

Tactile-Filter: Interactive Tactile Perception for Part Mating

Kei Ota
 Mitsubishi Electric

Devesh K. Jha
 Mitsubishi Electric Research Labs

Hsiao-Yu Tung
 MIT

Joshua B. Tenenbaum
 MIT

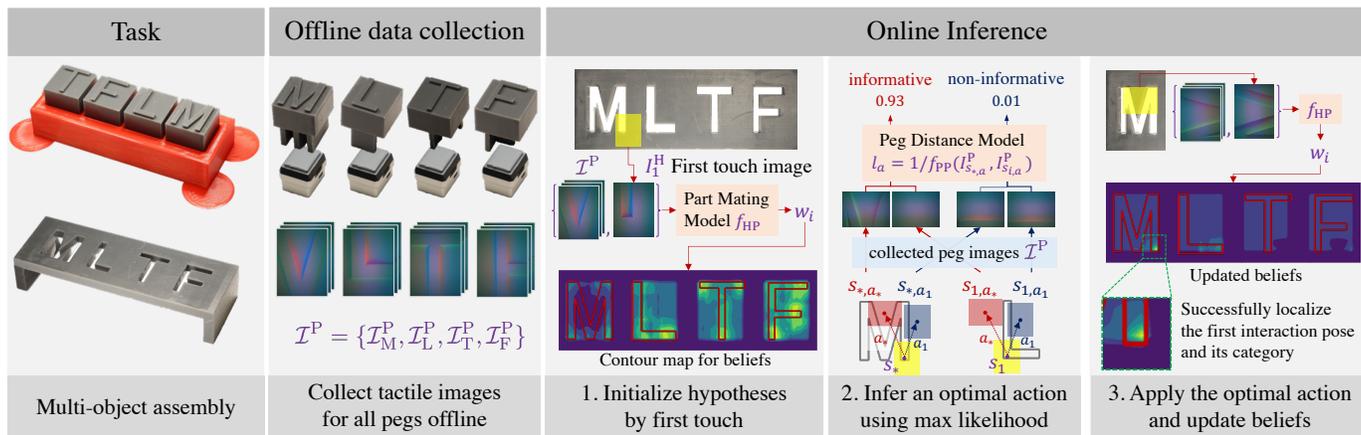


Fig. 1: Overview of the proposed *Tactile-Filter*. As shown in the figure, we consider the task of the part mating without any prior knowledge of 3D mesh of objects and which objects fit together. We assume that the robot has access to a collection of tactile images for the set of pegs as shown in figure (Offline data collection). During inference, the robot tries to identify which peg would fit into a given hole by the proposed *Tactile-Filter*. An initial set of hypotheses (denoted by $s \in \mathcal{S}$) is generated using the tactile image from the first touch and a trained part mating model, which predicts the correspondence between parts that fit together. We compute an optimal action for sampling the next image on the hole surface in order to minimize the uncertainty of the current estimate using a maximum likelihood approach. This is also illustrated in the figure, where given an initial touch, we can select an optimal action that results in maximum reduction of uncertainty. This method allows us to find the peg for the right fit as well as localize the hole (as we finally get the correct hypothesis) while minimizing the number of interactions during the task. (MLTF stands for Maximum Likelihood Tactile Filter). [Best viewed in color]

Abstract—Humans rely on touch and tactile sensing for a lot of dexterous manipulation tasks. Our tactile sensing provides us with a lot of information regarding contact formations as well as geometric information about objects during any interaction. With this motivation, vision-based tactile sensors are being widely used for various robotic perception and control tasks. In this paper, we present a method for interactive perception using vision-based tactile sensors for a part mating task, where a robot can use tactile sensors and a feedback mechanism using a particle filter to incrementally improve its estimate of objects (pegs and holes) that fit together. To do this, we first train a deep neural network that makes use of tactile images to predict the probabilistic correspondence between arbitrarily shaped objects that fit together. The trained model is used to design a particle filter which is used twofold. First, given one partial (or non-unique) observation of the hole, it incrementally improves the estimate of the correct peg by sampling more tactile observations. Second, it selects the next action for the robot to sample the next touch (and thus image) which results in maximum uncertainty reduction to minimize the number of interactions during the perception task. We evaluate our method on several part-mating tasks with novel objects using a robot equipped with a vision-based tactile sensor. We also show the efficiency of the proposed action selection method against a naive method. See supplementary video at https://www.youtube.com/watch?v=jMVBg_e3gLw.

I. INTRODUCTION

Humans rely on tactile sensing to monitor and control manipulation tasks during interactions with the environment. Tactile sensing allows us to monitor and respond to contact forces, adapt to object slip during grasping, and perform various perception tasks to build models of the environment. Even in situations where visual observation is not possible, tactile sensing enables us to interpret different types of interactions with the environment. For instance, we can locate objects in cluttered environments even in the absence of visual cues, identify the correct key for a lock by feeling the different options, or determine the type of video port (e.g., HDMI) on a monitor screen by touch alone. It has been the long-standing goal of robotics to imitate such intelligent behavior during manipulation tasks. Motivated by this goal, we present an interactive perception method for robots that utilize vision-based tactile sensors to construct reliable models for part mating.

A lot of manufacturing tasks could be decomposed into a sequence of insertion tasks. Object insertion is a well-studied contact-rich manipulation task in robotics [16, 36]. However,

the task becomes extremely challenging when the geometry of the mating objects is unknown. Also, this can make the task of part mating significantly complex as the uncertainty in geometry can limit the ability to understand any possible contact formation between the parts [16]. This complexity is further amplified in assembly tasks that require precise geometric information, as tolerances between mating parts become critical. Achieving the required level of precision in manufacturing tasks can be challenging when relying solely on vision-based algorithms.

To address these challenges, we propose leveraging vision-based tactile sensors located at the robot’s gripper for precise perception in these tasks. We present an interactive perception method, *Tactile-Filter*, using vision-based tactile sensors for estimating part correspondence for part mating, i.e., to estimate which parts fit into each other using tactile sensors in the absence of any vision sensing. We train a deep learning model to predict correspondence between the correct mating parts, observed using a tactile sensor. In the presence of partial observation or non-unique contact patch, we make use of a particle filter to aggregate the information from multiple touches and improve the estimate of the correct fit for a given hole. To minimize the number of interactions between the robot and the object, a maximum likelihood-based action selection method is used during the proposed interactive perception. The proposed method is tested on several different test environments with objects of different shapes and sizes that are not used in training the model to show the generalization of the proposed approach. Figure 1 shows the abstract idea of the proposed interactive perception method.

Contributions: This paper has the following contributions:

- 1) We present a part mating problem that deals with objects of unknown shapes, aiming to identify and estimate the pose of mating parts using minimal number of interactions. To address this problem, we present a novel approach called *Tactile-Filter*, which combines contrastive learning, particle filter, and tactile sensing for part mating.
- 2) Through our experiments conducted on novel objects, we demonstrate that our proposed method effectively resolves uncertainty by iteratively updating its belief during interactions. We demonstrate the ability to generalize to objects not encountered during training. Furthermore, we introduce an action selection method within our approach, which leads to significant improvements in efficiency.

It is noted that for the sake of brevity and clarity of presentation, we will use the word *peg* for the male part and *hole* for the female part in our paper.

II. RELATED WORK

Vision-based tactile sensors have attracted a lot of attention recently [44, 22]. These sensors provide a high-resolution capture of the contact patch during contact formation and thus can help in localization of contacts, detection of slip, etc [22, 29, 14]. Consequently, vision-based tactile sensors

have been used for a lot of control and perception tasks [34, 21, 16, 35, 25]. Most of these methods make use of displacement of the contact patch to recover a signal indicating slip which can be used for control of the manipulation task. Similarly, there has been prior work that uses these sensors for object classification and slip stabilization using feedback from these sensors [6, 45], or for learning visuotactile servoing [7, 18].

