Can Modular Finger Control for In-Hand Object Stabilization be accomplished by Independent Tactile Feedback Control Laws?

Filipe Veiga[1] and Jan Peters[1,2]

Abstract— **Currently grip control during in-hand manipulation is usually modeled as part of a monolithic task, yielding complex controllers based on force control specialized for their situations. Such non-modular and specialized control approaches render the generalization of these controllers to new in-hand manipulation tasks difficult. Clearly, a grip control approach that generalizes well between several tasks would be preferable. We propose a modular approach where each finger is controlled by an independent tactile grip controller. Using signals from the human-inspired biotac sensor, we can predict future slip – and prevent it by appropriate motor actions. This slip-preventing grip controller is first developed and trained during a single-finger stabilization task. Subsequently, we show that several independent slip-preventing grip controllers can be employed together without any form of central communication. The resulting approach works for two, three, four and five finger grip stabilization control. Such a modular grip control approach has the potential to generalize across a large variety of in-hand manipulation tasks, including grip change, finger gaiting, between-hands object transfer, and across multiple objects.**

I. INTRODUCTION

To date, most grasping and in-hand manipulation approaches are based on global monolithic planning and control strategies. These monolithic policies determine finger placement, grasping configuration, finger trajectories, contact forces and contact locations for the entire hand throughout the entire manipulation task [1], [2]. Such an approach requires highly accurate kinematic/dynamic models of both hand and object along with precise sensing of positions (of both hand and object) as well as interaction forces. In practice, uncertainty is frequently too high for jointly using all of these components and, thus, the focus of grasping has become largely data-driven [3].

Unfortunately, data-driven approaches do not come for free but either require a large training data set, retricting the tasks to sufficiently similar scenarios [3] or a low-dimensional representation that covers the considered manipulation tasks (such as synergies [4], motion primitives [5], etc.). All data-driven solutions are obviously task- and platform-specific. The learned polices inherently couple the employed degrees of freedom, rendering the incorporation of tactile feedback from multiple fingers difficult.

To enable a control policy to both deal with uncertainty (e.g., in contact locations and forces) and to generalize well beyond a limited set of cases, an approach needs to be both data-driven and modular. Human grasping and manipulation appears to largely data-driven [6] but there is strong evidence that the resulting control strategies are just based on local sensing [7], rendering the control of the fingers largely independent from each other [7] thus, being highly modular. The correlations in human movement appear to arise from higher-level planning instead of low-level control [8]. Clearly such an approach would be desirable in robotic grasping and manipulation.

Ensuring grip stability of the object in the hand is frequently central to both grip control (i.e., stabilizing an object) and in-hand object manipulation (i.e., moving an object between stable grip configurations). Classical robotics approaches often rely on measures such as form- or force-closure for assessing grip stability; clearly for imperfect models and imperfect contact/force sensing, the use of such measures is not straightforward. Many researchers have proposed alternative grasp stability measures [9], [10], [11], [12], [13] and developed accompanying control strategies. For unknown objects, it may make more sense to reactively prevent instability than to maximize a potentially arbitrary stability measure [14]. In grasping and in-hand manipulation, instability is synonymous with slip [15]. While slip cannot

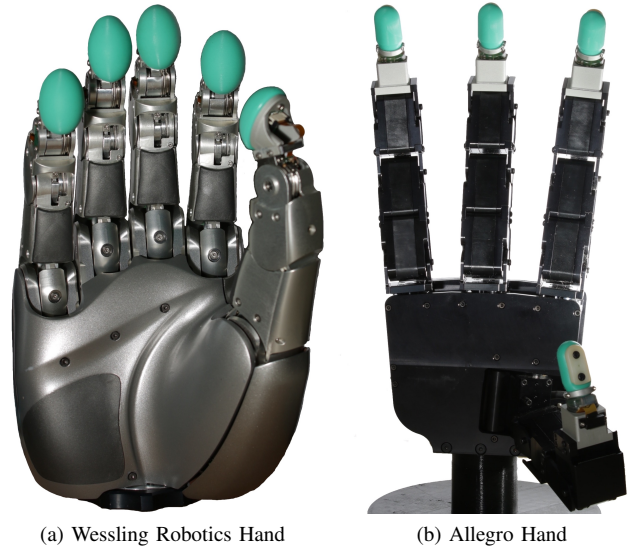(a) Wessling Robotics Hand   (b) Allegro Hand

Fig. 1. The controllers proposed in this paper were successfully evaluated both on the Wessling Robotics Five Finger Hand and the Allegro Four Finger Hand. Each hand only has identical fingers, which allows the learned models to be transferred easily across the four or five fingers of the respective hands. Different versions of the BioTac sensor are employed in both hands, i.e., the regular BioTac with 19 electrodes and the BioTac SP with 22 electrodes.

