

Placing by Touching: An empirical study on the importance of tactile sensing for precise object placing

Luca Lach^{*1,2}, Niklas Funk^{*3}, Robert Haschke¹, Séverin Lemaignan², Helge Joachim Ritter¹, Jan Peters^{3,4,5,6}, Georgia Chalvatzaki^{3,5}




Fig. 1: Stable object placing sequence with a *blind* robot, from left to right: (1) an object is handed to the robot in an unknown pose. Subsequently, by leveraging the tactile sensor readings inside the gripper, we estimate how the object needs to be reoriented for stable placement. (2) Given the estimate, a controller aligns the object with the placing surface. (3) We move the robot down to place the object on the table. Contact is again detected using tactile sensors. (4) The gripper is opened, and the robot retracts.

Abstract—This work deals with a practical everyday problem: stable object placement on flat surfaces starting from unknown initial poses. Common object-placing approaches require either complete scene specifications or extrinsic sensor measurements, e.g., cameras, that occasionally suffer from occlusions. We propose a novel approach for stable object placing that combines tactile feedback and proprioceptive sensing. We devise a neural architecture that estimates a rotation matrix, resulting in a corrective gripper movement that aligns the object with the placing surface for the subsequent object manipulation. We compare models with different sensing modalities, such as force-torque and an external motion capture system, in real-world object placing tasks with different objects. The experimental evaluation of our placing policies with a set of unseen everyday objects reveals significant generalization of our proposed pipeline, suggesting that tactile sensing plays a vital role in the intrinsic understanding of robotic dexterous object manipulation. Code, models, and supplementary videos are available on <https://sites.google.com/view/placing-by-touching>.

I. INTRODUCTION

Human dexterity is impeccable and is largely attributed to the human sense of touch. Tactile sensing enables precise, reliable, dexterous manipulation, a crucial component for the

vision of generalized autonomy and of more capable, intelligent autonomous robotic systems [1]. Integrating the sense of touch is a key research topic in robotics, with prominent works including tactile insertion [2]–[4], in-hand manipulation [5]–[8], assembly [9]–[11], human-robot-interaction [12], [13], and object pose estimation [14], [15].

In the past few years, significant effort has been put into designing sensitive and compact tactile sensors that allow easy integration with robotic systems. Compared to more traditional Force/Torque (F/T) sensors that only provide one 6-dimensional F/T measurement at the sensor’s location, tactile sensors offer a high spatial resolution and measurements directly at the points of contact. However, the high-dimensional tactile signals are usually unsuitable for direct integration in control loops and require additional preprocessing. Tactile sensors can be realized using a wide range of sensing principles [16]. Recently, tactile sensors based on piezo-resistive sensing distributed in an array of taxels [17]–[21], and those relying on cameras capturing a soft gel’s surface [22]–[27] have become increasingly popular amongst the robotics community.

Herein, we focus on piezo-resistive tactile sensors. In particular, we use a pair of Myrmex tactile sensors [18], which feature a 16×16 taxel array sampled at 1 kHz, i.e., matching the frequency of good F/T sensors, whilst visuotactile sensors can perform at best around 90 Hz. Regarding costs, tactile sensors like Myrmex or GelSight are considerably cheaper than F/T sensors.

This work studies the benefit of local tactile measurements between gripper and object for stable and reliable object placing. Stable object placement is an essential skill for any autonomous robotic system, particularly for capable assistive household robots. It forms the basis for many tasks, such as object rearrangement, assembly, sorting, and storing goods. While a large body of prior works exists on stable

* Authors contributed equally.

This work was supported by the European Union Horizon 2020 Marie Curie Actions under Grant 813713 NeuTouch, the DFG Emmy Noether Programme (CH 2676/1-1), the BMBF Project Aristotle, and the AICO grant by the Nexplore/Hochtief Collaboration with TU Darmstadt.

¹ Neuroinformatics Group, Technical Faculty, Bielefeld University {llach, rhaschke, helge}@techfak.uni-bielefeld.de

² PAL Robotics, Barcelona, Spain

³ Computer Science Department, Technical University Darmstadt,

⁴ German Research Center for AI (DFKI), Research Department: Systems AI for Robot Learning, ⁵ Hessian.AI, ⁶ Centre for Cognitive Science. {niklas.funk, jan.peters, georgia.chalvatzaki}@tu-darmstadt.de

This work has been submitted to the IEEE for possible publication. Copyright may be transferred without notice, after which this version may no longer be accessible.

object placing [28]–[34], none of those works investigates the contribution of tactile feedback in stable placing. Rather, they rely either on vision systems, which are prone to occlusions and require external sensors, or accurate scene descriptions, which demand cumbersome manual labor. We attempt to fill this gap by investigating the impact of tactile sensing in this simple yet challenging scenario.

We propose an effective pipeline for translating taxel-based measurements into useful features for learning a pose correction signal to ensure optimal object placement. A placing action is optimal if the object’s placing normal (orthogonal to its placing surface) is colinear with the normal of the placing surface, e.g., a table. Our method comprises a deep convolutional neural network that predicts a corrective rotation action for the gripper. Given the current tactile sensor readings and potentially adding other signals, e.g., F/T information, we predict a rotation matrix w.r.t. the current gripper frame (cf. Fig. 3(a)). The z-axis of this predicted frame corresponds to the object’s placing normal. This prediction is subsequently used to plan a hand movement to align the object’s placing normal with the table’s normal. After this single-step prediction and alignment, we attempt to place the object on the surface while keeping the previously determined orientation fixed. The major challenge here is to predict the object’s placing normal solely from tactile and proprioceptive sensors instead of employing traditional extrinsic vision-based methods. To assess the importance of learning-based placing policies for this problem, we compare our method to two classical baseline approaches.

Our main contributions are twofold; (i) the development and training of tactile-based policies for stable object placing without requiring any extrinsic visual feedback, and (ii) an open-source suite of our dataset, CADs, pretrained models, and the codebase of all methods (both classical and deep learning ones) from our extensive real-robot experiments.

Overall, our study confirms that tactile sensing can be a powerful and valuable low-cost addition to robotic manipulators: their signals provide features that increase reliability and robot dexterity.

