Pivoting with Tactile Feedback

by

Pepijn Bogaard

to obtain the degree of Master of Science
at the Delft University of Technology,
to be defended publicly on Friday April 17, 2026 at 10:00 AM.

Student number:	4217489
Programme:	MSc Robotics
Supervisors:	Dr. M. Wiertlewski, G. Vitrani, MSc
Examination committee:	Dr. M. Wiertlewski, Dr. C. Della Santina, G. Vitrani, MSc

An electronic version of this thesis is available at http://repository.tudelft.nl/.

TU Delft

Pivoting with Tactile Feedback

Pepijn Bogaard
Delft University of Technology, The Netherlands

***Abstract*—Reliable robotic manipulation in everyday environments remains difficult because robots must interact with objects that differ in shape, weight, and frictional properties. Manipulation tasks in such unstructured environments often require more than stable grasping, including reorientation of the object within the grasp. A common example is pivoting, a reorientation task in which the object rotates within the fingers while translation is prevented. This requires information about contact forces and proximity to slip at the fingertip, which can be provided by tactile sensing. Recent work has shown that tactile data can be used to predict a frictional safety margin for slip caused by shear force. However, in tasks involving rotation, such as pivoting, slip depends on both shear force and torque. In this thesis, a generalized safety margin is introduced for pivoting by extending the safety margin estimation to also account for torque. This is done using the Limit Surface, which describes the boundary between sticking and slipping under combined force and torque. An experimental setup was built to measure slip under different force and torque combinations, from which the Limit Surface curves were fitted. Using those curves and the displacement field from the tactile sensor, deep learning models were then trained to predict the contact forces and the generalized safety margin. Validation on unseen data and experiments on grasp control and pivoting showed that tactile sensing can be used to predict and regulate friction during pivoting. This thesis extends the frictional safety margin approaches of previous studies from translational loading to pivoting. Estimating and regulating friction during pivoting could help robots perform a wider range of manipulation tasks in unstructured environments.**

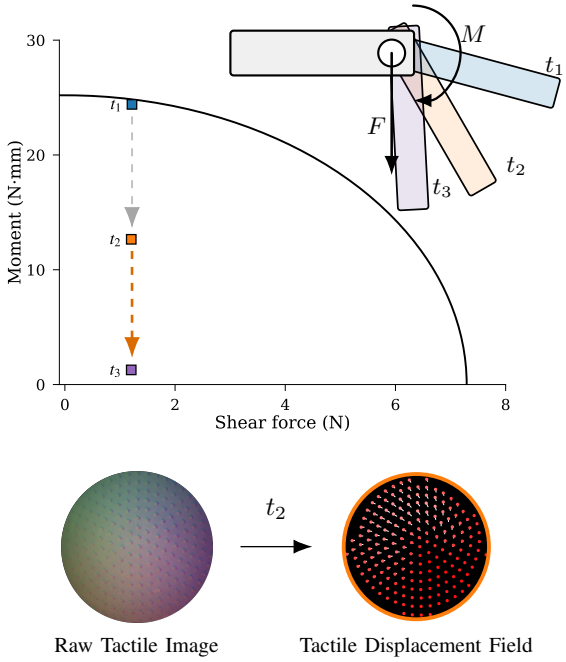

Fig. 1: Robotic gripper pivoting an object with gravity. The location on the Limit Surface is shown at t_1, t_2 and t_3. The raw tactile sensor image and the calculated displacement field is shown for t_2.

I. Introduction

Robots excel in structured environments such as manufacturing, where they execute repetitive tasks with high precision and reliability [1]. Their adoption has enabled more efficient production and reduced the need for humans to perform repetitive or dangerous work [2]. Despite this success, their deployment in the unstructured real world remains limited [3]. A key reason for this is that robots struggle to reliably handle diverse objects in such environments [1]. As a result, many manipulation tasks are still performed by humans. With growing workforce shortages in sectors such as healthcare and agriculture, this dependence on human labor is becoming increasingly difficult to sustain [4], [5]. Therefore, there is an increased need for robotic systems capable of reliable manipulation in everyday environments.

Achieving such reliable manipulation remains difficult because robots must interact with objects that vary in shape, weight, and frictional properties. In addition, many tasks require manipulation beyond stable grasping, such as reorientation. Consider a household robot loading a dishwasher. This mundane task for humans is quite challenging for a robot. It must handle delicate glasses, plates, or metal cutlery and place them in the correct position. This reorientation task involves controlled in-hand manipulation, where the object must move within the hand without being dropped. To achieve this, the robot must be able to sense and interpret what happens at the fingertip during manipulation. In particular, it must estimate contact forces and the margin to slip. Tactile sensing can provide this type of information. A promising solution is the use of vision-based tactile sensors. Using an embedded camera, these sensors observe the deformation of the soft surface of the robotic fingertip. By tracking markers on the surface, they can estimate contact forces and detect slip. Their relatively low cost and high resolution make them attractive for robotic manipulation [6].

Existing tactile sensing based solutions can already estimate contact forces, detect slip, and in recent work, predict a grasp safety margin [7], [8]. However, most of these methods focus on translational slip caused by shear force and do not model the role of torque. This becomes important for reorientation tasks such as pivoting, where slip is controlled to allow rotation but not translation. Although prior work has studied

pivoting with tactile sensing, a method to quantify grasp stability under a combination of shear force and torque is still missing [9], [10]. A formulation is therefore needed that can describe proximity to slip under these conditions.

This thesis addresses this problem by extending the concept of the Limit Surface to formulate a generalized safety margin for motions involving both translation and rotation. To achieve this, the ShadowTac tactile sensor is used because it can track in-plane deformation and also estimate the indentation caused by the grip force [8]. Furthermore, the dimple tracking method directly provides a displacement field without relying on computationally intensive convolutional neural networks. This displacement field is then processed using a PointNet-inspired learning model adapted for tactile data [11].

A dedicated experimental setup was developed to collect data under varying normal force, shear force, and torque, allowing slip conditions to be identified and the Limit Surface to be fitted. Based on this data, the adapted machine learning model was trained to predict a generalized safety margin that accounts for both translational and rotational slip. The approach was validated through experiments demonstrating stable grasping under torsional loading and controlled pivoting motions.

The trained model predicts both contact forces and the generalized safety margin directly from tactile data. The results show that tactile sensing can be used not only to detect slip, but also to regulate friction during pivoting tasks. This enables grasp control under varying loads during rotational motion and supports the use of controlled slip during manipulation.

The ability to estimate and regulate friction plays an important role in robotic manipulation. By allowing robots to sense how close a contact is to slipping under combined force and torque, this work enables more reliable and flexible object manipulation in unstructured environments. In particular, capabilities such as in-hand manipulation and pivoting can help robots perform a wider range of tasks in the real world. In this way, advances in tactile-based manipulation could help robots have a similar impact in the unstructured real world as they already have in structured industrial settings.

II. Related work

A. Safety margins in Tactile Sensing

When robots interact with objects, they must regulate their grip force carefully. The grip must be strong enough to prevent unintended slip, but not so large that it damages the object. Reliable manipulation therefore requires a measure of how close the contact is to slipping. The frictional safety margin provides such a measure.

Boonstra et al. [7] studied how tactile sensing can be used to estimate the frictional safety margin. For this purpose, they developed a vision-based tactile sensor called the Chroma-Touch. The sensor tracks dimple displacement to estimate in-plane deformation and uses color changes of the dimples to estimate out-of-plane deformation. Using the tactile images, they trained a convolutional neural network to predict the frictional safety margin directly. This prediction was combined with a grip-force controller to demonstrate stable grasping under increasing shear loads. The results show that tactile sensing can be used to regulate grasp stability through the estimated frictional safety margin.

