A Sense of Touch for the Shadow Modular Grasper

Nicholas Pestell, Luke Cramphorn, Fotios Papadopoulos and Nathan F. Lepora

***Abstract*— In this study, we have designed and build a set of tactile fingertips for integration with a commercial, three-fingered robot hand, the Shadow Modular Grasper. The fingertips are an evolution of an established optical, biomimetic tactile sensor, the TacTip. In developing the tactile fingertips, we have progressed the technology in areas such as miniaturisation, development of unconventional shaped finger-pads and integration of multiple sensors. From these fingertips, we extract a set of high-level features with intuitive relationships to tactile quantities such as contact location and pressure. We present a simple linear-regression method for predicting roll and pitch angle of the finger-pad relative to a surface normal and show that the method generalises to unknown depths. Finally, we apply this prediction to a grasp-control framework with the Modular Grasper and show that it can be used to adjust the grasp on three real-world objects from the YCB object set in order to attain a greater area of contact at each fingertip.**

I. INTRODUCTION

Robot hands have seen accelerated development in recent years [1], advancing attributes such as dexterity, grip strength and ease of use. Two-fingered grippers are deployed in large numbers for repetitive manufacturing tasks, whereas, more advanced, multi-fingered hands are yet to find applications outside of research. Thus, a gap persists for automation of small scale production, where robots are required to grasp and manipulate unknown objects [2]. This gap can only be filled by dexterous, multi-fingered robot hands.

Given the advances in the state-of-the-art of robot hands, it is surprising such hands have not yet found widespread application. One contributing factor may be a lack of sufficient tactile sensing capabilities. Indeed, it is known that humans rely heavily on their sense of touch to maintain a stable grasp [3]. Whilst there have been many attempts at improving grasp stability with the introduction of tactile sensing, primarily with data driven approaches [4]–[9], these methods are often impractical due to the large quantities of training data required and poor generalisability. In general, the tactile sensors are low resolution, array-based technologies, which do not allow for extraction of high-level features with a more direct relationship to the object being held.

NP, LC and NL were supported by an EPSRC IAA award on a 'Tactile smart grasping system'. NL was supported by a Leverhulme, Leadership Award on a 'A biomimetic forebrain for robot touch' (RL-2016-39).

NP, LC and NL are with the Department of Engineering Mathematics and Bristol Robotics Laboratory, University of Bristol, Bristol, U.K. FP is with the Shadow Robot Company, London, U.K.
NP, LC NL affiliation E-mail: {n.pestell, ll14468, n.lepora}@bristol.ac.uk, FP affiliation email: fotios@shadowrobot.com.

Fig. 1: Image of the developed tactile sensors integrated with the Shadow Modular Grasper. Base, proximal and distal joints are labelled **B**, **P** and **D** respectively.

This study presents the development of a high-definition, biomimetic tactile fingertip and its integration with a three-fingered, fully-actuated, robot hand: the Shadow Robot Company's Modular Grasper (Fig. 1). We extract a set of features from high-dimensional raw tactile images and infer information relevant to grasp quality using simple algorithms and small amounts of training data. The overall aim is to improve grasp quality in a way that is robust to variations such as object shape, orientation and weight.

II. BACKGROUND AND RELATED WORK

A seminal study by Kawasaki et al. [10] in 2002 set a high benchmark for forthcoming work into endowing robot hands with a sense of touch. The authors presented a sophisticated, 16-DOF, anthropomorphic hand, the Gifu hand II. The hand was equipped with 624 resistive tactile pads distributed across all five fingers and the palm.

In 2011, Romano et al. [11] conducted an innovative study into a grasp control framework using touch as an integral component, using comparatively rudimentary hardware: two 5x3 capacitive tactile sensors and an accelerometer integrated with the PR2 two-fingered gripper. Different states within the control system were triggered by hard-coded tactile signals.

Since Romano's influential paper, many researchers have integrated tactile sensors with dexterous robot hands [12]–[15], leveraging this sense of touch to predict grasp stability [4], [5] and/or perform grasp stabilisation [7], [9]. In contrast to the Romano's study, however, the trend has been towards data driven methods:

In [4], an under-actuated, Robotiq gripper was equipped with 6 array-based pressure sensors. A kernel logistic regression model, trained with tactile data from 192 grasps, predicted grasp success to 89% accuracy. Similarly, in [5], two 7x4 capacitive tactile arrays were integrated on the finger-pads of a Robotiq two-fingered, griper. 1000 grasps trained a CNN to achieving a prediction accuracy of 88.4%.

The authors of [7] propose a method for effective re-grasp: a three-fingered Barrett hand was equipped with array-based Biotac tactile sensors. Reinforcement learning was used to predict stable re-grasp of a single object, trained with a total of 50 hours of real robot data. Grasp success was improved from 42% to 97%.

The bias towards these data-driven methods is partially due to low resolution of the tactile sensors employed. Without sufficient acuity, sensors provide un-intuitive representations of tactile contact. In contrast, optical tactile sensors, e.g. the Gelsight [16] and the TacTip [17], provide high-resolution tactile images which ease interpretation.

