Learning the Tactile Signatures of Prototypical Object Parts for Robust Part-based Grasping of Novel Objects

Emil Hyttinen[1] Danica Kragic[1] Renaud Detry[2]

***Abstract*— We present a robotic agent that learns to derive object grasp stability from touch. The main contribution of our work is the use of a characterization of the shape of the part of the object that is enclosed by the gripper to condition the tactile-based stability model. As a result, the agent is able to express that a specific tactile signature may for instance indicate stability when grasping a cylinder, while cuing instability when grasping a box. We proceed by (1) discretizing the space of graspable object parts into a small set of prototypical shapes, via a data-driven clustering process, and (2) learning a touch-based stability classifier for each prototype. Classification is conducted through kernel logistic regression, applied to a low-dimensional approximation of the tactile data read from the robot's hand. We present an experiment that demonstrates the applicability of the method, yielding a success rate of 89%. Our experiment also shows that the distribution of tactile data differs substantially between grasps collected with different prototypes, supporting the use of shape cues in touch-based stability estimators.**

I. INTRODUCTION

It is well-accepted that touch plays a crucial role in human grasping [1]. Recent progress in artificial touch sensing hardware has allowed the robotics community to endow robots with touch capabilities and to show that tactile sensing plays an equally important role in robot grasping [2], [3], [4], [5].

This paper addresses the problem of grasping novel objects. A novel object is one that the robot sees for the first time, one for which the robot has no 3D shape model, and no previous grasping experience. Tactile sensing is particularly important in this situation: Because the visual data provided by the robot's camera show only one side of the object, the robot must apply at least one finger onto a surface that it cannot see, resulting in a relatively uncertain grasping plan. By confronting the tactile information gathered upon contact to tactile impressions gathered during previous successful and unsuccessful grasp trials, the robot can either strengthen its belief in the success of the grasp, or decide to apply caution and modify its plan.

The problem of assessing grasp stability from touch has been studied by several groups [6], [4]. Bekiroglu et al. [4] modeled success probabilities from touch exclusively, allowing for a model that generalizes to novel objects. In another paper, Bekiroglu et al. [6] showed that taking visual information into account increases prediction performances. However, in that paper [6], the model relies on full 3D object representations, which is not applicable to novel objects.

In this work, our aim is to learn a grasp quality model that exploits both tactile/proprioceptive information and the partial 3D data that the robot sees from where it stands. The most straightforward solution to this problem is to train a classifier on positive and negative grasp examples, with each example parametrized by a concatenation of the tactile data, finger joint angles, and a mix of 3D shape features extracted from the 3D image. This approach will however suffer from a high-dimensional input space, making the job of the classifier difficult. Instead, we build on the fact that many grasp planners that work with novel objects already do some sort of shape classification. Specifically, planners are often composed of a dictionary of prototypical shapes. To plan a grasp, the robot fits the prototypical shapes to a 3D image of the scene, and it executes the grasp that corresponds to the prototype that best fits the 3D image. A dictionary of prototypical shapes forms a discretized representation of the space of 3D shapes that fit well within the robot's hand. Therefore, the label of the best-fitting prototype can be used as a highly-compressed characterization of the shape of the object near the grasping point.

This paper evaluates a tactile-aware part-based grasp planner. The grasp planner is composed of n prototypical object parts. The planner encodes information on how to position, orient and preshape the gripper relative to each of the parts in order to produce a grasping configuration. The planner also encodes, for each part, a model of the tactile feedback and finger joint angles that are expected upon contact with an object. The learning of prototypical parts and how to grasp them is discussed in previous work [7]. In this paper, we focus on learning the tactile/joint models, and evaluating the model's ability to characterize the stability of grasps onto novel objects. We define a grasp as stable if the object remains rigidly attached to the hand when the hand lifts up the object, and when the operator applies to the object a force equivalent to the one he would apply to lift it up. The tactile/joint models encode the probability of success of a grasp given the tactile/joint data issued by the sensors upon contact – when the hand is fully closed. The ℓ^{th} model is constructed by executing a number of grasps on different objects using the ℓ^{th} prototype and training a kernel logistic regression classifier on the resulting tactile/joint data and success/failure labels. We present an extensive experiment