[1]Filipe Veiga and Jan Peters are with the Technische Universität Darmstadt, Darmstadt, Germany, FG Intelligent Autonomous Systems. e-mail: {veiga, peters}@ias.tu-darmstadt.de

[2]Jan Peters is also with the Max-Planck-Institut für Intelligente Systeme, Tübingen, Germany.

be directly measured, it can be detected and predicted based on high-frequency tactile sensing [14]. We propose that a grip stabilizing control law should aim to ensure that there will be no future slip.

Inspired by progression from one-finger over two-fingers to the whole hand proposed by [16] in the context of tactile object exploration and by the independent control hypothesis in human grip control by [7], [8], we aim to develop independent control policies based on tactile feedback for each finger that in conjunction generalize from one-finger to five-finger gripping and easy in-hand manipulation.

Edin/Burstedt's Hypothesis [7], [8]
While human grasp planning is global, human grib control is local in every finger – they "communicate" through the object. It does not really matter whether the fingers belong to different humans.

To accomplish this goal, we endow every finger with learned predictive models of future slip based on the high-frequency features of the tactile finger-object interaction. The local control law in the finger chooses the finger's movement such that probability of future slip is minimized. The resulting control law is capable stabilizing objects against other objects (such as a table or a wall), jointly stabilizing objects with more robotic fingers (as in in-hand object stabilization or gripping) or against the hand of a human operator (human-robot joint stabilization). It can also be employed for in-hand manipulation by stabilizing an object with several fingers while one or more fingers move the object within the stable grip. The coordination between modular finger controller occurs only indirectly through the tactile signals observed by each finger.

This modular approach enables a higher-level planning system to operate with less object knowledge while requiring fewer complex models for control than analytical approaches. The proposed reactive control framework generalizes potentially better than monolithic data-driven approaches across multiple tasks, a variety of objects and different robotic platforms. Our approach reproduces findings in human motor control where the absolute amount of force applied by single digits will always settle just above the minimal amount of forces to prevent slip [6], [7] in static setting.

Our Working Hypothesis
If a single finger controller can enable a human and a robot finger to hold objects well-together based only on actile sensing, then a modular *N*-finger grip control can be accomplished through *N* independent tactile sensing-based finger grip control.

II. MODULAR TACTILE SENSING-BASED IN-HAND OBJECT STABILIZATION

Classical approaches to sensing-based in-hand stabilization of objects have relied on accurate force-torque sensors, precise models of the held objects, optimal contact points, highly accurate finger control and the central computation of joint-level forces to be applied in order to accomplish stable force closure. Such requirements clearly cannot always be met and simpler approaches to sensing-based in-hand stabilization are desirable. In this paper, we propose a novel approach that goes into the opposite direction: We assume that single fingers can accomplish stabilization without central control purely relying on the prediction of slip based on tactile sensing and, thus, can be considered independent modules. As a first step, we introduce a single finger tactile control approach which works well for stabilizing objects pinned against other objects. Subsequently, in a second step, we discuss how this approach can be used in conjunction with more finger – from its own hand as well as from another agent; see Fig. 2 for an experiment where this method stabilized an object jointly with a human finger.

A. Single-Finger Stabilization based on Slip Prediction

Humans appear to have two components during tactile stabilization of objects, i.e., they assess the potential danger of slip and reactively adjust the applied forces until slip is no longer likely [7]. Thus, to obtain our single-finger tactile sensing-based object stabilization controller, we require two components, i.e., (1) a slip predictor and (2) a force adjustment method based on slip prediction.

1) Slip Prediction: We can formalize such an approach by treating slip prediction as a classification problem. Here, a classifier $f(\cdot)$ is employed to predict the state at time $t + \tau_f$ as slip or not slip (a preliminary version of this idea has been presented in [14]). To predict slip in such a way, features $\phi(\cdot)$ of the raw sensor values $\mathbf{x}_{(t-\tau_H):t}$ need to be extracted for a window in time (i.e., the last τ_H time steps). Thus, we need to train the slip predictors

$$c_{t+\tau_f} = f(\phi(\mathbf{x}_{(t-\tau_H):t})) \qquad (1)$$

using $c_{t+\tau_f} \in \{\texttt{slip}, \texttt{contact}, \neg\texttt{contact}\}$ as labels for the prediction at time-step $t + \tau_f$ where τ_f is the prediction horizon, i.e., it denotes the number steps into the future where slip is predicted.