Vitrani et al. [8] later introduced another vision-based tactile sensor, called the ShadowTac. Instead of relying on color changes, ShadowTac uses photometric stereo with three colored LEDs to reconstruct the displacement field. Using this displacement field, they trained a simpler MLP network to estimate the safety margin. Compared with raw tactile images, the displacement field allows simpler models than CNN-based approaches. They also demonstrated stable grasping using the predicted safety margin.

Together, both studies show that tactile sensing can be used to estimate the frictional safety margin and to support stable grasping. However, both works focus on grasp stability under translational loading. They do not address tasks in which slip depends on a combination of shear and torque, such as pivoting.

B. Pivoting with Friction

In addition to stable grasping, many manipulation tasks require the object to be reoriented within the hand. A common example is pivoting, where the grasped object rotates relative to the fingers while contact is maintained. Pivoting therefore depends strongly on friction, since the contact must allow controlled rotation, but prevent translation. Tactile sensing is well suited to such tasks, because it can provide information about the contact state during motion.

Viña et al. [9] studied the pivoting task using a two-finger robotic gripper with tactile sensors and visual feedback. They combined grip-force control with this feedback to regulate the pivoting motion. Using this approach, they achieved a desired pivot velocity while maintaining the grasp. However, their method does not use the concept of a frictional safety margin and still relies partly on vision.

Costanzo et al. [10] also studied the pivoting task using only tactile sensing and no visual feedback. To model the frictional contact during pivoting, they used the concept of the Limit Surface. This describes the slip boundary under a combination of force and torque. Based on the Limit Surface and additional modeling, they constructed an observer to estimate the pivot velocity from the contact forces. Using this observer, they controlled the pivoting motion while maintaining a stable grasp. Slip was prevented by regulating the slip velocity close to zero. However, their method does not define a frictional safety margin, and grasp stability is instead enforced indirectly through velocity control.

Together, these studies show that tactile sensing can be used to control pivoting tasks. However, neither approach defines a frictional safety margin for pivoting. As a result, they do not provide a direct measure of proximity to slip during such motions.

C. Concept of the Limit Surface

To characterize the slip boundary under combined force and torque, prior work introduced the concept of the Limit Surface.

The Limit Surface defines the boundary between stick and slip. It therefore provides a way to describe when a contact remains stable and when slip begins. This makes it relevant for manipulation tasks such as pivoting, where both force and torque act at the contact.

An early study of robotic finger contacts under combined force and torque was presented by Jameson [12]. This work was later refined by Howe et al. [13] in collaboration with Cutkosky. They studied the frictional behavior of a hemispherical finger and found that the slip boundary in the force-torque plane has an approximately elliptical shape. This boundary was later formalized in the literature as the Limit Surface.

In summary, prior work has shown that the Limit Surface is a suitable way to describe the slip boundary under combined force and torque. This makes it relevant for manipulation tasks that involve translation and rotation. However, prior work has not combined the concept with tactile safety margin estimation. A tactile method to quantify proximity to slip during such motions is still missing.

III. Methods

A. Generalized safety margin

When studying slip in a robotic gripper with two soft hemispherical fingertips, the most commonly used model is the Coulomb dry friction model. It states that the maximum tangential force the contact can withstand before slip occurs is given by [14]:

$$F_{xy}^{\text{slip}} = F_{xy}^{0} + \mu F_z \tag{1}$$

where μ is the dry friction coefficient, F_z is the applied grip force and F_{xy}^{0} captures the small amount of friction present at zero grip force. For this type of contact, which assumes pure translation, there is also a concept that quantifies how far the grasp operates from the onset of slip. This is called the frictional safety margin and is given by: [15]:

$$\text{SM}(t) = \frac{F_{xy}^{\text{slip}} - F_{xy}(t)}{F_{xy}^{\text{slip}}} \tag{2}$$

However, in real-world applications, combined tangential forces and moments are much more common than pure tangential loads or pure torques.

Assuming a Hertzian contact, given by [16]:

$$a = \left(\frac{3}{4}\frac{R}{E^*}F_z\right)^{\frac{1}{3}} \tag{3}$$

the maximum moment before slip occurs is given by [16]

$$M_z^{\text{slip}} = \mu\frac{3\pi}{16}aF_z \tag{4}$$

These two equations assume that the onset of slip in translation is independent of the applied moment, and that the onset of slip in rotation is independent of the applied tangential force. However, as shown by Jameson [12] and later by Howe [13], the onset of slip depends on a combination of torque and shear force.

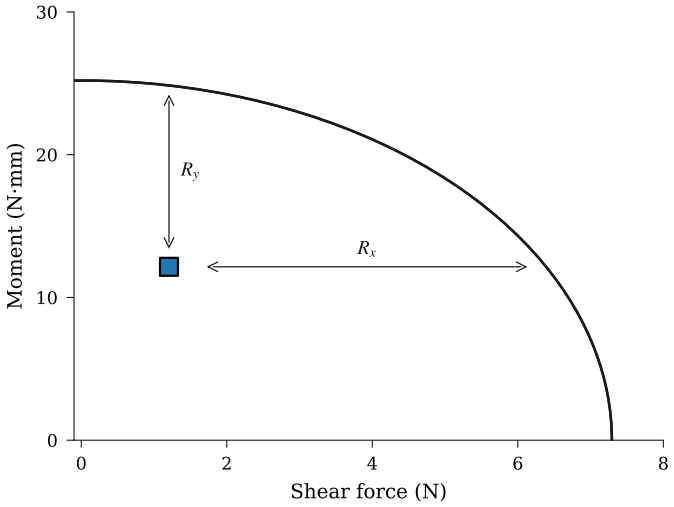

Fig. 2: Limit surface in the $(F_{xy}, |M_z|)$ plane at constant normal force F_z. Generalized safety margin denoted by R_y and R_x.

They introduced the concept of the Limit Surface, which defines the boundary between dry friction and slip (or viscous friction) for a given grip force F_z, plotted in the (F_{xy}, M_z) plane. Using the Limit Surface, it is possible to predict under which combination of grip force, tangential force, and moment a contact will transition from dry friction to viscous friction.

This thesis extends the concept of the frictional safety margin by adding the rotational component. Together, these form the generalized safety margin (GSM), which is given by

$$\text{GSM}(t) = \begin{bmatrix} R_y \\ R_x \end{bmatrix} = \begin{bmatrix} \left(M_z^{\max} - M(t)\right)/M_z^{\max} \\ \left(F_{xy}^{\max} - F_{xy}(t)\right)/F_{xy}^{\max} \end{bmatrix} \tag{5}$$

For a given contact, the GSM is shown in Fig. 2. Unlike the frictional safety margin, this extended version can also be used for contacts undergoing combined translational and rotational motion.

B. Experimental Setup

To use the generalized safety margin, the Limit Surface first had to be fitted experimentally. Therefore, an acquisition setup was created to measure slip under different combinations of normal force, tangential force, and torque. The acquisition setup used can be seen in Fig. 3.

The setup consists of a tactile sensor (Section III-E) mounted on a stepper motor. The stepper motor generates a rotational displacement, which results in a torque. It is mounted on a horizontal linear guide. Using a pulley, a weight can be attached to apply a shear force. Above the ShadowTac sensor, a 6-axis force sensor, the Mini40 is mounted. It is attached to two vertical linear guides. The indentor used was attached to the force sensor. For this study, a flat piece of Perspex plastic was used under dry conditions.