The Gelsight was integrated with the Weiss WSG-50 parallel gripper [9], where deep CNNs predicted the probability of successful re-grasp given a proposed action. An impressive re-grasp success rate of 83.8% using vision and tactile was achieved. However, this method required extreme quantities of data; 6,450 grasps on 65 objects.

Here our presented fingertip design is based on an established tactile sensing device: the TacTip [17], [18]. Originally developed in 2009, the TacTip is an optical tactile sensor with a biomimetic design based on human tactile sensing. Physical interactions cause deformation of a soft skin producing deflections of markers on the inside of the sensing surface. The inside surface is illuminated by an array of LEDs and imaged with a camera system. The markers are tracked by an image processing algorithm and their positions are mapped to a classification of the tactile stimulation [19].

The TacTip has previously shown potential for integration with robot hands, for example in [20] where a TacTip based tactile thumb (TacThumb) device was integrated with the Open-Hand model M2 gripper and in [21] where two TacTip devices were mounted as fingertips on the Open-Hand model GR2 gripper. In both of these studies, the authors utilise a supervised-learning method to achieve precise in-hand manipulation of custom-made objects.

Whilst the aforementioned literature shows clear benefits of tactile sensing for grasping, a common shortcoming is the nature of the tactile data available and/or the amount of data required to interpret it. Here we present a more flexible platform by integrating a highly sensitive, high-resolution, optical tactile sensor with a fully-actuated industrial robot hand. We demonstrate the capability of the developed system for improving grasps on unknown, real-world objects.

III. MATERIALS AND METHODS

A. Shadow Modular Grasper

The Shadow Modular Grasper is an extensive yet practical, three-fingered hand. By enabling a large set of functional configurations it simplifies the task of grasping real-world objects from small precision tools to large assembly parts.

The hand is fully actuate, with 9 degrees of freedom (three per finger). The system is fully integrated with ROS and the user can control both specific joints and whole-hand grasps. Each identical finger has base, proximal and distal joints with dedicated actuation and can be easily attached and detached resulting in a modular system. The full hand has a total mass of 2.7 kg and a payload of 2 kg. Each finger also features a back-drivable gearbox enabling inherent compliance, which is an essentially feature when working in unstructured environments. The whole unit requires only two connections: power and comms (EtherCAT) [22].

B. Tactile Sensing

1) Tactile Fingertip Design: Tactile sensing is enabled by replacing the fingertips of the Modular Grasper with three custom-built tactile sensors (Fig. 1). The fingertips are comprised of two main components: a compliant finger-pad which deforms when contacted and a rigid body which fixes the camera in place and houses other electronics (Fig. 2).

The finger-pad is fully 3D-printed as a single part with a multi-material 3D-printer (Stratasys Objet). A sensing surface 'skin' is printed in Tango Black+ (Shore A 26-28). Importantly, the skin is deformable, which both enables transduction of tactile stimulation and also prevents damage to grasped objects. The inside of skin is tessellated with a triangular pattern of 97 pins (Tango Black+, 3 mm length and 2 mm diameter). White markers on the end of these pins are printed in rigid Vero White and provide a visual representation of the tactile stimulation. A rim (Vero White) enables a press-fit connection to the fingertip body. A clear acrylic sheet (2 mm thick) is glued into the rim resulting in a small cavity between skin and the acrylic lens. This cavity is filled with a two-part cure, clear, silicone gel (RTV27905, Techsil UK (∼Shore OO 10)) using a manual injection method. The gel helps to reduce hysteresis whilst still enabling deformation. The finger-pad shape closely matches the shape of the Modular Grasper's original fingertip contact surface.

The finger-pad is press-fit into a hollow fingertip body. The body is 3D-printed in ABS. A 2.0 megapixel CMOS array USB web-cam (ELP cameras) is mounted on the back of the fingertip body via four M2 screws. The camera is used in HD mode (1920x1080). The shape and size of the fingertip body was designed to provide a full view of the markers whilst minimizing the overall depth of the fingertip. The markers are illuminated by four LEDs which are arranged on two narrow PCB strips of two LEDs each. The PCBs are glued to the inside of the body, close to the interface with the finger-pad (Fig. 3).

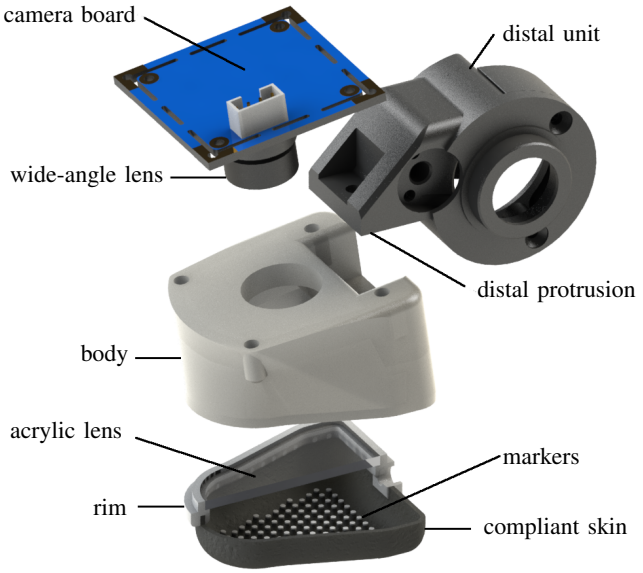

Fig. 2: Exploded CAD model of the tactile fingertip.