Several studies have focused on estimating object pose using tactile sensors which can be broadly categorized into two directions. The first group of studies addresses pose estimation from a single touch [27, 3, 17, 24], employing regression [27, 17] and contrastive learning [3]. However, these methods exhibit limitations when applied to large objects or objects with non-unique contact patches. The second group of studies focuses on pose estimation of large objects, often utilizing particle filtering techniques to narrow down the distribution of possible poses [28, 24, 38, 5]. However, these methods face challenges when applied to real objects due to factors such as low-dimensional sensors [28], the requirement of a large number of interactions for training interaction policies [24] or collecting samples [38], and the reliance on 3D models of objects and/or simulators [24, 38, 5]. In contrast to these two sets of works, we assume the shape of the objects is not known a priori, and allow only a limited number of interactions (a maximum of 10) by generating informative actions that effectively reduce ambiguity. Furthermore, our method not only estimates object pose but also identifies the object type from multiple candidates, adding an additional layer of complexity to the problem.

There exists another category of works that integrate vision and tactile sensors to estimate object pose [28, 12, 7, 17]. By leveraging the capabilities of vision sensors, these methods can reduce the number of interactions required to estimate object pose by narrowing down the distribution of possible poses. While our study specifically focuses on object pose estimation using vision-based tactile sensors, our method is not mutually exclusive with the methods that combine vision and tactile sensors. By integrating both modalities, we can leverage the strengths of each sensor type and potentially improve the accuracy and efficiency of object pose estimation.

Our problem setup is relevant to the interactive perception literature, where the goal is also to interact with the objects and update iteratively on the estimation of the states [4, 41, 37]. While previous work focuses mostly on state estimation from raw visual perception [19, 4], point clouds [33, 40, 30], or 3D states of the objects [31]. In contrast to these previous works, we consider state estimation using tactile sensors.

Our work is also related to perception during insertion or part mating. Vision is mostly insufficient to perform a lot of insertion tasks due to the precision required during these tasks [16, 26]. Consequently, there has been a lot of work making use of tactile and/or wrench measurements on contact formation between the mating parts. The idea behind most of these tasks is to make use of tactile measurements and a feedback mechanism to iteratively correct the pose error between the mating parts [16, 13, 23]. In all these methods,

geometric information is not explicitly used. Furthermore, they do not explicitly update the uncertainty in measurements. In contrast, we present a method where the robot can make use of the geometric information upon contact formation using tactile sensors to iteratively estimate the part correspondence as well as precise localization.

III. PROBLEM STATEMENT

In this section, we present a formal statement of the problem which is studied in this paper. We also motivate the problem by discussing some common scenarios where the proposed problem could arise and the proposed method could be useful.

We consider the task of perception during part mating while performing automated assembly. We consider scenarios where the robot can not use a vision sensor to perceive the target object to perform the desired mating task. Such situations could arise in tasks where a robot has to assemble a product where occlusions are created by other parts (e.g., consider the assembly of an electronic board). Apart from occlusion, these tasks could also require precision in pose estimates which might be very difficult to obtain using vision. To formalize the problem, we define the task for the robot as identifying the correct peg from a known set of possible choices by multiple observations of the hole using the vision-based tactile sensor(s). The goal here is to design algorithms that can identify the correct peg with minimum physical interaction with the hole using a tactile sensor. We make the following assumptions in the proposed study:

- 1) The possible number of pegs for the part mating task is fixed and known as apriori.
- 2) The rough location of the target hole is known so that the robot can establish initial contact with the part and it does not need to perform this rough localization using touch.
- 3) The robot can collect a dense set of tactile images for all the candidate pegs by touching each of the pegs at various different locations and orientations, prior to experiments.

The first assumption is not restrictive as we will generally have a limited number of parts to assemble. The second assumption is also very easily met using common vision methods with rough precision in localization. The third assumption would require that the robot has access to a detailed geometric model of the possible pegs observed using tactile sensors. This would require that the robot performs some exploration to collect this data. This assumption is required as the size of the tactile sensor may be small compared to the size of the objects that the robot is interacting with.

We focus on the setup where the target objects are larger than the size of a tactile sensor such that the entire object can not be observed using a single touch or a single touch can result in non-unique observations (consider when parts of the objects could be similar). Such problems would require multiple touches and a method to aggregate the data from multiple touches. For example, consider the task of inserting pegs into the shape of alphabet letters that are larger than

the sensor size (see Fig. 4). If the target hole is the letter "A", whether we can be certain that it is indeed an "A" with a single touch depends on where we land our finger. For instance, if the first touch is made on the horizontal line segment of the letter "A," this feature may also be present in other alphabets such as "B," "D," "E," and so on. However, by making contact with features unique to the letter "A," such as the lower left intersection, the probability of other candidates can be reduced. So the question is how do we aggregate information from multiple touches, and how do we select places to touch that will maximize the information so as to reduce required interactions?

IV. TACTILE FILTER

We propose *Tactile-Filter*, an uncertainty-aware interactive perception method to identify the correct peg from a candidate set that fits into a given hole for assembly using tactile sensors. At the core of the framework is a feature-matching model that computes the probability that a peg can pair with a hole by measuring the distance between the corresponding tactile images in a joint feature space. Using the feature matching model, we can construct the three critical steps in standard particle filtering: (1) we can initialize a set of possible hypotheses from the first touch by comparing the tactile image on the hole with the candidate images for the pegs collected before the experiment, (2) we then sample the most probable hypothesis and generate an action that maximizes uncertainty reduction, and (3) we apply the inferred action to the real system and update beliefs about the shape and pose of the target hole. Next, we detail each of these components.

A. Learning to find mating part

Given a pair of peg and hole images, we train a model that maximizes the similarity score if the image for the hole corresponds to the image for the peg. To this end, we use a contrastive learning framework [8, 20] to learn the feature space, similar to the work MoCo-v3 [9].

MoCo-v3 has two encoders, f_q and f_k , with output vectors q and k . The goal of learning is to find the key vectors k that correspond to the query vectors q . In our case, we consider the query vectors q to be the vectors from images of the holes and learn to maximize the similarity score between q and the vectors computed from the corresponding peg images k , while minimizing the similarity between the query vectors and a set of vectors from the negative peg images $\{k^-\}$. In MoCo-v3, this is formulated by minimizing the InfoNCE loss [32]:

$$\mathcal{L} = -\log \frac{\exp(q \cdot k^+/\tau)}{\exp(q \cdot k^+/\tau) + \sum_{k^-} \exp(q \cdot k^-/\tau)}, \quad (1)$$

where τ is the hyperparameter. More details can be found in [9], implementations in [11], and the training procedure is shown in Fig. 3. We denote the model used for measuring similarities between the hole images and peg images as f_{HP} , and name it *part mating model*. Additionally, we train a model, denoted as f_{PP} and referred to as the *peg distance model*, to calculate the similarity between peg images. This model

will be utilized to produce informative actions (as detailed in Section IV-C).

B. Generating hypotheses

In cases where only a partial shape of a hole can be observed using a tactile sensor, multiple corresponding choices for peg may exist. Thus it might be hard or impossible to determine the object with a deterministic approach. To address this uncertainty, we explicitly generate and maintain a set of hypotheses representing potential candidate choices using a particle filter.

Since the class of the target hole and its orientation and location is unknown (Sec. III), we generate a set of hypotheses \mathcal{S} where each hypothesis s_k is a quadruple: $s_k = (c_k, x_k, y_k, \theta_k)$, where c_k represents the possible object categories $c_k \in \mathcal{C}$, and x_k, y_k, θ_k are the SE(2) relative planar displacements from the center of the object. An example is shown in Fig. 1 and is also explained in Fig. 2.