2) Force adjustment: We require a control law that converts the predicted slip label $c_{t+\tau_f}$ into joint velocities $\dot{\mathbf{q}}_t$, as we only have velocity control access for both our robot hands (i.e., the Allegro hand and the Wessling five finger hand in Fig. 1) and for the PA-10 robot arm (in Fig. 2). Inspired by neuroscientific studies [7], our controller only maintains a sufficient statistic on the past slippage of the finger, i.e.,

$$l_{t+1} = \begin{cases} l_t + s_{\text{slip}} & \text{if } c_{t+\tau_f} = \texttt{slip}, \\ l_t - s_{\neg\text{slip}} & \text{if } c_{t+\tau_f} = \texttt{contact}, \\ l_t & \text{if } c_{t+\tau_f} = \neg\texttt{contact}, \end{cases} \qquad (2)$$

which basically is the weighted balance between the time where the object slip is predicted and where not. Here, adding a weight s_{slip} should result in the controller to apply more force to the object while subtracting $s_{\neg\text{slip}}$ will reduce the amount of applied forces. We employ the resulting statistic l_t in an exponential regulator that increases the desired velocity of the finger tip towards the contact direction if slip is
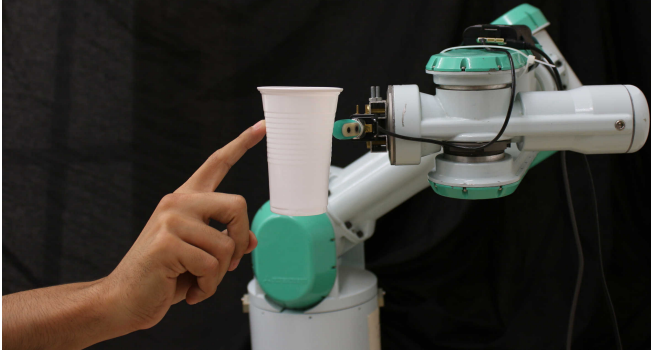
Fig. 2. Single finger stabilization (as described in Sec. II-A) can be decomposed into a slip prediction component and a control law that attempts to adjust all commands to obtain the minimum force where no slip occurs. Even when the slip predictors are trained by pinning objects against a fixed plane, these slip predictors generalize to multi-finger scenarios. Here, we depict a PA10 used as a single that predicts slip based on a single BioTac sensor but successfully stabilizes an object together with a human finger. The same predictors and controller are used in the multi-finger scenarios through-out this paper.

frequently predicted, implicitly increasing the the normal contact force. The velocity quickly goes to jittering around zero when no slip is predicted. These requirements result into the control law

$$\dot{\mathbf{q}}_t = \beta \mathbf{N}(\mathbf{x}_t) e^{\alpha l_t} \tag{3}$$

where $\beta > 0$ and $\alpha > 0$ are controller parameters and $\mathbf{N}(\mathbf{x}_t)$ denotes the estimated contact normal at time t.

B. Multi-Finger Gripping by Single-Finger Predicted Slip Control

When progressing towards in-hand stabilization and in-hand manipulation, more fingers are required and the complexity of the tasks quickly scales accordingly with the hand dexterity. Generally, a higher dimensionality can be coped with either by identifying a lower-dimensional manifold for the problem or by decomposing the problem. Following the core insight in [7] that human multi-finger grip stabilization appears to be accomplished by separate neural circuits that interact through the object instead of via the central nervous system. Following their insight, we hypothesize that *multi-finger robot gripping can be accomplished using the same single-finger predicted slip controller on all fingers independently from each other*. As a first scenario, we reproduce the scenario in [7], where two humans jointly hold an object using one finger each – with the same apparent ease as if one human was using one finger from each hand. The underlying neural control appears to be unaffected, no matter whether it interacts with finger from the same human or another human. We reproduce this effect replacing one human finger by a robot finger in a human-robot joint stabilization task as shown in Fig. 2 – and the single finger controllers of both the human and the natural finger directly work well together.

Thus, to fully utilize the dexterous capabilities of the hand, we propose that each hand should be seen as a set of independently controlled fingers[1]. A set of independent fingers, in contrast to a fully connected manipulator, allows decomposing the object stabilization control problem such that each finger separately predicts future slip based on tactile sensing and avoids it by independently adjusting the applied forces. While it may appear that synchronization only through the tactile feedback may appear counter intuitive, it actually not only greatly reduces the problem dimensionality but also ensures that the fingers only influence each other when it is needed for the stabilization of the object. As a result, it not only becomes more straightforward to design stabilizing control laws but the synchronization becomes more robust.

III. EXPERIMENTAL EVALUATION

To explore the capabilities of the concept of finger independence discussed in Sec.II-B when using the control strategy proposed in Sec.II-A.2, we perform several experiments. We attempt to stabilize the object with a varying number of fingers with both the Allegro and the Wessling hands, Sec. III-C. The showcase of aforementioned experiments is preceded by a detailed description of the two robotic platforms used, Sec.III-A.1, an account of the tactile sensors mounted on these platforms and their differences, Sec.III-A.2, and a detailed outline of the procedure used to acquire the ground truth data for the slip classifiers, Sec. III-B.

A. Experimental Setup: Testing Platforms & Tactile Sensors

We use two different robotic hands to assess the generalization capabilities of our approach across different robots. Each hand as sensorized fingertips, equipped with two different versions of the BioTac sensors. These two versions of the sensors, although similar, have significant differences with respect to shape and sampling frequencies. The Allegro and Wessling hands are introduced in Sec.III-A.1 and the two versions of the BioTac sensors are described in Sec.III-A.2.

