

General Object Tip Detection and Pose Estimation for Robot Manipulation

Dadhichi Shukla, Özgür Erkent, and Justus Piater

Intelligent and Interactive Systems, Institute of Computer Science, University of Innsbruck
Technikerstr. 21a, 6020 Innsbruck, Austria
dadhichi.shukla@uibk.ac.at

Abstract. Robot manipulation tasks like inserting screws and pegs into a hole or automatic screwing require precise tip pose estimation. We propose a novel method to detect and estimate the tip of elongated objects. We demonstrate that our method can estimate tip pose to millimeter-level accuracy. We adopt a probabilistic, appearance-based object detection framework to detect pegs and bits for electric screw drivers. Screws are difficult to detect with feature- or appearance-based methods due to their reflective characteristics. To overcome this we propose a novel adaptation of RANSAC with a parallel-line model. Subsequently, we employ image moments to detect the tip and its pose. We show that the proposed method allows a robot to perform object insertion with only two pairs of orthogonal views, without visual servoing.

Keywords: Pose estimation, Tool tip detection, Peg-in-hole insertion

1 Introduction and Related Work

A fundamental challenge for robots operating in environments not designed for robots lies in generic, vision-guided manipulation skills. While factory settings can control the robot’s environment to reduce variability, uncertainty and thus perceptual requirements to a minimum, in uncontrolled environments a robot must be able to interact in flexible, reactive perception-action loops.

One generic sensorimotor skill that arises in many different contexts is the insertion of elongated objects into their receptacle, such as screws or pegs into holes, drills onto pilot holes, or screwdriver bits onto the screw head. These operations require the pose of the insertable tip of the object to be estimated with sufficient precision.

Here we present a generic, robust and accurate method for detecting tips of generic, elongated objects for purposes such as robotic insertion. Examples of such objects are shown in Figs. 1a and 1b.

Many prominent strategies involve visual servoing methods [14,3,7] and rely on 3D models of the objects or on visual markers to solve peg-in-hole-style challenges. Hee-Chan Song et al. [9] recently proposed an automated peg-in-hole method for complex-shaped parts. However, such methods critically depend on CAD models of the objects involved. An important aspect in image-based visual servoing is to determine the set of visual features to be used in the control scheme. Visual servoing methods proposed by Chaumette [2] and Liu et al. [8] employed image moments as image features which




Fig. 1: (a)–(b) Elongated objects: Peg, variety of screws, screwdriver bits. (c) Robot inserting a screw into a hole.

can efficiently increase the stability of the system. Since the objects in this work are very small in size compared to the robot hand, they are often strongly occluded, making their visual features difficult to track. Recent work by Hoffman et al. [5] can be related to this study of estimating tool tip location. The method localizes the tool tip by using 3D depth information to segment the tool from the background and performing contour segmentation based on level sets [1]. The method proposed by Stückler et al. [10] also falls in this category. It would be difficult to adapt these methods for the objects in this study since there will be no or negligible depth information due to the minuscule size of the objects.

In this study, therefore, we propose a novel strategy independent of CAD models, fiducial markers or depth information that applies to different types of screws or other elongated objects. After tip pose estimation, it relies on a calibrated camera/robot setup to insert the tip into its receptacle. We do not use visual servoing; instead, we introduce a discrete visual feedback system capable of inserting the object with only two pairs of orthogonal views of the object in hand. Details are discussed in section 4. The experimental setup of this work is shown in Fig. 1c, where the left robot arm grasps the object, and the eye-in-hand system is mounted on the right robot arm.

The vision part of the system contains two steps: (1) Detection of the object held by the robot hand, and (2) Estimation of the tip of the object. The objects of interest in this work are textureless or reflective. Thus, conventional, appearance-based methods have trouble detecting the object. Therefore, we adopt a probabilistic object detection framework based on image gradients proposed by Teney et al. [12], which naturally accommodates variability in scale, shape and appearance of the objects. A brief description of the approach is presented in section 2.1.

However, screws are difficult to detect using this framework due to their reflective characteristics. At the high spatial resolution we require for precise tip estimation, these create spurious image features like edges and corners. We turn this problem to our advantage: Screw threadings exhibit conspicuous, linearly-aligned corner features which are conveniently detected using RANSAC [4]. We use a parallel-line model to detect parallel pairs of collinear points. We further discuss our RANSAC parallel-line (RPL) model in section 2.2.

Precise pose estimation normally requires accurate object models. In this work, however, we seek a generic, object-independent method. Precisely estimating tool tips without tool-specific models is a serious challenge. We succeed thanks to a combina-

tion of three methods: First, screws are detected using *RPL*, and all other objects using *Teney et al.'s method*. Then, building on the assumption that the object of interest exhibits suitable rotational or reflective symmetry about its principal axis, we compute *image moments* proposed by Hu [6] on the objects segmented by either of the first two methods, which greatly increases the localization accuracy of the estimated tip poses. A brief description of image moments and their use in this study is given in section 3.