Weights could be placed on top of the force sensor to apply a desired normal force. For the normal force, eight different

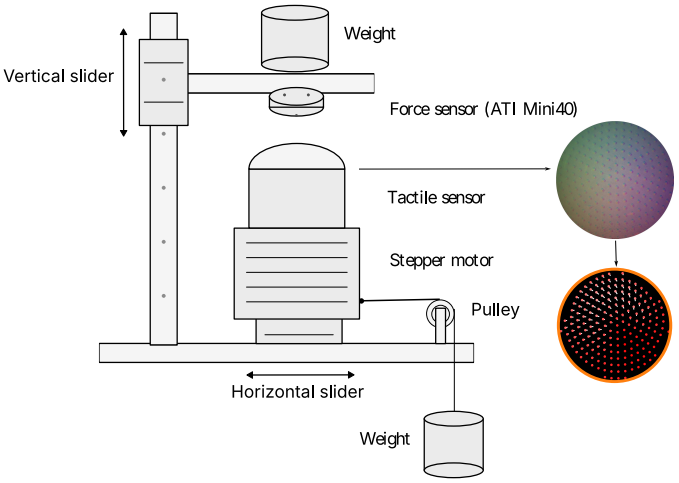

Fig. 3: Overview of the acquisition setup.

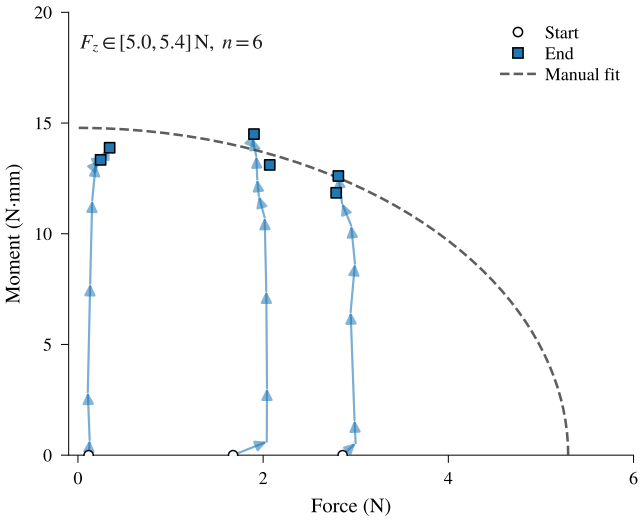

Fig. 4: Limit surface with fitted ellipse.

levels were applied, from 0 to 700 g, resulting in 2 N to 9 N when the weight of the force sensor and frame itself is included. For the tangential forces seven discrete levels from 0 to 600 g, in steps of 100 g, were used. This resulted in a force range of 0 to 6 N. For rotation, the stepper motor was driven at 7 steps per second, with a step size of 1/8 step, for 10 seconds. Assuming 200 steps per 360°, this resulted in a total rotation of 15.75°. The forces were recorded at 5000 Hz. The images were recorded at 50 Hz. Each acquisition lasted 15 seconds. In total, 120 acquisitions were performed. For each combination of F_z, F_{xy} and M_z, three acquisitions were performed: two in the counterclockwise direction and, due to time constraints, one in the clockwise direction.

C. Data Processing

The forces were plotted over time. Since F_z and F_{xy} were kept constant while the rotation was increased until slip occurred, M_z was used to identify the exact onset of slip. For each unique combination of F_z, F_xy, three data runs were recorded, referred to as triplets. The triplets were analyzed together to make it easier to identify whether one of the runs contained incorrect data or an outlier.

D. Limit Surface Fitting

Using the identified slip points, two methods were used to construct the Limit Surface for each grip force. The first method was to fit ellipses visually. For each discrete F_z, one ellipse was fitted. One challenge in the data was friction in the vertical sliders. As a result, F_z formed a distribution rather than discrete values. To address this, the forces were grouped into eight bins, corresponding to eight discrete F_z levels used in the experiments. For each discrete F_z, the identified slip points were plotted in the $(F_{xy}, |M_z|)$ plane, and an ellipse was fitted visually, as shown in Fig. 4. This figure also shows three acquisitions from start to slip, indicated by the arrows.

The second method was to fit a quadratic surface using linear regression. The fitted surface is given by:

$$z = c_0 + c_1 x^2 + c_2 y^2 \tag{6}$$

E. Tactile Image Processing

With the Limit Surfaces experimentally fitted and the GSM defined as a measure of slip safety, it becomes possible to train a model that predicts this safety margin from tactile input. For this thesis, a ShadowTac tactile sensor was manufactured using the procedure described by Vitrani et al [8]. ShadowTac is a vision-based tactile sensor that measures deformation by tracking a dimple pattern in a silicone dome. The sensor is shown in Fig. 5. The dimples are illuminated by three LEDs positioned on different sides, resulting in multicolored shadows. By tracking these shadows, a displacement field can be obtained. This displacement field changes with the deformation of the silicon dome. An additional benefit of using three colors is that photometric stereo can also be used to estimate the indentation of the silicon dome. The ability to track both the in-plane deformation caused by rototranslational motion and the indentation caused by grip force was the reason for choosing ShadowTac in this thesis. Furthermore, the dimple-tracking code developed by Vitrani et al. directly provides a displacement field without requiring high computational resources like convolutional neural networks.

The tactile images were processed using the dimple-tracking pipeline from the existing ShadowTac codebase. For each data run, the first captured image was used as the initialization image. At that moment, the sensor was not yet in contact with an object. This state was therefore used as the reference for zero displacement and zero indentation. Subsequent images were processed to track the dimples and reconstruct a displacement field and surface height using photometric stereo. For the sensor manufactured in this thesis, each image contains $N = 196$ trackable dimples. For each dimple i, the state at time t is defined as:

$$\mathbf{s}_i^t = \begin{bmatrix} x_i^0 & y_i^0 & x_i^t & y_i^t & \Delta h_i^t \end{bmatrix} \in \mathbb{R}^{1\times 5}. \tag{7}$$

where (x_i^0, y_i^0) is the initial coordinate of the i-th dimple,

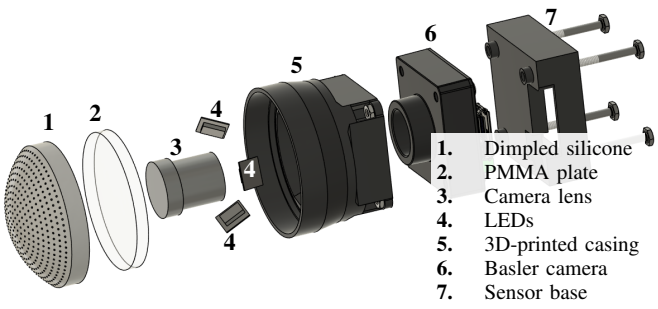

Fig. 5: Exploded view of the tactile sensor [8].

(x_i^t, y_i^t) represents the current coordinates at time t, and Δh_i^t denotes the relative normal displacement.

The tactile state of a single image is obtained by stacking the individual dimple states row-wise, resulting in the tactile state matrix at time t:

$$\mathbf{S}^t = \begin{bmatrix} \mathbf{s}_1^t \\ \mathbf{s}_2^t \\ \vdots \\ \mathbf{s}_{196}^t \end{bmatrix} \in \mathbb{R}^{196\times 5}. \tag{8}$$

A visualization of the processing pipeline is shown in Fig. 6. Encoding the deformation of the tactile sensor in this matrix form makes it possible to develop models that predict force and slip.

F. Force and Slip prediction

Using the processed tactile data, two machine learning models were trained. The first model predicts the contact forces: F_z, F_{xy}, and M_z. The second model predicts the generalized safety margin: R_y and R_x.

Why Two Separate Models: There are several reasons for training two independent models. First, the ground truth for the generalized safety margin is computed from the measured forces and the fitted Limit Surface. As a result, the safety margin labels already contain the modeling errors induced by the ellipse-fitting procedure. If only a GSM model were evaluated, its predictions would be compared against a quantity that already contains those errors. By training a separate force model, the PointNet-inspired architecture can also be evaluated directly against the raw force measurements. This makes the model performance easier to interpret, since the comparison is made with a known physical quantity.

Second, it is useful to have a force prediction model in addition to a slip prediction model. When slip occurs, the rotational equations of motion for the grasped object are [9]:

$$I\ddot{\theta} = \tau_g(\theta) + \tau_c \tag{9}$$

where I is the object inertia, τ_g the gravitational torque, and τ_c the contact torque generated by the gripper. From the acceleration, the velocity and displacement can be obtained. The force model is therefore useful for tasks like motion planning.

