

Enabling Robot Manipulation of Soft and Rigid Objects with Vision-based Tactile Sensors

Michael C. Welle*¹, Martina Lippi*², Haofei Lu¹, Jens Lundell¹, Andrea Gasparri², Danica Kragic¹

Abstract—Endowing robots with tactile capabilities opens up new possibilities for their interaction with the environment, including the ability to handle fragile and/or soft objects. In this work, we equip the robot gripper with low-cost vision-based tactile sensors and propose a manipulation algorithm that adapts to both rigid and soft objects without requiring any knowledge of their properties. The algorithm relies on a touch and slip detection method, which considers the variation in the tactile images with respect to reference ones. We validate the approach on seven different objects, with different properties in terms of rigidity and fragility, to perform unplugging and lifting tasks. Furthermore, to enhance applicability, we combine the manipulation algorithm with a grasp sampler for the task of finding and picking a grape from a bunch without damaging it.

I. INTRODUCTION

Humans have a remarkable ability to manipulate a wide range of objects in various situations. For instance, we can effortlessly pick up and hold objects of different shapes and sizes, adjust our grip, and move them around without much thought. Furthermore, we can manipulate *fragile* objects without damaging them. This level of flexibility and adaptability cannot be easily replicated in robotic systems [1]. A major challenge for current robotic systems is the deficiency in sensory capabilities when compared to humans: human fingers are soft, sensitive tactile sensors that enable humans to feel contact forces as well as detect slippage of objects with ease. Equipping robotic hands or grippers with similar tactile sensing capabilities has been a long-standing goal of the research community [2].

Recently, vision-based tactile sensors have received increased popularity in the community [3] thanks to their relatively low cost, high resolution, and ease of use. Examples of vision-based tactile sensors are GelSight [4], TacTip [5], and DIGIT [6]. The basic principle of these sensors is to use a camera to capture the deformation of an elastomer material occurring in case of contact with objects. By combining this vision-based tactile sensing with powerful machine learning techniques, researchers have managed to demonstrate impressive tasks such as in-hand manipulation [6] and ball rolling [7] (see Sec. II for more details). However, these solutions typically rely on large neural networks that may




Fig. 1: Sequence of states for picking a grape berry from a bunch, without damaging it, using visual-tactile sensors mounted on the robot gripper.

be tailored to specific sensors and objects and may require re-training to address different tasks. Additionally, many of the existing approaches are designed to handle rigid objects.

In this work, we aim to bridge this gap by presenting a simple learning-free touch and slip detection technique for vision-based tactile sensors that relies on a single hyperparameter, which can be easily tuned without gathering large datasets. We design an algorithm that combines such detection methods to handle a wide range of objects, from soft to rigid ones, without requiring any prior knowledge of the object properties. Finally, we integrate this algorithm into a grasping pipeline [8] that proposes a set of grasp poses from which we select one to execute for picking a grape from a bunch, as shown in Fig. 1. In detail, our contributions are:

- A simple-to-deploy touch and slip detection method;
- The design of an algorithm to manipulate soft or rigid objects, which can potentially be fragile;
- Extensive real-world experimental validation using low-cost DIGIT sensors [6] with seven different objects, such as tomatoes and ridged connectors (all videos on the project website¹);
- Deployment of the framework in a comprehensive grasping pipeline detaching table-grapes from a bunch with an off-the-shelf robotic platform.

II. RELATED WORK

Here, we review prior work in vision-based tactile manipulation. For an overview of non-vision-based tactile sensing and manipulation, we refer the reader to the following two reviews [9], [10].

¹ <https://vision-tactile-manip.github.io/exp/>

*Contributed equally.

¹KTH Royal Institute of Technology Stockholm, Sweden, {mwelle,haofeil,jelundel,dani}@kth.se

²Roma Tre University, Rome, Italy {martina.lippi,andrea.gasparri}@uniroma3.it

This work was supported by the Swedish Research Council, Knut and Alice Wallenberg Foundation, by the European Research Council (ERC-884807), and by the European Commission (Project CANOPIES-101016906).

The general interest in vision-based tactile manipulation took off with the invention of the GelSight sensing technology [4], and, in particular, its adaption to fit gripper form factors [11]–[14]. Since conception, these and other similar types of sensors have been used in a range of tasks, including pose estimation [15], [16], grasping [17]–[20], slip detection [21], shape reconstruction [22], [23], object tracking [24], and closed-loop control [7], [25]–[27].