[1]Emil Hyttinen and Danica Kragic are with the Centre for Autonomous Systems and the Computer Vision and Active Perception Lab, CSC, KTH Royal Institute of Technology, Stockholm, Sweden {emilhy,dani}@kth.se

[2]Renaud Detry is with the University of Liège, Belgium renaud.detry@ulg.ac.be

This work was supported by the Swedish Foundation for Strategic Research, the Swedish Research Council, and the Belgian National Fund for Scientific Research (FNRS).

that evaluates our model on a set of 192 grasps on 32 objects. We compare our shape-aware tactile model to a classifier that considers tactile and joint data only. While the results of both approaches are similar, we present a careful analysis of the data, that indicates that shape information is relevant to grasp stability estimation. Those results shed new light on grasp stability estimation and will help in further developing shape- and tactile-aware predictors. In summary, the main contributions of this work are (1) a model of grasp stability that uses touch, joint data, and local object shape, (2) results on a large data set of 32 objects, and (3) an analysis the data collected during the experiment supporting the usefulness of shape information in grasp quality prediction.

II. Related work

Before robot tactile sensing became viable and affordable, robot grasping was performed open-loop [8], [9]: A grasp pose was planned from vision, then executed by the robot. Grasp planning is however a difficult task, and the success of a grasping plan is difficult to guarantee. Authors have shown that grasps can be simulated prior to execution [10] in a virtual environment where so-called grasp quality measures [11] help evaluate stability. Unfortunately, this process relies on the same noisy perception as grasp planning, and it has been shown that the correlation between the predicted quality and real-world success is limited [12], [13].

To ensure stable grasps, sensory feedback collected after the application of a grasp can be used. In this case visual feedback alone does generally not perform well since it is bad at capturing the local properties of the contact areas, and because of occlusion caused by the hand or the object itself. Tactile sensors capture useful information about the interactions between the hand and the object and are suitable for feedback when the hand and the object is in contact.

Tactile sensing can be used to estimate different object parameters useful to grasping, such as object position [14], shape [15], [16], [17], [18] and object class [19], [20], [18].

Naturally, tactile can also be used directly to estimate stability. Multiple authors have shown that it is possible to model the shape of hand-object contacts that lead to successful or unsuccessful grasps [6], [21], [22], [23], [24], [25], thereby allowing the robot to assess stability before attempting to lift an object.

Dang et al. [21] have studied means of estimating grasp stability using the 3D locations of object-hand contacts. In later work, Dang et al. [22], [23] devised means of correcting grasps that are considered unstable. Miao et al. [24] used a similar approach as Dang et al. [23], but based their stability estimation on other tactile features. By contrast to our work, neither Dang et al. nor Miao et al. take object shape information into account when assessing stability.

Bekiroglu et al. [6] trained a robot to classify the stability of a grasp from joint angles, tactile feedback and visual feedback. This work was however limited to known object, as the pose and identity of the object was used as in put to the classifier. The aim of this paper is to devise a method that continues using both tactile and visual information while being applicable to novel objects.

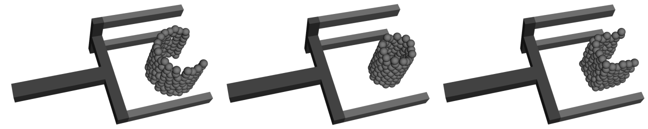

Fig. 1. The three prototypes used in our experiments. The left image shows the first prototype, the middle image shows the second prototype, the right image shows the third. The figure also shows the gripper poses associated with the prototypes. For clarity, the gripper shown here has a simplified shape, it is not representative of the gripper used in our experiments.

