2024 IEEE International Conference on Robotics and Automation (ICRA 2024)
May 13-17, 2024. Yokohama, Japan

Learning to estimate incipient slip with tactile sensing to gently grasp objects

Dirk-Jan Boonstra* Laurence Willemet[1]* Jelle Luijkx Michaël Wiertlewski
Dept. of Cognitive Robotics *Dept. of CSAIL* *Dept. of Cognitive Robotics* *Dept. of Cognitive Robotics*
TU Delft *MIT* *TU Delft* *TU Delft*
The Netherlands United States The Netherlands The Netherlands
ORCiD 0000-0002-9827-0636 lwilleme@mit.edu j.d.luijkx@tudelft.nl m.wiertlewski@tudelft.nl

***Abstract*— To gently grasp objects, robots need to balance generating enough friction yet avoiding too much force that could damage the object. In practice, the force regulation is challenging to implement since it requires knowledge of the friction coefficient, which can vary from object to object and even from grasp to grasp. Tactile sensing offers a window in the contact mechanics and provides information about friction. Notably touch can detect the precursor of the object slipping away from the grasp. To find this information, tactile sensors measure the deformation field of an artificial skin in both the normal and tangential direction. However, current approaches only react to slip and therefore react too late to perturbations. The object slips, inducing a failure of the grasp and damage. In this study, we introduce a method that uses machine-learning to anticipate slip by computing the so-called *safety margin* of the grasp. This safety margin represents the extra lateral force that maintains the contact away from the frictional limit. To find this value, we use a high-density camera-based tactile sensor to measure the 3D deformation of the surface via the movement of 82 colored markers. We trained a Convolutional Neural Network (CNN) to estimate the safety margin from the tactile images. Because it gives a distance to slip, the safety margin is a powerful metric for regulating grasp forces. As a testament of this effectiveness, we show that a simple proportional controller can robustly grasp a wide variety of objects. The results show that this control method outperforms slip detection methods, by reducing regrasp reaction times while decreasing grasping forces to 1-3 N.**

INTRODUCTION

When dynamically manipulating objects with a robotic gripper, the contact with fingers constantly evolves. During the movement, the pressure and traction distributions change in response to the dynamics and as a function of the friction and material properties. Consequently, it can be difficult to estimate and predict how the object will move within the grasp and whether or not the grasp will be stable. This prediction is crucial for grasping since the forces at the contact determine if the object can rotate, pivot, slide, or stay in place. Without the information regarding the frictional resistance, a controller cannot optimally determine the force that would maintain a stable grasp. Therefore friction-agnostic approaches generally overestimate the grip force to avoid a catastrophic loss of grip [1], [2]. Large forces prevent dropping objects, but also restrain manipulation flexibility [3].

Tactile sensing offers a promising avenue for capturing the mechanical interaction at the interface between the environment and the fingers. Robotic tactile sensors capture the deformation of an artificial skin from which they can infer high-order information, such as material properties (i.e. compliance, texture, curvature) or the contact state (i.e. distance to slip or effort). Tactile sensors discretize the mechanical interaction, represented by the pressure or deformation field, often using miniaturized high-resolution cameras pointing at the membrane [4], [5].

The pressure or deformation field can be processed to estimate contact shape and force [6], or to detect slip from physics-based models [7]. More complex mechanical interactions can be captured using machine learning approaches [8]–[10]. However, when tactile sensors are deployed for grasping regulation they are used to detect slip which makes the reaction to perturbation too slow and they often fails to regain stability after slip [11]–[13].

At a mechanical level, the transition from stick to slip for a soft fingertip occurs gradually. When the tangential force increases from a fully stuck contact, the outer edge of the contact area begins slipping while the center remains stuck. The slip region grows until the entire contact area is in the slip state and the object fully slips [14], [15]. It is postulated that humans use the ratio between the stick and the slip region inferred from the skin deformation to estimate the safety margin [16]. This distance from the onset of slip is believed to be ultimately used to regulate their grip force [17], [18].

In this work, we measured the pattern of deformation of the artificial fingertip before the onset of slip with an iterated version of our ChromaTouch tactile sensor [19], [20]. We trained a convolutional neural network (CNN) to estimate the frictional strength using the safety margin. The model performance is evaluated against an unseen dataset, which showed an average prediction accuracy of 98.2% from the ground truth, when computing the MSE loss over the entire range of safety margin predictions. The 50 Hz refresh rate and the accuracy of the estimation make it suitable for real-time grasping applications on soft and complex objects such as fruits and vegetables, which is demonstrated on three fragile fruits.