Furthermore, the GSM model serves a different purpose. The grip force must be adjusted to reach the state of full sliding. For this, it is useful to predict how close the contact is to slip, which is what the GSM model provides. Alternatively, the GSM could be computed from the predicted forces using the Limit Surface. However, for high-speed control loops, this would be computationally expensive. It is therefore preferable to predict the GSM directly.

Finally, the most important reason is generalization. Slip does not occur uniformly across the contact. From the literature, it is known that slip first occurs at the outer regions of the contact, before propagating inward. A model trained directly on the GSM has the potential to learn such local deformation patterns. With sufficient data across different materials and shapes, the model may learn to detect incipient slip beyond what is a combination of global forces with a manually fitted Limit Surface.

For these reasons, two models were trained independently. The force model was included to allow a direct physical evaluation against measured forces. The GSM model was trained to predict the onset of slip directly, with the potential for generalization.

Tactile Input: Both models use the same tactile input. The tactile sensor gives a set of tracked marker positions, also referred to as a point set. To process this set of points, a network inspired by PointNet was used [11]. PointNet is a deep neural network architecture designed for point clouds. It uses multiple multilayer perceptrons (MLPs) followed by max pooling. Initially, a standard MLP was tested in this thesis. However, unlike an MLP, PointNet is invariant to the order of the input points. For the displacement field used here, this property is beneficial because it makes data augmentation operations such as rotation easier.

Dataset Construction: In total, 40 different combinations of normal force and shear force were recorded. Each combination consisted of two counterclockwise rotations and one clockwise rotation. To obtain a balanced dataset and ensure unbiased validation, one complete counterclockwise rotation run from each combination was reserved exclusively for validation and was never seen during training.

The dataset contained $N = 23,901$ tactile images. For each image n, a feature matrix was constructed as defined in Eq. 7. Stacking all samples results in the input tensor

$$\mathbf{X} \in \mathbb{R}^{23901\times 196\times 5}. \tag{10}$$

Here, 196 corresponds to the maximum number of tracked markers, and each marker is represented by a 5-dimensional feature vector. Due to noise and deformation, not all images contained exactly 196 markers. Samples with fewer tracked markers were padded with zeros to maintain a consistent input size.

The same input tensor $\mathbf{X}$ was used for both prediction models.

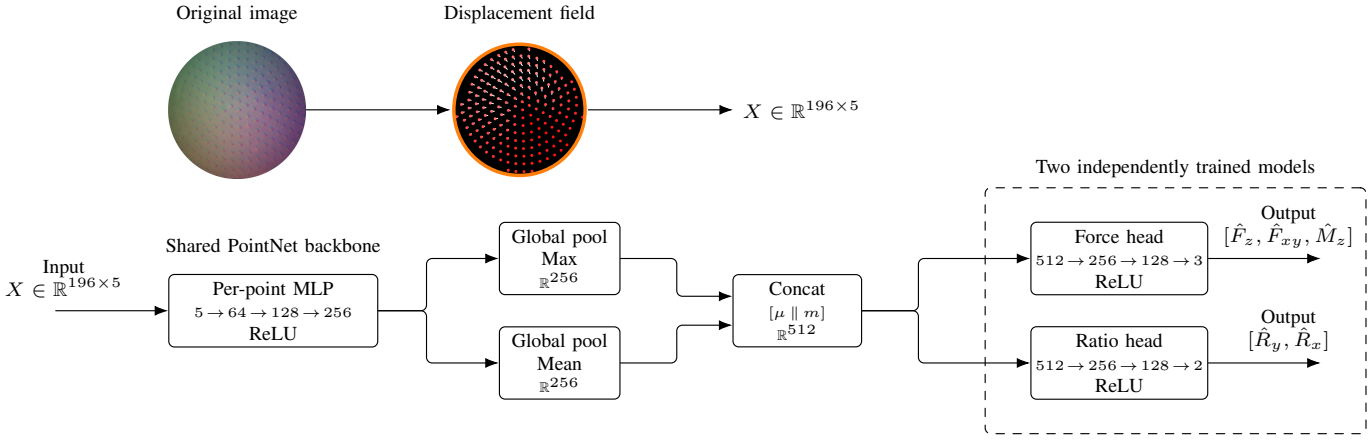

Fig. 6: Tactile processing and PointNet based ML architecture.

Ground Truth Construction: For the force prediction model, the ground-truth vector was constructed directly from the recorded force-torque measurements:

$$\mathbf{y} \in \mathbb{R}^{23901\times3}, \qquad \mathbf{y} = \begin{bmatrix} |F_z| & |F_{xy}| & |M_z| \end{bmatrix}. \tag{11}$$

For the safety margin prediction, the ground-truth values required an additional computation step. Using the fitted limit surface curves, the generalized safety margin was calculated for each sample. The resulting target vector is then given by:

$$\mathbf{y} \in \mathbb{R}^{23901\times2}, \qquad \mathbf{y} = \begin{bmatrix} R_y & R_x \end{bmatrix}. \tag{12}$$

The dataset was randomly split into a training and validation set using an 80/20 ratio.

Network Architecture: The network used in this work is inspired by the PointNet architecture. However, the original PointNet architecture was developed for point-cloud classification. To predict continuous forces and safety margins, the architecture was changed to use regression instead. Several other modifications were introduced:

- PointNet uses a mini-network called T-net to align the point clouds such that the network is invariant to geometric transformations. The dimples in the tactile sensor are already aligned on a fixed surface, therefore this component was not used.
- PointNet only uses max pooling. To capture both small localized peaks and overall contact deformations a combination of max pooling and average pooling was applied.
- PointNet uses 3D point coordinates. In the proposed model, however, each marker is represented by a 5-dimensional feature vector instead.

G. Experimental Evaluation

To evaluate the trained machine learning models, three experiments were done, as shown in Fig. 7.

The first experiment evaluated how well the slip prediction model could maintain a safe grasp under increased load.

The second experiment tested whether the trained model could be used to pivot an object without dropping it and how it performed across different objects. Four objects were selected: a protractor triangle, a box cutter, a plastic bottle, and a smartphone. The experiment started by gripping the object at 90° and controlling for a safe grasp using the predicted vertical safety margin, $\hat{R}_y$. After 10 seconds, the grip force was gradually decreased over 20 seconds, allowing the object to pivot.

The final experiment was to perform a demonstration of how

Fig. 7: Overview of the three experimental setups. (1) Object grasping and pivoting [17]. (2) Control stable grasp under increased weight. (3) Rotation experiment with free rotational degree of freedom.

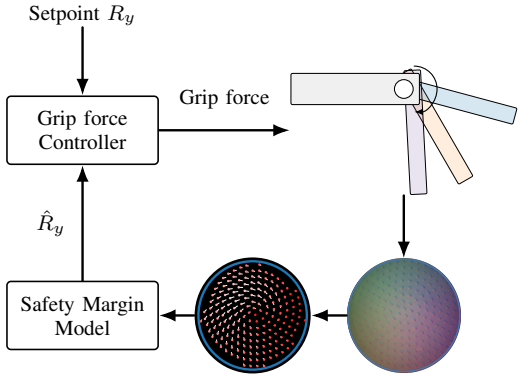

Fig. 8: Closed-loop grip force control using tactile sensing.