One can initialize the particle set by sampling from a uniform distribution within a reasonable range (e.g., c_k from candidate categories, x_k from $[-X_{\text{MAX}}, X_{\text{MAX}}]$, which is the lower and upper limits of reasonable sizes of the peg, etc.). However, it will be inefficient as the sampling space becomes huge as the target object becomes larger, and/or the number of candidate categories increases. As an alternative, we initialize the particle set after obtaining the first tactile image of the target hole $I_{t=1}^H$ (which we denote as I_1^H) by utilizing the previously collected set of peg images \mathcal{I}^P and pre-trained part mating model f_{HP} . Specifically, we compute the similarities, w_i , between the observed initial tactile image for the hole I_1^H and each tactile image of the peg I_i^P from the set of peg images \mathcal{I}^P as:

$$w_i = f_{\text{HP}}(I_1^H, I_i^P). \quad (2)$$

We sample a particle proportional to this likelihood. Therefore, the probability of a particle given the initial hole image can be written as:

$$p(s = i | I_1^H) = \frac{w_i}{\sum_{i \in \{1, \dots, |\mathcal{I}^P|\}} w_i} \quad (3)$$

where, i is the index of the set of previously collected peg images as $i \in \{1, \dots, |\mathcal{I}^P|\}$ (see Fig. 1). We then initialize the set of hypotheses by independently sampling K particles with the above distribution. It is noted that the above distribution is a categorical distribution over all peg images.

C. Selecting informative action

In order to efficiently determine the category and pose of the target hole, we aim to calculate an optimal action that can maximize uncertainty reduction. While it is possible to compute such an optimal action by maximizing information gain against all possible peg images, it necessitates the integration of all latent variables, making it computationally infeasible within a reasonable time frame. As an alternative, we utilize the existing hypothesis set \mathcal{S} to enhance the sample efficiency.

We sample the most probable hypothesis from the current set of particles $s_* = \arg \max_{s_k \in \mathcal{S}} w_k$ and determine the optimal

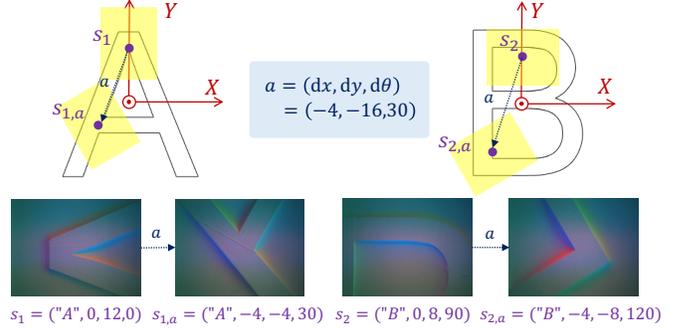


Fig. 2: This picture defines particles and actions available to the robot. A particle s is defined as a tuple consisting of the class of the object and a pose in SE(2) w.r.t. a frame attached to the center of the object. Each particle s_i is associated with the corresponding image $I_{s_i}^P$ observed via the tactile sensor at that location. An action a is equivalent to the transform in SE(2) applied to a particle s_i . The pose of a particle obtained by applying an action a can be obtained by applying the transform in SE(2) to the pose of the particle s , and we denote it as $s_{i,a}$. Thus, an action a will result in the observation of the contact patch $I_{s_{i,a}}^P$ at the new pose corresponding to the particle $s_{i,a}$. [Best viewed in color]

action by simulating it on a set of previously collected peg images \mathcal{I}^P (see Fig. 1). The action $a = (dx, dy, d\theta)$ is represented as a transform in SE(2) to the pose of the particles, and we denote the particle s_k with the updated pose by applying the action a as $s_{k,a}$. With this updated pose and the peg images \mathcal{I}^P , we can also obtain the peg image when applied to the action a , which we denote $I_{s_{k,a}}^P$. If there is no corresponding pose in the collected set of images, we assign an empty image which is a tactile image without any contact $I_{\text{NoContact}}$. We visually explain the definition of particles and actions in Fig. 2.

Given a current most-likely hypothesis, the next optimal action can be selected by finding the most informative action. To do that, we compute the distance between the tactile images obtained by applying any possible action to the most probable hypothesis and the remaining particles in set \mathcal{S} . The action that maximizes the sum of this distance over all the particles is selected as the optimal action. Such an action is favored only when the peg image of the sampled particle is close to the hole image, while other images have a greater distance when applying the same action to all the other particles. More concretely, we define the likelihood of an action as

$$l_a = \sum_{s_i \in \mathcal{S} \setminus \{s_*\}} 1/f_{\text{PP}}(I_{s_*,a}^P, I_{s_i,a}^P). \quad (4)$$

The optimal action can be then selected by maximizing the likelihood as:

$$a^* = \arg \max_{a \in \mathcal{A}} l_a, \quad (5)$$

where \mathcal{A} is a set of actions with which the tactile sensor can observe the peg after applying the action. This action set \mathcal{A}

Algorithm 1 *Tactile-Filter*

Input Number of candidate pegs N^P . A set of dense tactile images for each peg categories $\mathcal{I}^P = \{\mathcal{I}_1^P, \dots, \mathcal{I}_{N^P}^P\}$, number of particles K , maximum number of interactions N^{\max} , pre-trained part mating model f_{HP} and peg distance model f_{PP} (Sec. IV-A), threshold to stop iteration δ^{prob}

Output Mating peg category and its displacement in x, y, θ coordinate from the center of the peg

- 1: Touch and observe the first hole image I_1^H
 - 2: Compute similarity score w_i between the hole image I_1^H and a peg image $I_i^P \in \mathcal{I}^P$ as $w_i = f_{\text{HP}}(I_1^H, I_i^P)$
 - 3: Initialize a particle pool $\mathcal{S} = \emptyset$
 - 4: **for** $k \leftarrow 1$ to K **do**
 - 5: Sample a particle from the set of the peg images according to the categorical distribution defined in Eq.(3)
 - 6: Add the sampled particle $s_k = (c_k, x_k, y_k, \theta_k)$ to the particle pool: $\mathcal{S} \leftarrow \mathcal{S} \cup s_k$.
 - 7: **end for**
 - 8: **for** $t \leftarrow 2$ to N^{\max} **do**
 - 9: Select the most probable particle $s_* = \arg \max_{s_k \in \mathcal{S}} w_k$
 - 10: Infer the optimal action a_* according to Eq. (5) using s_* and \mathcal{S} .
 - 11: Move the robot with the inferred action a_* and get the tactile image I_t^H
 - 12: **for** $k \leftarrow 1$ to K **do**
 - 13: Compute importance weight for the particle $w_k = f_{\text{HP}}(I_t^H, I_{s_k}^P)$
 - 14: **end for**
 - 15: Compute the posterior distribution defined in Eq.(6)
 - 16: Re-sample K particles from \mathcal{S} using the posterior distribution
 - 17: Compute updated posterior $p_t(c|I_{1:t}^H)$ of each peg category from the particles as defined in Eq. (9)
 - 18: **if** $\max_{j \in \{1, \dots, N^P\}} p_t(c_j) > \delta^{\text{prob}}$ **then**
 - 19: Break the current loop
 - 20: **end if**
 - 21: **end for**
 - 22: Select the most probable particle $s_* = (c_*, x_*, y_*, \theta_*) = \arg \max_{s_k \in \mathcal{S}} w_k$ (to get the localization estimate, i.e., x_*, y_*, θ_*)
 - 23: **return** Object category c_* and its displacement x_*, y_*, θ_* .
-

can be obtained by calculating the L1 pixel distance between the tactile image of the peg after applying the action $I_{s_*, a}^P$ and the peg image that does not have contact $I_{\text{NoContact}}^P$, and see if it exceeds a threshold as $\mathbb{I}(\|I_{s_*, a}^P - I_{\text{NoContact}}^P\|_1 < \delta^{\text{act}})$.