We aim to create a tactile-enabled gripper that maintains a squeezing force on an arbitrary object so that the safety margin remains constant (Fig. 1A). To do so, we designed an impedance control gripper (Fig. 1B) which regulates

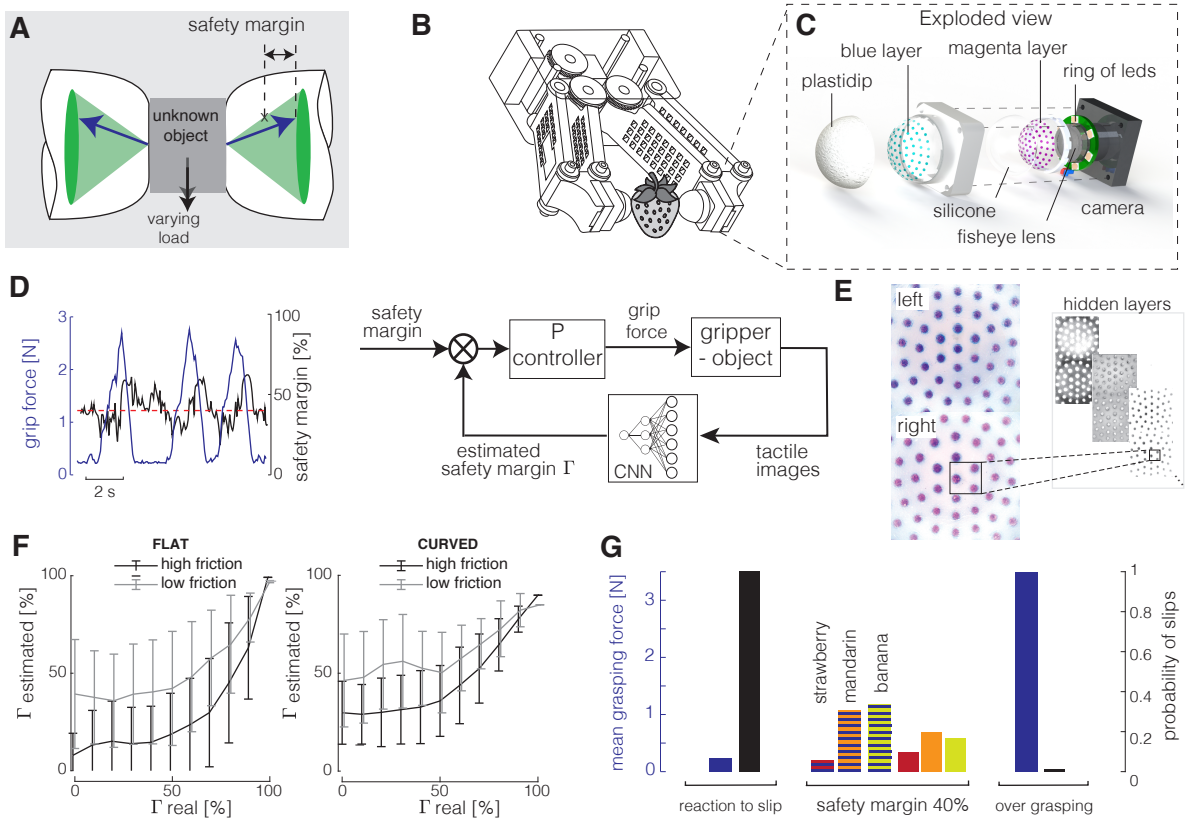

Fig. 1. **A**. Typical evolution of the interaction force when manipulating an object. The grip force is maintained with a safety margin Γ over the minimum required grip force defined by the friction cone. **B**. Render of the custom-made parallel gripper. **C**. Exploded view of the tactile sensor ChromaTouch. **D**. Grip force control to maintain a constant safety margin (red dotted line). **E**. Hidden layers of the convolutional neural network are used to estimate the safety margin. **F**. Deviation of the estimation compared to the real safety margin. **G**. Mean grasping force and probability of slips for three control strategies: reaction to slip [13], constant safety margin of 40%, and over grasping strategy with a fixed 3.5 N grasping force.

its grasping force in real-time. The gripper has two soft tactile sensing fingertips, capturing the 3D deformation of a membrane using an embedded camera (Fig. 1C). The images of the interaction are fed to a CNN to estimate the frictional safety margin Γ (Fig. 1E,F), which is used to adjust real-time grasping force (Fig. 1D), improving object manipulation and minimizing object slip (Fig. 1G). A desired target for Γ is set between 20-60%, resembling human behavior [18].

Materials and Methods

Tactile sensing gripper

The tactile gripper consists of two main components: a set of Chromatouch tactile sensors and a custom-made robotic gripper with force control (Fig. 2A). The Chromatouch tactile sensing mechanism relies on a color-mixing principle [19]. Two layers of colored-markers are first 3D-printed with a Stratasys J735 PolyJet printer in flexible transparent Agilus-Clear with a Shore hardness of 30A. These layers are bonded together using a 1.2 mm-thick elastic silicone (Smooth-On SortaClear 12A), cast between the two marker layers. Three layers of white pigmented silicone (PlastiDip) are sprayed over the outside of the dome, to block light from external sources and to help diffuse internal light. A per-finger embedded USB camera (Basler Dart daA1920-160uc with a Basler Evetar M13B02118W fisheye lens attached) acquires 896 px × 896 px tactile images at 100 Hz with a surface resolution close to 22 px/mm. More information on the design of these tactile fingertips can be found in our previous work [20].