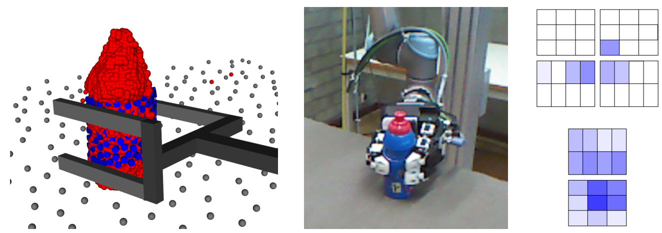

Fig. 2. Planning a grasp onto a new object: the gripper is brought to the pose given by the best-fitting prototype. The rightmost images show the data produced by the tactile sensors, where the darkness of each cell is proportional to the value computed by that cell. The figure shows the distal and intermediate pads only. To understand their arrangement, one can imaging looking at the fully open hand along an optical axis that is normal to the palm. In other words, the top and bottom arrays show the distal pads, and the middle arrays show the intermediate pads.

III. Method

The aim of this paper is to predict the probability of success of a grasp prior to lifting up the object. The information that is exploited to make this prediction consists of (1) the tactile data collected upon contact with the object by the sensors placed on the inner sides of the fingers, (2) the hand's configuration (joint angles) once fully closed around the object, and (3) a characterization of the 3D shape of the object in the vicinity of the grasping point. The characterization of the object's shape emerges from the algorithm that plans where to grasp the object. This algorithm is described in Section III-A. Section III-B describes how grasp success is predicted from hand readings. Section III-C details how the raw data issued by the hand is processed before it is submitted to the model of Section III-B.

A. Part-based Grasp Planning

Planning a grasp on a novel object seen from a single viewpoint is a difficult task, as at least one of the robot's fingers has to be applied to a side of the object that is not visible. This problem can be addressed by constructing a dictionary of grasping prototypes (also referred to as *templates* in the literature) that encode the shape of an object part and how to grasp it. By fitting the prototypes to a partial 3D view of an object and selecting the best fitting prototype, the robot implicitly postulates the shape of some of the object's faces that are not directly observable, which allows it to devise a workable multi-finger grasp. The dictionary of prototypes can be constructed by hand [26] or learned from experience [7]. This paper relies on the planner of Detry et al. [7], which learns prototypes from grasps demonstrated to the robot via teleoperation. Specifically, the planner is composed of three grasp prototypes illustrated in Figure 1. Figure 2 illustrates the planner in action: the best-fitting prototype provides the robot with a gripper pose. For further information on the planner and how it is learned, we refer the reader to the work of Detry et al. [7].

B. Learning Prototypical Tactile Signatures

The aim of this paper is to allow the robot to estimate the stability of a grasp planned with the algorithm discussed above, using both the shape of the object and the tactile imprints/hand configuration available upon contact with the object. Specifically, we define a model of the probability of success of a grasp given the shape of the tactile imprints and the hand's joint angles; the shape of the object around the grasping point is taken into account by learning a separate model for each prototype. In this way, we model the shape of object-gripper contacts that lead to a stable or unstable grasp *given* that we are expecting to be grasping, e.g., a box-like object, or a cylinder-like object. By proceeding this way, the robot should ultimately be able to identify stability cues within certain tactile or finger data. For instance, a fully closed finger, or a finger that reports no contact at all, is likely to indicate that the grasp will not work as intended. Conversely, regular contacts on all finger phalanges are a good indicator of success.

Defining success probabilities via hand-written rules such as the two examples given above would be a tedious task, that would need to be repeated each time a new prototype is created. Moreover, given the high dimensionality of the data that can potentially be produced by touch sensors, such rules would quickly become overly complex. Instead, we let the robot learn success probabilities from experience. We collect a dataset by executing several grasps with each prototype, onto objects of different sizes and shapes. We then train one model for each prototype ℓ, using the data generated with prototype ℓ, i.e., grasps planned onto novel objects via ℓ. When grasping a novel object, the probability of success of the grasp is computed by confronting the tactile and joint data read from the hand to the stability model of the prototype that was used for planning the grasp.

In machine learning, the probability of a binary variable taking one outcome or the other is often modeled via logistic regression. Logistic regression is a linear model that essentially plugs a linear regression model into the logistic function, to produce a model whose values are bounded to $[0, 1]$. Given the nature of tactile and proprioceptive information, a linear mapping from touch and joint angles to success probability is unlikely to emerge. We instead use kernel logistic regression (KLR), which trades the linear regression model of logistic regression for a nonlinear kernel regression model.