II. RELATED WORK

Object placing. Stable object placing is a crucial skill for autonomous robotic systems. Many prominent tasks in the robotic community, such as object rearrangement or assembly, require robotic pick & place sequences that heavily rely on this skill. The authors of [29] propose a model-based pointcloud-conditioned approach for stable object placing by matching polygon models of object and environment. Similarly, [28] uses pointcloud observations for extracting meaningful feature representations for learning to place new objects in stable configurations. More recently, [32] proposes to exploit learned keypoint representations from raw RGB-D images for solving category-level manipulation tasks. [31] also uses a combination of vision and learning for manipulating unknown objects. [33] presents a planning algorithm for stable object placement in cluttered scenes requiring a fully specified environment. Closely related to our work is

[34], which presents an iterative learning-based approach for placing household objects onto flat surfaces but using a system of three external depth cameras for input.

While most of these works deal with the problem of generating stable placing poses for unknown objects, the main difference to our work lies in the input observation—none of them considers tactile sensing. Instead, they all rely on single or even multiple depth/RGB images. Relying on image data might be problematic due to gripper-object occlusions in highly cluttered scenes, especially if the object ends up inside the gripper without any prior knowledge of its pose. Additionally, external sensing systems require careful and precise calibration w.r.t. the robot, which is often tedious, time-consuming, and error-prone. In contrast, tactile sensors directly provide the contact information between the object and gripper, independent of the surrounding environment.

In-hand object pose estimation. Due to the inherent difficulties of estimating a grasped object’s pose and due to its importance for tasks like pick & place or in-hand manipulation, multiple methods for object-in-hand pose estimation have been developed. The authors of [35] solely exploit tactile sensors and match their signal with a local patch of the object’s geometry, thereby estimating its pose. Other works [36], [37] make use of both visual and tactile inputs. While [36] only requires an initial visual input for initializing a particle filter, [37] employs an extended Kalman filter constantly using vision & touch. Recent progress in deep learning has fostered data-driven methods for in-hand object pose estimation. [38]–[40] present end-to-end approaches based on RGB images. While [38], [39] directly outputs pose predictions, [40] learns observation models that can later be exploited in optimization frameworks. [41] fuses vision and tactile in an approach that self-determines the reliability of each modality, while [42] exploits a learned tactile observation model in combination with a Bayes filter. Following the successes of recent deep learning approaches, we propose to learn an end-to-end direct mapping from tactile input for estimating the grasped object’s placing normal. We want to point out that we are not interested in estimating the object’s full 6D pose. Instead, we only focus on aligning the object with the placing surface. Moreover, our proposed method does not require any repetitive measurements or filtering and solely needs to be queried once. Finally, our method requires a very small training set.

Insertion. Stable object placing is also related to tactile insertion. Successful completion of both tasks requires suitable alignment between object & table, or peg & hole. Several works approach challenging insertion tasks using tactile sensors [2]–[4], [43], [44]. The authors of [43] leverage vision-based tactile sensors for precisely localizing small objects inside the gripper. This information is subsequently exploited for small-part insertions using classical control. [2], [3] also focus on solving tight insertion tasks using learned tactile representations. Both exploit the tactile measurements as a feedback signal to predict residual control commands. Recently, [4] demonstrated tactile insertion through active extrinsic sensing, i.e., explicitly estimating the contact between

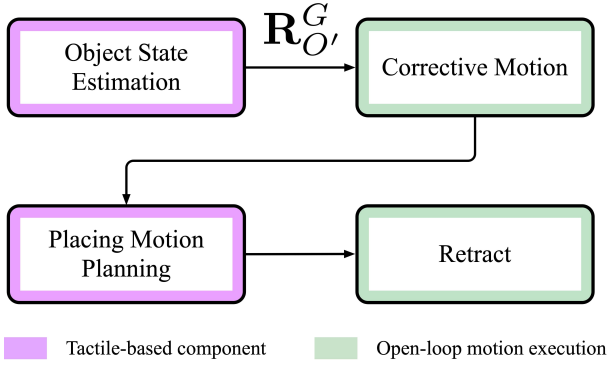


Fig. 2: Schematic overview of the four phases of stable object placing corresponding to the sequence shown in Fig. 1. Components relying on tactile data are highlighted, as well as those utilizing motion planning and execution.

a peg and a hole. Contrarily to tactile placing, tactile insertion requires environment interactions – typical strategies involve highly compliant policies for guiding the insertion [45]. In tactile-based object placing instead, it is crucial to estimate the properties of the currently held object for identifying stable placing poses. To the best of our knowledge, we are the first to investigate tactile sensing for the stable placing of objects onto flat surfaces, providing a thorough analysis of the different combinations of sensing modalities and comparisons with classical and deep approaches for concluding on the importance of tactile sensing when planning object-placing actions.

III. STABLE OBJECT PLACING

This work considers one of the most common household manipulation tasks: placing on planar surfaces like shelves or tables. For an object to be placed optimally, one of its (potentially many) placing faces need to be aligned with the placing surface prior to releasing from the hand. Which faces of an object are safe for placement is both object- and context-dependent. In the following, we will first describe the different phases of stable object placing, then provide a mathematical definition of the pose correction action necessary for aligning the object with the placing surface. Finally, we describe the controllers that execute the pose correction and object placing.