For this work, we custom-designed a parallel gripper where the force could be finely tuned and fit the need of delicate robot grasping. The design, publicly available[1], consists primarily of 3D-printed parts, with tactile sensors embedded at the fingertips. The fingers are driven using one servo-motor (Dynamixel XH430-W210-R), through a set of modified POM gears. The servo-motor enables our gripper with current control, which is calibrated to forces by grasping a regular force sensor (ATI Nano43). We confirmed a linear relation between motor current and grasping force (Fig. 2B).

Working principle of the tactile sensor

To validate the working principle of the tactile sensor, we calibrated the acquired images towards actual 3D deformation using an analytical model. We first acquired images during a normal indentation and a lateral slide (Fig. 3A). On Fig. 3B, images of the preload, and image differences of the initial contact, incipient slip and full slip are shown. We proceeded to find a robust transformation of the sub-image around each marker into the normal displacement u_z

[1] https://github.com/Dirrkk/fuse-gripper

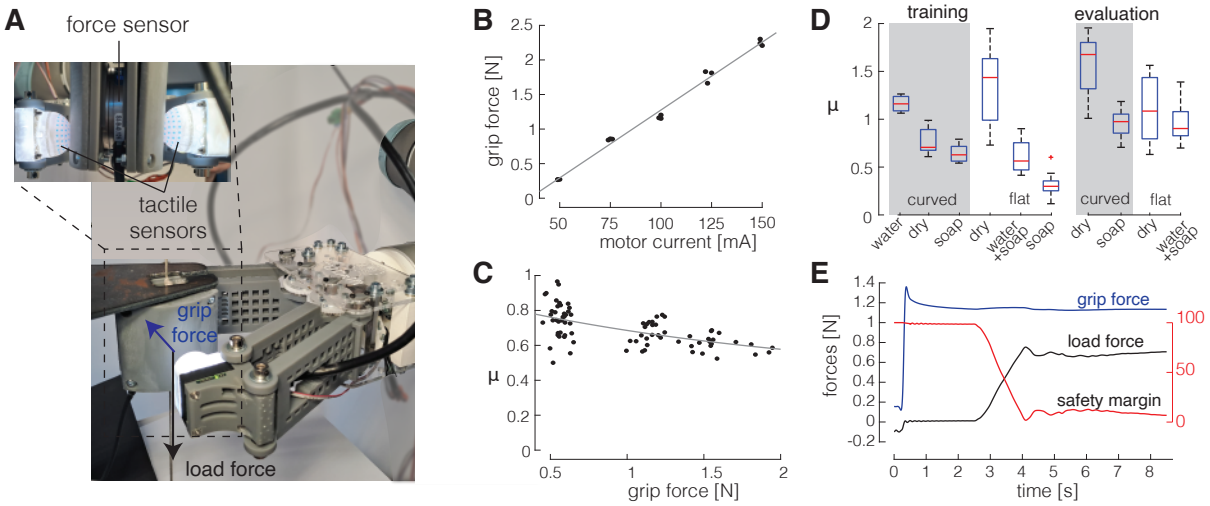

Fig. 2. **A**. Experimental setup for data collection. **B**. Calibration of the grasping force. **C**. Influence of the grip force on the friction coefficient. **D**. Mean and standard deviation of the friction coefficient estimated from the point of slip for each dataset (training and evaluation). **E**. Grip force and load force evolution along one trial. The gripper first closes its arms on the load-compensated object, the DC-motor applies a pulling force after 2 s. The safety margin is computed using an approximation of the friction coefficient.