Evolving the TacTip into a practical fingertip for robot hands requires a reduction of size which provides a set of challenges unique to this brand of optical tactile sensor. As the camera is brought closer to the sensing surface, edge markers move out of view and focus is lost. To overcome this, a number of lenses were considered for increasing view angle and close-up focus. The final system employs a wide-angle lens (2.1 mm focal length, 150° view angle), which enables focused images of the full marker array (Fig. 4a).

2) Hardware Integration: The aim is for modular tactile fingertips without modifying other parts of the hand. This introduces a set of major design challenges. Shown in Fig. 3, the distal unit fixes to the fingertip by way of a ∼15 mm protrusion dictated by the original hand design. To overcome occlusion of the markers by the distal protrusion, the camera is shifted forward and fixed off-parallel from the sensing surface by an angle of ∼8°. This geometry is unique from all previous versions of the TacTip, where the camera is mounted directly above and parallel to markers.

Integration of three sensors has, to date, not been achieved with TacTip-based sensing. A solution proposed here is to connect each tactile fingertip to its own dedicated USB-hub. With three dedicated hubs, the data transfer occurs in parallel without reducing the frame rates, ∼20 fps per camera.

3) Tactile Feature Extraction: Feature extraction is performed using a Voronoi method previously demonstrated to achieve direct inference of pressure and contact locations with the TacTip [23]. Marker positions are tracked with a simple blob detection algorithm implemented with OpenCV in Python. A Voronoi tessellation is created over the sensor skin, treating the marker positions as seeds (Python, SciPy) (Figs. 4a and b). The areas of each Voronoi cell are related to local skin deformation, where increased size corresponds to indentation and hence pressure.

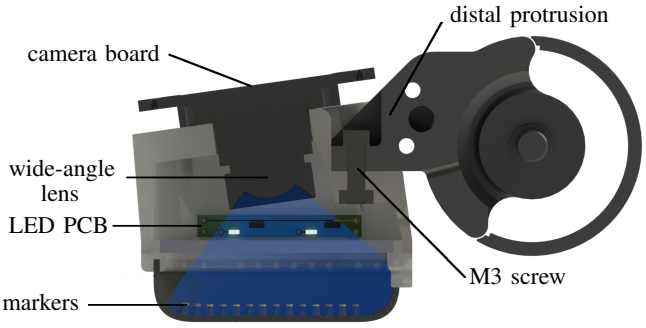

Fig. 3: Cross-section view of tactile fingertip and distal unit. The view angle of the camera is shown in blue.

Visual representations of the surface deformation are obtained by interpolating the change in Voronoi cell areas over the fingertip. A centre-of-pressure, a tactile analogue of centre-of-mass, is computed as an average of marker positions weighted by their corresponding cell area (Figs. 4c and d).

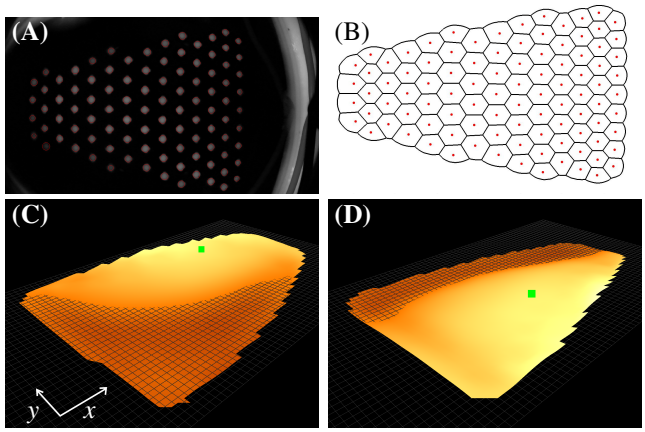

Fig. 4: **(A)**: View of markers imaged by the camera and tracked with OpenCV. **(B)**: Voronoi tessellation over markers. **(C)** and **(D)**: Visual representation of surface deformations with centre-of-pressure shown as a green spot.

C. Off-line Testing

Orientation of the fingertip relative to the contact surface may be of importance when grasping an object. For example, greater frictional forces are achieved with a larger contacting surface area, which is affected by relative angle between fingertip and object.

Here we examine the ability of the presented sensor and described feature extraction technique for perceiving roll and pitch relative to a flat surface.

1) Data collection: The fingertip is mounted as an end-effector on a six degree-of-freedom robot arm (UR5, Universal Robotics). The sensor maintains continual contact with a flat acrylic plate and the robot re-orients the sensor relative to the plate. Data is sampled randomly from a 2D grid of roll, ϕ, and pitch, θ, values,

$$-16^\circ \leq \phi \leq 16^\circ, \quad \text{and} \quad -11^\circ \leq \theta \leq 3^\circ, \qquad (1)$$

Experimental set-up is shown in Fig. 5. ϕ and θ angles are equally space by 2° and 1° respectively yielding, a total of $N_\phi N_\theta = 17 \times 15 = 255$ data points.