A. The four phases of object-placing by touching

We structure stable object placing into four phases, as shown in Figs. 2 & 1. At the beginning of the object-placing pipeline, we only assume the object to be inside the gripper. Yet, we do not have any prior information about its pose, as would be the case after a handover. In the first phase, given sensory information, we estimate the object’s placing normal, which allows quantifying the misalignment with the placing surface (see Sec. III-B). Based on the misalignment prediction, we compute and execute a corrective motion (see Sec. III-C). For the third phase, a placing motion is planned and executed that moves the object towards the placing surface linearly while maintaining the object’s orientation. Once the robot detects table-object contact in the tactile

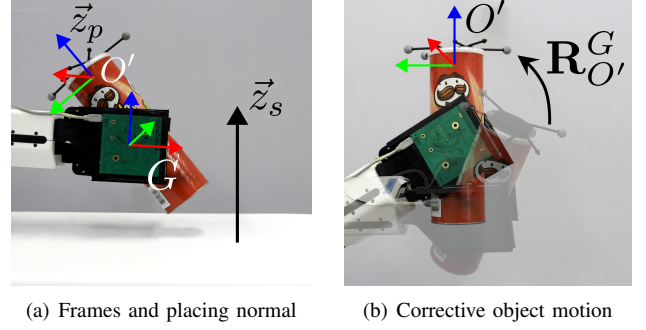


Fig. 3: Illustration of the problem setting. (a) shows the gripper frame G , a possible object placing frame O' and the surface’s placing normal \vec{z}_s . (b) Result of the corrective motion is to align the object with the placing surface based on the estimation of $\mathbf{R}_{O'}^G$.

responses, it transitions to the fourth phase, where it opens up the gripper and retracts.

B. Definition of pose correction action for stable placing

As illustrated in Fig. 3(a), we define the object’s placing normal $\vec{z}_p = \mathbf{R}_O^W \mathbf{R}_{O'}^O \vec{z} = \mathbf{R}_{O'}^W \vec{z}$ to lie along the \vec{z} -axis in the so-called local object placing frame O' with the rotation matrix from a reference frame, e.g. “world”, to the object placing frame given by $\mathbf{R}_{O'}^W \in SO(3)$. \mathbf{R}_O^W & $\mathbf{R}_{O'}^O$, thus, describe the rotation matrices from the world to the object frame, and from the object frame to the object’s placing frame, respectively. Note that we deliberately introduce this local object placing frame O' , in addition to the object’s pose frame O , for two reasons. On the one hand, there might exist multiple placing frames per object, and on the other hand, to highlight that knowing the object’s pose might not be informative enough for stable placing, for instance, in the case of lacking precise information about the object’s geometry. Since we only consider scenarios where the object was already grasped, $\mathbf{R}_{O'}^W$ can be decomposed into two distinct rotation matrices

$$\mathbf{R}_{O'}^W = \mathbf{R}_G^W \mathbf{R}_{O'}^G, \quad (1)$$

where $\mathbf{R}_G^W, \mathbf{R}_{O'}^G \in SO(3)$ describes the orientation of the gripper w.r.t. the world and the object’s placing frame within the gripper, respectively. While the former can be reliably estimated via forward kinematics from proprioceptive feedback, the latter is usually unknown. Here, we make the assumption, that one suitable placing face of the object is oriented toward the ground, which is generally the case after grasping or human-robot object handovers of various household objects. Thus, $\mathbf{R}_{O'}^G$ has to be estimated based on sensor measurements as it is an indispensable ingredient for generating a motion that corrects the object misalignment and enables appropriate placing.

C. Object reorientation and placing motion

Next, we require a placing controller that, given $\mathbf{R}_{O'}^G$, generates two movements: a corrective arm movement that aligns \vec{z}_p with \vec{z}_s (cf. Fig. 3(b)) and a downward placing motion that stops on table contact and releases the object

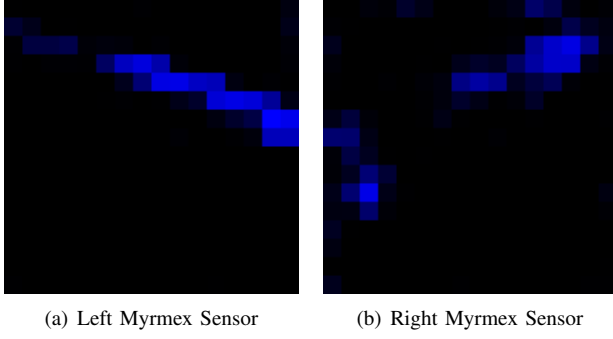


Fig. 4: Sensor response of both Myrmex sensors while holding the cylindrical object. Each reading of the 16×16 taxels corresponds to the currently measured normal force. The images were scaled up, and the force readings are amplified for visualization purposes.

afterward. For the corrective movement, we first project \bar{z}_s into the local gripper frame:

$$\bar{z}_s^G = \mathbf{R}_{W'}^G \bar{z} \quad (2)$$

Then, we calculate the rotation $\mathbf{R}_{G'}^G$ that rotates the current gripper frame such that in the resulting one (G'), $\bar{z}_p^{G'}$ aligns with $\bar{z}_s^{G'}$. This is achieved by finding the rotation axis with $\bar{a} = \bar{z}_p^G \times \bar{z}_s^G$ and the rotation angle $\theta = \cos^{-1}(\bar{z}_p^G \cdot \bar{z}_s^G)$, with the vector cross-product \times . Using the rotation $\mathbf{R}_{G'}^G$, we can try to find an inverse kinematics (IK) solution that realizes the object reorientation.

An IK solution that leads to a secure placing configuration should satisfy more constraints than the object reorientation alone. After reorientation, we will execute a linear downward movement in Cartesian space to acquire table-object contact, and we need to ensure that the object will be the first point of contact with the table. Arm configurations suitable for placing thus need to ensure that the object frame (located in the gripper) has a lower z -coordinate in the world frame than the wrist. The commonly used hierarchy-of-tasks approach [46] allows us to formulate such additional constraints for the IK. Since many IK solvers are sensitive to the initial arm configuration, we search for solutions starting from 20 different initial poses, which are sensible placing configurations and chose the one with the lowest error that is most similar to our current arm pose. The motion to reach this solution is generated by linear interpolation in joint space starting from the robot's current arm position.

For the placing motion, a trajectory is generated that moves the gripper linearly downward. This is realized using TIAGo's torso joint. During this motion, tactile measurements are continuously monitored. We interpret a spike in tactile sensation as making contact with the table. The torso controller is then signaled to halt execution. Finally, the object is released by opening the gripper and moving the torso upward again.