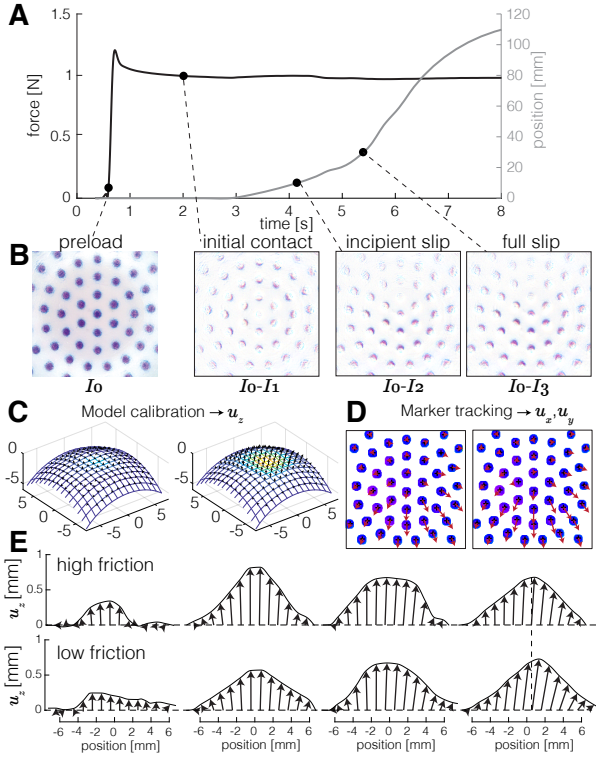

Fig. 3. **A**. Grasping force and vertical position of an object being gripped. **B**. Tactile images obtained during 4 phases. **C**. Interpolated normal displacement u_z of the grid of markers calibrated using a filtered version of Hertz contact model. **D**. Lateral displacement u_x and u_y obtained with markers centroid tracking. **E**. Reconstruction of the 3D displacement of each marker along the direction of slippage for a high and a low friction condition. The influence of friction is indicated by the vertical line.

(Fig. 3C). To find this transformation, we used as ground-truth a filtered version of the Hertzian contact theory. This model predicts a parabolic displacement of the surface of the membrane. For further explanation about the method, refer to our previous work [20]. The lateral displacements u_x, u_y are computed using a marker tracking method (Fig. 3D). The effective 3D displacements are shown at different time courses during slip, for a high and a low friction condition in Fig. 3E. The output displacements have a sub-millimeter precision in the normal and tangential plane.

Experimental setup

We designed an experimental setup to impose arbitrary tangential and normal forces on the tactile sensors (Fig. 2A). During a trial, the tactile gripper first grasps a stationary suspended object with a predefined grasping force. The apparent weight of the object was programmed at 0g by constantly pulling the object against the direction of the gravitational field using two DC-motors (Faulhaber, 2642 012 CR) and a capstan transmission. A force sensor (ATI Nano43) is embedded inside this object, which measures the 3D force interaction at the gripper fingertips. A second-long ramp-pulling force is applied to the suspended object two seconds after the object has been grasped, trying to pull it away from the fingertips using the DC motors. The pulling force produced by the DC motors is controlled using a linear servo amplifier (Maxon, LSC 30/2).

The position of the object was measured with an encoder (Baumer, BTIV 24S 16.24K 1024 G4 5) and the safety margin Γ was estimated using the force measurement and an approximation of the friction coefficient (Fig. 2E):

$$\Gamma(t) = \frac{f_f^* - f_f(t)}{f_f^*}, \tag{1}$$

with f_f the current frictional force, and f_f^* the maximum applicable force to overcome the frictional strength of the grasp, after which the object starts to slip. This frictional force was computed using an estimation of the friction coefficient as the average force ratio when the object is

slipping (Fig. 2D). The influence of the grip force on the friction coefficient has been neglected (Fig. 2C).

During data collection, a randomized experimental plan containing the grasping force and the load force and its rate was followed. The grasping forces vary from 1.0 to 2.5 N in 0.5 N intervals, the load force was uniformly chosen between 2 and 3 N and the force rate was controlled at 2 different levels (0.5 N/s or 1 N/s). The training dataset consisted of a flat and a curved object (radius of curvature $R = \infty$ and $R = 45$ mm) at 3 different frictional conditions (high, medium, low). These friction conditions were obtained by adding water or soap on the surface of the objects, the corresponding friction coefficients are reported Fig. 2D (left side). Every condition has been cycled 16 times using the experimental plan, resulting in a total of 96 trials, or 177.000 (left+right) images. We used 80% for training and 20% for testing. This procedure has been repeated 4 days later (the temperature was $6°C$ lower and the humidity decreased by 3%), to collect a validation dataset consisting of the same 2 objects, on 2 frictional conditions (high and low), resulting in 64 trials and 118.000 images. The data analysis is performed on this validation dataset.

Model training

The images from the tactile gripper and the above-computed safety margin Γ are linked together using a CNN ShuffleNetV2 [21]. This network is lightweight and mobile, and showed promising results in related tactile sensing studies [3]. Our application is suited for a lightweight network since the tactile images represent close-contact information with a limited pixel size, and do not need to contain complete environmental scenes. Furthermore, the lightweightness gives tactile grasping demonstrations in a portable setting.

Inputs of the network are the tactile images from both fingertips. The raw images are resized to 224×224 px, after which they are concatenated vertically and fed into the network, ensuring both images have the same timestamp. ColorJitter and GaussianBlur data augmentation techniques are used to increase the generalization capabilities of the model.

The output of the ShuffleNetV2 network is adapted to a single floating point value, which equals the estimated Γ for the input image. This reduction in output space is obtained by combining a set of linear layers with LeakyReLU activation functions and Batch Normalization (1D) layers, decreasing the output space from 1024 nodes to 1.