1) Testing Platforms: Two hands are used to test our stabilization controllers as we which to show that the proposed approach is platform independent. Both hands are depicted in Fig. 1

The *Allegro Hand*, produced by Wonik Robotics, is a lightweight four fingered hand with four joints per finger, for a total of 16 actuated degrees of freedom. The thumb has an abduction joint, two metacarpal joints (rotation and flexing) and a proximal joint. The remaining fingers do not have metacarpal rotation joints and instead have a distal joint. Each finger is equipped with a BioTac sensor on the fingertip.

The *Wessling Robotic Hand*, produced by Wessling Robotics, is composed of five modular designed robotic fingers. Each finger consists of three actuated degrees of freedom and coupled joint, offering 15 actuated degrees of freedom and 20 joints in total. The hand's fingertips are equipped with BioTac SP sensors.

---

[1]Note that this does not imply independent planning, no independent making and breaking contact but only the stabilization phase.
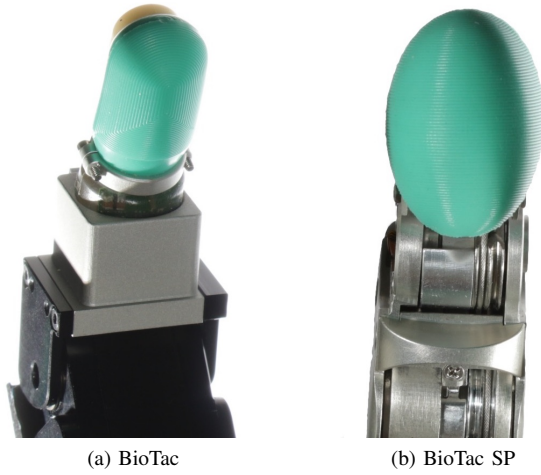
(a) BioTac   (b) BioTac SP

Fig. 3. Syntouch's BioTac and BioTac SP fingertip sensors respectively equipping the Allegro and Wessling hands fingers. These multi-model sensors i. provide information on finger pressure using pressure transducers, ii. temperature and temperature flow via a set of heaters in conjunction with a thermistor, and iii. local deformation through impedance-sensing electrodes distributed on the fingers surface.

2) Tactile Sensors: The raw tactile data is extracted from the fingertip sensors equipped on each hand. The BioTac and BioTac SP, respectively for the Allegro and Wessling Hands, produced by Syntouch, LLC. Both sensors are multi-channel tactile sensors inspired by the human finger and provide multi-modal responses composed of low frequency P_{dc} and high frequency P_{ac} pressure, temperature T_{dc} and temperature flow T_{ac} and local skin deformation E. The sensors can be seen on the respective fingertips in Fig. 3

The BioTac, acquires its pressure signals P_{dc} and P_{ac} using a pressure transducer. Nineteen impedance-sensing electrodes $E \in \mathbb{R}^{19}$ are spread across the finger surface measuring local deformation. Finally, T_{dc} and T_{ac} are measured using a set of heaters coupled with a thermistor. All channels are sampled at a rate of 100 Hz. The P_{ac} is acquired by the sensor at a rate of 2.2 kHz, but is still sampled at 100 Hz, producing batches of 22 values every 10 ms.

In comparison, the BioTac SP differs from the BioTac in two aspects. Firstly, the sensor is designed with an elliptical shape influencing the number of electrodes that can be distributed on the finger surface. The SP's shape allows twenty two electrodes $E_{SP} \in \mathbb{R}^{22}$ to be distributed on it's surface, in contrast to the nineteen in the BioTac. The SP then as a denser sensing of local finger deformation. The other difference is the sampling frequency of the sensor. For the SP, the high frequency pressure, P_{ac}, is sampled at 4.545KHz, the electrodes E_{SP} and the pressure P_{dc} are sampled at 1KHz and the temperatures T_{ac} and T_{dc} are sampled at 200Hz. This allows for a much faster feedback loop on than the one available when use the BioTac.

B. Tactile Training Data

For training the slip classifiers, data has to be collected on both platforms as they are equipped with similar yet different
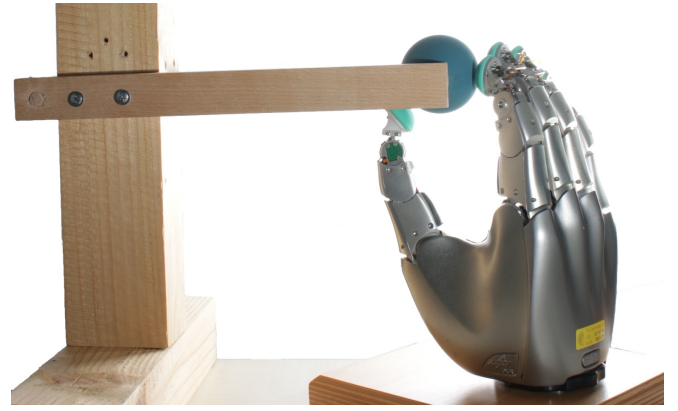

Fig. 4. The data collection setup used for acquiring the necessary training data for the slip classifiers. In this particular example, the Wessling Hand was used. An object is fixed by two wooden plates while in the hand's work space. After being initialized in a predefined position, three fingers are flexed until each becomes in contact with the object, having their pressure immediately adjusted to a target pressure. In a final step, the three finger survey the object's surface with random direction and velocity.

sensors. The data collection procedure is identical for both platforms.