Let us denote the dataset collected with one of the grasp prototypes by

$$Z = \{(x_i, y_i)\}_{i=1,\dots,n} \tag{1}$$

where each pair (x_i, y_i) is composed of tactile readings and joint angles $x_i \in \mathbb{R}^d$ (defined in Section III-C below), and a binary stability label $y_i \in \{\text{stable, unstable}\}$. KLR models the stability probability of a grasp characterized by a perceptual vector x with the help of a weighted sum of the similarities between x and each vector in the training dataset Z. The weights associated to stable grasps will generally be positive, while those associated to unstable grasps will be negative. If x resembles percepts of Z that lead to stable grasps, its probability of stability will thus be high. In order to restrict values to the $[0, 1]$ interval, KLR models probabilities by plugging the weighted sum described above into the logistic function $f(z) = \frac{1}{1+e^{-z}}$, which smoothly grows from 0 to 1 as its argument varies from minus infinity to infinity. Weights are usually chosen to maximize the probability of the training set.

Formally, we model the probability of pose- and touch-conditional grasp stability as

$$p(y = \text{stable}|x; v) = \frac{1}{1 + \exp\{-\sum_{i=1}^{n} v_i \mathcal{K}(x, x_i)\}} \tag{2}$$

where $p(y = \text{stable}|x)$ is the probability of success of a grasp characterized by the tactile and pose vector x (see Section III-C), $\mathcal{K}$ is a kernel function that models the similarity between two perceptual readings (defined in Section III-C below) and v is a weight vector chosen to maximize the regularized stability probability of the data

$$-\sum_{i=1}^{n} \log p(y_i|x_i; v) + c\,\text{trace}(vKv^T) \tag{3}$$

where K is the kernel Gram matrix, with $K_{ij} = \mathcal{K}(x_i, x_j)$, and c is a constant. This problem can be solved, e.g., with Newton's method. For more details, we refer the reader to the work of Yamada et al. [27].

We train one classifier per prototype. The probability of success of a grasp x planned with prototype ℓ is modeled with

$$p(y = \text{stable}|x, \ell) = p_\ell(y = \text{stable}|x), \tag{4}$$

where p_ℓ corresponds to the probability given by the classifier learned for prototype ℓ.

C. Tactile and Proprioceptive Data

Our experimental platform is composed of a Robotiq three-finger gripper equipped with tactile sensors. Six sensors are fixed on the inside of distal and intermediate finger links, the seventh is attached to the palm (see Figures 2 and 3).

The data we use in our model is acquired from the tactile sensors and finger encoders at the time a grasp is executed. Each tactile sensor pad gives 8 to 12 spatially distributed pressure measurements (see Figure 2). Classifying grasps based on raw tactile data would be a difficult task given their high dimensionality. Instead, we consider our tactile pads to be two-dimensional images and we apply ideas from image processing. We represent each pad using image moments,

Fig. 3. Robotiq three-finger gripper equipped with tactile sensors. The proximal sensors are currently not functional.

as suggested in [28], [29]. The general parametrization of image moments for one tactile pad f is given by

$$m_{p,q} = \sum_a \sum_b a^p b^q f(a,b) \tag{5}$$

where p and q represent the order of the moment, and a and b are the horizontal and vertical indices of tactile cells. We compute moments up to order one, $(p+q) \in \{0,1\}$ for the fingertips and only the basic moment given by $p,q=0$ for the other pads. The basic moment $p,q=0$ corresponds to the total pressure measured by a pad, and moments of order 1 correspond to the coordinate of the center of the contact between a pad and the object. The image moments of the seven tactile pads yield a total of 13 values characterizing the contacts between the hand and the object.

The Robotiq gripper has ten degrees of freedom. However, we can only control (and measure) one degree of freedom for each finger, plus the lateral tilt of the two side-by-side fingers. The fingers are designed to mechanically comply to the shape of the enclosed object. As a result, we cannot directly compute the configuration of the hand. Nevertheless, the information provided by the finger encoders is useful and can help disambiguating grasp stability.

