “Benchmarking protocols for evaluating small parts robotic assembly systems,” vol. 5(2), 883–889, 2020.
- [13] D. P. Kingma and J. Ba, “Adam: A method for stochastic optimization,” 2015.
- [14] Y. Koga, H. Kerrick, and S. Chitta, “On cad informed adaptive robotic assembly,” *arXiv preprint arXiv:2208.01773*, 2022.
- [15] M. Lambeta, P.-W. Chou, S. Tian, B. Yang, B. Maloon, V. R. Most, D. Stroud, R. Santos, A. Byagowi, G. Kammerer, *et al.*, “Digit: A novel design for a low-cost compact high-resolution tactile sensor with application to in-hand manipulation,” *IEEE Robotics and Automation Letters*, 2020.
- [16] R. Li, R. W. Platt, W. Yuan, A. ten Pas, N. Roscup, M. A. Srinivasan, and E. H. Adelson, “Localization and manipulation of small parts using gelsight tactile sensing,” *Proc. IEEE/RSJ Int. Conf. on Intelligent Robots and Systems (IROS)*, 2014.
- [17] W. Lian, T. Kelch, D. Holz, A. Norton, and S. Schaal, “Benchmarking off-the-shelf solutions to robotic assembly tasks,” *Proc. IEEE/RSJ Int. Conf. on Intelligent Robots and Systems (IROS)*, 2021.
- [18] T. Lozano-Pérez, “Motion planning and the design of orienting devices for vibratory part feeders,” in *IEEE Journal Of Robotics And Automation*. MIT AI Laboratory, 1986.
- [19] T. Lozano-Perez, M. T. Mason, and R. H. Taylor, “Automatic synthesis of fine-motion strategies for robots,” *The International Journal of Robotics Research*, vol. 3, no. 1, pp. 3–24, 1984.
- [20] J. Luo, O. Sushkov, R. Pevceviciute, W. Lian, C. Su, M. Vecerik, N. Ye, S. Schaal, and J. Scholz, “Robust multi-modal policies for industrial assembly via reinforcement learning and demonstrations: A large-scale study,” *Robotics: Science and Systems (RSS)*, 2021.
- [21] K. McKee, “Automatic assembly by g. boothroyd, c. poli and le murch, marcel dekker, new york, 378 pp., 1982,” *Robotica*, vol. 3, no. 3, pp. 195–196, 1985.
- [22] P. Moritz, R. Nishihara, S. Wang, A. Tumanov, R. Liaw, E. Liang, M. Elibol, Z. Yang, W. Paul, M. I. Jordan, *et al.*, “Ray: A distributed framework for emerging ai applications,” in *13th USENIX Symposium on Operating Systems Design and Implementation (OSDI 18)*, 2018, pp. 561–577.
- [23] Y. Narang, K. Storey, I. Akinola, M. Macklin, P. Reist, L. Wawrzyniak, Y. Guo, A. Moravanszky, G. State, M. Lu, *et al.*, “Factory: Fast contact for robotic assembly,” *Robotics: Science and Systems (RSS)*, 2022.
- [24] B. K. Natarajan, “Some paradigms for the automated design of parts feeders,” *The International journal of robotics research*, vol. 8, no. 6, pp. 98–109, 1989.
- [25] R. Okumura, N. Nishio, and T. Taniguchi, “Tactile-sensitive newtonianvae for high-accuracy industrial connector-socket insertion,” *arXiv preprint arXiv:2203.05955*, 2022.
- [26] H. Qiao, B. Dalay, and R. Parkin, “Fine motion strategies for robotic peg-hole insertion,” *Proceedings of the Institution of Mechanical Engineers, Part C: Journal of Mechanical Engineering Science*, vol. 209, no. 6, pp. 429–448, 1995.
- [27] I. Radosavovic, R. P. Kosaraju, R. Girshick, K. He, and P. Dollár, “Designing network design spaces,” in *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, 2020, pp. 10 428–10 436.
- [28] G. Schoettler, A. Nair, J. Luo, S. Bahl, J. A. Ojea, E. Solowjow, and S. Levine, “Deep reinforcement learning for industrial insertion tasks with visual inputs and natural rewards,” in *2020 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, IEEE, 2020, pp. 5548–5555.
- [29] G. Schoettler, A. Nair, J. A. Ojea, S. Levine, and E. Solowjow, “Meta-reinforcement learning for robotic industrial insertion tasks,” in *2020 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, IEEE, 2020, pp. 9728–9735.
- [30] O. Spector and D. D. Castro, “Insertionnet - A scalable solution for insertion,” *IEEE Robotics and Automation Letters*, 2021.
- [31] O. Spector, V. Tchuiev, and D. D. Castro, “Insertionnet 2.0: Minimal contact multi-step insertion using multimodal multiview sensory input,” *Proc. IEEE Int. Conf. Robotics and Automation (ICRA)*, 2022.
- [32] S. Wang, M. Lambeta, P.-W. Chou, and R. Calandra, “TACTO: A fast, flexible, and open-source simulator for high-resolution vision-based tactile sensors,” *IEEE Robotics and Automation Letters (RA-L)*, vol. 7, no. 2, pp. 3930–3937, 2022.
- [33] Y. Wang, W. Huang, B. Fang, F. Sun, and C. Li, “Elastic tactile simulation towards tactile-visual perception,” in *Proceedings of the 29th ACM International Conference on Multimedia*, 2021, pp. 2690–2698.
- [34] B. Wen, W. Lian, K. E. Bekris, and S. Schaal, “You only demonstrate once: Category-level manipulation from single visual demonstration,” *Robotics: Science and Systems (RSS)*, 2022.
- [35] W. Yuan, S. Dong, and E. H. Adelson, “Gelsight: Highresolution robot tactile sensors for estimating geometry and force,” *Sensors* 17, 2017.
- [36] T. Z. Zhao, J. Luo, O. Sushkov, R. Pevceviciute, N. Heess, J. Scholz, S. Schaal, and S. Levine, “Offline meta-reinforcement learning for industrial insertion,” in *2022 International Conference on Robotics and Automation (ICRA)*, IEEE, 2022, pp. 6386–6393.