Training was done using PyTorch on an Ubuntu 20.04 machine with an Intel Xeon CPU and using CUDA on a Nvidia RTX3060Ti GPU with 8GB of video memory. We used 50 epochs with the batch size set to 64. The MSEloss function was used for backpropagation. For the optimizer, we made use of the Ranger21 framework [22], which is built around the AdamW optimizer, while also providing several techniques to further increase performance and prevent influences from local minima. We used Ranger21 default parameters and set the number of iterations to 100 (number of epochs) x 141600 (length of the training dataset).

Controller design

The trained CNN described above outputs one floating point value (Γ), based on a set of images from the left and right fingertips combined. A simple proportional controller is deployed on the gripper, which takes the estimated Γ value as input, and outputs the grip force required to maintain a target value for Γ. By adapting the target value, the distance to slip can be varied on a per-object basis. The K_P gain was set to 2, and the Γ difference is in the range 0-100%. The minimum applied grip force is set to 0.25 N, to keep the object in a force closure grasp. To prevent damage to the tactile sensor's soft silicone layer, the maximum applied grip force was limited to 3.2 N.

RESULTS

Safety margin estimation

To show the accuracy of the trained CNN, we compared the model output, which only sees the tactile images with ground truth data extracted from the force measurement collected during the experiments. We estimated the safety margin Γ on a validation dataset of unseen images when the gripper interacts with a flat or a curved object and with a high or a low friction coefficient. The ground-truth safety margin was recorded using a force sensor embedded into the object (see Materials & Methods for further details). To compare the predictions of Γ with the real values, we compute the Mean Squared Error (MSE) loss over the validation datasets, the trial average can be found in Table I. Taking the average over the four datasets yields a combined MSE loss of 0.01821, giving the total model an accuracy of 98.2% when predicting Γ.

TABLE I

MEAN SQUARED ERROR (MSE) LOSS (MEAN $\pm$ STD) COMPUTED OVER THE 4 VALIDATION DATASETS: A FLAT AND A CURVED OBJECT EVALUATED ON BOTH HIGH AND LOW FRICTION CONDITIONS. THE MSE SHOWN IS THE COMBINED AVERAGE OVER ALL 16 TRIALS PER CURVATURE/FRICTION CONDITION.

MSE loss [-]	Flat	Curved
high friction	0.03005 $\pm$ 0.04232	0.01357 $\pm$ 0.02482
low friction	0.01256 $\pm$ 0.02480	0.01665 $\pm$ 0.01941

The confusion matrices are shown in Fig. 4A with the safety margin Γ divided in 10% bins to evaluate model performance on the full 0 to 100% range. The model showed similar performance on the flat and the curved object (Fig. 1F). Thus, to show the influence of friction on the safety margin estimation, the flat and curved object datasets are averaged in a single confusion matrix. From the diagonal trend in the figure, we can see that the estimated safety margin is positively correlated with the ground truth safety margin in both friction conditions. However, friction has a significant influence on the safety margin prediction accuracy (Anova, $F(1, 63) = 4.27$, p= 0.043). We observed similar performance for both friction conditions when the safety margin is higher than 40%. However, when the safety margin drops below 10%, the model performs better in the high

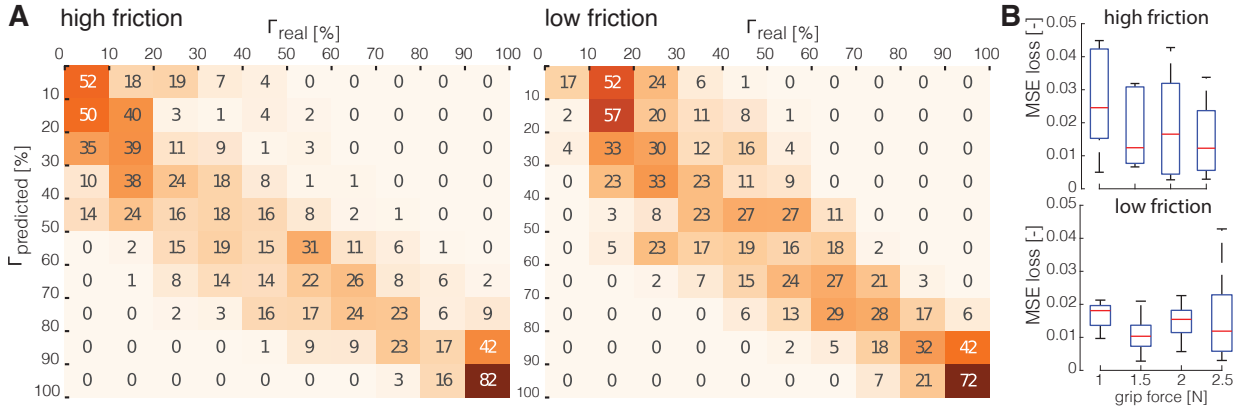

Fig. 4. **A**. Confusion matrix of Γ real versus Γ estimated, with experiments for both flat and curved objects combined. The experiment is done at high and low friction conditions. A perfectly trained model would result in only predictions on the matrix's diagonal. **B**. MSE loss (mean ± std) of the safety margin prediction for each applied grip force for high and low friction condition.

friction condition. In the low friction case, we can see that the prediction is underestimating the ground-truth with errors up to 0.4 when the real safety margin is equal to 0 (Fig. 1F).