This procedure is visually depicted in Fig. 1, where the optimal action a^* is shown to generate a more distinct contact patch (observable through touch) when applied to all particles (s_* and s_1 in the figure). On the other hand, the non-informative action a_1 produces similar contact patches that would not effectively disambiguate the current belief.

D. Update hypotheses

After obtaining the optimal action a_t at time step t , we apply it on the real robot and observe the tactile image of the hole I_t^H . We then update the probability of each hypothesis by comparing the observed hole image I_t^H and the peg images from the current hypotheses I_{s_k, a_t}^P :

$$p(s_t | I_{1:t}^H, a_{1:t-1}) \propto p(I_t^H | s_t) p(s_t | I_{1:t-1}^H, a_{1:t-1}) \approx \sum_{k=1}^K w_k p(s_t | I_{1:t-1}^H, a_{1:t-1}), \quad (6)$$

where $w_k = \frac{f_{\text{HP}}(I_t^H, I_{s_k, a_t}^P)}{\sum_{k=1}^K f_{\text{HP}}(I_t^H, I_{s_k, a_t}^P)}$ is the likelihood (weight) for the hole image to match with the peg image of the k^{th} particle.

The second term can be written as:

$$p(s_t | I_{1:t-1}^H, a_{1:t-1}) = \sum_{s_{t-1} \in \mathcal{S}} \underbrace{p(s_t | s_{t-1}, a_{t-1})}_{\text{forward dynamics}} \underbrace{p(s_{t-1} | I_{1:t-1}^H, a_{1:t-1})}_{\text{obtain through recursion}}, \quad (7)$$

where the first term represents the deterministic forward dynamics. Given the particle is in state s_k at time step t , the dynamics can be expressed as:

$$p(s_{t+1} | s_t, a_t) = \begin{cases} 1 & \text{if } s_{t+1} = s_{k, a_t} \\ 0 & \text{otherwise.} \end{cases} \quad (8)$$

The second term in Eq.(7) is initialized with the prior defined in Eq. (3) and can be obtained through recursion. We update the distribution of the particles by regenerating a new set of particles through weighted sampling based on w . The full algorithm is shown in Alg. 1.

E. Terminal condition

After updating the posterior of the particles with Eq. (6), we compute the posterior probability of each peg category as follows:

$$p_t(c | I_{1:t}^H) = \frac{\sum_{s_k \in \mathcal{S}} \mathbb{I}(s_k \in c)}{|\mathcal{S}|}, \quad (9)$$

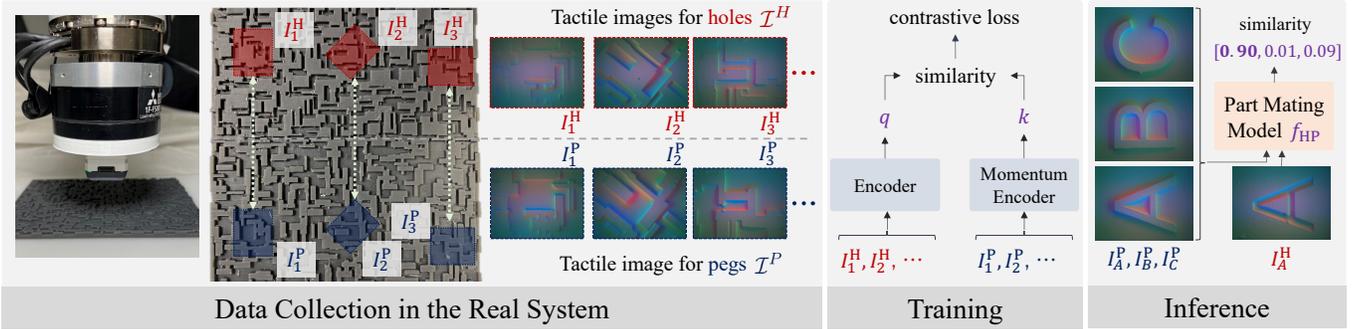


Fig. 3: Data Collection, Training, and Inference of the Part Mating Model: The left block depicts the data collection process using the *MAZE* board that features various shapes, including hole and peg shapes in the upper and lower halves, respectively. This board is placed on the robot platform and the robot arm equipped with a tactile sensor at the tip of the wrist makes contact with the board to collect data denoted as \mathcal{I}^H and \mathcal{I}^P , which corresponds a set of images for pegs and holes, respectively. The middle block illustrates the training procedure for the part mating model. It is trained in a self-supervised manner using a contrastive loss that encourages the model to produce high scores only when images corresponding to true mating parts are provided. The right block demonstrates the model’s generalization to different shapes after training. [Best viewed in color]

where $\mathbb{I}(s_k \in c)$ is an indicator function that returns 1 only when the category of the particle c_k belongs to the category c . The algorithm terminates when the majority of particles belong to a specific class, indicated by $\max_{j \in \{1, \dots, N^P\}} p_t(c_j) > \delta^{\text{prob}}$, which is a user-specified parameter for termination.

V. EXPERIMENTS

In this section, we evaluate the performance of the *Tactile-Filter* algorithm in two different test scenarios. The first scenario, referred to as the **small objects**, involves a collection of small objects that can be fully captured by a single touch of the tactile sensor, thereby making the estimation problem relatively simpler to solve. The second scenario, referred to as the **large objects**, involves objects that are larger than the size of the tactile sensor, requiring multiple touch measurements to accurately estimate their shape. All the test objects used in the pose estimation experiments are novel and are not used for training the contrastive learning model.

A. Training the part mating model

Tactile sensor. We use a commercially available GelSight Mini [2] tactile sensor, which provides 320×240 compressed RGB images through the Robot Operating System (ROS) at a rate of approximately 25 Hz, with a field of view of 18.6×14.3 millimeters.

Robot platform. The MELFA ASSISTA robot [1], a collaborative robot with 6 DoF, is used in this study. The tactile sensor is mounted on the robot’s wrist during data collection (see Fig. 3). It is noted that we do not use the force torque sensor mounted at the wrist of the robot as shown in the Fig. 3.

Data collection. In order to train a model that is capable of generalizing to a diverse set of shapes, we designed a board for data collection so that it features random polygonal shapes to simulate pegs and holes of arbitrary shapes. The shapes were generated through a process that involved creating a maze (we name it *MAZE* board), adding random perturbations to the

position and size of the walls that make up the maze, and then exporting the result for 3D printing. This board was designed such that any arbitrary hole patch sampled from the upper half has a corresponding mating peg patch in the lower half (see Fig. 3). To collect data for training, we sampled several different locations and orientations on the upper half *MAZE* board from a high-resolution grid to collect the hole images, and then collect the corresponding peg images from the lower half. This resulted in a total of approximately 23,000 pairs of images of pegs and holes which perfectly fit with each other.

Preprocessing. In this study, the tactile sensor used has RGB LEDs with different colors on each of the three surfaces [2]. As a result, even when the same object is in contact, the color may differ depending on the position of the image captured. To mitigate the potential impact on generalization performance, we obtained an image of a non-contact situation $I_{\text{NoContact}}$ during data collection, reducing the impact by subtracting the image. Then, the average and variance of each RGB channel were calculated for all images, and the images were normalized before being input into the model.

Training. As described in Sec. IV-A, we use MoCo-v3 for our part mating model f_{HP} and peg distance model f_{PP} . We train the models with the collected images using the *MAZE* board for 500 epochs. To improve generalization capability, we augment the data by using random cropping and horizontal or vertical flips, which will be applied to the pairs of images inputted to the model during training.