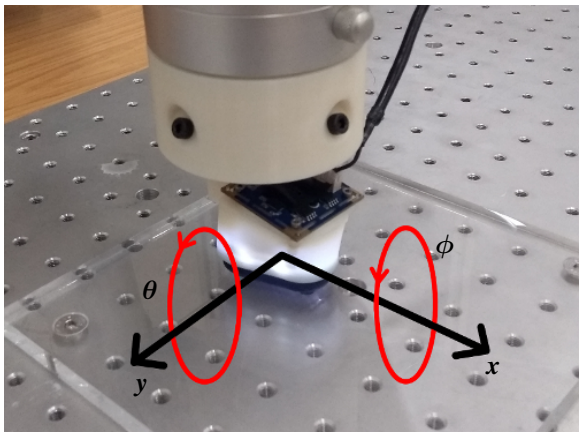

Fig. 5: Data collection set-up with tactile fingertip mounted as an end-effector on a UR5 robot arm. Showing roll, ϕ, and pitch, θ orientation relative to the sensor.

Three seconds of training data (~ 60 frames) are collected for each sample. A sample consists of a time series of centre-of-pressure vales r_{ki}, where $1 \leq k \leq N_{\text{frames}}$ and $1 \leq i \leq N_{\text{dims}}$; $N_{\text{frames}} \simeq 60$ and $N_{\text{dims}} = 2$ for x and y positions.

Three separate test sets are collected of 200 data points each, sampled at random from a continuous space within the same range of ϕ and θ. Each set is collected at a different depth since we are interested in assessing the degree to which our methods are pressure invariant. The depths, -0.5 mm $\leq z_l \leq$ 0.5 mm, are equally spaced, where $z_l = 0$ mm corresponds to the depth used for training.

The data collection procedure for training and testing is repeated three times, once for each fingertip.

2) Perception: Prior to training, data is averaged across frames, so each sample has $N_{\text{dims}} = 2$ features. We map centre-of-pressure-xy position to ϕ and θ via three separate multivariate, linear models: A simple linear model (1st-order polynomial), 2nd- and 3rd-order polynomials.

D. System Integration - Improving Grasp Stability

For the purpose of this study, we intend to use the predicted ϕ and θ (Section III-C) to adjust a grasp. Three Python drivers, one for each sensor, run on the host PC and interact with the grasp controller (C++) via a ROS-network.

The grasp controller builds upon the existing method for the Shadow Modular Grasper: The high-level strategy consists of independent control of each joint through its dedicated motor (full actuation). Each joint can be controlled in either position or torque mode. A grasp is comprised of two distinct phases: *(i)* closing phase and *(ii)* adjustment phase. These two phases are implemented within an update loop running at 1 KHz which iteratively updates joint modes and targets according to sensor data.

i) During the closing phase all joints are controlled in position mode and commanded to a set of target angles via a PID which is an implementation of the roscontrol ROS-package. Throughout this phase, the hand controller is listening to a set of 'contact-detection' ROS-topics, published by each driver. Contact is determined by upward threshold crossing on the surface deformation (Fig. 4). Upon contact-detection each respective finger stops moving.

TABLE I: R^2 scores for 1st- (simple), 2nd- and 3rd-order polynomial linear regression for tips A, B and C, calculated for test sets at contact depths of -0.5, 0 and 0.5 mm.

	1st order			2nd order			3rd order		
Depth (mm)	**-0.5**	**0**	**0.5**	**-0.5**	**0**	**0.5**	**-0.5**	**0**	**0.5**
tip-A	0.91	0.90	0.91	0.92	0.91	0.92	0.84	0.95	0.67
tip-B	0.87	0.89	0.72	0.89	0.92	0.71	0.85	0.94	0.28
tip-C	0.88	0.88	0.81	0.90	0.90	0.82	0.76	0.93	0.40

ii) After all three sensors have detected contact the adjustment phase is entered: the controller switches all proximal joints to torque mode and applies a fixed squeezing torque to the grasped object. Base and distal joints remain in position mode and are servoed with the same PID controller as used in phase-*i)*. Here the PID input is ϕ and θ predictions for base and distal joints respectively. These predictions are again received from the sensors over a set of ROS-topics. To steady the grasp, this adjustment phase has a time-out period, the length of which, is obtained from experimentation (Section IV-B). see Fig. 1 for reference of joint names.

IV. RESULTS

A. Off-line Testing

Here we compare three competing model types (simple, 2nd- and 3rd-order linear models) for predicting ϕ and θ. Our model should accurately approximate the training data whilst generalising to unobserved samples. This latter criterion is particularly important when working with flexible platforms such as robot hands. Table I shows R^2 scores for three competing model types, for all three tips at each test depth.

Where test sets were collected at training depth (0 mm), the 3rd-order polynomial model achieved the highest R^2 scores across all three tips (0.95, 0.94, and 0.93). These were also the highest scores over the whole set. This model, however, under-performed at the two previously unseen depths, particularly for the highest position (0.5 mm) where it achieved the lowest three R^2 scores of the whole set (0.67, 0.28 and 0.40).