IV. OBJECT POSE CORRECTION ESTIMATION WITH TACTILE SENSING

As motivated previously, the key component for stable object placing is the estimation of the object placing frame

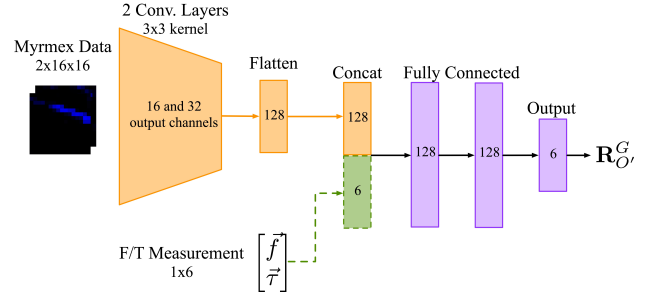


Fig. 5: Overview of our general neural network architecture used in all models. Tactile measurements from both sides of the gripper are stacked along the channel dimension of the convolutional layers. Depending on the model type, either the tactile feature embedding, the raw F/T data, or both concatenated are fed to the MLP. Finally, the network estimates the corrective rotation $\mathbf{R}_{O'}^G$.

w.r.t. the gripper frame ($\mathbf{R}_{O'}^G$). According to this transformation, the object is re-oriented prior to placing. However, determining this quantity is difficult, as the object is handed over in an unknown pose, it is occluded by the gripper, and herein we also do not assume any prior knowledge about the object type. We, therefore, propose estimating $\mathbf{R}_{O'}^G$ from the signals of the tactile sensors inside the gripper. The Myrmex tactile data is represented as a 16×16 matrix of normal force readings that are normalized in $[0, 1]$. Fig. 4 shows a visualization of the tactile readings from the right and left sensors while holding an object. Next, we, first, introduce our proposed Neural Network, and, subsequently, explain line-fitting baselines for recovering $\mathbf{R}_{O'}^G$ from the tactile readings.

Since we require a solution that is flexible enough to deal with different objects, that can handle the sensors' noise, and convert the high-dimensional readings into a signal suitable for reorienting the objects, we propose to employ a neural network. The general network architecture is visualized in Fig. 5. To process the tactile data, we first use two convolutional layers with a 3×3 kernel each, and 16 and 32 output channels respectively. The output of the last convolutional layer is then fed to a Multilayer Perceptron (MLP) consisting of two hidden layers with 128 neurons each and ReLU activation functions, followed by a dropout layer with a dropout probability of $p = 0.2$. F/T data can be optionally fed into the MLP as an additional input signal, which is concatenated with the tactile features.

To smoothly represent the rotation matrix in the output layer, we use the 6D representation comprising the first two columns of $\mathbf{R}_{O'}^G$, as introduced in [47] and has been shown to exhibit superior properties for learning in $SO(3)$. Each estimate is then converted into an $SO(3)$ rotation matrix for the computation of the loss, which is defined as

$$\mathcal{L}(\mathbf{R}, \bar{z}_p^{gt}) = \cos^{-1}(\mathbf{R}_G^W \mathbf{R} \bar{z} \cdot \bar{z}_p^{gt}), \quad (3)$$

where $\mathbf{R} = \mathbf{R}_{O'}^G$ is the quantity of interest and the prediction of the network, and \bar{z}_p^{gt} is the ground truth measurement of the object's placing normal in the world frame that is obtained through an OptiTrack motion capture system. By taking \cos^{-1} , the loss lies in the interval $[0, \pi]$, and can be interpreted as the angular distance error between predicting \bar{z}_p using the network's output \mathbf{R} and the ground truth.



(a) Objects used during data collection (left) and (b) Variations of in-hand object poses used during evaluation and out-of-distribution evaluation (right).

Fig. 6: Experiment objects and object poses used for evaluation and data collection.

TABLE I: Training results of networks with different input sensor modalities. We report the lowest test loss averaged over 10 batches.

	Tactile	F/T	Tactile + F/T
Test Loss (rad)	0.03	0.43	0.05

V. EXPERIMENTAL EVALUATION

For the experimental evaluation of all models, we use the TIAGo robot from PAL Robotics, equipped with a parallel-jaw gripper. To obtain tactile data, we exchange TIAGo’s non-sensorized fingers with Myrmex sensors that are fixed to the gripper via a 3D-printed adapter. This piezo-resistive sensor produces 16×16 normal force measurements with a readout frequency of 1kHz. Each sample is represented as a 2D matrix of normalized force measurements, where we transform the raw sensor readings from $[0, 4095]$ to lie in $[0, 1]$, where values of 1 refer to the maximum force that can be measured in a sensor cell (cf. Fig. 4).

A. Data collection and training

We used two primitive 3D-printed objects for data collection, i.e., a cylinder (2.25cm radius, 15cm length) and a cuboid ($5 \times 5 \times 19$ cm) as shown on the left in Fig. 6(a). To collect the ground-truth orientation that is required to train our networks, we use OptiTrack, an external, infrared camera-based marker tracking system. Thus, during data collection, markers were attached to both, robot and object.

For training-data generation, the objects were randomly positioned in the gripper and then grasped using the force controller from [48]. The robot’s arm was then manually brought into a realistic initial pose using gravity compensation mode. Then, a sample was collected by storing tactile data, F/T data, $\mathbf{R}_{O'}^G$, \mathbf{R}_G^W and \hat{z}_p^{gt} . For each arm pose, the object’s in-hand pose was varied 10 times within a range of 160° corresponding to plausible grasp poses. Afterward, the arm’s pose was changed, and the process was repeated. Using this method, we collected 800 samples per object and 1600 samples in total. We trained every neural net architecture for 40 epochs and reserved 20% of the data for testing. The network parameters were stored every time the running average of the test loss over the last 10 batches was lower than before.

Table I shows the lowest average test loss for all three architectures. The tactile-only model achieved the best result, although only marginally outperforming the tactile + F/T model with its test performance being only lower by 1.1° .

The F/T-only model performed significantly worse with an average angular error of 0.43 rad, indicating that it might not have learned the task sufficiently well.