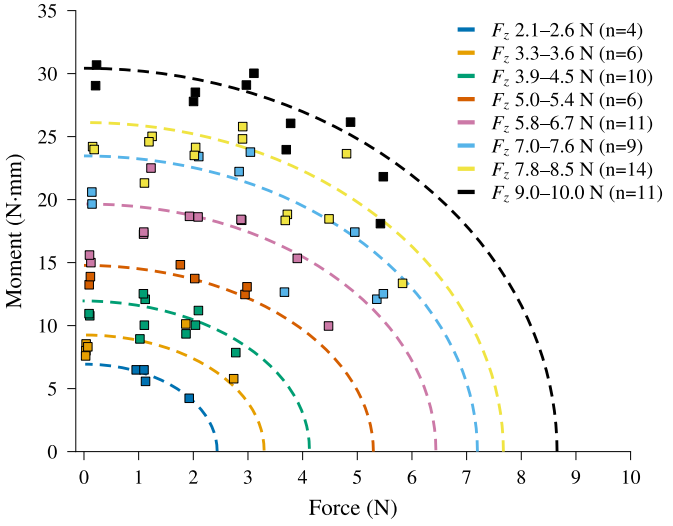

Fig. 9: Visually fitted limit surface curves using identified slip points (n = 71) visualized in the $(F_{xy}, |M_z|)$ plane. Achieved results for $|M_z|$ are an MAE of 1.77 Nmm, RMSE of 2.32 Nmm and mean absolute relative error of 10.20% ($P_{90} = 21.39\%$).

the model could be used in a more practical setting. When a robot picks up an object, it can be useful for one degree of freedom to remain free, while another is constrained. This was demonstrated in the final experiment. The robot grasps an object and then keeps the rotational degree of freedom free by rotating the robot gripper while keeping the object vertical.

H. Grip controller

Each experiment required a controller to maintain a setpoint for the vertical safety margin, R_y. The control diagram is shown in Fig. 8. The sensor hardware and software convert the tactile information into a displacement field. The ML model uses the displacement field to predict the vertical safety margin, $\hat{R}_y$. The grip force controller compares this prediction with the setpoint and adjusts the grip force accordingly. The controller used was a PI controller with a gain, K_p of 2500, an integrator, K_i of 1000, a dead-band of 0.01, an integrator clamp of 3.0, and a low-pass filter with a time constant of 0.02 on the input. The control parameters were determined heuristically.

IV. Results

A. Limit Surface Fitting Results

Of the 120 runs, 71 were used to fit the Limit Surface curves, because runs with clockwise rotation showed very different slip points.

The visually fitted curves are shown in Fig. 9. The plot also shows a clear gap in the data. The points with low torque are missing. The curves were also fitted using a quadratic surface, as shown in Fig. 14, resulting in a mean absolute relative error of 11.81% and P_{90} of 22.49%. This is slightly worse than the visually fitted curves.

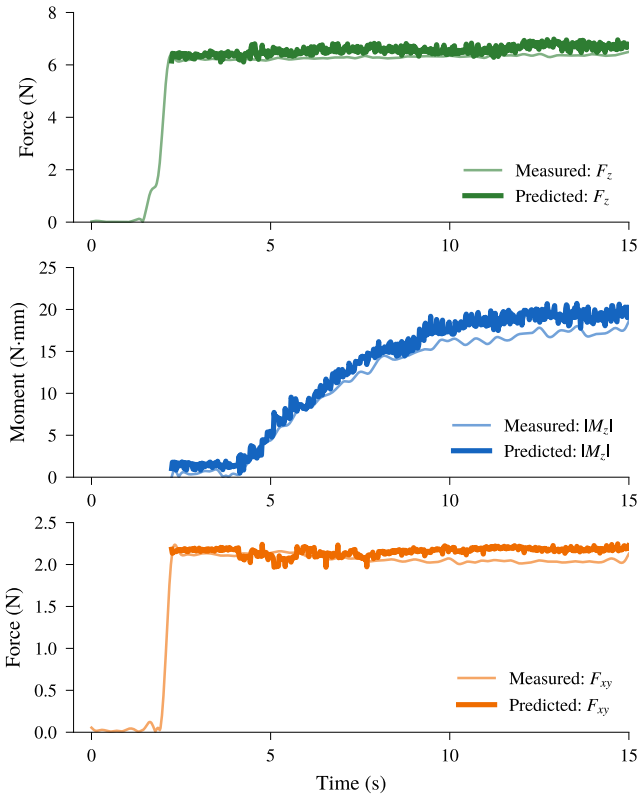

Fig. 10: Measured and predicted contact forces over time for an unseen validation trial. Using the displacement field from the tactile sensor, the model predicts the normal force F_z, shear force F_{xy}, and torque $|M_z|$.

B. Force and Slip prediction Performance

The force and slip prediction models were evaluated on $n = 39$ unseen validation runs. These runs used the same flat Perspex indentor and measurement setup as the training dataset. A representative time-series prediction for the contact forces is shown in Fig. 10. The first plot shows the measured normal force F_z over time together with the predicted F_z. The second subplot shows the moment, and the third shows the shear force. Across all unseen validation runs, the models achieved a root mean squared error (RMSE) of 0.59 N for F_z, 0.44 N for F_{xy}, and 2.46 N mm for M_z. The corresponding mean absolute errors (MAE) were 0.46 N, 0.29 N, and 1.78 N mm.

Fig. 11 shows a representative time-series prediction of the generalized safety margin. The first plot shows the vertical safety margin, and the second subplot the horizontal safety margin. Across all unseen validation runs, the RMSE for R_y was 0.11 and the MAE was 0.086. For R_x, the RMSE was 0.13 and the MAE was 0.091.

C. Safe Grasp Under Increased Load

A single trial of the safe grasp experiment is shown in Fig. 12. First, the cup is grasped with a minimum grip force

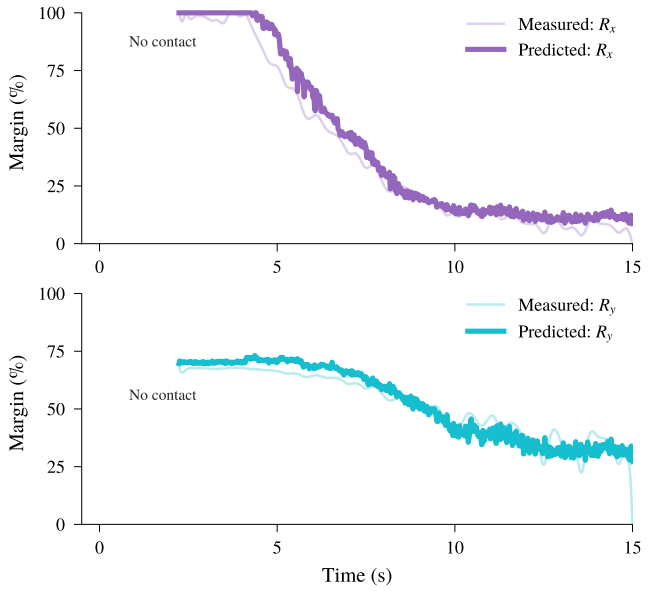

Fig. 11: Measured and predicted safety margins over time for an unseen validation trial. Using the displacement field from the tactile sensor, the model predicts the vertical safety margin R_y and the horizontal safety margin R_x.