Most of the above works [7], [15], [17]–[19], [22], [23], [25]–[27], tackles the problem with data-driven methods. For instance, in grasping, Calandra et al. [17] used a Convolutional Neural Network (CNN) to classify grasp success given an RGB and tactile image of the grasped object, while the authors of [18], [19] used a CNN to suggest local gripper adjustments based on the current tactile image for improving grasps. Other example works in vision-based tactile deformable object manipulation have explored closed-loop control of the gripper width for cable manipulation [25] and pivoting [26]. From an application perspective, the main limitation of the data-driven methods is gathering datasets diverse enough for training models that can generalize to many unseen objects, which is both time-consuming and non-trivial to realize.

To mitigate the problem inherent to data-driven vision-based tactile manipulation methods, a few works have explored learning-free approaches [16], [21], [24]. The earliest of these works, [16], did tactile pose estimation of a USB stick and used the estimated pose for inserting the USB into a port. Similar to [16], Izatt et al. [24] proposed an Extended Kalman Filter (EKF)-based method for tracking the pose of a manipulated object using point clouds from an RGB-D camera and the tactile image from a GelSight sensor.

Out of the learning-free methods, the one most similar to ours is [21], where the authors also proposed an algorithm for adjusting the grasp force to avoid object slippage when grasping. That work, however, required markers on the GelSight sensor to detect slippage and mostly evaluated the method on rigid objects. Our method, on the other hand, is marker-free and can handle rigid, deformable, and *fragile* objects. Moreover, it is the first vision-based tactile manipulation method to successfully realize cable unplugging.

III. METHOD

As mentioned in the introduction, we consider a robot equipped with two vision-based tactile sensors mounted on its gripper, as shown in Fig. 1. Our focus is on designing a manipulation strategy that adaptively modulates the gripper closure in order to manipulate an object of interest, which might either be soft or rigid and potentially fragile. Specifically, the objective is to ensure that the gripper can securely grasp the object without losing it during the manipulation task and, at the same time, it does not compress the object more than necessary to hold it. The latter aspect is fundamental in the case of fragile objects, such as table-grapes, where over-compression may result in damage to the objects. Examples of manipulation tasks that require such ability include unplugging or detaching objects, which

require a certain amount of force to be exerted without over-compressing the object, and moving fragile items, where it is crucial to modulate the gripper closure carefully, as mentioned above. As tactile sensors, we use DIGIT sensors, but the approach can be generalized to any vision-based tactile sensor.

In order to successfully modulate the gripper closure and perform a manipulation task on both soft and rigid objects, we identify two main capabilities that the system must provide:

- 1) Touch detection, which establishes whether contact with an external object has occurred, regardless of the texture, shape, rigidity, and other physical properties of the object;
- 2) Slip detection, which establishes, after a contact detection, whether the object is sliding out of the gripper.

Based on these capabilities, a manipulation algorithm is designed which establishes the robot reaction in terms of motion and modulation of the gripper closure in response to either touch or slip detection. Desirable features to achieve such capabilities are sensitiveness and ease of tuning. Specifically, for touch detection, we want the system to be sensitive enough to detect contact, even with soft objects, without over-compressing them. Nevertheless, a balance must exist that allows the system to be robust with respect to noisy data of the tactile sensor, and avoid misclassifying such noise as a contact state. Regarding slip detection, it should be sensitive enough to prevent the object from completely slipping out of the gripper, but at the same time, it should not be overly sensitive and detect slip for minor changes in the object’s positioning. Our approach addresses both capabilities with the same algorithm and only relies on a single hyperparameter that is easy to tune. In the following, we first describe the touch and slip detection method and then describe how to exploit it into a manipulation algorithm.

A. Touch and Slip Detection

The basic principle behind the touch and slip detection methods is the same. Given a reference image of a tactile sensor, which is updated over time according to the manipulation algorithm, we evaluate the pixel-wise variation of consecutive frames with respect to the reference image. If this variation is above a certain threshold, a *change* in the tactile sensor state is recorded. Next, the change detection is *categorized* as either touch or slip detection. In the following, firstly, we focus on how to recognize a change in the sensor state and, secondly, we describe how the classification into touch or slip is made. Note that these detections are carried out independently for each sensor and then combined in the manipulation algorithm.

Let $I_t^s \in \mathbb{R}^{h \times w \times 3}$ be the normalized tactile input image (i.e., values in the range $[0, 1]$), with height h and width w , of the sensor s at time step t , where $s \in \{1, 2\}$, and let $R^s \in \mathbb{R}^{h \times w \times 3}$ be the normalized reference image related to the same sensor. Then, we calculate the RGB difference image $D_t^s \in \mathbb{R}^{h \times w \times 3}$ as

$$D_t^s = |I_t^s - R^s|, \quad (1)$$




Fig. 2: Examples of touch detection on a grape (left column) and slip detection on an AUX connector (right column)