The 3rd order model is expected to perform well on data close to the training set, being the most flexible. However, this flexibility is likely to be the cause of poor performance seen at other depths, as the model responds to details which are specific to the training data.

In contrast, both 1st- and 2nd-order polynomials perform more consistently across all depths: The lowest R^2 scores were 0.72 and 0.71 for the 1st- and 2nd-order polynomials respectively. This improved consistency is likely due to a lower variance in the models, giving better generalisation.

The 2nd-order polynomial slightly outperforms the 1st-order model at all but one tip-depth combination. The mean R^2 scores are 0.86 and 0.88 for 1st- and 2nd-order polynomials respectively. We conclude, therefore, that the 2nd-order model is the most suitable of the presented methods for predicting ϕ and θ for application with robot hands. This is a due to a suitable balance between bias and variance such that model performs well across all of the presented test sets.

Visualisations of the 2nd-order polynomial model for tip-B and the data are shown in Fig. 6. The data is well ordered

and the model appears a suitable fit. We observe strong correlations of centre-of-pressure-x and -y positions with θ and ϕ respectively. This is expected since the x- and y-axis align with θ and ϕ respectively (see Fig. 5). We also observe some correlation of centre-of-pressure-y and -x positions with and θ ϕ respectively, particularly when the centre-of-pressure is at extremes of the y-axis. This suggests that both features are useful predictors, particularly when ϕ is large, justifying our use of a multiple regression model.

B. Improving Grasp Stability

Here we investigate the capabilities of the integrated Shadow Robot Modular Grasper with tactile fingertips. Specifically, we look at the capacity for judging quality of grasp based on metrics used throughout this paper (ϕ and θ), which are predicted using a 2$^{\text{nd}}$-order polynomial regression model (Section IV-A). Then we explore the potential for adjusting the hand pose according to the control framework described in section III-D.

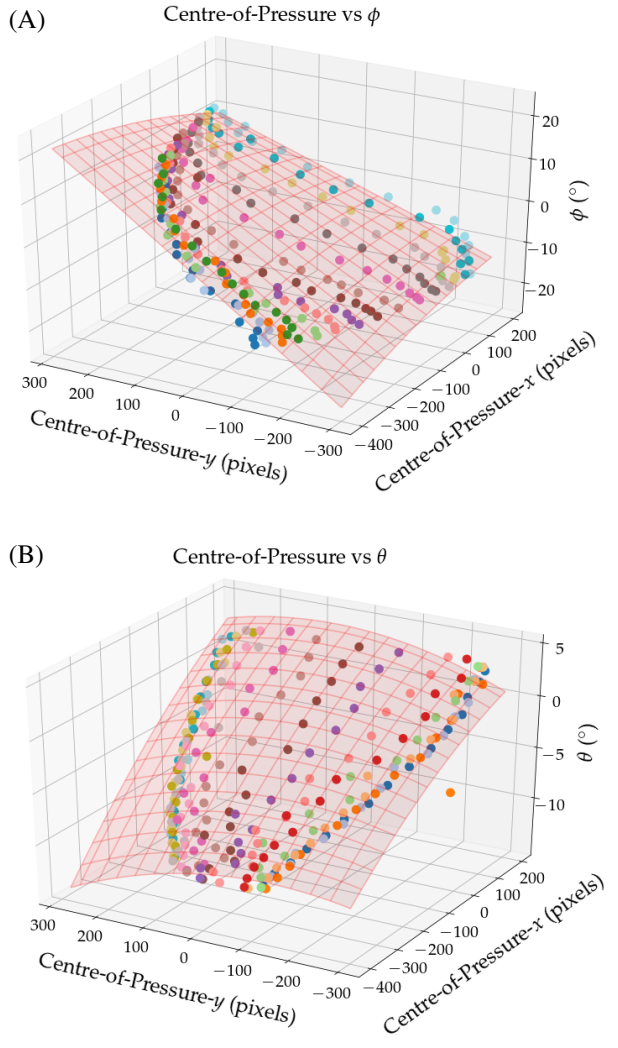

Fig. 6: **(A)**: Scatter plot of xy-centre-of-pressure vs ϕ. Each colour represents a constant θ. **(B)**: Scatter plot of xy-centre-of-pressure vs θ angle. Each colour represents a constant ϕ. In both plots the surface shows a 2$^{\text{nd}}$ degree polynomial fit.

Fig. 7 shows images of the hand grasping three objects from the YCB object set: a Rubik's cube (left), Pringles can (centre) and mustard bottle (right). Alongside each image are tactile visualisations from each fingertip associated with that grasp. The objects are initially held in place by a human participant before passing over to the robot when all three fingers have made contact. The top row shows images at initial contact detection (prior to tactile adjustment) and the bottom row shows images ∼10 s later after the base and distal joints of each finger have been adjusted according to tactile sensing. We also provide a supplementary video of the robot completing these three grasps.