In summary, the hand provides us with four joint values and thirteen values extracted from the tactile data. A grasp x_i in Eq. 1 is thus parametrized by

$$x_i = (J_i, M_i), \tag{6}$$

where we denote by J_i the vector of four joint values of grasp i, and by M_i the thirteen tactile values. In the text below, J_i is referred to as a *joint feature*, and M_i is referred to as a *tactile feature*. The vector x_i is referred to as *grasp feature*.

Grasp features are normalized to zero mean and unit variance prior to learning. The kernel $\mathcal{K}$ modeling the similarity between two grasp features is implemented with an isotropic Gaussian of variance σ, optimized via cross-validation.

IV. Experiments

In this section, we evaluate the applicability of our grasp stability predictor. We first compute the success rate of the prototype-aware tactile model discussed above. We then plot and discuss the distribution of successful and unsuccessful grasps in grasp feature space, and compare our classifier to a simpler classifier that does not take prototype data into account.

Fig. 4. Set of objects used for evaluation.

We have collected a total of 192 grasps, with 32 stable grasps and 32 unstable grasps for each prototype. Grasps were collected by repeatedly placing one of the objects of Figure 4 in front of the robot, grasping the object via the part-based planner discussed above, and assessing whether the grasp is stable. Stability was assessed by lifting the object with the robot, and slightly pushing the object by hand to verify that it is firmly bound to the gripper. Before lifting up the object, joint and tactile information were recovered and stored along with the ID of the prototype used to plan the grasp. We limited our experiments to pinch grasps because the Robotiq hand's ability to apply power grasps was beyond our expectations and it proved difficult to gather a significant number of failed power grasps. Figure 5 shows examples of successful and unsuccessful grasps along with tactile readings.

Let us denote by s_{all} the set of 192 grasps discussed above, and by s_1, s_2, and s_3 the three subsets of s_{all} that contain grasps collected with prototype 1, 2, and 3, respectively. A KLR classifier was trained on each dataset, yielding classifiers C_1 (from s_1), C_2 (from s_2), C_3 (from s_3), and $C_{\text{no shape}}$ (from s_{all}). As explained above, C_1, C_2 and C_3 form our shape-aware model C (see Eq. 4). $C_{\text{no shape}}$ is a classifier that differs from C by ignoring shape information – it was trained on all grasps, using tactile and joint data, but not the prototype IDs. Classifiers were trained using cross validation, optimizing the parameter c of Eq. 3, and the variance σ of the grasp kernel. The cross validation used all samples belonging to one object as test set and all other samples as training set.

In Figure 6 we illustrate the success rates of the different classifiers. The green dashed and red solid lines plot the success rates of C and $C_{\text{no shape}}$ as functions of the number of dimensions of the input grasp features made available to the classifiers. As explained above, a grasp feature contains 17 elements, four from joint angles and 13 computed from

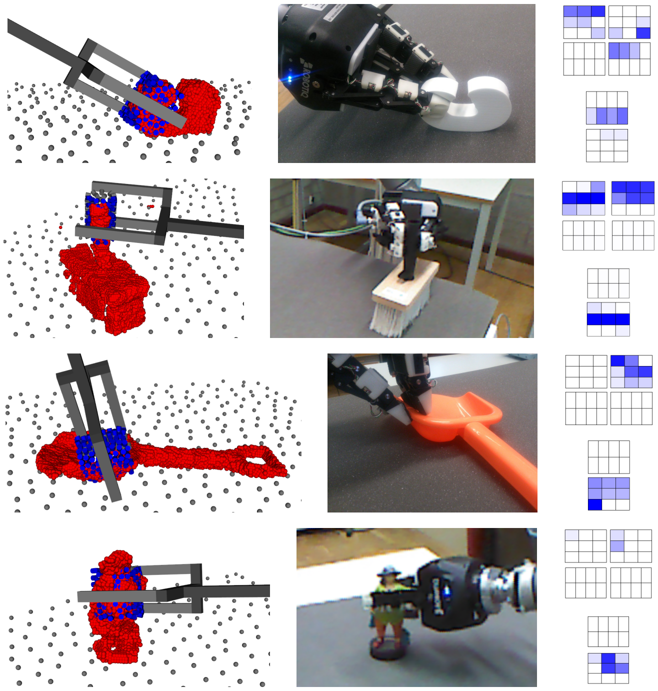

Fig. 5. Examples of successful and unsuccessful grasps, and tactile readings. The first two grasps are successful, while the last two fail when the robot attempts to lift the objects up. See Figure 2 for more details on the arrangement of tactile pads.