Firstly an object is fixed to a support that holds it on the hand's work space. All fingers are positioned in an initial configuration and are subsequently flexed until each becomes in contact with the object. After the fingers are in contact with the object, the pressure applied by each finger is adjusted by a PID controller until a target pressure is reached on all fingers. Finally, small surveying motions are performed on the object surface by three fingers simultaneously. Acquiring data from three sensors simultaneously reduces the necessary number of training trials. All data from each of the fingers is concatenated into a single data set that will be used to train each of the individual slip predictors. The data collection setup is exemplified in Fig. 4 for the Wessling hand data acquisition.

The described procedure relies on randomly selected velocities in task space for the object surface servoing. Target pressures are selected from three possible values from the set $P* = [20, 50, 80]$, in sensor grounded pressure units. Spanning the data across multiple pressures in conjunction with randomly selected velocity parameters and the distinct contact locations observed across the three fingers, serve for training slip classifiers that are not characterized only by applied pressure, contact location or fingertip velocities. Additionally, all sensor values concerning either pressure or finger deformation are grounded before training, preventing parametric differences in the sensors (for example nominal fluid pressure) from correlating to slip.

Four objects are used in our experiments and are depicted in Fig. 6. Note that from these for objects, only two of them (ball and cardboard box) are used for training the classifiers with the two remaining object only being used for testing.