In Fig. 8, we show the base and distal joint angles and centre-of-pressure-x and -y positions plotted against time when grasping the Rubik's cube, for fingers A, B and C, (as labelled in Fig. 7). All three fingers detect contact at approximately the same time, as shown by the green lines. Subsequently, both joint angles are adjusted before reaching the time-out, observed in Fig. 8 as a flattening of the blue curves. We chose a time-out of 10 s which gives the hand suitable time to re-adjust. During the same period, the centre-of-pressure position was adjusted towards 0 for both x and y dimensions. The shapes of these curves suggests that hand control performed as designed: to servo base and distal joints in order to shift the centre-of-pressure to the middle of each finger-pad. Furthermore, the shape of base and distal joint angle curves are roughly mirrored by the centre-of-pressure-y and x positions respectively, suggesting that there is a strong relationship between these variables.

Fig. 7 shows that the hand made initial contact (top row) with each object on all three fingers, with tactile visualisations showing clear deformation on all fingertips. The object is maintained in elevation, indicating that the initial grasp was successful in all cases.

The top and bottom rows of Fig. 7 show noticeable differences in both the grasp images and the tactile visualisations. In general, the grasp images show that the fingertips rotated around each object to minimise ϕ and θ as described in Fig. 5. Inspection of the tactile visualisations suggests that overall deformation of each fingertip increased subsequent to adjustments in all cases. Again, this suggests that the grasp controller performed as designed: to increase contact surface area and therefore improve grasp quality.

The hand made a good initial grasp of the Pringles can owing to its symmetrical shape, with ϕ for all three fingertips close to zero prior to adjustment. We observe a slight modification of θ for each fingertip which in-turn increased the contact surface area.

The mustard bottle is the most irregular of the three presented objects, as confirmed in the variance of tactile visualisations prior to adjustment. Despite this, the hand was able to maintain a stable grasp throughout the adjustment phase and, greatly increase the contact surface area on all three fingertips.

The Rubik's cube, produced a particularly interesting result: fingertips A and C, rotated from glancing edge contacts to contacting on individual faces, naturally increasing the

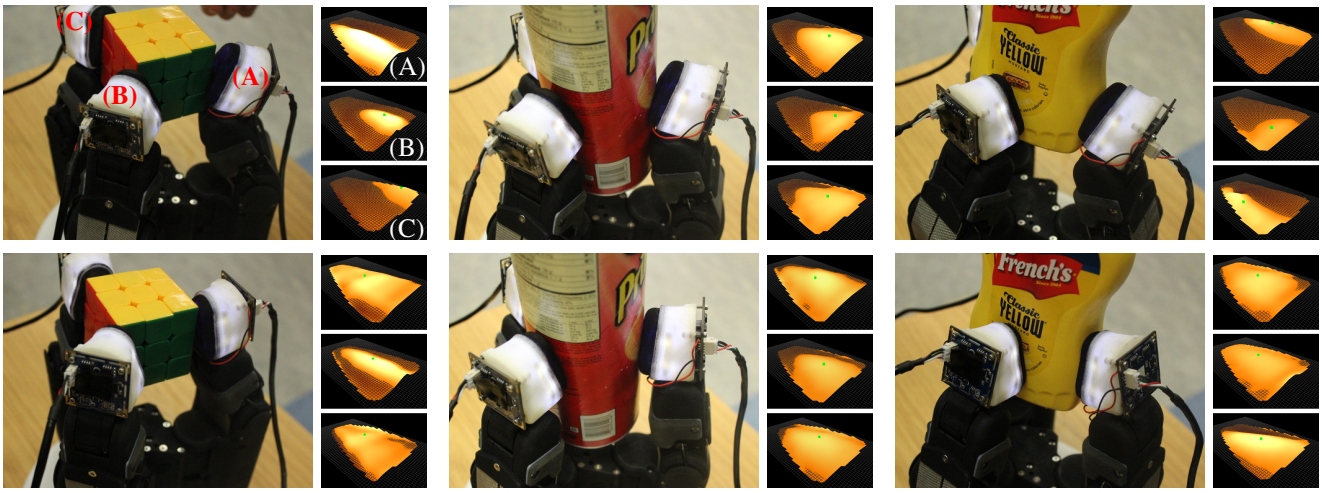

Fig. 7: Images of the performed grasps on the Rubik's cube, Pringles can and mustard bottle, before and after tactile adjustment, top and bottom rows respectively. Tactile visualisations for the three fingertips are displayed to the right of each grasp image. The top left image and visualisations are labelled to show how the visualisations relate to the image. Hand orientation and visualisation order is consistent throughout the figure.