which is converted into a mono-dimensional image $\bar{D}_t^s \in \mathbb{R}^{h \times w}$ by averaging the three (RGB) channels. At this point, a thresholding operation is applied to mitigate the noise in the tactile sensor data. Specifically, we assign zero value to pixels with intensity below a noise threshold $\tau_n^s \in [0, 1]$, and a value of 1 otherwise. The result is a binary difference image $\tilde{D}_t^s \in \mathbb{B}^{h \times w}$, with $\mathbb{B} = \{0, 1\}$, capturing the relevant difference of the current tactile image compared to the reference image. Such a binary image is combined with the one at the previous time step to detect a change in the sensor state. Specifically, we define the change image $C_t^s \in \mathbb{B}^{h \times w}$ at time t as

$$C_t^s = \tilde{D}_{t-1}^s \odot \tilde{D}_t^s, \quad (2)$$

where \odot denotes the element-wise multiplication (Hadamard product). Then, we rely on the fulfillment of the following condition to determine whether a change has occurred

$$\frac{1}{wh} \sum_{i=1}^h \sum_{j=1}^w C_t^s(i, j) \geq \tau_d, \quad (3)$$

with $\tau_d \in [0, 1]$ a detection threshold, i.e., a change occurs if the ratio of ones in C_t^s to the total number of pixels exceeds a specified threshold. Note that the use of consecutive images in C_t^s enhances the system robustness since it enables disregarding pixels that are equal to one for a single time

step. Such occurrences can arise due to the presence of noise in the vision-based tactile sensor.

After detecting a change, we classify it by simply analyzing when it has occurred. Specifically, since at the start, we assume no contact with the object, when the first change is recognized, this is classified as touch detection. On the other hand, if the change occurs after a touch has been established, this is classified as slip detection. Note that in the first case, the reference image is related to the “empty” sensor, i.e., no object in contact, while in the second case, the reference image represents the image obtained when holding the object. More details on the reference image update are provided in the next section, along with the manipulation algorithm. Fig. 2 (left column) shows an example of images obtained in case of touch detection with a grape, where Fig. 2 (top left) represents the reference (“empty”) image, Fig. 2 (middle left) shows the current tactile image and Fig. 2 (bottom left) reports the change image. Similarly, Fig. 2 (right column) shows examples of images obtained in case of slip detection of an AUX connector.

Note that two thresholds are used in the above for detection, i.e., the noise threshold τ_n^s and the detection threshold τ_d . To determine the first threshold τ_n^s , we propose an easy automatic calibration procedure for each sensor s to be run when there is no contact. In detail, we record the maximum values in the averaged difference image for K time steps and then compute the mean value, i.e.,

$$\tau_n^s = \frac{1}{K} \sum_{t=1}^K \max(\bar{D}_t^s). \quad (4)$$

It follows that the only hyperparameter to select for the touch/slip detection is τ_d as detailed in the following.

B. Manipulation Algorithm

Fig. 3 presents an overview of the proposed manipulation algorithm. Briefly, the basic idea is that once touch detection is complete, the robot can initiate the manipulation motion. If any slip occurs, the gripper closure is adjusted to ensure a firm grasp and the manipulation motion can be continued.

At the start, we assume that the robot is in a pre-grasping pose, wherein the closure of the gripper will result in the grasping of the object. Additionally, we assume that the noise thresholds τ_n^s have been calculated according to the calibration procedure described earlier. No knowledge is assumed regarding the object to manipulate in terms of texture, shape, rigidity, and other physical properties.

The first objective is to retrieve the initial reference images R^s of both the tactile sensors, i.e., $s \in \{1, 2\}$ (state 1 in the Fig. 3). Such reference images are representative of the tactile images when no contact is happening, as shown for instance in Fig. 2 (top left). These are obtained by collecting and averaging N tactile images for each sensor, i.e.,

$$R^s = \frac{1}{N} \sum_{t=1}^N I_t^s, \quad s \in \{1, 2\}. \quad (5)$$

Once the reference images are defined, touch detection is activated on both sensors (indicated with blue color), and the

