Citation for published version:

Corradi, T, Hall, P & Iravani, P 2015, Bayesian tactile object recognition: learning and recognising objects using a new inexpensive tactile sensor. in *2015 IEEE International Conference on Robotics and Automation (ICRA)*. Proceedings / IEEE International Conference on Robotics and Automation., IEEE, pp. 3909-3914, IEEE International Conference on Robotics and Automation (ICRA) 2015, Seattle, USA United States, 26/05/15. <https://doi.org/10.1109/ICRA.2015.7139744>

DOI:

[10.1109/ICRA.2015.7139744](https://doi.org/10.1109/ICRA.2015.7139744)

Publication date:

2015

Document Version

Early version, also known as pre-print

[Link to publication](#)

University of Bath

Alternative formats

If you require this document in an alternative format, please contact:
openaccess@bath.ac.uk

General rights

Copyright and moral rights for the publications made accessible in the public portal are retained by the authors and/or other copyright owners and it is a condition of accessing publications that users recognise and abide by the legal requirements associated with these rights.

Take down policy

If you believe that this document breaches copyright please contact us providing details, and we will remove access to the work immediately and investigate your claim.

Bayesian Tactile Object Recognition: learning and recognising objects using a new inexpensive tactile sensor*

Tadeo Corradi¹, Peter Hall² and Pejman Iravani¹

Abstract—We present a Bayesian approach to tactile object recognition that improves on state-of-the-art in using single-touch events in two ways. First by improving recognition accuracy from about 90% to about 95%, using about half the number of touches. Second by reducing the number of touches needed for training from about 200 to about 60. In addition, we use a new tactile sensor that is less than one tenth of the cost of widely available sensors. The paper describes the sensor, the likelihood function used with the Naive Bayes classifier, and experiments on a set of ten real objects. We also provide preliminary results to test our approach for its ability to generalise to previously unencountered objects.

I. INTRODUCTION

A Bayesian system for object learning and recognition using purely tactile, orientation independent information is presented. A novel, inexpensive sensor is used, mounted on a robotic arm which learns in an automatic manner to recognise objects outperforming state-of-the-art. We also provide some evidence that the system can recognise previously unseen objects.

The system learns and recognises objects from single-touch events using a newly developed sensor [1]. Tactile sensations are encoded using Zernike Moments and objects are modeled by a sum of Gaussian distributions. The approach presented does not use the orientation information of the objects and requires only a very limited number of training samples, making a substantial improvement over previous work. A fully automated robot system (depicted in Fig. 1) was constructed to learn the tactile appearance of 10 household objects and to recognise these with an accuracy of 87% after 15 touches and 95% after 30 touches.

II. RELATED WORK

A. Tactile sensors

Tactile sensors can be designed using a variety of techniques, the most common being piezo-resistive sensors, conductive polymers, or capacitive sensors [2]. The most widely used in robotics include the impedance based BioTac [3], the Weiss tactile array [4], and the capacitive array based DigiTact [5], all of which have a price tag exceeding USD 700. Recently, efforts have been made at creating cheaper and more accessible sensors. The TakkTile TakkArray [6] is an open source and open hardware sensor based on an array of

Fig. 1. The new tactile sensor mounted on a KUKA KR5-siXX-R650 robotic arm, currently exploring the tactile appearance of a mug.

MEMS barometers, it has a retail price of USD 500, and their material cost is approximately USD 200. The TacTip [7] aims to provide higher resolution whilst remaining inexpensive as they can be non-professionally manufactured (material cost is approximately USD 200). It is a biologically inspired tactile sensor based on the deformation of a silicone rubber hemispherical surface and the consequential displacement of a number of internal papillae. A digital camera is used to observe this displacement.

B. Recognition by grasping

Recently, there have been several projects involving recognition by grasping using machine learning techniques. Principal Component Analysis (PCA), Self Organizing Maps and Artificial Neural Networks have been combined to process the output of Weiss tactile sensory arrays attached to a number of robotic end-effectors, to recognize household objects [8]. Novel recursive Gaussian kernels have been designed to encode the various stages of contact during grasping leading to a robust on-line system capable of learning new models and classifying objects in real time [9]. The most accurate system, to the best of our knowledge, is the one developed by [10]. They extends HMP (Hierarchical Matching Pursuit, a multi-layer hierarchical feature learning system) to include temporal information. They test their method on 6 tactile databases and produce an accuracy of between 80% and 100%. Whilst it is evident that combining proprioceptive with tactile information is likely to yield better results than either modality alone [11], [12], using grasp limits the size

*This work was funded by the Engineering and Physical Sciences Council (EPSRC), UK

¹Department of Mechanical Engineering, University of Bath, Claverton Down, Bath, BA27AY, UK t.m.corradi@bath.ac.uk

²Department of Computer Science, University of Bath, Claverton Down, Bath, BA27AY, UK

of the object to be identified, requires a robotic hand, and requires a grasp to be achieved.