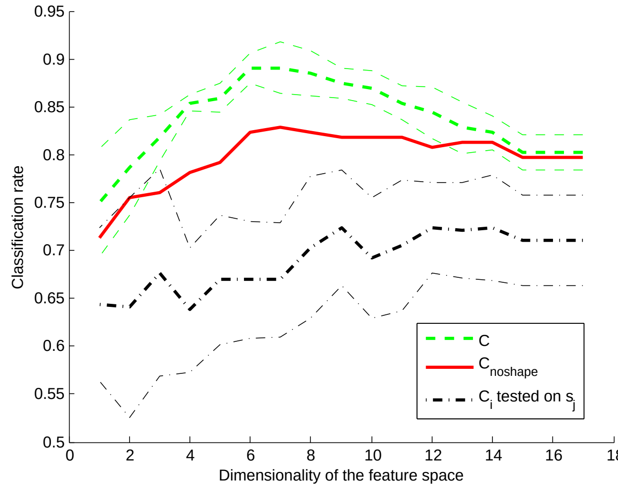

Fig. 6. Classification rates as a function of the number of dimensions of features made available to the classifier. The thick lines show the average success rates for a classifier or group of classifiers. The thin lines show one standard deviation when computing this statistic is possible. The green dashed lines show the shape-aware classifiers discussed in this paper. The red solid line shows the classifier that does not take shape information into account. The black dash-dot lines show the results of classifiers that are trained on grasp examples collected with a different prototype than the one the classifiers are tested with.

the tactile pads. Figure 6 shows the success rate of classifiers trained using subsets of the 17 dimensions. We started with all the 17 dimensions of the features and then we removed the dimensions one by one in order of importance, with the dimension least important for classification being removed first. The overall shapes of the curves shown in the figure indicate that successful/unsuccessful grasps can be discriminated by using only 6 to 8 of the dimensions of the grasp features. The last two dimensions are irrelevant – they correspond to the palm sensor and the finger tilt joint. The finger tilt is left unchanged in this work, and no palm contacts have been observed.

Figure 6 shows that, using an optimal number of features, our method manages to correctly classify 89% of the grasps in our database. Figure 6 also shows that the difference in success rate between C and $C_{\text{no shape}}$ is of limited statistical significance, and, at first glance, one may conclude from this result that using shape information does not help in classifying stability. However, further analysis of our data and the classifiers tends to show otherwise.

Figure 7 shows the 192 grasp features projected to a 2D space via multidimensional scaling. In effect, the 2D Euclidean distance between two points in the plots of Figure 7 is an approximation of the actual Euclidean distance between the two corresponding grasp features in 17-dimensional space. For clarity, features are shown on three separate plots that correspond to the three prototypes. Despite the low dimensionality it represents, Figure 7 can bring insight on several aspects of the data. It appears rather clear that grasp features emerging from different prototypes follow overlapping but distinct distributions. The difference in distribution of successful and unsuccessful grasps is even clearer. These observations tend to indicate that while KLR is powerful enough to discriminate successful/unsuccessful grasps when all three prototypes are mixed together, conditioning on shape has the potential of easing the burden of the classifier.

The black dash-dot line of Figure 6 illustrates the same observation quantitatively. This line has been computed by averaging the success rates of classifiers that were trained on grasp examples collected with a different prototype than the one the classifier was tested with. In other words, the rightmost point of the black dash-dot line corresponds to the average of the success rate of C_1 tested on s_2 and s_3, C_2 tested on s_1 and s_2, and C_3 tested on s_1 and s_2, using all 17 dimensions of grasp features. The rates obtained in this way are substantially lower than those provided by C or $C_{\text{no shape}}$. This means that a classifier learned, for instance, from s_1, will miss-classify some grasps of s_2 that would be correctly classified by C_2. We believe that this result is a strong indication that shape information is relevant to stability estimation, and that this problem should be studied further.