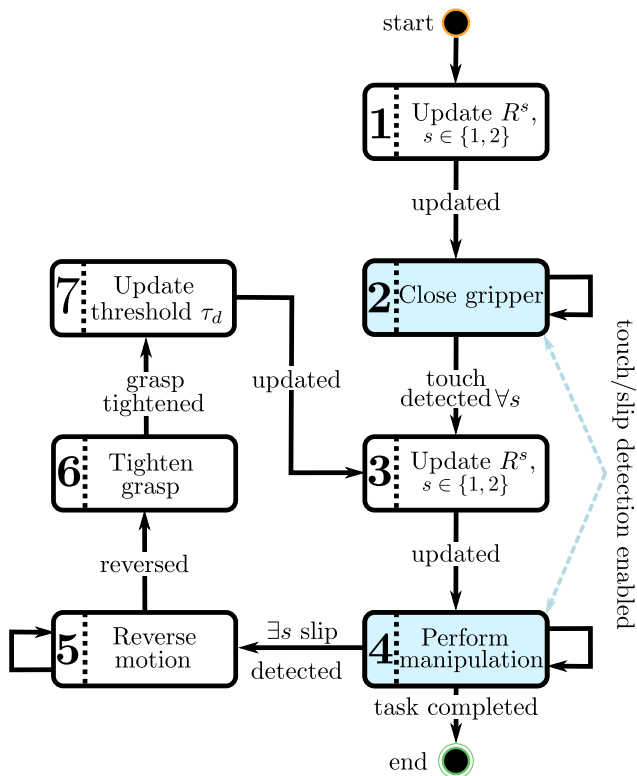


Fig. 3: Overview of the manipulation algorithm. The states where the touch/slip detection is active are highlighted in blue color.

gripper begins to close gradually (state 2 in the figure). The gripper keeps closing until both sensors detect the occurrence of touch, in order to ensure that the object can be held by the two fingers of the gripper.

At this point, the reference images $R^s, \forall s$, are updated again to enable slip detection. As in (5), we gather the next N tactile images and average them for each sensor (state 3 in Fig. 3). In this case, the reference images are representative of the tactile images when the robot is holding the object, as shown for instance in Fig. 2 (top right). Therefore, a variation with respect to the reference image would indicate a slip of the object.

After completing the update of the reference images, the robot starts the manipulation and moves as required (state 4 of Fig. 3). For instance, in an unplugging task, the robot would begin moving in the unplugging direction. While the manipulation is in progress, slip detection is activated.

If slip is detected by *either* of the tactile sensors, we apply a reaction strategy. Our first step is to reverse the manipulation motion, meaning that the robot returns to the configuration before slip detection (state 5 of the figure). In contact-rich manipulation tasks, such as unplugging objects, this action helps to prevent the object from slipping out of the gripper completely in the event that slip conditions persist over multiple time steps. However, in non-contact-rich tasks, such as object transportation, this phase might be unnecessary.

Once the reverse motion is complete, the gripper closure is modulated and the grasp is tightened (state 6 of Fig. 3). To this aim, the gripper is closed by a small amount depending on the gripper resolution. In our case, we utilize displacements of 0.001 m. The last step of the slip reaction strategy is updating the detection threshold τ_d (state 7 of the figure). Specifically, we express the threshold as

$$\tau_d = k_s \tau_{d,ini}$$

where $\tau_{d,ini}$ is the initial value of the threshold that the user has to select, while $k_s \in \mathbb{R}^+$ is a positive gain which is doubled every time a slip is detected, i.e., $k_s = 2 k_{s,prev}$ where $k_{s,prev}$ is the previous value of k_s and is initialized to 1. This exponential increase of the threshold allows us to adapt the slip detection sensitiveness over time: the tighter the object is grasped, the higher the tolerance for slip detection must be; otherwise, the system would keep detecting slippage for any minor change, preventing the successful manipulation of the object.

At the end of the slip reaction strategy, the reference images are updated (state 3 of the figure) and the manipulation is recovered (state 4 of the figure) until task completion.

C. Integration into Grasping Pipeline

Our manipulation algorithm assumes that a pre-grasping pose is given. In simple scenarios, such as picking objects from a table, using predefined grasp poses is reasonable. However, manually defining successful grasp poses in more challenging scenarios, such as picking grapes from a bunch, is non-trivial. Therefore, in more challenging scenarios, we use the grasp sampling method in [8] to suggest good grasp poses for the manipulation algorithm.

In short, the grasp sampler from [8] is a conditional generative neural network that takes as conditional inputs an incomplete point cloud of the object to grasp and a reference grasp approach direction. From these inputs, the sampler then generates multiple grasps whose approach direction is close but not exactly the reference one. To simplify the task, we select a region of interest given the point cloud and feed it to the grasp sampler. In addition, we also manually select, from all generated grasps, the grasp the robot should reach, which is provided to the low-level Cartesian controller. Once the grasp pose is reached, the manipulation algorithm described above is executed.