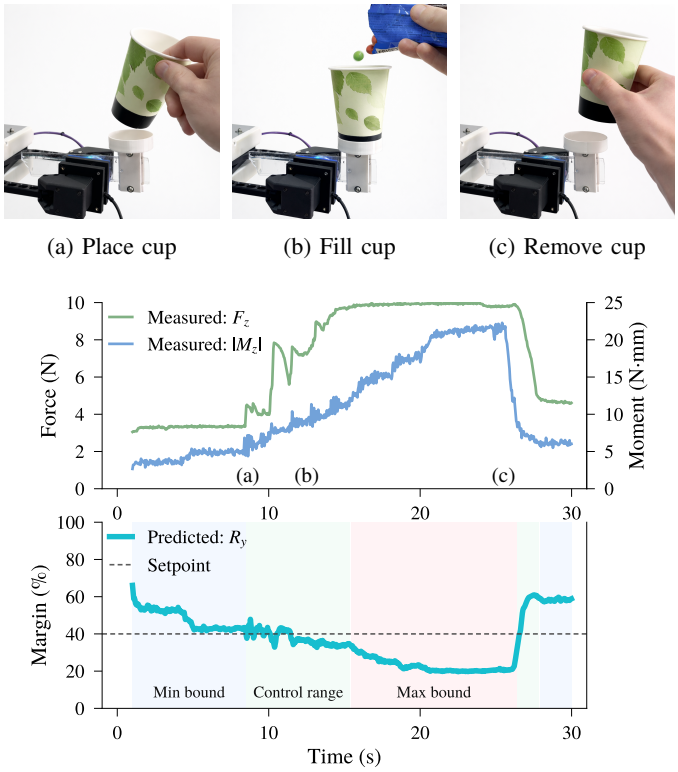

Fig. 12: Measured forces and predicted safety margin during the safe grasp under increased load.

of around $F_z = 3\,\mathrm{N}$. The cup is then placed and filled with chocolate candies, increasing the contact moment from around $M_z = 2.5\,\mathrm{Nmm}$ to $M_z = 20\,\mathrm{Nmm}$. Initially the grip force controller tries to maintain a setpoint of 40% for the vertical safety margin, R_y, but it quickly saturates and reaches the maximum grip force for this experiment, which is around $F_z = 10\,\mathrm{N}$. As the load continues to increase, the predicted safety margin drops below the setpoint to around 20%. At the end of the experiment, the cup is removed, and the grip force controller decreases the grasp force.

D. Pivoting Results for Unseen Objects

The next experiment evaluated pivoting with different objects. A representative data run with a smartphone as the grasped object is shown in Fig. 13. First, the gripper grasps the object and controls for a safe grasp. Then, after $t = 10\,\mathrm{s}$, the controller slowly releases the grip force until the object starts to pivot. The top subplot shows the measured grip force and moment. The bottom subplot shows the vertical safety margin during the pivot motion. In total, 12 runs were done with 4 different objects. The results are shown in Table I. The table shows the measured forces from the force sensor and the predicted safety margins from the model at which pivot occurs. The most important metric is the vertical safety margin. This metric shows high repeatability for each object, since the margin at which pivot occurs is consistent across repeated runs of the same object. However, it differs substantially between objects. For example, the bottle starts to pivot at around $\hat{R}_y = 0.5$, whereas the box cutter starts to pivot at around $\hat{R}_y = 0.3$.

TABLE I: Per run measured forces and torques by force sensor and predicted safety margins from the model at the onset of pivoting. Each object was evaluated over three independent runs (R1–R3).

Object	Run	$\lvert F_z \rvert$ (N)	$\lvert M_z \rvert$ (N mm)	$\hat{R}_y$	$\hat{R}_x$
Protractor	R1	3.27	7.41	0.412	0.932
	R2	3.27	8.53	0.421	0.931
	R3	3.33	7.89	0.464	0.937
Box Cutter	R1	7.64	26.10	0.290	0.728
	R2	9.31	27.55	0.289	0.750
	R3	9.57	25.05	0.309	0.760
Bottle	R1	4.53	14.48	0.484	0.741
	R2	4.89	14.82	0.485	0.705
	R3	5.03	14.34	0.493	0.704
Phone	R1	5.88	17.93	0.338	0.781
	R2	6.12	18.47	0.352	0.768
	R3	5.95	17.64	0.341	0.775

V. Discussion

This thesis studied whether tactile sensing can be used to estimate proximity to slip during pivoting. The approach used combines tactile sensing, Limit Surface based friction modeling, and machine learning to make this possible. The first question is how well the Limit Surface curves were fitted from the measured slip data.

Fig. 9 shows that the visually fitted curves give a mean

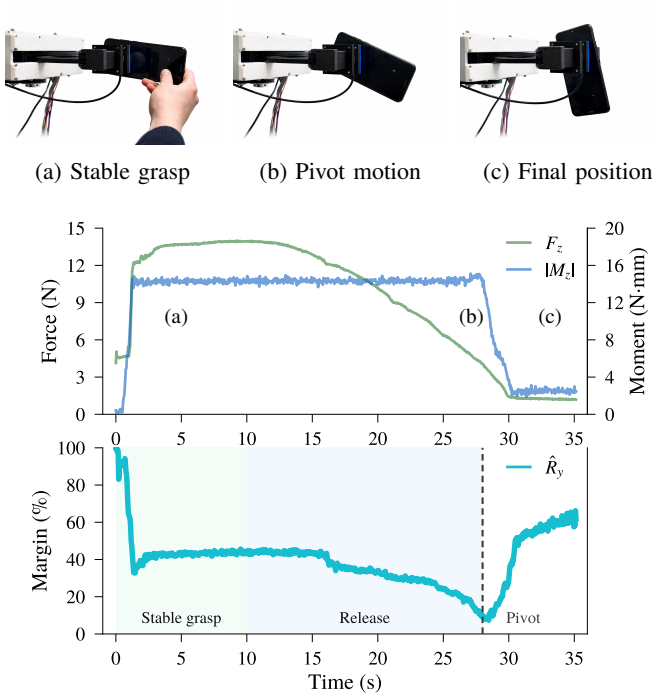

Fig. 13: Measured forces and predicted safety margins during the pivot motion.

absolute relative error of around 10% for the moment. This is slightly better than the quadratic surface fit shown in Fig. 14, while both methods achieve a P_{90} error of around 20%. Given that the test setup was assembled from aluminum extrusion profiles, this level of error is expected. Nevertheless, the achieved accuracy appears sufficient for safe grasp control as demonstrated later with the first experiment.

With the fitted Limit Surfaces, the next step is to use the tactile data to predict the underlying contact state. First the model was trained to predict the contact forces. The time-series in Fig. 10 show that the tactile sensor and trained model can estimate the contact forces accurately on unseen data from the same experimental setup. The low error metrics indicate that the tactile data contains enough information to accurately describe force and torque at the contact.

The prediction of the generalized safety margin shows similar performance, as shown in Fig. 11. This is expected, because the ground-truth safety margin was constructed from the measured contact forces and the fitted Limit Surfaces. The safety margin model therefore behaves much like a force predictor with the fitted Limit Surfaces embedded. These results show that the model learned to accurately predict the safety margins for unseen data within the same experimental setup.

Before testing generalization, the first experiment evaluated whether the predicted safety margin could already be used for safe grasp control. In this experiment, the safety margin prediction model was used together with the gripper force controller. The results in Fig. 12 show that the controller increases the grip force as the torque increases and the vertical safety margin decreases. This shows that the predicted safety margin can be used in closed-loop control to maintain a stable grasp under increasing torsional load.

The second experiment then evaluated whether the model could also be used for pivoting across different objects. As shown in Fig. 13, the object was first grasped safely and then gradually released until pivoting occurred. During this experiment, it was found that the robotic gripper could handle more current than the limit used in the first experiment, so the maximum grasp force was increased to 15 N. However, the model was only trained with grasp forces up to 10 N, which explains why the vertical safety margin is less sensitive during the initial grasp decrease between 10 and 15 s. After that, the vertical safety margin follows the decrease in grasp force until the slip point at 27 s. Once the object starts to pivot, the rotation brings its center of mass horizontally closer to the contact, so the measured torque decreases. Counterintuitively, the vertical safety margin also increases, even though the pivoting motion is still ongoing. A likely explanation is that, once slip starts, the deformed sensor surface is free to move back toward its initial position. The model may then interpret this as a lower contact deformation, and therefore as a higher safety margin.

Table I shows the predicted vertical safety margins at which the objects begin to pivot. The values are repeatable across repeated runs with the same object, but they differ substantially between objects. Importantly, pivoting would be expected to occur near a safety margin of 0%, but the data shows that it occurred much earlier, at around 30% to 50%. Together, this shows that the model does not generalize well across objects and that its predictions are overly optimistic. The model predicts the grasp to be safer than it really is. The limited generalization is likely caused by training with only a single indentor. The too optimistic predictions are likely caused by errors in the Limit Surface fitting, and by the different conditions used during training and during the experiments. The training data was collected with a single indentor in the test setup. The experiments were performed with real objects using a two-finger robotic gripper mounted on a robotic arm.