C. Single contact tactile recognition

Recognition using a single touch at a time is a possible solution which remains relatively unexplored. As far as we know, the best results so far are achieved by [13], requiring 60 touches to converge to 90% recognition accuracy, using 200 touches for training, over a set of 5 objects.

The most common approaches for single contact tactile object recognition are voxel based or point clouds [14], [15], [16]. Recently, a very efficient and accurate combination of both was developed [17], which is able to model the object shape and the uncertainty about occupied space. They achieve above 80% accuracy in recognition over a set of 45 objects, and from only 10 touches; however, object 3D models are required in advance. Voxel representations and point-clouds provide a natural way of representing tactile information about objects, but they can be cumbersome in terms of computational power for recognition, as they usually comprise a large number of points/voxels whose matching to a database can be complex, and are prone to noise which is difficult to model. Attempts to address these problems include merging points that are close into a probability point modelled by a Kalman filter [18], and clustering to subdivide the point cloud into regions which are then encoded as features [19].

D. Appearance based tactile-only recognition

One of the first attempts at a tactile-only recognition is [20], which uses geometric features such as lines and points and their evolution over time. Their accuracy recognising objects is high (83%), however the number of shapes is only 6 and they are very basic predefined geometric solids (cylinder, cone, etc.). The two notable recent pieces of research which most closely relate to our study are the work of Schneider et al. [21], and the work of Pezzementi et al. [13].

The first [21], involves the repeated application of a two fingered grasps using a gripper equipped with Weiss tactile array sensors. Features are extracted, then a bag-of-features approach is used to recognise household and industrial objects. They use an information theoretic approach for maximum expected information gain to inform grasping position. They obtain an accuracy of 84.6% in recognition, using 830 tactile images for training and 16 to 20 tactile images in the testing set. The object pose is strictly known and fixed (small translation variance is tolerated). It could be argued that this work uses proprioception (they know the height of the gripper) and thus is not purely appearance based.

Pezzementi et al. [13] use simulations to compare various methods of feature extraction, and create clusters of these features to compile feature histograms to be compared for object recognition. Most of their testing is performed in simulation using 3D models of objects. The physical testing was done using DigitTact sensors over a set of 5

objects (the context was recognition of plastic letters) using a predefined exploring routine. They use 200 samples for training and 100 for testing. The accuracy in these physical experiments reaches 90% for one of their feature choices after approximately 60 touches. It would be interesting to see this system tested on a larger set of objects, since its simulated performance is quite good.

III. SENSOR AND TACTILE DATA REPRESENTATION

The new sensor [1] used in this paper is based on the same principle as the TacTip. However, it has neither papillae nor internal gel. Instead it has a plain black smooth opaque silicone rubber hemispherical membrane of radius 40mm and thickness 1mm, mounted at the end of a rigid opaque encasing for the digital camera, 3D printed in ABS¹. The camera has a resolution of 640 by 480 pixels, and incorporates a set of 8 white LEDs. The shading pattern of light is used as input. When the sensor is in contact with an object, the shading pattern on the membrane changes accordingly (see Fig. 2). In recent work, it was shown to recognise seven basic shapes with over 95% accuracy [1].

Due to the circular geometry of the sensor image, a rotationally invariant representation was required. In previous work, a number of encoding methods were compared and it was suggested that Zernike Moments together with PCA achieved the best performance [1]. Zernike Moments have been shown to be useful when scale, rotation and translation invariances are sought [22], and have been successfully used for basic shape recognition [23]. Zernike moments here refers to the absolute value of the inner product of a vectorised image with a vectorised Zernike polynomial, a set of radial complex polynomials defined on the unit disk (see Fig. 3).