V. Conclusion

We presented a method for estimating the stability of grasps planned onto novel objects. Grasp planning is achieved via a part-based planner, composed of grasp prototypes learned from experience. Grasp stability is estimated via kernel logistic regression, the stability model is trained on hand joint angles, a low-dimensional approximation of tactile readings, and the ID of the prototype a grasp is planned

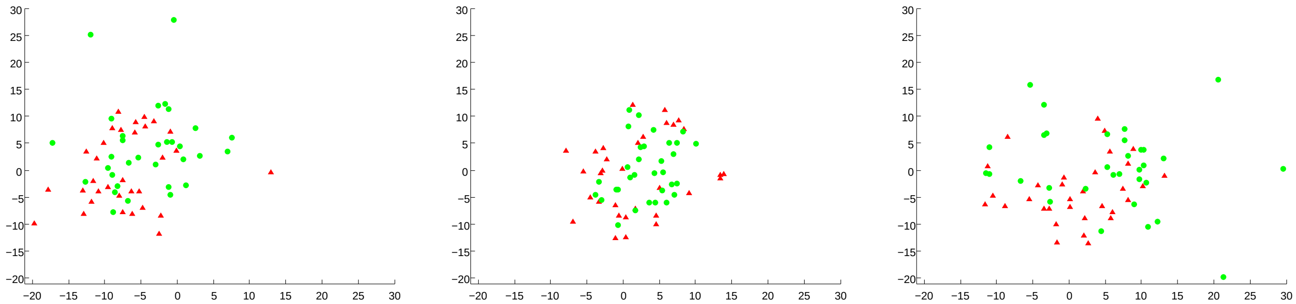

Fig. 7. Plots of all grasp features projected to a 2D plane via MDS. The three plots respectively correspond to grasps planned with prototype 1, 2 and 3. Green circles correspond to successful grasps and red triangles to failed grasps.

with. This way, the stability estimator only uses information about the approximate shape of the object near the grasping point, which makes it transferable to novel objects that partly resemble one of the training objects.

Our experimental results support the applicability of our method, showing that grasp outcomes are correctly predicted in 89% of the test cases. We have compared our model to a model that only relies on tactile and joint data. Our model slightly outperforms the latter. We have additionally shown that the distribution of tactile and joint data is dependent on the approximate shape of the object near the grasping point, and that predicting the outcome of a grasp planned with one prototype is difficult to do based on the information gathered with a different prototype. We believe that those results show evidence of the relevance of shape information in stability estimation, and indicate that shape-aware stability is a problem that deserves further attention.

In future work, we plan to conduct experiments on a larger scale, to provide further insight on the relevance of different types of tactile and shape features on grasp stability. We also plan to investigate different ways of adapting grasps based on tactile sensing.

References

[1] R. Johansson, "Sensory input and control of grip," in *Novartis Foundation Symposium*, pp. 45–59, 1998.

[2] R. Howe, N. Popp, P. Akella, I. Kao, and M. Cutkosky, "Grasping, manipulation, and control with tactile sensing," in *IEEE International Conference on Robotics and Automation*, 1990.

[3] A. Bicchi, J. Salisbury, and P. Dario, "Augmentation of grasp robustness using intrinsic tactile sensing," in *IEEE International Conference on Robotics and Automation*, 1989.

[4] 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*, submitted.

[5] L. Jentoft, Q. Wan, and R. Howe, "Limits to compliance and the role of tactile sensing in grasping," in *Proceedings of International Conference on Robotics and Automation (ICRA), 2014.*

[6] Y. Bekiroglu, R. Detry, and D. Kragic, "Learning tactile characterizations of object- and pose-specific grasps," in *IROS*, pp. 1554–1560, 2011.

[7] R. Detry, C. H. Ek, M. Madry, and D. Kragic, "Learning a dictionary of prototypical grasp-predicting parts from grasping experience," in *IEEE International Conference on Robotics and Automation*, 2013.

[8] J. Coelho, J. Piater, and R. Grupen, "Developing haptic and visual perceptual categories for reaching and grasping with a humanoid robot," in *Robotics and Autonomous Systems*, vol. 37, pp. 7–8, 2000.