Finally, the grip force applied has no significant influence on the MSE, although we can see that lower grip force results in higher errors in the high friction condition (Fig. 4B).

Controller validation

To evaluate the safety margin framework in a real time grasping task, the CNN model output is used with a Proportional controller to control the grip force of the gripper. Experiments are performed on a set of delicate soft fruits: a strawberry (20 g), a mandarin (90 g) and a banana (180 g). In a normal force-controlled environment, assuming $\mu = 1$ in dry circumstances, a desired grasping force would need to exceed respectively 0.2, 0.9 and 1.8 N at least. A 50 g weight was attached with a wire from the grasped object to induce a sudden increase in load force on the object. Fig. 5A shows the grip force required to hold the objects with a given target safety margin of 20%, 40%, and 60%, which is around the desired safety margin target used by humans as discussed before. The boxplots represent the average and standard deviation over 5 trials. We can see a significant increase in grip force when increasing the safety margin target for all fruits ($p< 0.001$ and $F(2, 31) = 189.7$, $F(2, 33) = 19.6$, $F(2, 29) = 80.74$ for the strawberry, the mandarin and the banana respectively). To measure the control reaction time, the attached weight is dropped, causing a sudden spike in load force. The gray zones from Fig. 5A show that, on average, this increase in load force resulted in an increased grip force, especially for the 60% safety margin target, which allows for higher grasp forces.

The probability of object slip decreases with the increase in safety margin, as shown in Fig. 5B. We can see that for the 20% target, the gripper managed to keep the relatively light strawberry between its fingertips for most cases, while the heavier banana and mandarin experienced more slips. Increasing the safety margin target to 40% and 60% caused fewer slips for the heavier objects. We can even see that for the 60% case, the strawberry experienced zero cases of full slip. During the ≤40% safety margin targets, the grasping force stays below what is needed for a force-grasp strategy, demonstrating the benefit for fragile objects.

We also performed an experiment with slowly increasing load force, when grasping an empty cup which was slowly manually filled with rice. The results can be seen in Fig. 5C. At around 10 seconds, as weight is slowly added, we recorded an increase in grip force to maintain the set control target. We can see that over the whole range, a higher grip force is required to maintain the 60% target, while the 40% target only increases grip force to around 1 N when maintaining the grasp.

Conclusion and Discussion

We introduced a new method for controlling the grasping force of a gripper based on a new metric called the frictional safety margin. The frictional safety margin Γ was extracted from tactile sensor images using a convolutional neural network and its prediction shows an average accuracy of 98.2% compared to the safety margin found with force measurements. The trained network was evaluated on a validation dataset recorded on a different day. The robustness of the results shows the model capabilities in expanding to other friction conditions, as these vary on a daily basis and are highly sensitive to environmental conditions like humidity.

Despite its excellent performance, the main limitation of the proposed approach is the prediction of Γ when the frictional contact is either small or slippery. To overcome the first issue, we decided to run our experiments within a limited range of grasping forces, between 0.25 and 3.2 N. Limiting the range kept results accurate while preventing damage of the tactile sensors. On slippery surfaces, we noticed that the accuracy of the estimate decreases for lower safety margins, especially in low friction conditions. The decrease in accuracy can be caused by the sparsity of data used for training in the lower range of safety margin. During data acquisition, the low safety margin, which corresponds to gross slippage, was not always reached because of limitations in the experimental setup in which maximum applicable load