Prior work has shown that tactile sensing can be used to estimate frictional safety margins during grasping. Boonstra et al. and Vitrani et al. showed that these predictions can be used to enable stable grasp control. However, these studies focused on translational loading and did not consider motions in which rotation also plays a role. The results of this thesis build on their work by showing that tactile data can also be used for safety margin estimation during pivoting.

Tactile sensing has also been used for pivoting in prior work. Viña et al. controlled pivoting using tactile sensing together with visual feedback, while Costanzo et al. used tactile sensing and the Limit Surface to model the contact friction. A key difference is that this thesis focuses on predicting contact forces and estimating proximity to slip, rather than on estimating motion variables such as velocity or position. It also uses the Limit Surface to define a safety margin for pivoting.

In this way, this thesis combines tactile sensing, safety margin estimation, pivoting, and Limit Surface modeling in a single approach. This extends tactile safety margin estimation from grasping under translational loading to tasks in which rotation also plays a role.

However, the current approach has several limitations.

A first limitation of the Limit Surface fitting is the lack of data points with low moment and high shear force. These missing points are important because they determine where the ellipse crosses the x-axis. The initial experiments were designed to force rotational slip, but not horizontal slip, in order to prevent misalignment and potential mechanical damage. A few months later, part of the experiment was repeated to specifically measure horizontal slip. However, the acquisition setup had to be reassembled, which led to different results. Due to time constraints, and to maintain a consistent dataset, the analysis was continued with the data from the initial acquisition.

The partial retrial also led to a reconsideration of how the slip point was defined. In the initial setup, slip was defined as the maximum reached torque. However, slip already starts below that, when the contact transitions from dry friction to viscous friction and reaches a certain slip speed. Under these conditions, the viscous friction force can become larger than the dry friction force. This means that the slip points used to fit the Limit Surface were defined too late, which is another limitation of the approach.

The main challenge with using tactile sensing for safety margin prediction is generalization across different materials and shapes of objects. This is also the main limitation of this thesis. Training was done with a single indentor, which does not allow the model to learn to generalize across different objects and materials. As a result, the trained model behaves much like a calibrated force sensor with the fitted Limit Surfaces built in for the specific indentor used during acquisition.

Another limitation of this thesis is its focus on rotation. The experiments were designed around pivoting and torque, while horizontal sliding and variation in shear force were not evaluated. This is especially relevant for the grasp control experiment, in which the added chocolate candies increased the moment up to 20 Nmm but added only about 70 g of weight. As a result, they contributed very little to the shear force at the contact. A more complete evaluation of the controller should therefore also include experiments in which the shear force is increased up to a few newtons.

A further limitation is the limited operating range of the tactile sensor and gripper. A minimum grip force is needed to produce enough indentation, because otherwise no dimples move and the tactile sensor becomes effectively blind. The maximum grip force is limited by the robotic gripper and also by the silicone structure of the tactile sensor.

Finally, no higher-level motion planning tasks were performed. The predicted contact forces could in principle be used to estimate quantities such as acceleration and displacement. It would therefore have been interesting to evaluate this in practice.

VI. Conclusion

This thesis presented an approach that combines tactile sensing, Limit Surface modeling, and machine learning to estimate proximity to slip during pivoting. The goal was to extend earlier work on tactile safety margin estimation for translational loading to tasks that also involve rotation. This is important for pivoting, where the object must rotate relative to the fingers while contact is maintained.

Using data acquired from the test setup, it was possible to fit the Limit Surface and to train two deep learning models for contact force prediction and generalized safety margin prediction. Validation on unseen data from the same experimental setup showed that the models can predict these quantities accurately for this setup. The first experiment then showed that the predicted safety margin can be used for safe grasp control under increased torsional load. The second experiment, which evaluated pivoting across different objects, showed that the approach is repeatable for repeated trials with the same object, but does not yet generalize well across different objects. Furthermore, the current model is not yet accurate enough to predict the exact onset of pivoting, since pivoting should happen near a safety margin of 0%, but in practice occurred earlier, at around 30% to 50%.

The main finding of this thesis is that tactile sensing can be used to predict the contact state reliably, as long as the conditions are similar to those used during training. The Limit Surface proved to be a useful concept for extending safety margin estimation to motions that involve both translation and rotation. At the same time, the results also showed that fitting the Limit Surface accurately is difficult, since the onset of slip must be identified precisely and limitations of the experimental setup can affect the measured slip points.

The combination of the Limit Surface, the generalized safety margin, and the PointNet-inspired model has clear potential for future work on generalization across different objects and materials. To get there, much more and more diverse data is needed than the single indentor used in this thesis.

It would also help to collect more data in the region with high shear force and low torque, since this part was missing in the current dataset. If the model can eventually learn to generalize across objects and materials, tactile sensing could support more reliable friction regulation during manipulation, including controlled sliding and pivoting. This could help enable more capable in-hand manipulation in unstructured environments.

VII. Acknowledgments

This work reflects the input and ideas of many people in the group.

First of all, I would like to thank the professor leading the group, Michaël Wiertlewski. With his scientific insight and his eye for the human side of research, he continuously introduced new ideas throughout this thesis. The environment he created, with smart and motivated postdocs, PhD students, and master's students, was a great place to work during this project.

I would also like to thank my daily supervisor, Giuseppe Vitrani, who was highly involved throughout the project. He contributed many good ideas while maintaining a clear overview of the work. His academic experience helped me understand how this thesis fits into a larger context.

I would also like to thank Laurence Willemet for the good ideas and the feedback on the introduction.

Lukas Stracovsky helped with creating the videos, building the acquisition setup, and staying late several times so the experiments could continue into the evening.

Overall, I am grateful to everyone in the group for the discussions and collaboration.

REFERENCES

[1] H. Kasaei and M. Kasaei, "MVGrasp: Real-time multi-view 3D object grasping in highly cluttered environments," *Robotics and Autonomous Systems*, vol. 160, p. 104313, Feb. 2023. [Online]. Available: https://www.sciencedirect.com/science/article/pii/S0921889022002020

[2] Z. Liu, Q. Liu, W. Xu, L. Wang, and Z. Zhou, "Robot learning towards smart robotic manufacturing: A review," *Robotics and Computer-Integrated Manufacturing*, vol. 77, p. 102360, Oct. 2022. [Online]. Available: https://www.sciencedirect.com/science/article/pii/S0736584522000485

[3] C. C. Kemp, A. Edsinger, and E. Torres-Jara, "Challenges for Robot Manipulation in Human Environments."

[4] Ministerie van Volksgezondheid, Welzijn en Sport, "Rijksbegroting 2025: Volksgezondheid, welzijn en sport (xvi)," Rijksoverheid, Tech. Rep., 2024, accessed: 2026-04-03. [Online]. Available: https://www.rijksoverheid.nl/binaries/rijksoverheid/documenten/begrotingen/2024/09/17/xvi-volksgezondheid-welzijn-en-sport-rijksbegroting-2025/XVI%2BVolksgezondheid%2C%2BWelzijn%2Ben%2BSport%2BRijksbegroting%2B2025.pdf

[5] Sociaal-Economische Raad, "Werken aan veranderkracht," Sociaal-Economische Raad, Tech. Rep., 2024, accessed: 2026-04-03. [Online]. Available: https://www.ser.nl/-/media/ser/downloads/adviezen/2024/werken-aan-veranderkracht.pdf

[6] A. Solayman, H. Sajwani, O. Abdul Hay, R. Chang, L. AbuAssi, A. Ayyad, Y. Zweiri, and Y. Abdul Samad, "PiezoSight: Coupling vision based tactile sensor neural network processing and piezoresistive stimulation for enhanced piezo-vision hybrid sensing," *Materials & Design*, vol. 253, p. 113899, May 2025. [Online]. Available: https://www.sciencedirect.com/science/article/pii/S0264127525003193

[7] D.-J. Boonstra, L. Willemet, J. Luijkx, and M. Wiertlewski, "Learning to estimate incipient slip with tactile sensing to gently grasp objects," in *2024 IEEE International Conference on Robotics and Automation (ICRA)*, May 2024, pp. 16 118–16 124. [Online]. Available: https://ieeexplore.ieee.org/abstract/document/10611517

[8] G. Vitrani, B. Pasquale, and M. Wiertlewski, "Shadowtac: Dense measurement of shear and normal deformation of a tactile membrane from colored shadows," in *2025 IEEE International Conference on Robotics and Automation (ICRA)*, 2025, pp. 5004–5010.