Three trials are done for each value of $P*$ on the two separate objects for a total of 18 trials. Considering that three fingers are used per trial, the resulting data set correspond
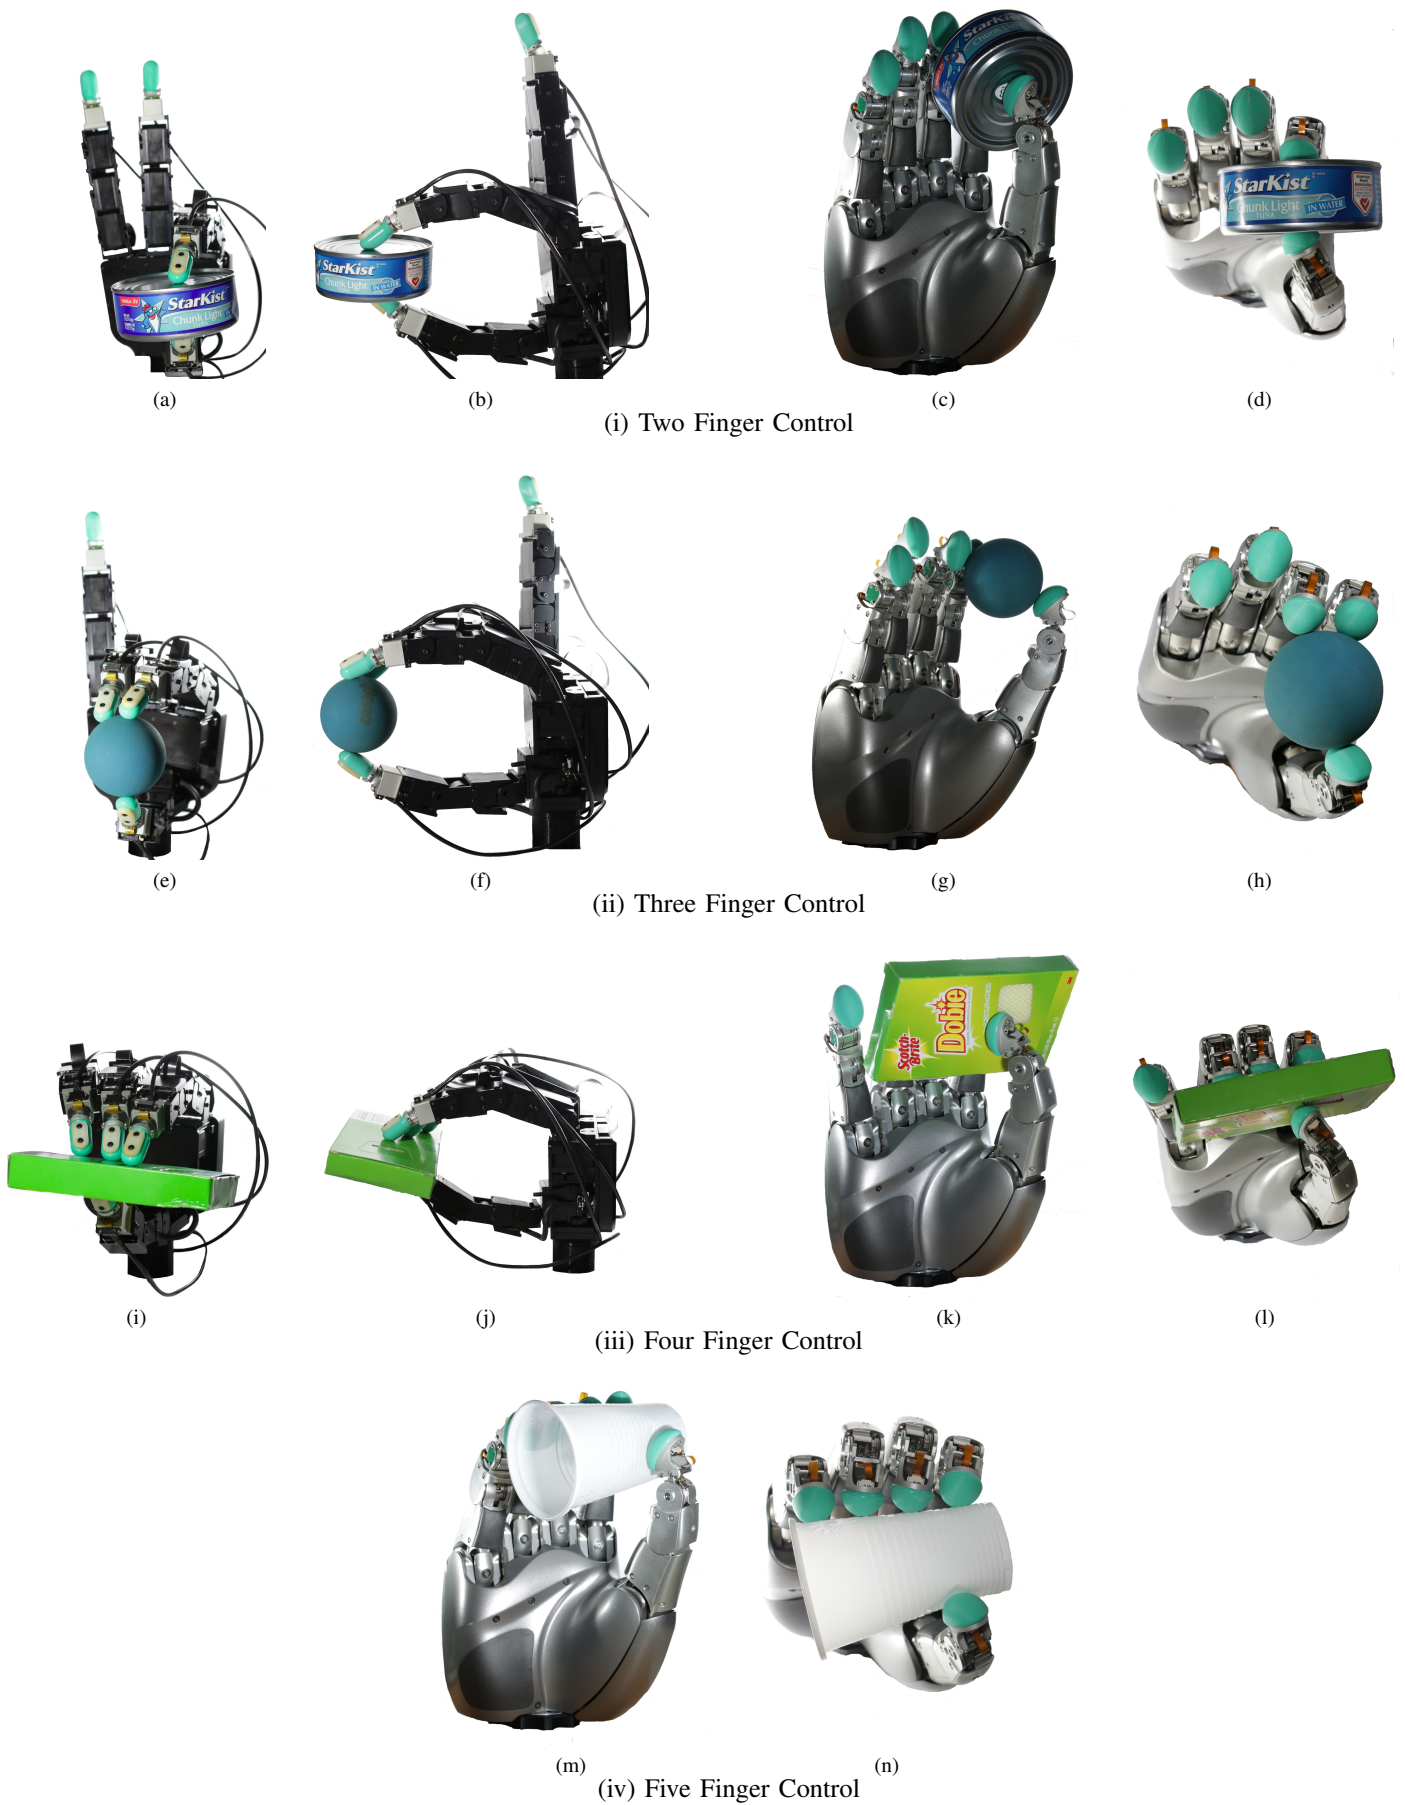
Fig. 5. Stabilization of several objects using different grasp configurations with varying number of fingers. Both hands are controlled using independent controllers for each finger. In (i), a tuna can is stabilized using two fingers of the (a,b) Allegro and (c,d) Wessling Hands. In (ii), a tripod grasp is used to stabilize a ball using the (e,f) Allegro and (g,h) Wessling Hands. In (iii), four fingers are used to stabilize a cardboard box with the (i,j) Allegro and (k,l) Wessling Hands. Finally in (iv) the full five fingers of the Wessling hand are used to stabilize an empty plastic glass (m,n). For all grasps, the stabilization controllers of each finger were only trained with data collected from the ball and the cardboard box.