IV. EXPERIMENTS

We validated the effectiveness of the proposed approach with a Franka Emika Panda robot, equipped with two DIGIT sensors on the gripper, performing various manipulation tasks, as shown in Fig. 4. In detail, we performed an extensive validation of the manipulation algorithm on 7 different objects, including rigid, soft, and fragile ones. Finally, inspired by the needs of the European Project CANOPIES, focusing on the precision agriculture in table-grape vineyards, we tested the overall grasping pipeline for detaching a (soft and fragile) table-grape berry from a bunch. As mentioned earlier, the challenge in manipulating these

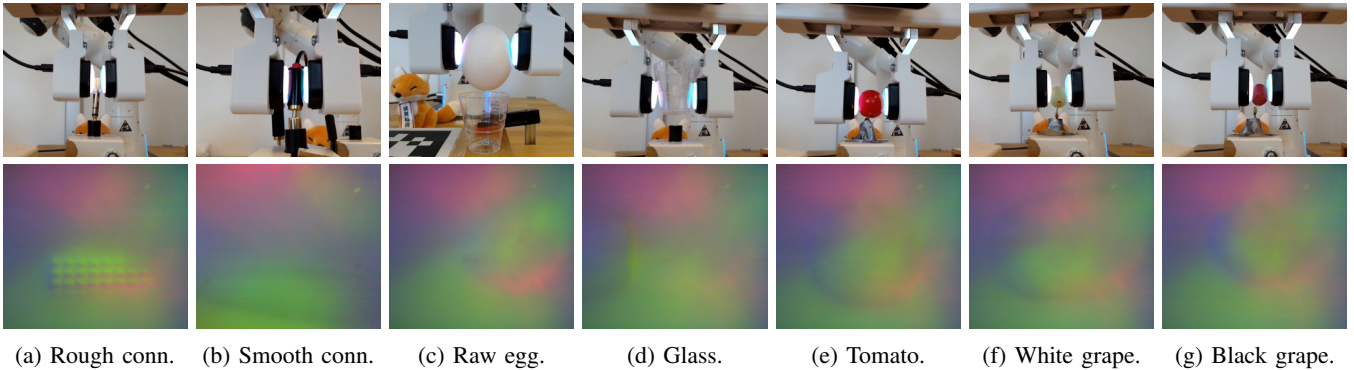


Fig. 4: Illustrations of the considered objects to manipulate (top row) and the respective tactile sensor images (bottom row).

soft objects is to apply the necessary force to detach them, which typically varies based on the object, while avoiding any damage to the object. Examples of all manipulation tasks and of the integrated pipeline are provided in the supplementary video.

For all experiments, we set the initial detection threshold equal to $\tau_{d,ini} = 0.01$ and considered displacements of 0.001 m to tighten the gripper in state 6. The noise thresholds τ_n^s were collected before each experiment using $K = 100$ frames for each sensor, and the reference images R^s were updated using $N = 10$ frames. We set the closing gripper speed to 0.0015 m/s. All experimental videos and tactile sensor images are available on the project website¹.

A. Manipulation Algorithm Validation

To validate the manipulation algorithm, we considered seven objects shown in Fig. 4. In the figure, for each object (top row), an example of the respective tactile image is shown (bottom row). We performed two unplugging tasks of rigid connectors with different forces needed to unplug from their sockets and different textures, one with a rough surface (Fig. 4a) and one with a smooth surface (Fig. 4b). Next, we carried out lifting tasks on two fragile objects, a plastic glass (Fig. 4d) and a raw egg (Fig. 4c). Then, we realized detaching tasks of three soft fragile objects, that are tomatoes (Fig. 4e) and white (Fig. 4f), and black (Fig. 4g) grapes. For these soft objects, we attached the respective connecting branch to the experimental setup in order to generate a realistic detaching force. In all tests, we started from a pre-grasping pose and required the robot to move up to perform the unplugging, lifting, or detaching task. Each manipulation task of an object was repeated three times. The results for all objects are provided in Table I. Specifically, we report the average and standard deviation of duration, compression percentage with respect to the gripper width at touch time, and the number of slip detections.

We can observe an obvious correlation between the number of slippages and the duration and compression needed to manipulate a given object: the higher the number of slippages, the longer the task duration and the higher the compression. The rough connector is by far the object that needs the most force to be unplugged, and therefore is the

Experiment	Duration [s]	% Compression	# Slippages
Rough conn.	108.92 ± 1.43	45.45 ± 1.27	11.00 ± 0.00
Smooth conn.	79.88 ± 3.77	10.47 ± 1.28	5.00 ± 0.82
Egg	68.37 ± 2.20	3.69 ± 0.32	1.00 ± 0.00
Glass	51.52 ± 1.51	1.67 ± 0.74	0.67 ± 0.47
Tomato	43.50 ± 2.09	4.78 ± 2.31	0.67 ± 0.47
White grape	55.45 ± 1.15	13.79 ± 3.75	1.33 ± 0.47
Black grape	58.38 ± 1.52	18.45 ± 2.28	2.67 ± 0.47

TABLE I: Results in terms of duration, % Compression, and # Slippages. Each manipulation is repeated three times, and average and standard deviation values are reported.

one resulting in the highest number of slippages (equal to 11) since a strong grasp must be achieved to succeed in the manipulation. Similarly, the highest object compression compared to the touch configuration (equal to 45% on average), and longest duration (equal to 109 s on average) are recorded. A much easier unplugging task is obtained with the smooth connector, resulting in 5 slippages on average. This shows that objects of the same category can yield very different results based on their specific features and properties.