Let $m \geq n$ be non-negative integers, and let $0 \leq \phi \leq 2\pi, 0 \leq \rho \leq 1$ define a polar coordinate system. Then the

¹3D model of the tactile sensor encasing, and links to the other components are available at: <https://github.com/Exhor/bathtip>

Fig. 2. The new tactile sensor design (left). The main body is 3D printed in ABS. The tip is a 1mm thick silicone rubber hemisphere. At the base (not visible) there is a USB eSecure web-cam (running at 640 by 480 pixels) with 8 LEDs illuminating the inside of the silicone hemisphere. As the tip makes contact with an object, it deforms resulting in a specific shading pattern (right).

Fig. 3. The first Zernike polynomials evaluated on a unit disk. Here depicted as modulus (red) and phase (blue).

even and odd Zernike polynomials are defined as:

$$Z_n^m(\rho, \varphi) = R_n^m(\rho) \cos(m\varphi)$$

$$Z_n^{-m}(\rho, \varphi) = R_n^m(\rho) \sin(m\varphi),$$

Which can be indexed by:

$$Z_j = Z_{n(j)}^{m(j)}$$

Where $m(j), n(j)$ are Noll's indices of Zernike polynomials [24], and

$$R_n^m(\rho) = \sum_{k=0}^{(n-m)/2} \frac{(-1)^k (n-k)!}{k! \left(\frac{n+m}{2-k}\right)! \left(\frac{n-m}{2-k}\right)!} \rho^{n-2k}$$

Then, the d^{th} Zernike Moment of an image M is given by:

$$Zer_d(M) = \left| \sum_{i,j \in \{i^2+j^2 \leq n^2/2\}} M(i,j) Z_d'(i,j) \right|$$

Where,

$$Z_d'(i,j) := Z_j \left(\frac{\sqrt{i^2+j^2}}{\frac{\sqrt{2}}{2}n}, \arctan \left(\frac{j-n/2}{i-n/2} \right) \right)$$

Once the Zernike moments are obtained from the entire training set, PCA is performed. The Zernike moments of images obtained during validation/testing are multiplied by the PCA dimensionality reduction matrix obtained during training. This process is hereafter referred to as “finding the Zernike-PCA moments”. The number of components to be used is decided by inspecting the eigenvalues and retaining sufficiently many principal components so as to explain 95% of the variance in the training data.

IV. OBJECT LEARNING AND RECOGNITION

The proposed model stores the Zernike-PCA moments of all tactile images and their corresponding object labels given during training. During testing, the Zernike-PCA moments of each new tactile image is compared against those stored values, and the likelihood of the new image, given each learnt object, is computed. This likelihood is defined as the normalized sum of n_C Normal probability density functions, where n_C is the number of training images used for object C . Each one of these is evaluated at the sensed image's Zernike-PCA value, centered at one of the training points,

and with covariance given by the covariance matrix of all training points². The process is depicted in Fig. 4.

Formally, let the training set be $X_C = \{X_{C,i}, i = 1, \dots, n_C\}$, where $X_{C,i}$ is the Zernike-PCA moment vector corresponding to the i^{th} tactile image of object C , which was observed n_C times during training. Let W be the covariance matrix of X_C . Let $Y = \{Y_j, j = 1, \dots, m\}$ be the sequence of Zernike-PCA moments (PCA reduction is performed using the dimensionality reduction matrix obtained from the training data), where Y_j represents the Zernike-PCA moments of the j^{th} tactile image of the object being sensed for recognition. Then the likelihood of Y_j for a given object class C is defined as:

$$P(Y_j|C) = \frac{1}{n_C} \sum_{i=1}^{n_C} \mathcal{N}(Y_i|X_{C,i}, W) \quad (1)$$

Where,

$$\mathcal{N}(Y_i|X_{C,i}, W) = \frac{e^{-\frac{1}{2}(Y_i - X_{C,i})^T W^{-1}(Y_i - X_{C,i})}}{\sqrt{\|W\|(2\pi)^d}}$$

Here, d is the dimensionality of the feature vector. Using this likelihood function a Naive Bayes classifier was implemented. This assumes that observed Zernike-PCA moments are statically independent. Note that PCA projection here helps to mitigate against correlations between features.

$$P(C|Y) = \alpha \prod_{j=1}^m P(Y_j|C)P(C)$$

Where α is just a normalizing constant, and $P(C)$ can be estimated from the number of times each object is observed during training, which in our case forms a uniform prior distribution. Therefore object recognition can be performed using maximum a posteriori:

$$C_{pred} = \operatorname{argmax}_C P(Y_j|C)$$

The computational complexity arises from Equation 1. Assuming there are n observations times during training, the complexity is $O(dn^2)$ during training and $O(d^2n)$ during testing.