B. Line-fitting baselines

To assess whether a data-driven model is required to solve this task robustly and to gauge its performance compared to other approaches, we introduce two baseline models for comparison. Given the nature of our sensory data and the goal to find the object’s main axis within it, line-fitting methods naturally come to mind. We chose two popular methods from that field, Principal Component Analysis (PCA) [49] and Hough transforms [50]. As both methods work on individual images, we combine the two sensor readings into one frame by flipping one of the sensor images to account for symmetry. Therefore, the input to the baselines contains the information of both sensors. This should increase robustness as the sensors might be differently affected by noise.

1) *PCA baseline*: We treat the force readings as a bi-variate, uni-modal Gaussian and estimate its mean, standard deviations, and covariance matrix \mathbf{C} . To obtain the orientation of the object’s main axis, we calculate the first principal component of \mathbf{C} using PCA. Assuming that an object’s main axis lies along its largest grasping surface, the first principal component should constitute a decent estimation for said axis. We then calculate the angle of the line relative to the sensor and transform it into the rotation $\mathbf{R}_{O'}^G$, allowing us to generate a corrective motion that aligns the objects.

2) *Hough transform baseline*: Hough transform is a common tool in image processing for finding lines in images. Typically, a raw input image undergoes some preprocessing where edges are extracted, and finally, the Hough transform is applied to the resulting binary image. Empirically, we have found the Hough transform to show better performance if we simply create a binary image by assigning a value of 1 to taxels with a force above a noise threshold and 0 otherwise. We only consider the resulting line with the most votes (the most confident estimate). Lines are parametrized by the angle ψ between the x -axis and a line normal that intersects the origin and the distance to the origin of said normal. From ψ , we calculate the angle between the x -axis and the line itself $\phi = \pi - \psi$ and again calculate $\mathbf{R}_{O'}^G$, as with the other baseline.

C. Experimental setup

For the real-world experimental evaluation of placing success, we let each method place an object from 5 different arm poses and 4 different in-hand object poses, yielding 20 samples per object for each method. Fig. 6(b) shows the Pringles object in four exemplary in-hand object poses that were used during evaluation. During the evaluation, we again attach markers to the objects to assess each method’s performance w.r.t. ground truth. Additionally, we execute the placing controller (cf. Sec. III-C) for each trial and note whether the object was placed successfully, i.e. did not fall or tip over after opening the gripper. Situations, where the object tilts but is still supported by the gripper, are also considered a failure.

TABLE II: Comparison of different placing methods on the two 3D-printed training objects (cf. Fig. 6(a)). The angular error is the angle between the predicted \hat{z}_p and the one measured with OptiTrack.

Method	Metric	Cylinder	Cuboid	Average
OptiTrack	% Suc. Ang.Err.	90% -	95% -	92.5% -
Tactile	% Suc. Ang.Err.	95% 0.06 ± 0.03	85% 0.17 ± 0.12	90.0% 0.11 ± 0.08
Tactile + F/T	% Suc. Ang.Err.	90% 0.08 ± 0.04	75% 0.16 ± 0.19	87.5% 0.12 ± 0.08
F/T	% Suc. Ang.Err.	15% 0.38 ± 0.26	25% 0.39 ± 0.32	20.0% 0.38 ± 0.29
PCA	% Suc. Ang.Err.	90% 0.07 ± 0.02	10% 0.83 ± 0.42	50% 0.45 ± 0.22
Hough	% Suc. Ang.Err.	80% 0.09 ± 0.04	10% 0.76 ± 0.37	45% 0.42 ± 0.21

D. Ablation over sensor modalities with seen objects

We first compare the OptiTrack baseline with all three networks and the two line-fitting baselines on the objects used for training. Thus, we evaluate 6 methods on the two 3D-printed training objects and conduct 20 trials per object, resulting in 240 placing trials in total. Table II shows the results. Note, that we consider the OptiTrack measurements as the baseline since it was not possible to obtain more precise object state estimations. We can compare the OptiTrack baseline to the other methods based on success rate and the angular error between their predictions of the object normal and OptiTrack’s ground-truth measurement for \hat{z}_p^G (c.f. eq. 3).

A surprising result is that the tactile-only model performed better in terms of a success rate than the OptiTrack baseline. This can be explained by OptiTrack not measuring the marker positions perfectly and accurately all the time. When the robot is blocking the sight of some cameras on the markers, this leads to a degradation in measurement accuracy or even an interruption of tracking. Additionally, the cylinder was quite top-heavy, which requires a relatively precise alignment with the table for it to not topple. For the cuboid object, which is not as challenging to place safely, the OptiTrack baseline performed better than the tactile model, with 95% vs. 85% success rate, respectively.

Among the neural networks, the tactile-only model performed best in almost all metrics. Only in terms of estimation accuracy for the cuboid object did the tactile+F/T model achieve higher scores, however only negligibly so by 0.01 rad. It also exhibited a higher variance in accuracy as compared to the purely tactile model, resulting in less reliable results in terms of success rate. The evaluation also confirmed another indication from training, namely that the F/T-only model did not succeed in estimating the object’s placing normal sufficiently well for stable placing. In most cases, it simply predicted the identity matrix resulting in no corrective movements. Following a similar argumentation as in [51], we posit that the object’s mass might be negligible compared to the actual weight of the end-effector. Thus, F/T signals alone do not seem to carry enough information for assessing the relative transform between gripper and the object’s placing normal.

PCA has comparable performance to the tactile-based neural networks for the cylindrical object and showed the lowest

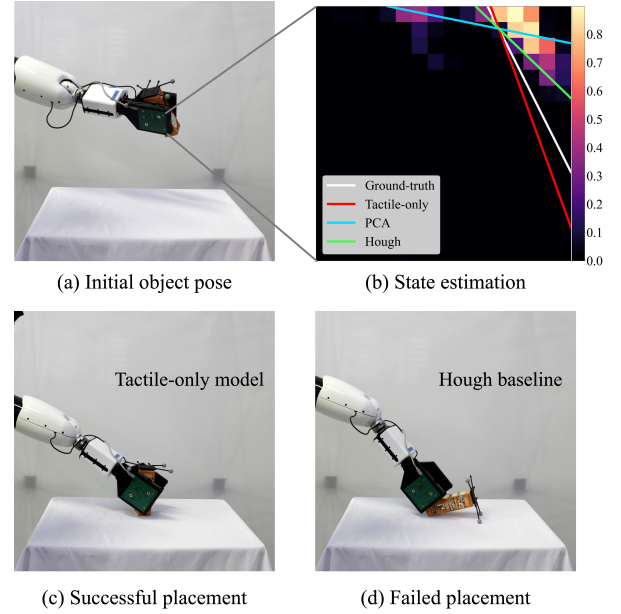


Fig. 7: Placement sequence comparing the tactile-only model to the two baselines. The network performed best with an angular error of 0.11 and placed the Tabasco correctly, while Hough and PCA both failed with estimation errors of 0.34 and 0.91 respectively. (b) depicts the force intensities for each taxel along with each model’s prediction of the object’s placing normal and the ground -truth obtained from OptiTrack.