The plastic glass and the tomatoes, on the other hand, led to the least amount of slippage detected (below 0.7 on average). Some tests were successful without requiring any adjustments to the gripper closure. This result was expected as, for the glass manipulation, the glass surface is not highly slippery, and there are no significant opposing forces that need to be overcome during the lifting process; for the tomato detaching, the connection of the fruit to the branch is relatively weak and, again, does not need to counteract high forces. However, in both cases, the proposed algorithm successfully accomplished the tasks without deforming or damaging the objects. A slightly higher number of slippages (equal to 1) was recorded with the raw egg lifting, due to the higher weight and a more slippery surface compared to the plastic glass.

The grapes were the most challenging objects we tested, as they are significantly softer than the elastomer material of the DIGIT sensor. As a result, minimal deformation of the sensor material occurs during grasping, while the grape itself can be significantly deformed. Additionally, grapes are connected

to the bunch by pedicels, which need to be pulled apart to detach a grape berry. This process requires exerting a certain amount of force, which must be carefully controlled not to cause the grape to get crushed and burst open. On average, slippage was detected 2.67 times for black grapes and 1.33 times for white grapes, due to the different properties of the specific grape varieties. However, in all tests of both grape varieties, the system was able to successfully pick the grape.

The norm of the estimated force at the end effector [28], denoted by $\|f_{ee}\|$ and based on the joint torque sensor measures, is shown in Fig. 5 for manipulating the rough connector (top) and a black grape (bottom). The states of the manipulation algorithm are also highlighted with different background colors. Some states are not visible due to short duration. The figure shows how the end effector force varies according to the specific phase of the manipulation algorithm. Specifically, for the rough connector, every time a slip is detected (beginning of yellow background blocks), the grasp is tightened, leading to an increasing force generated to unplug the object during subsequent manipulation phases (light red background). This trend is observed until the connector is finally unplugged at about $t = 80$ s. Note that during the tightening phases, we stop pulling the object, which is why a drop in the force norm is observed. Regarding the black grape manipulation, we observe much fewer slippages compared to the previous case. However, the same trend is preserved: at the beginning of the first two manipulation phases, an increase in the force norm is observed to detach the grape until about $t = 35$ s when the force drops since the grape is successfully detached.

B. Integrated Pipeline Validation

We now present the results obtained with the entire pipeline exploiting the suggestion of the grasping poses as described in Sec. III-C. The objective is to detach black grapes from a bunch, as shown in Fig. 1 and in the accompanying video. This application might be relevant, for instance, for assistive settings, and for quality inspection applications. In this case, during the manipulation phase, the robot was required to move horizontally in the opposite direction with respect to the bunch. For each grape detaching task, the user cropped the point cloud to mainly contain the object and selected the reference grasp approach direction the sampled grasps should be close to. Given these inputs, the grasp sampler first generated grasp poses, and then a human manually selected the pose the robot should reach. We detached three consecutive grapes from a small bunch with an average duration of 111.93 ± 12.14 for the manipulation algorithm. The compression of the black grape in this setting was 22.23 ± 7.72 , which is slightly higher than the compression observed in the single grape setting. Similarly, we also recorded a higher number of slippage occurrences (4.67 ± 3.77), which we attribute to the fact that the achieved grasps are not as stable as in the single grape setting. Nevertheless, our method was able to detach the grapes without damaging them in all cases¹.

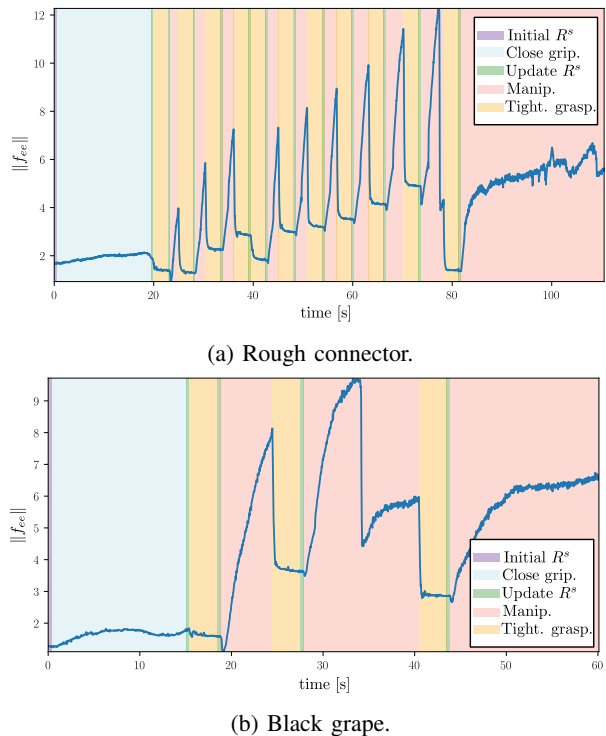


Fig. 5: Examples of the evolution of the force norm for manipulating the rough connector (top) and a black grape (bottom). The states of the manipulation algorithm are also highlighted with different background colors.