B. Small objects

We first evaluate the performance of the *TactileFilter* when applied to objects that fit in the size of the sensor.

Baselines. To understand the challenges encountered when identifying objects that might not be fully captured through a single touch, we compare our method against two methods that only use the initial image. The first baseline, referred to as *Pixel*, computes the L1 distance between the peg and hole

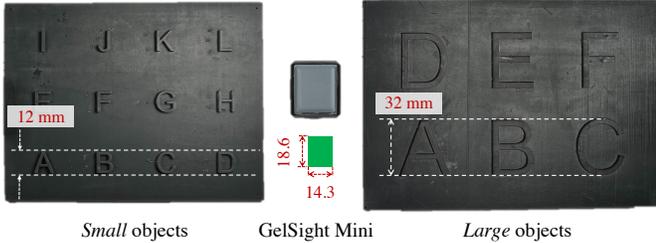


Fig. 4: Alphabet boards for our experiments. The left board contains small characters, each with a length of 12 mm and a maximum width of 16 mm, to fit within the size of the sensor if the robot makes contact with the center position. The sensor size is shown in the middle image. The board on the right has large characters with a length of 32 mm and a maximum width of 40 mm, requiring multiple interactions with the tactile sensor to obtain complete geometry for the object.

images and returns the index of the nearest neighbor image. The second baseline, *MoCo*, utilizes the pre-trained MoCo-v3 model to calculate the distance (negative of the MoCo-v3’s output) based solely on the first tactile image and without incorporating any subsequent interactions. The results of our method are denoted as *Ours* (n), where n indicates the number of interactions. It is important to note that the value of n includes the initial contact, therefore, *Ours* ($n = 1$) represents the results obtained without any additional interactions.

Settings. For this experiment, we have designed an evaluation board consisting of 12 alphabet characters (ranging from “A” to “L”), each with a maximum width of 16 mm and height of 12 mm, so the characters fit within a single touch. Since we would like to evaluate the model in situations where the pose of the object is unknown, resulting in only partial observation of the object and requiring multiple touches for accurate estimations, we collect data with displacements in $X, Y \in \{-8, -4, 0, 4, 8\}$ millimeter and $\theta \in \{-90, -60, \dots, 90\}$ degree from their center position. This results in $12 \times 5 \times 5 \times 7 = 2100$ images. Figure 4 shows examples of the characters we used for the experiment. The hyperparameter we used for our algorithm is the number of particles $K = 100$, the maximum number of iterations $N^{\max} = 10$, and the threshold to stop the iteration $\delta^{\text{prob}} = 0.95$.

Metrics. The performance of the results is assessed through two metrics. Firstly, we evaluate the accuracy in classifying the objects. For the baseline calculation, we calculate the distance between a hole image and all previously gathered peg images, select the image with the minimum distance, and consider the prediction to be accurate when the predicted image’s class matches the class of the inputted hole image. Additionally, the distance between the predicted pose and the ground truth pose is quantitatively measured.

With regard to the evaluation of the proposed method, we utilize the likelihood used for updating the particles to weight the prediction. The object with the highest weighted

TABLE I: Quantitative evaluation of single touch experiments with **small** objects on the alphabet board.

	Pixel	MoCo	Ours		
			$n = 3$	$n = 5$	$n = 10$
Accuracy [%]	0.0	39.6	81.8	90.7	95.0
Error XY [mm]	-	0.7	0.2	0.1	0.1
Error θ [deg]	-	5.4	0.9	0.3	0.1

TABLE II: Quantitative evaluation of multiple touch experiments with **large** objects on the alphabet board.

	Pixel	MoCo	Ours		
			$n = 3$	$n = 5$	$n = 10$
Accuracy [%]	11.9	41.9	58.7	72.3	85.0
Error XY [mm]	6.9	4.4	1.3	1.0	0.7
Error θ [deg]	16.6	15.5	4.4	2.9	1.5

probability is then evaluated with the target object. We also use weighted error between the particles and the ground truth image to compute the quantitative error.

Results and Analysis. The results are presented in Table I. A comparison between the two baselines, *Pixel* and *MoCo*, reveals that correspondences between parts cannot be obtained simply by comparing pixel values. The contrastive framework captures the features of mating parts, resulting in improved performance. However, the results using only the first contact are still not sufficiently accurate as the tactile sensor only observes a partial view of the object. In contrast, our method demonstrates a gradual improvement in performance as interactions are added. Additionally, as we can see from Table I, our method is able to achieve good localization accuracy both in position and orientation. In particular, we are able to achieve a submillimeter average error in localization which might be required for industrial insertion tasks.

C. Large objects

Settings. In the next set of experiments, we evaluate the performance of the proposed method when applied to objects that are larger than the size of the sensor. This scenario requires the robot to interact multiple times with the object to gain a comprehensive understanding of its shape. To this end, we have designed an evaluation board consisting of twelve alphabet characters (ranging from “A” to “L”), each with a maximum width of 40 mm and height of 32 mm. We tested the method on the location and orientation of the robot from $X, Y \in \{-20, -16, \dots, 20\}$ mm and $\theta \in \{-90, -60, \dots, 90\}$ with respect to the center position of each character. Figure 4 presents examples of the characters utilized in the experimental setup. As for the baselines, we compare the method against the same baselines as the previous experiment on *small* objects.

Results and Analysis. Table II shows the results on the large objects. Similar to the results obtained in the setting of small objects, our proposed model demonstrates improved performance compared to the baselines. However, it is also observed that a larger object size requires a greater number of interactions in order to achieve comparable accuracy. In

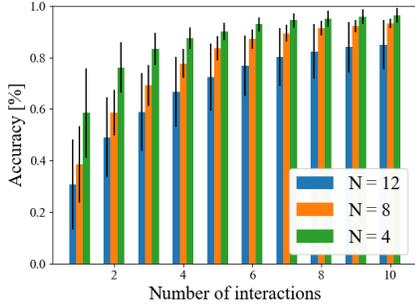


Fig. 5: Classification accuracy of the proposed method with a different number of classes evaluated on the *Large* objects. The result shows our method can quickly identify the correct class if the number of classes (shown by N) is limited.

Figure 5, we show the classification accuracy with respect to the number of interactions, and the results with randomly sampled smaller sets of 4 and 8 characters to evaluate the performance with a smaller number of possible candidates. The bar plots demonstrate that the proposed method can quickly identify the correct class if the number of classes is small.

D. Ablation on action selection strategy

Settings. To assess the effectiveness of the proposed action selection strategy, we compare the proposed method with a random action selection method (which we call *Random*). We evaluate the two methods on the *Small* and *Large* objects settings described earlier.

Results and Analysis. The results in Fig. 6 indicate the proposed maximum likelihood action selection approach demonstrates significant improvement in comparison to the method with regard to classification accuracy.

E. Evaluation on industrial connectors

Settings. To further evaluate the performance of the trained part mating model in an industrial setting, we collect tactile images of connectors and sockets from a Raspberry-Pi board, as depicted in Fig. 7.