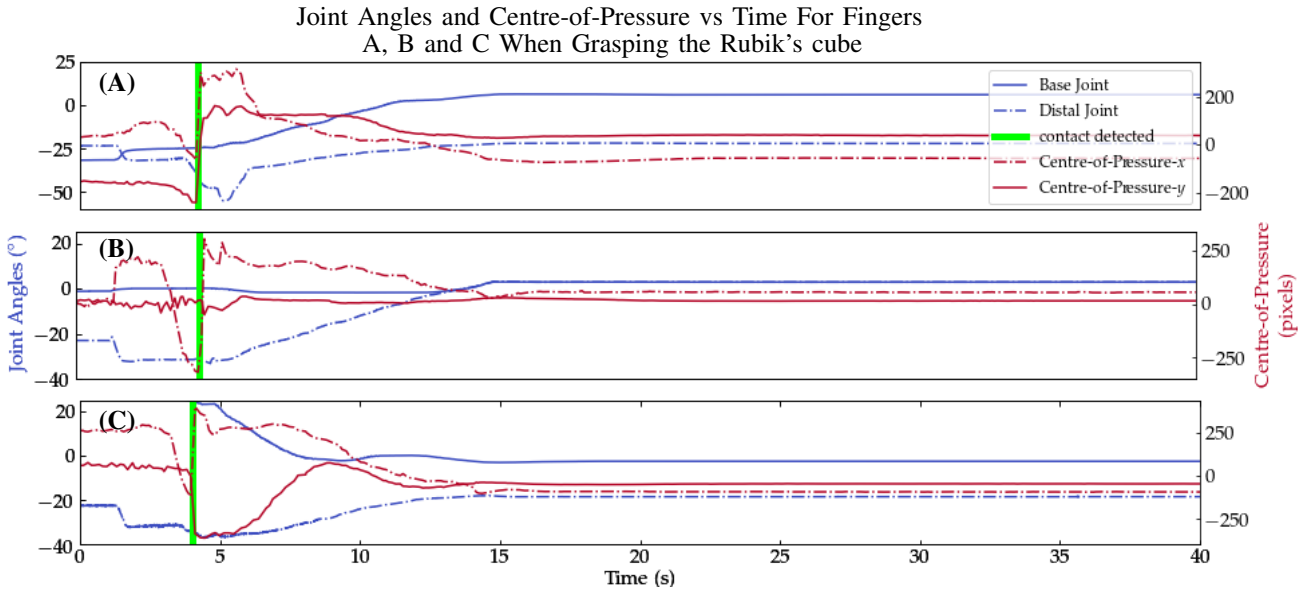

Fig. 8: Plots of base and distal joint angles (blue) and *xy* centre-of-pressure (red) versus time, for fingers A, B and C, whilst grasping the Rubik's cube. Vertical green lines shows when each finger detected contact.

surface area of contact. In contrast, fingertip B maintained its location on the edge of the cube. Whilst this grasp may not maximise the overall contact surface area, it is a stable configuration given its starting point. This suggests that the method is able to produce useful results even when presented with shapes not seen in training.

V. DISCUSSION

In this study, we presented the integration of an established optical tactile sensing technology, the TacTip, with a three-fingered, commercial robot hand: the Shadow Modular Grasper. The sensors themselves were tested for their capacity to predict roll and pitch angle relative to a flat surface. Finally, we integrated tactile output with the hand control and demonstrated grasps on three objects from the YCB set, using predicted roll and pitch angles to adjust the grasp for attaining greater contact surface areas at each finger-pad.

A set of specific challenges were presented throughout the design process. Notably, parallelisation of the three tactile sensors, which was overcome by utilising three dedicated USB-hubs. Furthermore, the sensors were miniaturised to a suitable fingertip size for the Modular Grasper, which required experimentation with camera location and lens types.

For predicting roll and pitch angles, three linear regression models were compared; simple, 2nd- and 3rd-order polynomials. The 2nd-order model was the most suitable due to its balance of accuracy and generalisation.

The Modular Grasper, equipped with the designed tactile fingertips, was able to successfully grasp and hold a Rubik's cube, Pringles can and mustard bottle. Additionally, the

hand maintained a grasp whilst adjusting all three base and distal joints according to the perceived roll and pitch angles. With this tactile adjustment, the finger-pads obtained greater surface area of contact and a more centrally located pressure, both of which are features that improve grasp quality.

Recently, OpenAI achieved in-hand manipulation with impressive levels of dexterity using deep reinforcement learning. To aid training, simulation was used, however, this is a challenge for touch, hence, this modality was omitted from their research. This highlights a need for learning-free approaches to using touch for application with any hand, without the need for re-training or simulation.

A key benefit of the presented approach to grasp stabilisation is that, whilst it requires high-resolution tactile sensing, it can be easily applied to any robotic hand with the required degrees of freedom. This is owing to the intuitive output from the sensors (roll and pitch angle) which can be directly applied to joint angles of the gripper. As a consequence no prior training with the robot hand is required.

REFERENCES

[1] Z. Xu and E. Todorov, "Design of a highly biomimetic anthropomorphic robotic hand towards artificial limb regeneration," in *2016 IEEE International Conference on Robotics and Automation (ICRA)*, May 2016, pp. 3485–3492.