V. DISCUSSION

In this work, we introduced a simple-to-deploy algorithm detecting touch and slippage using the vision-based DIGIT sensors for a wide range of soft and rigid objects. This capability is essential for various manipulation tasks, as highlighted in our demonstration of unplugging, lifting, and detaching tasks.

The use of camera-based vision sensors, like the DIGIT, generates images that can be processed by advanced machine learning techniques, as discussed in related work and demonstrated in [6], to enable touch and slip detection capabilities. Specifically, the Pytouch library [29] that is recommended for DIGIT use provides pre-trained machine learning models which can be utilized for touch detection. However, in our test, we could not reliably detect the object as shown at the link². We hypothesize that this is related to the fact that pre-trained models might be hardware specific. This motivates why we opted for a learning-free approach to touch and slip detection which can be easily adapted to different settings.

Another challenge that we encountered is that the quality of the tactile images, in terms of sharpness of the manipulated object, increases as the contact force increases, i.e., the harder we press, the more details are visible. This is not an issue for rigid objects, as impressively demonstrated in [6] (and shown in this work on the rough AUX connector), where we

²<https://vision-tactile-manip.github.io/exp/#pytouch>

can apply relatively high forces without damaging the object. However, it represents a major limitation when dealing with soft objects. In cases where the object is significantly softer than the sensor elastomer material, we might not be able to observe any deformation through the images, even if the object is being compressed. This motivates why we designed an increasing detection threshold which allows us to detect small changes at first (which are suitable for soft objects), but if needed it also allows us to apply significant forces after few iterations (which are suitable for rigid objects).

However, as a drawback of the proposed approach, we identify the fact that it may result in a rather slow manipulation process, since we need to iteratively detect slip occurrences in order to modulate the gripper closure. In contrast, humans can easily identify in advance properties of the objects, which enable their manipulation without the need of multiple slip occurrences. Therefore, we aim to improve this aspect in future work by inferring relevant object properties while performing manipulation and exploiting them to achieve a more proactive behavior.

VI. CONCLUSION

In this work, we proposed a manipulation algorithm for soft and rigid objects using vision-based tactile sensors mounted on the robot gripper. The algorithm is based on a touch and slip detection method which evaluates the variations in the sensor images and relies on a single hyperparameter. No knowledge about the shape, texture, and physical properties of the object to manipulate was required. We showed the effectiveness of the approach on seven real-world objects, including rigid, soft, and fragile ones, using two DIGIT sensors for unplugging, lifting, and detaching tasks. Finally, we combined the manipulation strategy with a grasping proposal system to provide a comprehensive framework and validated it on a grape detaching task.