V. EXPERIMENTS AND RESULTS

Two experiments were performed to test the accuracy of the object recognition method outlined above: one to recognise objects seen before within a fixed collection, the other to test generalisation to unseen objects. Both experiments were carried out under the same setup.

A. Experimental setup

The system consisted of a 6 degrees of freedom (DOF) KUKA KR5-sixx-R650 robotic arm, a 6 DOF force-torque sensor mounted on its end effector, and the new tactile sensor mounted on the force-torque sensor (see Fig. 1). The force-torque sensor was used to detect touch events and to ensure the safety of the robot-object interaction.

²In practice, this is the diagonal matrix of variances, since X_C is the scores matrix resulting from PCA.

Fig. 4. The recognition process: from touch to object posterior probabilities.

The initial location of the object is assumed to be known, but its orientation is unknown. Limited unintentional pose alteration (less than 5% of object size) does occur during the experiments, as a consequence of contact. The aim is to have the robotic arm move the sensor to various point on the object surface and collect the tactile information autonomously. Each object was manually placed and secured in this location. The robotic arms is programmed to perform the following exploration procedure:

- 1) Define a “safety hemisphere” of radius 30cm about the assumed object centre. The hemisphere occupies the space above the object.
- 2) Generate a set of random points on that hemisphere.
- 3) Take the sensor to the next unvisited position in the list, facing inwards towards the centre point.
- 4) Move the sensor linearly inwards, until a normal force of 75 grams is detected.
- 5) Record the tactile image.
- 6) Retract the sensor linearly away from the object back to the imaginary sphere.
- 7) Back to step 3.

B. Object recognition

The objective of the first experiment was to automatically explore, learn and recognise objects from a set of 10 household objects (see Fig. 5): stapler, toothbrush, porridge pot, mug, shampoo bottle, box, pen, ball, textbook, water bottle (empty).

A total of 120 tactile images were collected for each object. These were split into 60 for training, 30 for validation and 30 for testing. A number of tests were attempted using the validation data set for testing. Initially, a Naive Bayes classifier using clustering was implemented, which resulted in approximately 70% accuracy after 30 touches, using k-means. Alternative clustering methods were tested, but did not improve performance. In particular Gaussian Mixture Models seemed suitable due to the natural representation of the likelihood function for observed data, but the parameter estimation led to an under-determined system for such a small data set. The final choice of inference system is non-parametric, and as such there is no need for a validation data set for parameter estimation. Of the 90 samples (training and

testing) for each object, 100 different partitions (60 training images and 30 testing images) were made, the accuracy reported is the percentage of correct recognitions, averaged of these 100 iterations. Fig. 6 shows the confusion matrix after 5 and 15 touches.

After 15 touches the overall accuracy is 87% ; however, there is still a marked (approx. 19%) confusion between the toothbrush and the pen. These objects are very similar to touch in many of their local patches. This confusion represents 2.7% of the inaccurate predictions after 30 touches. There is high uncertainty about the stapler in the first 5 touches, perhaps reflecting the varied tactile features of its surface.

Fig. 7 shows the average accuracy for all objects, over 100 trials. As a comparison, best previous results (averaged over 7 trials) are shown [13]. The recognition accuracy follows a similar pattern in all methods, however our system gains a clear advantage from the start, and it stabilizes after about 25 touches.

C. Classifying unseen objects

In the second experiment, the potential for classification of previously unseen objects was preliminarily tested. The aim was to discern if the system had potential to classify objects that had not been used in training. Five previously

Fig. 5. The objects to be learnt and recognised.

Fig. 6. Confusion matrices showing recognition errors for the 10 household objects, after 5 and 15 touches.

Fig. 7. Accuracy as a function of the number of touches, averaged over 100 trials. Comparison to previous work by [13], showing their best performing features: “Polar Fourier” (PF) and “Moment Normalized Translation Invariant” (MNTI).

untouched objects were sensed and attempted to be classified using the system outlined above. The objects used were: a plastic card, a different mug, a different pen, a smaller and harder ball, and another textbook (soft-back). This time the full data set for the 10 known objects was used for training, and 120 images of the unseen object were used in testing. Fig. 8 shows the posterior probabilities of each of the known 10 objects, assigned to each of the new objects, against the number of touches.