variation there, with Hough transforms performing slightly worse, yet satisfactory. For the cuboid object, however, the performance decreased sharply for both classical methods. While cylindrical objects generate an easily identifiable, unique line in the sensor images, cuboid ones cover more contact areas and create less distinct patterns. The neural networks still appear to identify the relevant patterns robustly, while PCA and Hough heavily suffer from this fact. Different from the F/T-only network, the classical approaches do not involve a training phase and, thus, they cannot learn some structure of the data and generalize over them, as effectively done by our neural network approach. While the network internalized a range of plausible angle offsets, while learning to filter-out noise, PCA and Hough often produce severely wrong estimates, despite our preprocessing efforts, resulting in large errors and only a few successful placing trials.

E. Evaluation with previously unseen objects

After confirming the training results on the real robot in our first experiment, we conducted a second evaluation on objects that were not present in the training data. During this evaluation, we also attached markers to the objects to validate the prediction accuracy. We do not report prediction accuracy for the lipstick as attaching a marker ensemble would substantially alter its placing dynamics and, thus, report solely the success rate. We only evaluated the two best neural models from the previous experiment, namely the tactile-only and the tactile with F/T neural nets, along with the two classical approaches. We evaluated the 4 methods on 7 different household objects for 20 trials each, hence performing 560 placing trials in total. Two of the objects

TABLE III: Experimental results for five household testing objects (cf. Fig. 6(a)). All objects were unknown to the models before this evaluation, i.e., they were not present in the training set. Angular errors are reported in radians.

Model	Metric	Pringles	Glue Bottle	Tabasco	Mallow Pop	Cheez It	Shampoo	Lipstick	Average
Tactile	% Suc.	90%	85%	85%	80%	80%	85%	90%	85%
	Ang.Err.	0.07 ± 0.03	0.08 ± 0.04	0.10 ± 0.13	0.16 ± 0.10	0.10 ± 0.06	0.06 ± 0.03	-	0.09 ± 0.07
Tactile + F/T	% Suc.	85%	80%	70%	70%	85%	75%	80%	77%
	Ang.Err.	0.13 ± 0.20	0.09 ± 0.04	0.16 ± 0.10	0.14 ± 0.10	0.05 ± 0.06	0.10 ± 0.06	-	0.12 ± 0.10
PCA	% Suc.	90%	90%	15%	20%	80%	75%	85%	65%
	Ang.Err.	0.08 ± 0.06	0.07 ± 0.03	0.79 ± 0.61	0.70 ± 0.45	0.09 ± 0.05	0.08 ± 0.0	-	0.30 ± 0.24
Hough	% Suc.	85%	80%	60%	50%	70%	70%	80%	70%
	Ang.Err.	0.11 ± 0.06	0.10 ± 0.03	0.20 ± 0.30	0.35 ± 0.37	0.12 ± 0.08	0.16 ± 0.32	-	0.17 ± 0.17

were cylindrical (Glue Bottle & Pringles), three objects were box-like (Mallow Pop, Tabasco & Cheez-It), and the Lipstick was a small, elongated, rectangular object with rounded edges (see Fig. 6(a) on the right side).

The results of this evaluation are given in Table III. We, again, report the success rate of correct placements and the angular error between the model’s predictions and OptiTrack where possible. The network models perform very well on most unknown objects, indicating that our method generalizes across object primitives of unknown dimensions. Similar to the results from the previous evaluation on known objects, the tactile-only model showed superior performance in most cases, as it performed best on average in both metrics with a low variance in results. The tactile + F/T model only performed better in terms of both metrics when placing the Cheez-It box. We hypothesize that the slightly worse performance might result from the network receiving rather non-informative F/T signals alongside the tactile data.

The PCA baseline showed similar performance to the neural networks on cylindrical objects. As in the first experiment, its performance dropped sharply on box-like objects, which can likely be attributed to a less pronounced distribution of forces around the object’s main axis. Lastly, the Hough model performed well to mediocre on all evaluation objects. While it did not exhibit such a drastic performance drop on box-like objects as PCA, it did not manage to achieve satisfactory results consistently. Nevertheless, it performed better on box-like objects than on the cuboid test object. Paper boxes tend to create more distinct lines (one at each edge of the object) which are easier to detect with Hough transforms, whereas the much more rigid 3D-printed cuboid creates a larger contact patch, making it more difficult to detect lines. Fig. 7 shows placing trials of the box-like Tabasco object, where the tactile-only model successfully placed the object whereas both classical models’ predictions led to failed placements. Finally, all models showed reasonable to good performance on the Lipstick, an object that is difficult to place due to its small placing faces. The tactile-only network performed best among all methods, further underlining its generalizability.

F. Discussion

Our results show that tactile data is a crucial source of information for predicting the object’s placing normals, whereas F/T data has not been proven to be as informative for this task. Furthermore, our evaluation has shown that neural network models generalize to unknown objects with

high success rates and precision. It also revealed that classical approaches can work reliably on a set of objects with specific characteristics (cylinders) while not being accurate enough on objects that lack these (boxes). This leads us to the conclusion that, despite the effort of data collection, neural networks are the most sensible model choice for this task due to their high accuracy and robustness in real robotic experiments with unseen objects, of different sizes and weights than the ones used during training.