REFERENCES

- [1] A. Billard and D. Kragic, "Trends and challenges in robot manipulation," *Science*, vol. 364, no. 6446, p. eaat8414, 2019.
- [2] C. Chi, X. Sun, N. Xue, T. Li, and C. Liu, "Recent progress in technologies for tactile sensors," *Sensors*, vol. 18, no. 4, 2018.
- [3] K. Shimonomura, "Tactile image sensors employing camera: A review," *Sensors*, vol. 19, no. 18, 2019.
- [4] W. Yuan, S. Dong, and E. H. Adelson, "Gelsight: High-resolution robot tactile sensors for estimating geometry and force," *Sensors*, vol. 17, no. 12, p. 2762, 2017.
- [5] 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, pp. 216–227, 2018.
- [6] 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 Robot. Autom. Letters*, vol. 5, no. 3, pp. 3838–3845, 2020.
- [7] S. Tian, F. Ebert, D. Jayaraman, M. Mudigonda, C. Finn, R. Calandra, and S. Levine, "Manipulation by feel: Touch-based control with deep predictive models," in *IEEE Int. Conf. Robot. Autom.*, 2019, pp. 818–824.
- [8] Z. Weng, H. Lu, J. Lundell, and D. Kragic, "Gonet: An approach-constrained generative grasp sampling network," 2023. [Online]. Available: <https://arxiv.org/abs/2303.07972>
- [9] 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 Trans. Robot.*, vol. 36, no. 6, pp. 1619–1634, 2020.
- [10] Z. Kappassov, J.-A. Corrales, and V. Perdereau, "Tactile sensing in dexterous robot hands," *Robot. Auton. Syst.*, vol. 74, pp. 195–220, 2015.
- [11] S. Wang, Y. She, B. Romero, and E. Adelson, "Gelsight wedge: Measuring high-resolution 3d contact geometry with a compact robot finger," in *IEEE Int. Conf. Robot. Autom.*, 2021, pp. 6468–6475.
- [12] 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 *IEEE/RSJ Int. Conf. on Intell. Robots and Syst.*, 2018, pp. 1927–1934.
- [13] I. H. Taylor, S. Dong, and A. Rodriguez, "Gelslim 3.0: High-resolution measurement of shape, force and slip in a compact tactile-sensing finger," in *IEEE Int. Conf. Robot. Autom.*, 2022, pp. 10781–10787.
- [14] D. Ma, E. Donlon, S. Dong, and A. Rodriguez, "Dense tactile force estimation using gelslim and inverse fem," in *IEEE Int. Conf. Robot. Autom.*, 2019, pp. 5418–5424.
- [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 *Conf. Robot Learn.*, 2021, pp. 1015–1029.
- [16] 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 *IEEE/RSJ Int. Conf. on Intell. Robots and Syst.*, 2014, pp. 3988–3993.
- [17] R. Calandra, A. Owens, M. Upadhyaya, W. Yuan, J. Lin, E. H. Adelson, and S. Levine, "The feeling of success: Does touch sensing help predict grasp outcomes?" in *Conf. Robot Learn.*, vol. 78, 2017, pp. 314–323.
- [18] R. Calandra, A. Owens, D. Jayaraman, J. Lin, W. Yuan, J. Malik, E. H. Adelson, and S. Levine, "More than a feeling: Learning to grasp and regrasp using vision and touch," *IEEE Robot. Autom. Letters*, vol. 3, no. 4, pp. 3300–3307, 2018.
- [19] F. R. Hogan, M. Bauza, O. Canal, E. Donlon, and A. Rodriguez, "Tactile regrasp: Grasp adjustments via simulated tactile transformations," in *IEEE/RSJ Int. Conf. on Intell. Robots and Syst.*, 2018, pp. 2963–2970.
- [20] A. Yamaguchi and C. G. Atkeson, "Implementing tactile behaviors using fingervision," in *IEEE-RAS Int. Conf. Humanoid Robot.*, 2017, pp. 241–248.
- [21] S. Dong, W. Yuan, and E. H. Adelson, "Improved gelsight tactile sensor for measuring geometry and slip," in *IEEE/RSJ Int. Conf. on Intell. Robots and Syst.*, 2017, pp. 137–144.
- [22] M. Bauza, O. Canal, and A. Rodriguez, "Tactile mapping and localization from high-resolution tactile imprints," in *IEEE Int. Conf. Robot. Autom.*, 2019, pp. 3811–3817.
- [23] S. Suresh, Z. Si, J. G. Mangelson, W. Yuan, and M. Kaess, "Shapemap 3-D: Efficient shape mapping through dense touch and vision," in *IEEE Int. Conf. Robot. Autom.*, 2022, pp. 7073–7080.
- [24] G. Izatt, G. Mirano, E. Adelson, and R. Tedrake, "Tracking objects with point clouds from vision and touch," in *IEEE Int. Conf. Robot. Autom.*, 2017, pp. 4000–4007.
- [25] Y. She, S. Wang, S. Dong, N. Sunil, A. Rodriguez, and E. Adelson, "Cable manipulation with a tactile-reactive gripper," *Int. J. Robot. Res.*, vol. 40, no. 12-14, pp. 1385–1401, 2021.
- [26] C. Wang, S. Wang, B. Romero, F. Veiga, and E. Adelson, "Swingbot: Learning physical features from in-hand tactile exploration for dynamic swing-up manipulation," in *IEEE/RSJ Int. Conf. on Intell. Robots and Syst.*, 2020, pp. 5633–5640.
- [27] S. Dong and A. Rodriguez, "Tactile-based insertion for dense box-packing," in *IEEE/RSJ Int. Conf. on Intell. Robots and Syst.*, 2019, pp. 7953–7960.
- [28] S. Haddadin, S. Parusel, L. Johannsmeier, S. Golz, S. Gabl, F. Walch, M. Sabaghian, C. Jähne, L. Hausperger, and S. Haddadin, "The franka emika robot: A reference platform for robotics research and education," *IEEE Robot. Autom. Magazine*, vol. 29, no. 2, pp. 46–64, 2022.
- [29] M. Lambeta, H. Xu, J. Xu, P.-W. Chou, S. Wang, T. Darrell, and R. Calandra, "PyTouch: A machine learning library for touch processing," *IEEE Int. Conf. Robot. Autom.*, 2021.