Results and Analysis. The results of the evaluation of the Pixel baseline and our part mating model for the classification of connectors and sockets from the Raspberry-Pi board are presented in Table III. The Pixel baseline demonstrates improved performance in comparison to the small and large object experiments, due to the reduced number of classes in this setting and the unique size of each connector/socket, which simplifies the classification through the use of only L1 pixel distance. Although the part mating model outperforms the Pixel baseline, it misclassifies the female HDMI connector as the male USB-A connector. This is attributed to the significant distribution shift between the training set and the test set, where the pins on the surface of the male part are not present in the training data. To address this issue, future work can

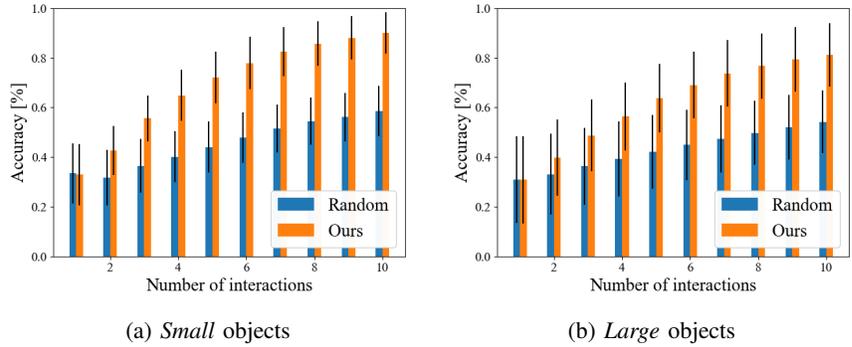


Fig. 6: Classification accuracy for *Small* and *Large* objects with different action strategies. As could be seen from these bar plots, our proposed method demonstrates significant improvement in comparison to the random action selection with regard to classification accuracy.

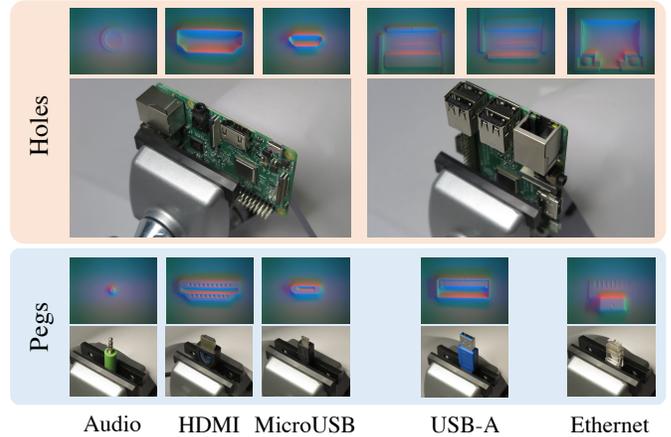


Fig. 7: Experimental setup for industrial connector identification on a Raspberry-Pi board. This image shows the observations of the six pegs and holes using the GelSight Mini sensor. Table III shows the classification results obtained by our model. [Best viewed in color]

TABLE III: Classification accuracy on the Raspberry Pi Board.

	Pixel	MoCo
Accuracy [%]	50.0	83.3

focus on enhancing the generalization capabilities of the part mating model.

F. Application to multi-object assembly

Settings. Once Tactile-Filter localizes the hole and identifies the corresponding peg, the robot can successfully insert the peg into the right hole. This is facilitated by the algorithm’s ability to estimate the pose with high precision, as demonstrated by the submillimeter average prediction error (refer to Table I and Table II). Consequently, we assess the proposed method in a real multi-object assembly scenario during the final experiment.

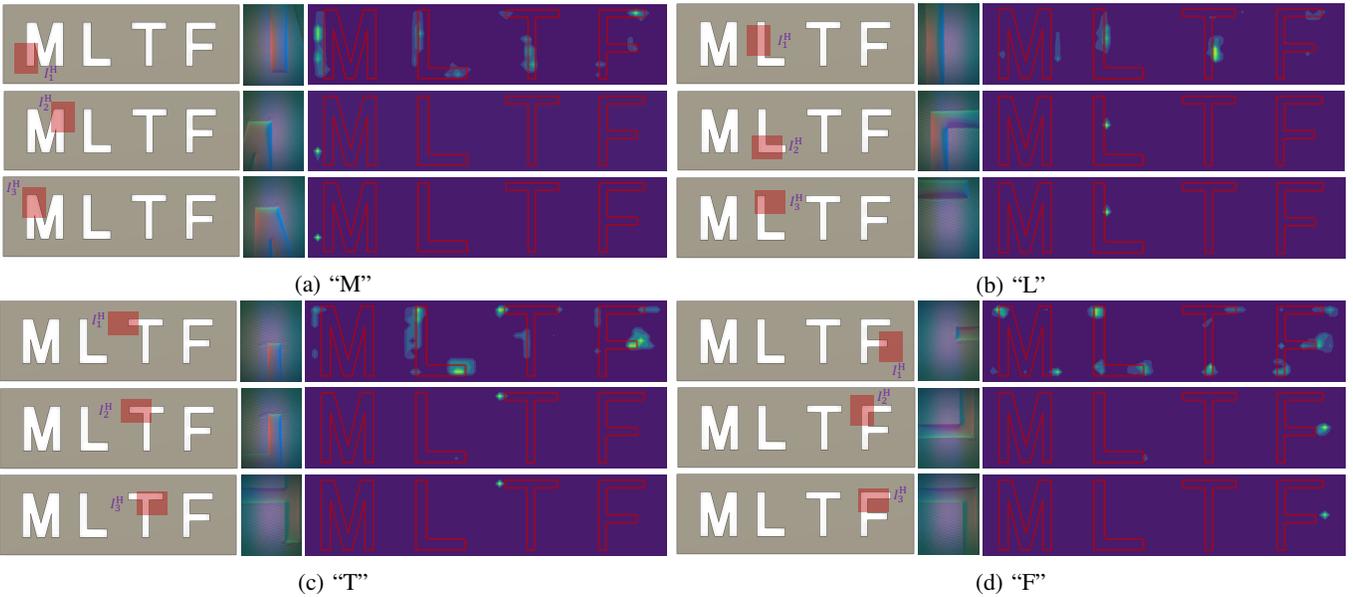


Fig. 8: Visualization of belief maps for holes categorized as “M”, “L”, “T”, and “F”. In each figure, the red regions in the left column show the spatial and temporal region captured by the tactile sensor during interactions, represented as I_t^H , where t indicates the interaction number or timestep. The center image shows the hole image captured by the tactile sensor, and the right figure illustrates the belief map generated by the current particles. Across all hole types, the results indicate a rapid convergence of the distribution of the hole’s initial contact pose and its corresponding category.

The task involves identifying and localizing four pegs and holes shaped like the alphabet characters “M”, “L”, “T”, and “F” (Maximum Likelihood Tactile Filter). Each peg has a length of 32 mm and a maximum width of 28 mm, as depicted in Fig. 1 (leftmost picture). Although the algorithm achieves accurate pose estimation for the holes, the robot still fails in insertion due to errors during grasping. To account for this, we design holes with a tolerance of 2 mm and treat insertion as a simple pick-and-place operation using an impedance controller. For more precise assembly, we can combine our method with prior work, such as [16, 15].

Results and Analysis. The qualitative results are available in the video accessible at https://www.youtube.com/watch?v=jMVBg_e3gLw. The video demonstrates the algorithm successfully identifies the correct peg and pose of the hole, and the robot successfully inserts the pegs. Moreover, the visualization generated from this experiment as shown in Fig. 8 demonstrates that our method is iteratively corrects its belief during the interactive perception. Finally, in terms of computational time, the most time-consuming steps of the algorithm involve updating importance weights using the part mating model f_{HP} in Eq. (2) and the peg distance model f_{PP} in Eq. (4). However, each of these steps takes approximately 0.3 seconds on a single GPU, which is significantly shorter than other robot operations. Thus, our algorithm is suitable for online control.

VI. CONCLUSIONS AND FUTURE WORK

Tactile sensing can allow robots to build reliable models of their environment to perform precise manipulation tasks.