On the other hand, our evaluation also revealed some limitations of our approach. Our method cannot deal well with objects that uniformly cover the complete sensor surface. To circumvent this problem, the ability of in-hand object manipulation would be required to again recover more distinctive tactile feedback. Yet, such maneuvers are severely limited by our parallel gripper. Thus, we hypothesize that future work should aim at executing a larger variety of exploratory object interactions to acquire better tactile signals and enable robust object placing with an even wider range of objects. Another limitation arises from our assumption of downward-pointing placing faces. In situations where this assumption is likely to be violated, tactile data could be used to fine-tune higher-level object state predictions, e.g., from a vision system with less accuracy.

VI. CONCLUSIONS

In this work, we investigated the potential of employing tactile sensing for obtaining stable object placing policies. This task is especially challenging as we neither assume any prior knowledge about the object’s initial pose, nor any access to external camera-based sensing. We presented an approach based on deep learning that takes as input tactile or F/T data, or a combination of both, and predicts a rotation matrix that is supposed to align the object’s placing normal with that of the placing surface. We trained a representation through supervised learning with a small dataset of primitive objects. In the first part of the experimental evaluation, we validated our approach using the training objects and concluded that F/T sensing alone is not sufficient for this task. Our subsequent evaluation of previously unseen household objects showed an impressive generalization of our approach. We also compared the neural network models to classical, learning-free methods, and found that our proposed tactile-based neural policy shows higher performance in general and better generalization properties, justifying the choice of a learning-based method. In the future, we intend to investigate a more interactive method for stable object placing through active touch.