The plastic card is very different to any known objects and as such causes high confusion initially. The system finally settles for classifying it as a mug or a pot. The new pen is initially very confidently classified as a pen, but after 10 touches there is growing confusion with the pot model. The other three objects are on average “correctly” classified. There is some confusion between the mug and the pot when classifying the new mug, which is understandable due to the similarity between the two known objects. These preliminary results show promise that the system may be generalisable to unseen objects, but are modest in scale and as such not conclusive: further research is required. It seems that objects very similar to the known ones (new book, new ball, new mug, new pen) are classified “correctly” very quickly, and as such the level of uncertainty at the beginning of the exploration could be used to inform a system that predicts new classes.

Fig. 8. Attempting to classify previously untouched objects. Posterior distribution over the known object classes, when testing is performed on five objects not sensed during training. Averaged over 100 trials.

D. Timings

All timings provided are for single-threaded, unoptimized, MATLAB code, running on a Core i7-4700MQ 2.4Ghz with 8Gb DDR3-1600 RAM. Zernike moment calculation took on an average of 3.7×10^{-3} s per tactile image. Feature dimensionality was always 21 or 22. For the first experiment (600 images in training), training took an average of 1.7×10^{-8} s, and testing 8.6×10^{-4} s per tactile image. For the second experiment (1200 images in training), training took an average of 1.7×10^{-8} s, and testing 1.2×10^{-3} s per tactile image. All these timings are substantially lower than the average time it takes the robotic arm to take a reading (approximately 30 seconds).

VI. CONCLUSION

A new inexpensive tactile sensor combined with an automated simple Bayesian object identity inference system were presented. They were shown to achieve accuracy in recognition outperforming state-of-the-art, for single contact, local appearance based tactile object recognition. The sensor

was made open source and can has a total material cost of approximately USD 30, substantially less than any other commercial or open source tactile sensor available, making it widely available to experts and hobbyists. A system was designed to autonomously collect tactile information from a range of household objects, using this new sensor, mounted on a robotic arm and aided by a force-torque sensor. These results are obtained using a very limited number of training, validation and testing images, about a third of previous similar work. In addition, preliminary results show potential for unseen object classification, yet more research is needed. Recognition is performed in real time.

Inference is performed using a Naive Bayes classifier. As such, there is an assumption of independence between observed features. This assumption is potentially limiting and a more sophisticated probabilistic model may be needed as the number of classes grows larger.

At present, exploration takes approximately 30 seconds per reading, 30 minutes to learn an object's representation and 15 minutes to recognise it with 95% confidence. Whilst attempts were made to create a reactive system, robot control is relatively rigid. It would be interesting to explore ways of using machine learning to make the robot control more efficient and self-adapting. Future work will also include sensor fusion, attempting to harness the potential shown here to complement active vision systems.

ACKNOWLEDGMENT

We would like to thank the Bristol Robotics Laboratory for lending us the TacTip sensor.