In this paper, we presented a novel method called *Tactile-Filter*. We presented an interactive perception method where a robot can improve its estimate for the perception task using tactile sensors while minimizing the number of interactions required with its environment. We considered the design of the method in the context of the task of part mating. In the absence of any vision input, we described a method where the robot could incrementally improve its estimate of correspondence between parts within a fixed number of available choices. We also proposed a maximum likelihood-based approach to select future actions to minimize the number of interactions during the perception task. The proposed method was verified using a vision-based tactile sensor and a physical robot on several tasks of part mating. The generalization of the proposed method to previously unseen scenarios was also illustrated.

As our method was trained and evaluated on a physical system, data collection was performed using a real robot. However, in future work, we aim to explore the possibility of utilizing an appropriate simulation environment [10, 43, 42, 39] to simulate contact patches for various object geometries. This approach would enable us to reduce reliance on physical robots for data collection and instead leverage simulations to acquire data. By doing so, we can learn and develop complex perception techniques with minimal usage of real-world data.

Another limitation of the proposed method is that we assume that tactile images consist of only the peg and hole parts. Therefore, the underlying deep learning model can get easily confused if distractors are present in the image (such as attachments to connectors in Fig. 7). In future research, we can work on this limitation by localizing the mating parts

from the tactile image to make the method more robust to such distractors.

REFERENCES

- [1] MELFA ASSISTA. <https://www.mitsubishielectric.com/fa/products/rbt/robot/items/assista/index.html>. Accessed: 2023-01-16.
- [2] GelSight Mini. <https://www.gelsight.com/gelsightmini/>. Accessed: 2023-01-16.
- [3] Maria Bauza, Antonia Bronars, and Alberto Rodriguez. Tac2pose: Tactile object pose estimation from the first touch. *arXiv preprint arXiv:2204.11701*, 2022.
- [4] Jeannette Bohg, Karol Hausman, Bharath Sankaran, Oliver Brock, Danica Kragic, Stefan Schaal, and Gaurav S. Sukhatme. Interactive perception: Leveraging action in perception and perception in action. *IEEE Transactions on Robotics*, 33(6):1273–1291, 2017. doi: 10.1109/TRO.2017.2721939.
- [5] Gabriele M. Caddeo, Nicola A. Piga, Fabrizio Bottarel, and Lorenzo Natale. Collision-aware in-hand 6d object pose estimation using multiple vision-based tactile sensors. In *Proceedings of IEEE International Conference on Robotics and Automation (ICRA)*, 2023.
- [6] Roberto Calandra, Andrew Owens, Dinesh Jayaraman, Justin Lin, Wenzhen Yuan, Jitendra Malik, Edward H Adelson, and Sergey Levine. More than a feeling: Learning to grasp and regrasp using vision and touch. *IEEE Robotics and Automation Letters*, 3(4):3300–3307, 2018.
- [7] Arkadeep Narayan Chaudhury, Timothy Man, Wenzhen Yuan, and Christopher G. Atkeson. Using collocated vision and tactile sensors for visual servoing and localization. *IEEE Robotics and Automation Letters*, 7(2): 3427–3434, 2022. doi: 10.1109/LRA.2022.3146565.
- [8] Ting Chen, Simon Kornblith, Mohammad Norouzi, and Geoffrey Hinton. A simple framework for contrastive learning of visual representations. In Hal Daumé III and Aarti Singh, editors, *Proceedings of the 37th International Conference on Machine Learning*, volume 119 of *Proceedings of Machine Learning Research*, pages 1597–1607. PMLR, 13–18 Jul 2020. URL <https://proceedings.mlr.press/v119/chen20j.html>.
- [9] X. Chen, S. Xie, and K. He. An empirical study of training self-supervised vision transformers. In *2021 IEEE/CVF International Conference on Computer Vision (ICCV)*, pages 9620–9629, Los Alamitos, CA, USA, oct 2021. IEEE Computer Society. doi: 10.1109/ICCV48922.2021.00950. URL <https://doi.ieeecomputersociety.org/10.1109/ICCV48922.2021.00950>.
- [10] Alex Church, John Lloyd, Raia Hadsell, and Nathan F. Lepora. Tactile sim-to-real policy transfer via real-to-sim image translation. In Aleksandra Faust, David Hsu, and Gerhard Neumann, editors, *Proceedings of the 5th Conference on Robot Learning*, volume 164 of *Proceedings of Machine Learning Research*. PMLR, 08–11 Nov 2022. URL <https://proceedings.mlr.press/v164/church22a.html>.
- [11] MMSelfSup Contributors. MMSelfSup: Openmmlab self-supervised learning toolbox and benchmark. <https://github.com/open-mmlab/mmselfsup>, 2021.
- [12] Snehal Dikhale, Karankumar Patel, Daksh Dhingra, Itoshi Naramura, Akinobu Hayashi, Soshi Iba, and Nawid Jamali. Visuotactile 6d pose estimation of an in-hand object using vision and tactile sensor data. *IEEE Robotics and Automation Letters*, 7(2):2148–2155, 2022. doi: 10.1109/LRA.2022.3143289.
- [13] Siyuan Dong and Alberto Rodriguez. Tactile-based insertion for dense box-packing. In *2019 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, pages 7953–7960, 2019. doi: 10.1109/IROS40897.2019.8968204.
- [14] Siyuan Dong, Daolin Ma, Elliott Donlon, and Alberto Rodriguez. Maintaining grasps within slipping bounds by monitoring incipient slip. In *2019 International Conference on Robotics and Automation (ICRA)*, pages 3818–3824. IEEE, 2019.
- [15] Siyuan Dong, Devesh K. Jha, Diego Romeres, Sangwoon Kim, Daniel Nikovski, and Alberto Rodriguez. Tactile-rl for insertion: Generalization to objects of unknown geometry. In *2021 IEEE International Conference on Robotics and Automation (ICRA)*, pages 6437–6443, 2021. doi: 10.1109/ICRA48506.2021.9561646.
- [16] Siyuan Dong, Devesh K Jha, Diego Romeres, Sangwoon Kim, Daniel Nikovski, and Alberto Rodriguez. Tactile-rl for insertion: Generalization to objects of unknown geometry. In *2021 IEEE International Conference on Robotics and Automation (ICRA)*, pages 6437–6443. IEEE, 2021.
- [17] Letian Fu, Huang Huang, Lars Berscheid, Hui Li, Ken Goldberg, and Sachin Chitta. Safely learning visuotactile feedback policies in real for industrial insertion. *arXiv preprint arXiv:2210.01340*, 2022.
- [18] Johanna Hansen, Francois Hogan, Dmitriy Rivkin, David Meger, Michael Jenkin, and Gregory Dudek. Visuotactile-rl: Learning multimodal manipulation policies with deep reinforcement learning. In *2022 International Conference on Robotics and Automation (ICRA)*, pages 8298–8304, 2022. doi: 10.1109/ICRA46639.2022.9812019.
- [19] Karol Hausman, Scott Niekum, Sarah Osentoski, and Gaurav S. Sukhatme. Active articulation model estimation through interactive perception. In *2015 IEEE International Conference on Robotics and Automation (ICRA)*, pages 3305–3312, 2015. doi: 10.1109/ICRA.2015.7139655.
- [20] Kaiming He, Haoqi Fan, Yuxin Wu, Saining Xie, and Ross Girshick. Momentum contrast for unsupervised visual representation learning. In *2020 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 9726–9735, 2020. doi: 10.1109/CVPR42600.2020.00975.
- [21] Francois R. Hogan, Jose Ballester, Siyuan Dong, and Alberto Rodriguez. Tactile dexterity: Manipulation prim-