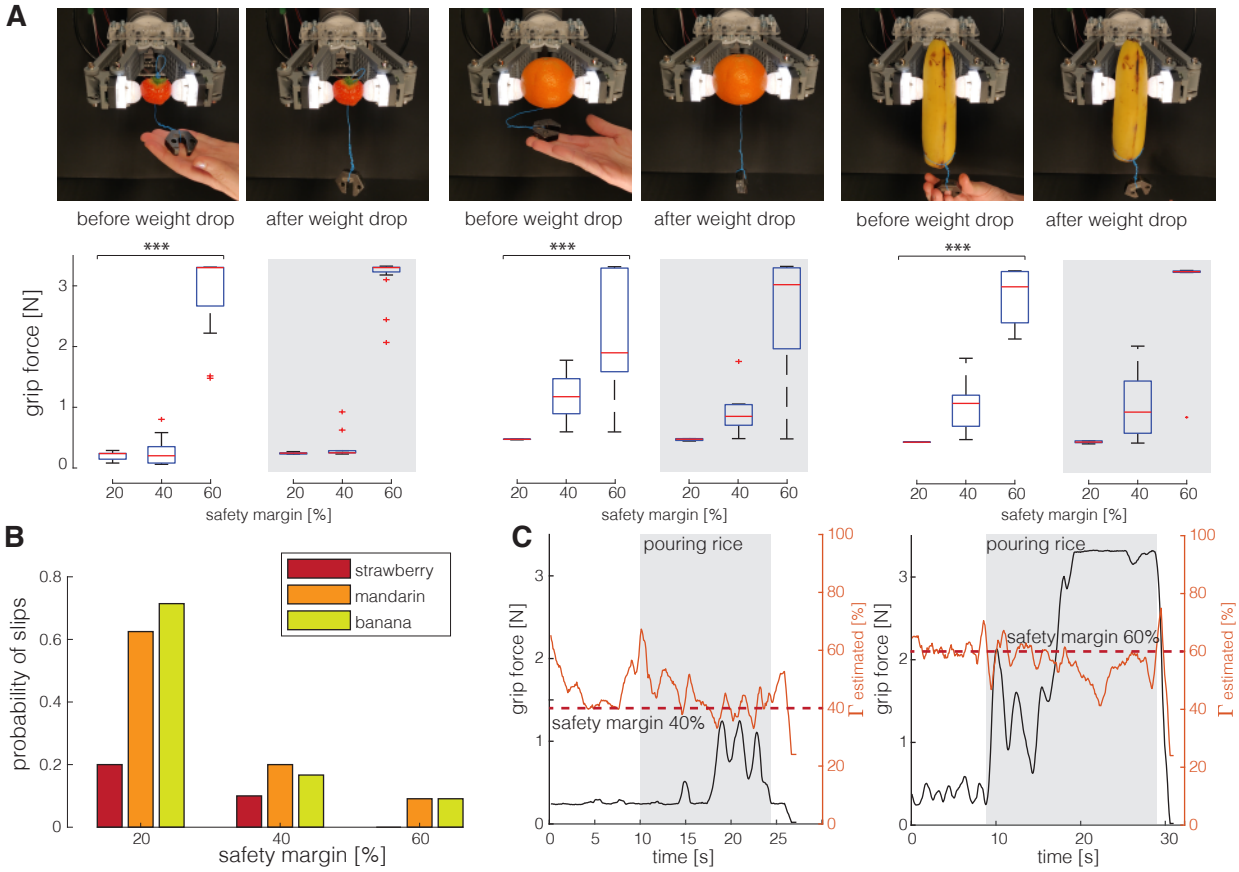

Fig. 5. **A**. Mean grasping force for 3 different fruits and 3 safety margin commands. The gray zone represents the period after the weight is dropped. **B**. Probabilities of slip as a function of the safety margin for the 3 fruits. **C**. Force and estimated safety margin when the gripper is grasping a cup being filled with rice. The safety margin was controlled at 40% and 60% respectively on the left and on the right.

force was limited to 3 N. This low accuracy can result in difficulties in handling fragile and slippery objects, as an underestimation of the safety margin will result in a higher grasping force than necessary. However, the study shows that the optimal target safety margin is around 40% so low values of safety margin will be rarely reached. In cases where the predictions are less accurate, the control is flexible and can quickly regulate a higher safety margin setpoint. The error in the estimation is also illustrated by a large standard deviation of the estimated safety margin. This uncertainty can result in fluctuations in grasp force.

A second validation has been done by evaluating the grasping performance of fragile real-life objects. These experiments show that, although the model has been trained on two object shapes, it can generalize to more complex scenarios. The frictional safety margin can also be used to increase grasping performance while reducing object damage. The trained network predicts Γ at 50 Hz on a desktop CPU, making it fast enough for real time control. The reaction time of the gripper after a weight drop has been measured at approximately 100 ms, which is in the same order of magnitude as the human reaction times to an external perturbation [23].

ACKNOWLEDGMENTS

This work was supported by the 4TU Soft Robotics program. J.L. acknowledges the EU's H2020 OpenDR project (grant No 871449). We would like to thank Jens Kober for insightful comments, Mostafa Atalla, Marlies Popken and Max Polak for technical assistance.

REFERENCES

[1] A. Rodriguez, M. T. Mason, and S. Ferry, "From caging to grasping," *The International Journal of Robotics Research*, vol. 31, no. 7, pp. 886–900, 2012.

[2] L. Zhang and J. C. Trinkle, "The application of particle filtering to grasping acquisition with visual occlusion and tactile sensing," in *2012 IEEE International Conference on Robotics and Automation*. IEEE, 2012, pp. 3805–3812.