[9] F. E. Viña B., Y. Karayiannidis, C. Smith, and D. Kragic, "Adaptive control for pivoting with visual and tactile feedback," in *2016 IEEE International Conference on Robotics and Automation (ICRA)*, May 2016, pp. 399–406.

[10] M. Costanzo, "Control of robotic object pivoting based on tactile sensing," *Mechatronics*, vol. 76, p. 102545, Jun. 2021. [Online]. Available: https://www.sciencedirect.com/science/article/pii/S095741582100043X

[11] C. R. Qi, H. Su, K. Mo, and L. J. Guibas, "Pointnet: Deep learning on point sets for 3d classification and segmentation," in *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, July 2017.

[12] J. Jameson, "Analytic Techniques for Automated Grasp," Ph.D. dissertation, May 1985.

[13] R. Howe, I. Kao, and M. Cutkosky, "The sliding of robot fingers under combined torsion and shear loading," in *1988 IEEE International Conference on Robotics and Automation Proceedings*, Apr. 1988, pp. 103–105 vol.1. [Online]. Available: https://ieeexplore.ieee.org/document/12032

[14] E. Popova and V. Popov, "The research works of Coulomb and Amontons and generalized laws of friction," *Friction*, vol. 3, pp. 183–190, May 2015.

[15] L. Willemet, N. Huloux, and M. Wiertlewski, "Efficient tactile encoding of object slippage," *Scientific Reports*, vol. 12, no. 1, p. 13192, Aug. 2022. [Online]. Available: https://www.nature.com/articles/s41598-022-16938-1

[16] M. M. Ghazaei Ardakani, J. Bimbo, and D. Prattichizzo, "Quasi-static analysis of planar sliding using friction patches," *The International Journal of Robotics Research*, vol. 39, no. 14, pp. 1775–1795, Dec. 2020. [Online]. Available: https://doi.org/10.1177/0278364920929082

[17] M. Costanzo, G. D. Maria, C. Natale, M. Costanzo, G. D. Maria, and C. Natale, "Detecting and Controlling Slip through Estimation and Control of the Sliding Velocity," *Applied Sciences*, vol. 13, no. 2, Jan. 2023. [Online]. Available: https://www.mdpi.com/2076-3417/13/2/921

APPENDIX A
LIMIT SURFACE FIT USING LINEAR REGRESSION

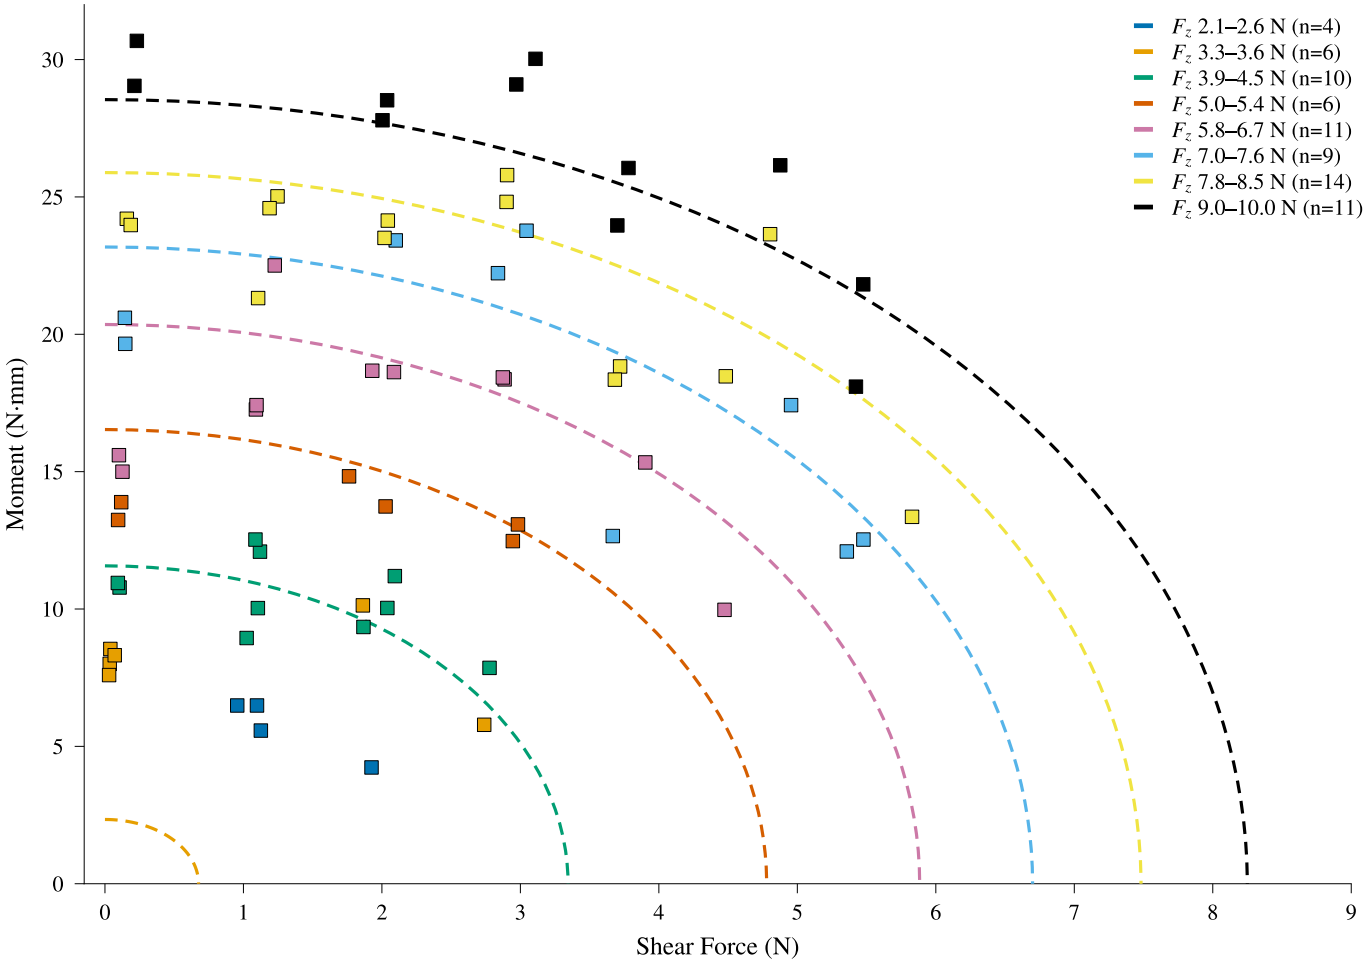

Fig. 14: Limit Surface curves fitted using least squares to the identified slip points ($n = 62$) in the $(F_{xy}, |M_z|)$ plane. For $|M_z|$, the fit achieved an MAE of $1.84\,\mathrm{Nmm}$, an RMSE of $2.35\,\mathrm{Nmm}$, and a mean absolute relative error of 11.81% ($P_{90} = 22.49\%$).