REFERENCES

- [1] E. Prassler, R. Bischoff, W. Burgard, R. Haschke, M. Hägele, G. Lawitzky, and M. Zöllner, "Towards service robots for everyday environments," *Springer Tracts in Advanced Robotics* 2012.
- [2] S. Dong and A. Rodriguez, "Tactile-based insertion for dense box-packing," in *Proc. IROS*. IEEE, 2019.
- [3] S. Dong, D. K. Jha, D. Romeres, S. Kim, D. Nikovski, and A. Rodriguez, "Tactile-rl for insertion: Generalization to objects of unknown geometry," in *Proc. ICRA*. IEEE, 2021.
- [4] S. Kim and A. Rodriguez, "Active extrinsic contact sensing: Application to general peg-in-hole insertion," in *Proc. ICRA*. IEEE, 2022.
- [5] H. Yousef, M. Boukallel, and K. Althoefer, "Tactile sensing for dexterous in-hand manipulation in robotics—a review," *Sensors and Actuators A: physical*, vol. 167, no. 2, 2011.
- [6] H. Van Hoof, T. Hermans, G. Neumann, and J. Peters, "Learning robot in-hand manipulation with tactile features," in *Proc. Humanoids*. IEEE, 2015.
- [7] B. Ward-Cherrier, N. Rojas, and N. F. Lepora, "Model-free precise in-hand manipulation with a 3d-printed tactile gripper," *IEEE Robotics and Automation Letters*, vol. 2, no. 4, 2017.
- [8] N. Funk, C. Schaff, R. Madan, T. Yoneda, J. U. De Jesus, J. Watson, E. K. Gordon, F. Widmaier, S. Bauer, S. S. Srinivasa *et al.*, "Benchmarking structured policies and policy optimization for real-world dexterous object manipulation," *IEEE Robotics and Automation Letters*, vol. 7, no. 1, pp. 478–485, 2021.
- [9] B. Belousov, B. Wibranek, J. Schneider, T. Schneider, G. Chaltatzaki, J. Peters, and O. Tessmann, "Robotic architectural assembly with tactile skills: Simulation and optimization," *Automation in Construction*, vol. 133, 2022.
- [10] N. Funk, G. Chaltatzaki, B. Belousov, and J. Peters, "Learn2assemble with structured representations and search for robotic architectural construction," in *Conference on Robot Learning*. PMLR, 2022, pp. 1401–1411.
- [11] N. Funk, S. Menzenbach, G. Chaltatzaki, and J. Peters, "Graph-based reinforcement learning meets mixed integer programs: An application to 3d robot assembly discovery," in *2022 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*. IEEE, 2022, pp. 10 215–10 222.
- [12] T. Bhattacharjee, A. Jain, S. Vaish, M. D. Killpack, and C. C. Kemp, "Tactile sensing over articulated joints with stretchable sensors," in *World Haptics Conference*. IEEE, 2013.
- [13] E. Dean-Leon, B. Pierce, F. Bergner, P. Mittendorfer, K. Ramirez-Amaro, W. Burger, and G. Cheng, "TOMM: Tactile omnidirectional mobile manipulator," in *Proc. ICRA*. IEEE, 2017.
- [14] D. Álvarez, M. A. Roa, and L. Moreno, "Visual and tactile fusion for estimating the pose of a grasped object," in *Iberian Robotics conference*. Springer, 2019.
- [15] M. B. Villalonga, A. Rodriguez, B. Lim, E. Valls, and T. Sechopoulos, "Tactile object pose estimation from the first touch with geometric contact rendering," in *Proc. CoRL*. 2021.
- [16] Z. Kappasov, J.-A. Corrales, and V. Perdereau, "Tactile sensing in dexterous robot hands — review," *RAS*, vol. 74.
- [17] B. Tise, "A compact high resolution piezoresistive digital tactile sensor," in *Proc. ICRA*, 1988.
- [18] C. Schürmann, R. Kōiva, and R. Haschke, "A Modular High-Speed Tactile Sensor for Human Manipulation Research." 2011 IEEE World Haptics Conference, 2011.
- [19] G. Büscher, R. Kōiva, C. Schürmann, R. Haschke, and H. Ritter, "Tactile dataglove with fabric-based sensors," in *Proc. Humanoids*, 2012.
- [20] F. Vidal-Verdú, Ó. Oballe-Peinado, J. A. Sánchez-Durán, J. Castellanos-Ramos, and R. Navas-González, "Three realizations and comparison of hardware for piezoresistive tactile sensors," *Sensors*.
- [21] O. Kerpa, K. Weiss, and H. Worn, "Development of a flexible tactile sensor system for a humanoid robot," in *Proc. IROS* 2003.
- [22] K. Kamiyama, H. Kajimoto, N. Kawakami, and S. Tachi, "Evaluation of a vision-based tactile sensor," in *Proc. ICRA* 2004.
- [23] B. Ward-Cherrier, N. Pestell, L. Cramphorn, B. Winstone, M. E. Giannaccini, J. Rossiter, and N. F. Lepora, "The tactip family: Soft optical tactile sensors with 3d-printed biomimetic morphologies," *Soft robotics*, vol. 5, no. 2, 2018.
- [24] 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 RAL*, vol. 5, no. 3, 2020.
- [25] E. Donlon, S. Dong, M. Liu, J. Li, E. Adelson, and A. Rodriguez, "Gelslim: A high-resolution, compact, robust, and calibrated tactile-sensing finger," in *Proc. IROS*. IEEE, 2018, pp. 1927–1934.
- [26] C. Sferrazza and R. D'Andrea, "Design, motivation and evaluation of a full-resolution optical tactile sensor," *Sensors*, 2019.
- [27] B. Romero, F. Veiga, and E. Adelson, "Soft, round, high resolution tactile fingertip sensors for dexterous robotic manipulation," in *Proc. ICRA*. IEEE, 2020.
- [28] Y. Jiang, M. Lim, C. Zheng, and A. Saxena, "Learning to place new objects in a scene," *The International Journal of Robotics Research*, vol. 31, no. 9, 2012.
- [29] K. Harada, T. Tsuji, K. Nagata, N. Yamanobe, H. Onda, T. Yoshimi, and Y. Kawai, "Object placement planner for robotic pick and place tasks," in *Proc. IROS*, 2012.
- [30] J. Ma, W. Wan, K. Harada, Q. Zhu, and H. Liu, "Regrasp planning using stable object poses supported by complex structures," *IEEE TR-CDS*, vol. 11, no. 2, 2018.
- [31] C. Mitash, R. Shome, B. Wen, A. Boularias, and K. Bekris, "Task-driven perception and manipulation for constrained placement of unknown objects," *IEEE RAL*, vol. 5, no. 4, 2020.
- [32] L. Manuelli, W. Gao, P. Florence, and R. Tedrake, "kpam: Keypoint affordances for category-level robotic manipulation," in *The International Symposium of Robotics Research*. Springer, 2019.
- [33] J. A. Haustein, K. Hang, J. Stork, and D. Kragic, "Object placement planning and optimization for robot manipulators," in *Proc. IROS*. IEEE.
- [34] R. Newbury, K. He, A. Cosgun, and T. Drummond, "Learning to place objects onto flat surfaces in upright orientations," *IEEE RAL*, vol. 6, no. 3, 2021.
- [35] J. Bimbo, S. Luo, K. Althoefer, and H. Liu, "In-hand object pose estimation using covariance-based tactile to geometry matching," *IEEE RAL*, vol. 1, no. 1, 2016.
- [36] D. Álvarez, M. A. Roa, and L. Moreno, "Tactile-based in-hand object pose estimation," in *Iberian Robotics conference*. Springer, 2017.
- [37] M. Pfanne, M. Chalon, F. Stulp, and A. Albu-Schäffer, "Fusing joint measurements and visual features for in-hand object pose estimation," *IEEE RAL*, vol. 3, no. 4, 2018.
- [38] S. Liu, H. Jiang, J. Xu, S. Liu, and X. Wang, "Semi-supervised 3d hand-object poses estimation with interactions in time," in *Proc. CVF*, 2021.
- [39] B. Doosti, S. Naha, M. Mirbagheri, and D. J. Crandall, "Hope-net: A graph-based model for hand-object pose estimation," in *Proc. CVF*, 2020.
- [40] P. Sodhi, E. Dexheimer, M. Mukadam, S. Anderson, and M. Kaess, "Leo: Learning energy-based models in factor graph optimization," in *CoRL*, 2022.
- [41] T. Anzai and K. Takahashi, "Deep gated multi-modal learning: In-hand object pose changes estimation using tactile and image data," in *IROS*. IEEE, 2020.
- [42] T. Kelestemur, R. Platt, and T. Padir, "Tactile pose estimation and policy learning for unknown object manipulation," *arXiv:2203.10685*, 2022.
- [43] R. Li, R. Platt, W. Yuan, A. ten Pas, N. Roscup, M. A. Srinivasan, and E. Adelson, "Localization and manipulation of small parts using gelsight tactile sensing," in *Proc. IROS*. IEEE, 2014.
- [44] S. Pirozzi and C. Natale, "Tactile-based manipulation of wires for switchgear assembly," *IEEE/ASME TR-M*, vol. 23, no. 6, 2018.
- [45] C. C. Beltran-Hernandez, D. Petit, I. G. Ramirez-Alpizar, and K. Harada, "Variable compliance control for robotic peg-in-hole assembly: A deep-reinforcement-learning approach," *Applied Sciences*, vol. 10, no. 19, p. 6923, 2020.
- [46] S. B. Slotine and B. Siciliano, "A general framework for managing multiple tasks in highly redundant robotic systems," in *proceeding of 5th International Conference on Advanced Robotics*.
- [47] Y. Zhou, C. Barnes, J. Lu, J. Yang, and H. Li, "On the continuity of rotation representations in neural networks," in *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 2019.
- [48] L. Lach, S. Lemaignan, F. Ferro, H. J. Ritter, and R. Haschke, "Bio-inspired grasping controller for sensorized 2-dof grippers," in *Proc. IROS* 2022, vol. 1. IEEE, 2022.

- [49] K. Pearson, "On lines and planes of closest fit to systems of points in space," *Philosophical Magazine*, vol. 2, no. 11, 1901.
- [50] R. O. Duda and P. E. Hart, "Use of the hough transformation to detect lines and curves in pictures," *Communications of the ACM*, vol. 15, no. 1, 1972.
- [51] Q. Li, O. Kroemer, Z. Su, F. F. Veiga, M. Kaboli, and H. J. Ritter, "A review of tactile information: Perception and action through touch," *IEEE T-RO*, vol. 36, no. 6, 2020.