[9] I. Kamon, T. Flash, and S. Edelman, "Learning to grasp using visual information," in *IEEE International Conference on Robotics and Automation*, vol. 3, pp. 2470–2476, 1996.

[10] C. Goldfeder and P. K. Allen, "Data-driven grasping," *Auton. Robots*, vol. 31, no. 1, pp. 1–20, 2011.

[11] C. Ferrari and J. Canny, "Planning optimal grasps," in *IEEE International Conference on Robotics and Automation*, pp. 2290–2295, 1992.

[12] J. Kim, K. Iwamoto, J. J. Kuffner, Y. Ota, and N. S. Pollard, "Physically based grasp quality evaluation under pose uncertainty.," *IEEE Transactions on Robotics*, vol. 29, no. 6, pp. 1424–1439, 2013.

[13] J. Weisz and P. K. Allen, "Pose error robust grasping from contact wrench space metrics," in *ICRA*, pp. 557–562, 2012.

[14] A. Petrovskaya and O. Khatib, "Global localization of objects via touch," *Trans. Rob.*, vol. 27, pp. 569–585, June 2011.

[15] J. B. Jarmo Ilonen and V. Kyrki, "Fusing visual and tactile sensing for 3-d object reconstruction," in *IEEE International Conference on Robotics and Automation*, 2013.

[16] A. Balestrino and A. Bicchi, "Adaptive surface following and reconstruction using intrinsic tactile sensing," in *Proc. IEEE Int. Works. on Sensorial Integration for Industrial Robots*, (Zaragoza, Spain), 1989.

[17] M. Moll and M. A. Erdmann, "Dynamic shape reconstruction using tactile sensors," in *In Proceedings of the 2002 IEEE International Conference on Robotics and Automation*, pp. 1636–1641, 2002.

[18] N. Sommer, L. Miao, and A. Billard, "Bimanual compliant tactile exploration for grasping unknown objects," in *Proceedings of International Conference on Robotics and Automation (ICRA), 2014.*

[19] S. Caselli, C. Magnanini, and F. Zanichelli, "Haptic object recognition with a dextrous hand based on volumetric shape representations," in *IEEE Int. Conf. on Multisensor Fusion and Integration, Las Vegas, NV*, pp. 2–5, 1994.

[20] P. Allen and R. Bajcsy, "Robotic object recognition using vision and touch," tech. rep., Proceedings of the 9th International Joint Conference on Artificial Intelligence, 1987.

[21] H. Dang and P. K. Allen, "Learning grasp stability," in *ICRA*, pp. 2392–2397, 2012.

[22] H. Dang and P. K. Allen, "Grasp adjustment on novel objects using tactile experience from similar local geometry," in *IROS*, pp. 4007–4012, 2013.

[23] H. Dang and P. K. Allen, "Stable grasping under pose uncertainty using tactile feedback," *Auton. Robots*, vol. 36, no. 4, pp. 309–330, 2014.

[24] L. Miao, Y. Bekiroglu, D. Kragic, and A. Billard, "Learning of grasp adaptation through experience and tactile sensing," in *Proceedings of International Conference on Intelligent Robots and Systems (IROS), 2014.*

[25] M. Madry, L. Bo, D. Kragic, and D. Fox, "St-hmp: Unsupervised spatio-temporal feature learning for tactile data," in *Proceedings of International Conference on Robotics and Automation (ICRA), 2014.*

[26] A. T. Miller, S. Knoop, H. I. Christensen, and P. K. Allen, "Automatic grasp planning using shape primitives," in *ICRA*, pp. 1824–1829, 2003.

[27] M. Yamada, M. Sugiyama, and T. Matsui, "Semi-supervised speaker identification under covariate shift," *Signal Processing*, pp. 2353–2361, 2010.

[28] J. Laaksonen, V. Kyrki, and D. Kragic, "Evaluation of feature representation and machine learning methods in grasp stability learning," in *Humanoid Robots (Humanoids), 2010 10th IEEE-RAS International Conference on*, pp. 112–117, Dec 2010.

[29] H. Zhang and N. Chen, "Control of contact via tactile sensing," *Robotics and Automation, IEEE Transactions on*, vol. 16, pp. 482–495, Oct 2000.