to 54 single finger trials.

Labeling the data is done automatically by using the finger's end-effector location and the pressure values observed on each finger. Since the object is fixed, if the finger is in contact (observed pressure above a certain threshold T_{Contact}) and the finger is moving then we assume slip is occurring. A minimum deviation $\Delta \mathbf{x}$ in finger position is enforced so that the skin deformation is taken into account.

C. Grip Stabilization Evaluation

In the following section, results all possible stabilization configurations will be shown. The configurations include single finger human-robot stabilization and all possible multi-finger grip (two, three, four and five finger grips). For each configuration we will show the finger pressure behavior of each finger during the stabilization procedure.

External disturbances are applied on the object by a human subject in order to test grip robustness.

1) One Finger Control: The single finger object stabilization work that serves as base for this work was performed on a PA10 robot arm in [14]. Stabilization of an object jointly with a human can be seen in Fig. 2. A extensive study of the slip predictors performance and of the stabilization performance was performed.

2) Two, Three and Four Finger Control: For grip configurations involving two to four fingers, we attempted to stabilzise each of the objects while only having the slip predictors trainned with a subset of the objects. The predictors were trained with two of the objects (the ball and the box), having the remaining objects (the tuna can and the plastic cup) as unknown objects. These unknown objects were chosen based on their distinc properties with respect to the objects in the training set. The tuna can is the heaviest of the objects while the plastic cup is the most deformable.

All objects were successfully stabilized in hand with each of the grip configurations. Some examples of these grasps are shown in Fig. 5.

3) Five Finger Control: For five fingered object stabilization, only the Wessling Hand can be used as the Allegro hand only possesses 4 digits. For this grip configuration we wish to particularly highlight the stabilization of the

Fig. 6.   The four objects used in our experiments. While the ball and the box are used for training and testing, the plastic cup and the tuna can are only used for testing.

plastic cup. The cup is highly deformable and, due to the kinematics of the Wessling hand, there is very few oposition between the thumb, the ring and little fingers, increasing the difficulty of the five fingered grip. This lack of finger oposition, requires the little and ring fingers to apply a mimimum amount of force, otherwise causing the object to shift on the remaining fingers, rendering the grip unstable. Even under these constraints, the controllers are able to successfully stabilize the plastic cup, using the five fingers in the Wessling hand. The stabilization is not only successfull but is also performed without deforming the object's surface. The resulting five fingered grip is depicted in Fig. 5.

IV. CONCLUSION AND DISCUSSION

In this section we will provide an overview of the contributions of this paper, Sec IV-A, describe the current short comings of the approach, Sec IV-B and outlining possible future directions for our work, Sec IV-C.

A. Summary of the Contribution

In this paper we presented a modular control approach capable of stabilizing objects in-hand using any number of fingers. Each finger was controlled independently, having the synchronization between fingers emerge from the tactile feedback of each finger controller. The tactile feedback allows for disturbances caused by poor contact distribution on the fingertip surface or introduced by other fingers action on the object to detected. Each finger will then automatically compensate for perceived changes that might compromise object stability, independently of these changes being cause by external sources, poor finger configurations or other fingers.

Our modular control approach was shown to be generalizable across multiple objects, even objects that are fairly different to the objects in the training set. It is also effective across different platforms, has shown by the experiments performed on the two robot hands that have relevant differences in their kinematic chains.

B. Current Shortcomings

In this work, we wished to assess if our independence assumption was valid and how much it is possible to do with designed controllers enriched with the tactile signals. Using the low dimensional slip signals that where acquired in previous work, allowed for the design of the controller used in this paper. As the full tactile state is much richer than the slip signals, we are possibly discarding relevant information.