REFERENCES

- [1] T. Corradi, P. Hall, and P. Iravani, "Tactile Features: Recognising Touch Sensations with a Novel and Inexpensive Tactile Sensor," in *Advances in Autonomous Robotics Systems*, ser. Lecture Notes in Computer Science, M. Mistry, A. Leonardis, M. Witkowski, and C. Melhuish, Eds. Springer International Publishing, Jan. 2014, no. 8717, pp. 163–172.
- [2] R. Dahiya, P. Mitterdorfer, M. Valle, G. Cheng, and V. Lumelsky, "Directions Toward Effective Utilization of Tactile Skin: A Review," *IEEE Sensors Journal*, vol. 13, no. 11, pp. 4121–4138, 2013.
- [3] N. Wettels, V. J. Santos, R. S. Johansson, and G. E. Loeb, "Biomimetic Tactile Sensor Array," *Advanced Robotics*, vol. 22, no. 8, pp. 829–849, 2008.
- [4] K. Weiss and H. Worn, "The working principle of resistive tactile sensor cells," in *Mechatronics and Automation, 2005 IEEE International Conference*, vol. 1, 2005, pp. 471–476 Vol. 1.
- [5] Pressure Profile Systems, "DigiTact II." [Online]. Available: <http://www.pressureprofile.com/digitacts-sensors>
- [6] L. Jentoft, Y. Tenzer, D. Vogt, J. Liu, R. Wood, and R. Howe, "Flexible, stretchable tactile arrays from MEMS barometers," in *2013 16th International Conference on Advanced Robotics (ICAR)*, Nov. 2013, pp. 1–6.
- [7] C. Chorley, C. Melhuish, T. Pipe, and J. Rossiter, "Development of a tactile sensor based on biologically inspired edge encoding," in *International Conference on Advanced Robotics, 2009. ICAR 2009*, 2009, pp. 1–6.
- [8] S. Navarro, N. Gorges, H. Worn, J. Schill, T. Asfour, and R. Dillmann, "Haptic object recognition for multi-fingered robot hands," in *2012 IEEE Haptics Symposium (HAPTICS)*, 2012, pp. 497–502.
- [9] H. Soh, Y. Su, and Y. Demir, "Online spatio-temporal Gaussian process experts with application to tactile classification," in *2012 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, 2012, pp. 4489–4496.
- [10] 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)(to appear)*, 2014.
- [11] J. K. Kim, J. W. Wee, and C. H. Lee, "Sensor fusion system for improving the recognition of 3D object," in *2004 IEEE Conference on Cybernetics and Intelligent Systems*, vol. 2, 2004, pp. 1207–1212.
- [12] N. Gorges, S. Navarro, D. Göger, and H. Worn, "Haptic object recognition using passive joints and haptic key features," in *2010 IEEE International Conference on Robotics and Automation (ICRA)*, 2010, pp. 2349–2355.
- [13] Z. Pezzementi, E. Plaku, C. Reyda, and G. Hager, "Tactile-Object Recognition From Appearance Information," *IEEE Transactions on Robotics*, vol. 27, no. 3, pp. 473–487, 2011.
- [14] A. Bierbaum, K. Welke, D. Burger, T. Asfour, and R. Dillmann, "Haptic exploration for 3D shape reconstruction using five-finger hands," in *2007 7th IEEE-RAS International Conference on Humanoid Robots*, 2007, pp. 616–621.
- [15] N. Gorges, P. Fritz, and H. Wörn, "Haptic Object Exploration Using Attention Cubes," in *KI 2010: Advances in Artificial Intelligence*, ser. Lecture Notes in Computer Science, R. Dillmann, J. Beyerer, U. D. Hanebeck, and T. Schultz, Eds. Springer Berlin Heidelberg, Jan. 2010, no. 6359, pp. 349–357.
- [16] N. Gorges, S. Navarro, and H. Worn, "Haptic object recognition using statistical point cloud features," in *2011 15th International Conference on Advanced Robotics (ICAR)*, 2011, pp. 15–20.
- [17] A. Aggarwal, P. Kampmann, J. Lemburg, and F. Kirchner, "Haptic Object Recognition in Underwater and Deep-sea Environments," *Journal of Field Robotics*, Aug. 2014.
- [18] M. Meier, M. Schopfer, R. Haschke, and H. Ritter, "A Probabilistic Approach to Tactile Shape Reconstruction," *IEEE Transactions on Robotics*, vol. 27, no. 3, pp. 630–635, June 2011.
- [19] M. Jin, H. Gu, S. Fan, Y. Zhang, and H. Liu, "Object shape recognition approach for sparse point clouds from tactile exploration," in *2013 IEEE International Conference on Robotics and Biomimetics (ROBIO)*, Dec. 2013, pp. 558–562.
- [20] R. A. Russell, "Object recognition by a 'smart' tactile sensor," in *Proceedings of the Australian Conference on Robotics and Automation*, 2000, pp. 93–8.
- [21] A. Schneider, J. Sturm, C. Stachniss, M. Reiser, H. Burkhardt, and W. Burgard, "Object identification with tactile sensors using bag-of-features," in *IEEE/RSJ International Conference on Intelligent Robots and Systems, 2009. IROS 2009*, 2009, pp. 243–248.
- [22] A. Khotanzad and Y. H. Hong, "Invariant image recognition by Zernike moments," *IEEE Transactions on Pattern Analysis and Machine Intelligence*, vol. 12, no. 5, pp. 489–497, May 1990.
- [23] Q. Wu and P. Hall, "Prime Shapes in Natural Images." British Machine Vision Association, 2012, pp. 45.1–45.12.
- [24] R. J. Noll, "Zernike polynomials and atmospheric turbulence," *Journal of the Optical Society of America*, vol. 66, no. 3, pp. 207–211, Mar. 1976.