- itives with tactile feedback. In *2020 IEEE International Conference on Robotics and Automation (ICRA)*, pages 8863–8869, 2020. doi: 10.1109/ICRA40945.2020.9196976.
- [22] S. Dong I. Taylor and A. Rodriguez. Gelslim 3.0: High-resolution measurement of shape, force and slip in a compact tactile-sensing finger. In *ICRA*, 2022.
- [23] Devesh K. Jha, Diego Romeres, William Yezazunis, and Daniel Nikovski. Imitation and supervised learning of compliance for robotic assembly. In *2022 European Control Conference (ECC)*, pages 1882–1889, 2022. doi: 10.23919/ECC55457.2022.9838102.
- [24] Tarik Kelestemur, Robert Platt, and Taskin Padir. Tactile pose estimation and policy learning for unknown object manipulation. In *Proceedings of the 21st International Conference on Autonomous Agents and Multiagent Systems*, AAMAS ’22, page 742–750, Richland, SC, 2022. International Foundation for Autonomous Agents and Multiagent Systems. ISBN 9781450392136.
- [25] Sangwoon Kim, Devesh K Jha, Diego Romeres, Parag Patre, and Alberto Rodriguez. Simultaneous tactile estimation and control of extrinsic contact. *arXiv preprint arXiv:2303.03385*, 2023.
- [26] Michelle A Lee, Yuke Zhu, Krishnan Srinivasan, Parth Shah, Silvio Savarese, Li Fei-Fei, Animesh Garg, and Jeannette 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)*, pages 8943–8950. IEEE, 2019.
- [27] Rui Li, Robert Platt, Wenzhen Yuan, Andreas ten Pas, Nathan Roscup, Mandayam A. Srinivasan, and Edward Adelson. Localization and manipulation of small parts using gelsight tactile sensing. In *2014 IEEE/RSJ International Conference on Intelligent Robots and Systems*, pages 3988–3993, 2014. doi: 10.1109/IROS.2014.6943123.
- [28] Shan Luo, Wenxuan Mou, Kaspar Althoefer, and Hongbin Liu. Localizing the object contact through matching tactile features with visual map. In *2015 IEEE International Conference on Robotics and Automation (ICRA)*, pages 3903–3908, 2015. doi: 10.1109/ICRA.2015.7139743.
- [29] Daolin Ma, Siyuan Dong, and Alberto Rodriguez. Extrinsic contact sensing with relative-motion tracking from distributed tactile measurements. In *2021 IEEE international conference on robotics and automation (ICRA)*, pages 11262–11268. IEEE, 2021.
- [30] Kaichun Mo, Leonidas J. Guibas, Mustafa Mukadam, Abhinav Gupta, and Shubham Tulsiani. Where2act: From pixels to actions for articulated 3d objects. In *Proceedings of International Conference on Computer Vision (ICCV)*, pages 6813–6823, October 2021.
- [31] Tonci Novkovic, Remi Pautrat, Fadri Furrer, Michel Breyer, Roland Siegwart, and Juan Nieto. Object finding in cluttered scenes using interactive perception. In *2020 IEEE International Conference on Robotics and Automation (ICRA)*, pages 8338–8344, 2020. doi: 10.1109/ICRA40945.2020.9197101.
- [32] Aaron van den Oord, Yazhe Li, and Oriol Vinyals. Representation learning with contrastive predictive coding. *arXiv preprint arXiv:1807.03748*, 2018.
- [33] Kei Ota, Hsiao-Yu Tung, Kevin A Smith, Anoop Cherian, Tim K Marks, Alan Sullivan, Asako Kanezaki, and Joshua B Tenenbaum. H-SAUR: Hypothesize, Simulate, Act, Update, and Repeat for Understanding Object Articulations from Interactions. *arXiv preprint arXiv:2210.12521*, 2022.
- [34] Yu She, Shaoxiong Wang, Siyuan Dong, Neha Sunil, Alberto Rodriguez, and Edward Adelson. Cable manipulation with a tactile-reactive gripper. *The International Journal of Robotics Research*, 40(12-14):1385–1401, 2021.
- [35] Yuki Shirai, Devesh K Jha, Arvind U Raghunathan, and Dennis Hong. Tactile tool manipulation. *arXiv preprint arXiv:2301.06698*, 2023.
- [36] Bruno Siciliano, Oussama Khatib, and Torsten Kröger. *Springer handbook of robotics*, volume 200. Springer, 2008.
- [37] Sanjana Srivastava, Chengshu Li, Michael Lingelbach, Roberto Martín-Martín, Fei Xia, Kent Elliott Vainio, Zheng Lian, Cem Gokmen, Shyamal Buch, Karen Liu, Silvio Savarese, Hyowon Gweon, Jiajun Wu, and Li Fei-Fei. Behavior: Benchmark for everyday household activities in virtual, interactive, and ecological environments. In Aleksandra Faust, David Hsu, and Gerhard Neumann, editors, *Proceedings of the 5th Conference on Robot Learning*, volume 164 of *Proceedings of Machine Learning Research*, pages 477–490. PMLR, 08–11 Nov 2022. URL <https://proceedings.mlr.press/v164/srivastava22a.html>.
- [38] Sudharshan Suresh, Zilin Si, Stuart Anderson, Michael Kaess, and Mustafa Mukadam. Midastouch: Monte-carlo inference over distributions across sliding touch. In *6th Annual Conference on Robot Learning*, 2022. URL <https://openreview.net/forum?id=JWR0nOf4w-K>.
- [39] Shaoxiong Wang, Mike Lambeta, Po-Wei Chou, and Roberto Calandra. Tacto: A fast, flexible, and open-source simulator for high-resolution vision-based tactile sensors. *IEEE Robotics and Automation Letters*, 7(2): 3930–3937, 2022. doi: 10.1109/LRA.2022.3146945.
- [40] Ruihai Wu, Yan Zhao, Kaichun Mo, Zizheng Guo, Yan Wang, Tianhao Wu, Qingnan Fan, Xuelin Chen, Leonidas Guibas, and Hao Dong. VAT-Mart: Learning visual action trajectory proposals for manipulating 3d articulated objects. In *International Conference on Learning Representations (ICLR)*, 2022.
- [41] Fei Xia, William B. Shen, Chengshu Li, Priya Kasimbeg, Micael Edmond Tchappmi, Alexander Toshev, Roberto Martín-Martín, and Silvio Savarese. Interactive gibbon benchmark: A benchmark for interactive navigation in cluttered environments. *IEEE Robotics and Automation*

Letters, 5(2):713–720, 2020. doi: 10.1109/LRA.2020.2965078.

- [42] Jie Xu, Tao Chen, Lara Zlokapa, Michael Foshey, Wojciech Matusik, Shinjiro Sueda, and Pulkit Agrawal. An End-to-End Differentiable Framework for Contact-Aware Robot Design. In *Proceedings of Robotics: Science and Systems*, Virtual, July 2021. doi: 10.15607/RSS.2021.XVII.008.
- [43] Jie Xu, Sangwoon Kim, Tao Chen, Alberto Rodriguez Garcia, Pulkit Agrawal, Wojciech Matusik, and Shinjiro Sueda. Efficient tactile simulation with differentiability for robotic manipulation. In *6th Annual Conference on Robot Learning*, 2022. URL <https://openreview.net/forum?id=6BIfCl6gsM>.
- [44] Wenzhen Yuan, Siyuan Dong, and Edward H. Adelson. Gelsight: High-resolution robot tactile sensors for estimating geometry and force. *Sensors*, 17(12), 2017. ISSN 1424-8220. doi: 10.3390/s17122762. URL <https://www.mdpi.com/1424-8220/17/12/2762>.
- [45] Wenzhen Yuan, Yuchen Mo, Shaoxiong Wang, and Edward H Adelson. Active clothing material perception using tactile sensing and deep learning. In *2018 IEEE International Conference on Robotics and Automation (ICRA)*, pages 4842–4849. IEEE, 2018.