Additionally, in this work we assumed that the tested objects were provided to the hand in configurations where the stabilization of said object would be possible, not requiring finger gating/re-positioning.

C. Future Work

Partitioning the hand into a set of independent fingers allows the manipulation problem to be viewed as a distributed problem where each finger solves the problem locally, with simpler models than when considering a complete model

for the full hand. This partitioning allows the models to be learned efficiently using data driven approaches that consider the the full sensor state, as the dimensionality of the problem is distributed across fingers. Our future work will focus on exploring the high dimensionality of the feedback signals and learning stabilization controllers using reinforcement learning approaches in these high dimensional spaces.

As each finger has only the goal of stabilizing the object while remaining in contact with it, manipulating an object would simply require one of the fingers to introduce the desired perturbation to the object, while the remaining fingers keep it stable. Exploring the assumption that small manipulations can be achieved by this kind of master slave paradigms is another interesting future work direction.

Finally, for complex manipulations, we believe that independently controlling the fingers will be necessary but not sufficient to achieve robust performance. Nonetheless, using the independent control as the base level in a hierarchical control framework might enable higher level control policies to perform these manipulations, effectively creating a robust control hierarchy, where the task complexity is distributed across the several levels of the hierarchy. Building such a hierarchy is also viewed as potentially interesting future work.

V. Acknowledgments

References

[1] K. Hertkorn, M. A. Roa, and C. Borst, "Planning in-hand object manipulation with multifingered hands considering task constraints," in *Robotics and Automation (ICRA), 2013 IEEE International Conference on*. IEEE, 2013, pp. 617–624.

[2] H. Van Hoof, T. Hermans, G. Neumann, and J. Peters, "Learning robot in-hand manipulation with tactile features," in *Humanoid Robots (Humanoids), 2015 IEEE-RAS 15th International Conference on*. IEEE, 2015, pp. 121–127.

[3] J. Bohg, A. Morales, T. Asfour, and D. Kragic, "Data-Driven Grasp Synthesis—A Survey," pp. 1–21, 2013. [Online]. Available: http://ieeexplore.ieee.org/xpls/abs{_}all.jsp?arnumber=6672028

[4] D. Prattichizzo, M. Malvezzi, and A. Bicchi, "On motion and force control of grasping hands with postural synergies," 2010.

[5] M. Kazemi, J.-S. Valois, J. A. Bagnell, and N. Pollard, "Robust object grasping using force compliant motion primitives," 2012.

[6] R. S. Johansson and J. R. Flanagan, "Coding and use of tactile signals from the fingertips in object manipulation tasks." *Nature reviews. Neuroscience*, vol. 10, no. 5, pp. 345–59, May 2009.

[7] B. B. Edin, G. Westling, and R. S. Johansson, "Independent control of human finger-tip forces at individual digits during precision lifting." *The Journal of physiology*, vol. 450, p. 547, 1992.

[8] M. Burstedt, B. B. Edin, and R. S. Johansson, "Coordination of fingertip forces during human manipulation can emerge from independent neural networks controlling each engaged digit," *Experimental brain research*, vol. 117, no. 1, pp. 67–79, 1997.

[9] Y. Bekiroglu, J. Laaksonen, J. A. Jørgensen, V. Kyrki, and D. Kragic, "Assessing grasp stability based on learning and haptic data," *IEEE Transactions on Robotics*, vol. 27, no. 3, pp. 616–629, 2011.

[10] H. Dang and P. K. Allen, "Stable grasping under pose uncertainty using tactile feedback," *Autonomous Robots*, pp. 1–22, 2013.

[11] M. Madry, L. Bo, D. Kragic, and D. Fox, "ST-HMP: Unsupervised spatio-temporal feature learning for tactile data," in *IEEE International Conference on Robotics and Automation (ICRA)*, 2014.

[12] M. Li, Y. Bekiroglu, D. Kragic, and A. Billard, "Learning of grasp adaptation through experience and tactile sensing," in *IEEE/RSJ International Conference on Intelligent Robotics and Systems (IROS)*, 2014.

[13] J. M. Romano, K. Hsiao, G. Niemeyer, S. Chitta, and K. Kuchenbecker, "Human-Inspired Robotic Grasp Control With Tactile Sensing," *IEEE Transactions on Robotics*, vol. 27, no. 6, pp. 1067–1079, 2011.

[14] F. Veiga, H. van Hoof, J. Peters, and T. Hermans, "Stabilizing novel objects by learning to predict tactile slip," in *IEEE/RSJ International Conference on Intelligent Robotics and Systems (IROS)*, 2015.

[15] M. Tremblay and M. Cutkosky, "Estimating friction using incipient slip sensing during a manipulation task," in *IEEE International Conference on Robotics and Automation (ICRA)*, 1993, pp. 429–434 vol.1.

[16] S. L. R. Klatzky, R. Bajcsy, "Object exploration in one and two fingered robots," in *IEEE International Conference on Robotics and Automation (ICRA)*, 1987.