[3] T. Bi, C. Sferrazza, and R. D'Andrea, "Zero-shot sim-to-real transfer of tactile control policies for aggressive swing-up manipulation," *IEEE Robotics and Automation Letters*, vol. 6, no. 3, pp. 5761–5768, 2021.

[4] N. F. Lepora, "Biomimetic active touch with fingertips and whiskers," *IEEE transactions on haptics*, vol. 9, no. 2, pp. 170–183, 2016.

[5] C. Sferrazza and R. D'Andrea, "Design, motivation and evaluation of a full-resolution optical tactile sensor," *Sensors*, vol. 19, no. 4, p. 928, 2019.

[6] 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.

[7] P. Griffa, C. Sferrazza, and R. D'Andrea, "Leveraging distributed contact force measurements for slip detection: a physics-based approach enabled by a data-driven tactile sensor," in *2022 International*

Conference on Robotics and Automation (ICRA). IEEE, 2022, pp. 4826–4832.

[8] F. Visentin, F. Castellini, and R. Muradore, "A soft, sensorized gripper for delicate harvesting of small fruits," *Computers and Electronics in Agriculture*, vol. 213, p. 108202, 2023. [Online]. Available: https://www.sciencedirect.com/science/article/pii/S0168169923005902

[9] S. Dong, D. Ma, E. Donlon, and A. Rodriguez, "Maintaining grasps within slipping bounds by monitoring incipient slip," in *2019 International Conference on Robotics and Automation (ICRA)*. IEEE, 2019, pp. 3818–3824.

[10] Z. Su, K. Hausman, Y. Chebotar, A. Molchanov, G. E. Loeb, G. S. Sukhatme, and S. Schaal, "Force estimation and slip detection/classification for grip control using a biomimetic tactile sensor," in *2015 IEEE-RAS 15th International Conference on Humanoid Robots (Humanoids)*. IEEE, 2015, pp. 297–303.

[11] J. Gao, Z. Huang, Z. Tang, H. Song, and W. Liang, "Visuo-tactile-based slip detection using a multi-scale temporal convolution network," 2023.

[12] J. W. James and N. F. Lepora, "Slip detection for grasp stabilization with a multifingered tactile robot hand," *IEEE Transactions on Robotics*, vol. 37, no. 2, pp. 506–519, 2020.

[13] F. Veiga, B. Edin, and J. Peters, "Grip stabilization through independent finger tactile feedback control," *Sensors*, vol. 20, no. 6, p. 1748, 2020.

[14] B. Delhaye, A. Barrea, B. B. Edin, P. Lefevre, and J.-L. Thonnard, "Surface strain measurements of fingertip skin under shearing," *Journal of The Royal Society Interface*, vol. 13, no. 115, p. 20150874, 2016.

[15] M. Tada, "How does a fingertip slip?-visualizing partial slippage for modeling of contact mechanics," in *2006 Proceedings of Eurohaptics*, 2006, pp. 415–420.

[16] L. Willemet, N. Huloux, and M. Wiertlewski, "Efficient tactile encoding of object slippage," *Scientific Reports*, vol. 12, no. 13192, 2022.

[17] F. Schiltz, B. P. Delhaye, J.-L. Thonnard, and P. Lefèvre, "Grip force is adjusted at a level that maintains an upper bound on partial slip across friction conditions during object manipulation," *IEEE Transactions on Haptics*, vol. 15, no. 1, pp. 2–7, 2021.

[18] A. M. Hadjiosif and M. A. Smith, "Flexible control of safety margins for action based on environmental variability," *Journal of Neuroscience*, vol. 35, no. 24, pp. 9106–9121, 2015.

[19] X. Lin and M. Wiertlewski, "Sensing the frictional state of a robotic skin via subtractive color mixing," *IEEE Robotics and Automation Letters*, vol. 4, no. 3, pp. 2386–2392, 2019.

[20] R. B. Scharff, D.-J. Boonstra, L. Willemet, X. Lin, and M. Wiertlewski, "Rapid manufacturing of color-based hemispherical soft tactile fingertips," in *2022 IEEE 5th International Conference on Soft Robotics (RoboSoft)*. IEEE, 2022, pp. 896–902.

[21] N. Ma, X. Zhang, H.-T. Zheng, and J. Sun, "Shufflenet v2: Practical guidelines for efficient cnn architecture design," in *Proceedings of the European conference on computer vision (ECCV)*, 2018, pp. 116–131.

[22] L. Wright and N. Demeure, "Ranger21: a synergistic deep learning optimizer," *arXiv preprint arXiv:2106.13731*, 2021.

[23] K. J. Cole and J. H. Abbs, "Grip force adjustments evoked by load force perturbations of a grasped object," *Journal of neurophysiology*, vol. 60, no. 4, pp. 1513–1522, 1988.