[2] Z. Kappassov, J.-A. Corrales, and V. Perdereau, "Tactile sensing in dexterous robot hands review," *Robotics and Autonomous Systems*, vol. 74, pp. 195 – 220, 2015. [Online]. Available: http://www.sciencedirect.com/science/article/pii/S0921889015001621

[3] R. S. Johansson and G. Westling, "Roles of glabrous skin receptors and sensorimotor memory in automatic control of precision grip when lifting rougher or more slippery objects," *Experimental Brain Research*, vol. 56, no. 3, pp. 550–564, Oct 1984. [Online]. Available: https://doi.org/10.1007/BF00237997

[4] E. Hyttinen, D. Kragic, and R. Detry, "Learning the tactile signatures of prototypical object parts for robust part-based grasping of novel objects," in *2015 IEEE International Conference on Robotics and Automation (ICRA)*, May 2015, pp. 4927–4932.

[5] J. Kwiatkowski, D. Cockburn, and V. Duchaine, "Grasp stability assessment through the fusion of proprioception and tactile signals using convolutional neural networks," in *2017 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, Sept 2017, pp. 286–292.

[6] L. Pinto and A. Gupta, "Supersizing self-supervision: Learning to grasp from 50k tries and 700 robot hours," in *2016 IEEE International Conference on Robotics and Automation (ICRA)*, May 2016, pp. 3406–3413.

[7] Y. Chebotar, K. Hausman, Z. Su, G. S. Sukhatme, and S. Schaal, "Self-supervised regrasping using spatio-temporal tactile features and reinforcement learning," in *2016 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, Oct 2016, pp. 1960–1966.

[8] M. Li, K. Hang, D. Kragic, and A. Billard, "Dexterous grasping under shape uncertainty," *Robotics and Autonomous Systems*, vol. 75, pp. 352 – 364, 2016. [Online]. Available: http://www.sciencedirect.com/science/article/pii/S0921889015001967

[9] 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," *CoRR*, vol. abs/1805.11085, 2018. [Online]. Available: http://arxiv.org/abs/1805.11085

[10] H. Kawasaki, T. Komatsu, K. Uchiyama, and T. Kurimoto, "Dexterous anthropomorphic robot hand with distributed tactile sensor: Gifu hand ii," in *IEEE SMC'99 Conference Proceedings. 1999 IEEE International Conference on Systems, Man, and Cybernetics (Cat. No.99CH37028)*, vol. 2, Oct 1999, pp. 782–787 vol.2.

[11] J. M. Romano, K. Hsiao, G. Niemeyer, S. Chitta, and K. J. Kuchenbecker, "Human-inspired robotic grasp control with tactile sensing," *IEEE Transactions on Robotics*, vol. 27, no. 6, pp. 1067–1079, Dec 2011.

[12] A. J. Spiers, M. V. Liarokapis, B. Calli, and A. M. Dollar, "Single-grasp object classification and feature extraction with simple robot hands and tactile sensors," *IEEE Transactions on Haptics*, vol. 9, no. 2, pp. 207–220, April 2016.

[13] M. Kaboli, A. D. L. R. T, R. Walker, and G. Cheng, "In-hand object recognition via texture properties with robotic hands, artificial skin, and novel tactile descriptors," in *2015 IEEE-RAS 15th International Conference on Humanoid Robots (Humanoids)*, Nov 2015, pp. 1155–1160.

[14] A. J. Spiers, M. V. Liarokapis, B. Calli, and A. M. Dollar, "Single-grasp object classification and feature extraction with simple robot hands and tactile sensors," *IEEE Transactions on Haptics*, vol. 9, no. 2, pp. 207–220, April 2016.

[15] H. Soh and Y. Demiris, "Incrementally learning objects by touch: Online discriminative and generative models for tactile-based recognition," *IEEE Transactions on Haptics*, vol. 7, no. 4, pp. 512–525, Oct 2014.

[16] R. Li and E. H. Adelson, "Sensing and recognizing surface textures using a gelsight sensor," in *2013 IEEE Conference on Computer Vision and Pattern Recognition*, June 2013, pp. 1241–1247.

[17] 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*, 2018.

[18] C. Chorley, C. Melhuish, T. Pipe, and J. Rossiter, "Development of a tactile sensor based on biologically inspired edge encoding," in *2009 International Conference on Advanced Robotics*, 2009, pp. 1–6.

[19] N. F. Lepora and B. Ward-Cherrier, "Superresolution with an optical tactile sensor," in *2015 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, 2015, pp. 2686–2691.

[20] B. Ward-Cherrier, L. Cramphorn, and N. F. Lepora, "Tactile manipulation with a tacthumb integrated on the open-hand m2 gripper," *IEEE Robotics and Automation Letters*, vol. 1, no. 1, pp. 169–175, Jan 2016.

[21] B. Ward-Cherrier, N. Rojas, and N. F. Lepora, "Model-free precise in-hand manipulation with a 3d-printed tactile gripper," *IEEE Robotics and Automation Letters*, vol. 2, no. 4, pp. 2056–2063, 2017.

[22] "Shadow Robot Company agile grasper documentation," https://agile-grasper.readthedocs.io/en/latest/, accessed: 2018-09-07.

[23] L. Cramphorn, J. Lloyd, and N. F. Lepora, "Voronoi features for tactile sensing: Direct inference of pressure, shear, and contact locations," in *2018 IEEE International Conference on Robotics and Automation (ICRA)*, 2018.