After estimating the tip, we plan a trajectory for the insertion movement using the KOMO motion planner proposed by Toussaint [13]. This motion planner produces trajectories that avoid self-collisions and obstacles in the scene. We illustrate in section 5 that the fully integrated system can achieve object insertion at millimeter-level accuracy.

2 In-Hand Object Detection

Partial occlusions are prevalent in in-hand object detection. The objects of interest like screws are highly occluded due to their tiny size relative to the hand. Features detected on the hand are confounded with the object features, which further adds to the detection problem. Note that peg-shaped objects can be detected by either of the strategies discussed in the two sections below.

2.1 Probabilistic Appearance-based Model

Learning object models: Pose-Appearance space The probabilistic appearance-based model proposed by Teney et al. [12] performs object detection, recognition and pose estimation in 2D images. The framework is trained on only three real images of the screw-driver bits as shown in Fig. 1b. Edge features of each image are defined on $\mathbb{R}^2 \times S_1^+$, accounting for the position and orientation defined as the appearance space A . Edge points x are then associated with the respective pose w to create *pose/appearance* pairs (x_i, w_i) . These pairs from all training images are concatenated to form the training set $T = (x_i, w_i)_{i=1}^M$. The training set is then used to define a continuous probability distribution ψ on the *pose/appearance* space given by

$$\psi(w, x) = \frac{1}{M} \sum_{(w_i, x_i) \in T} K_1(w, w_i) K_2(x, x_i), \quad (1)$$

where $w \in SE(3)$ and $x \in A$. The use of kernels K_1 and K_2 can be seen here as a Gaussian smoothing over the available training edge points. In a similar fashion the edge features of a test image are stored as the *observations* $O = \{x\}_{i=1}^N$, where $x_i \in A$. These are then used to define a probability density ϕ on A as

$$\phi(x) = \frac{1}{N} \sum_{x_i \in O} K_2(x, x_i). \quad (2)$$

Pose Inference The observation density ϕ is now matched to the learned model ψ in order to find the pose of the object in the test image. The pose w is modelled as random variable $W \in SE(3)$, and its distribution is given by

$$p(w) = \int_A \psi(w, x) \phi(x) dx. \quad (3)$$

The above expression measures the compatibility of the training data at pose w with the distribution of features observed in the test image. It essentially computes the cross-correlation of the distribution ϕ of observations with the model distribution $\psi(w, \cdot)$ at a given pose. Estimating the pose w then involves maximizing $p(w)$ over w , which is done using an efficient scheme [12,11]. An advantage of this framework is that the method generates one flexible model of a generic bit that can be applied to detect a variety of distinct screw-driver bits. Illustrative bit detection results are shown in Fig. 2a.




Fig. 2: Tip detection comparison between appearance-based method [12] and image moments.

2.2 RANSAC Parallel-Line Model

The position of the screw can be detected by fitting a pair of parallel lines passing through a pair of collinear corner features. To understand the proposed RANSAC Parallel-Line (RPL) model let us take n input data points from a test image I . For screws, corner features from the threadings are given as the input data points to RPL, as shown in Fig. 4a. In the case of a peg, locations of the edge features are given as the input data. A pair of parallel lines can be represented by three points. The goal of RPL is to best fit the parallel lines l_1 and l_2 to the observed data. Importantly, the inlier sets \mathcal{L}_1 and \mathcal{L}_2 associated with the lines l_1 and l_2 , respectively should not overlap, i.e., $\mathcal{L}_1 \cap \mathcal{L}_2 = \emptyset$.

To find the location of the screw, we propose a scoring scheme based on three objectives: (1) to maximize total number of inliers $|\mathcal{L}| = |\mathcal{L}_1| + |\mathcal{L}_2|$, (2) to maximize the orthogonal distance $d_{l_1 l_2}$ between the parallel lines, and (3) to minimize the longitudinal distance $d_{c_1 c_2}$ between the centroids of the two inlier sets, i.e. the distance between the centroid c_1 and the projection of the centroid c_2 onto the line l_1 . These variables are normalized by scaling them to the unit interval as

$$\hat{L} = \frac{|\mathcal{L}|}{n}, \hat{d}_{l_1 l_2} = \frac{d_{l_1 l_2}}{I_{diag}}, \hat{d}_{c_1 c_2} = \frac{d_{c_1 c_2}}{I_{diag}},$$

where, I_{diag} is the length of the diagonal of the image.

For a given model l_1, l_2 which fits the observed data, the variables \hat{L} , $\hat{d}_{l_1 l_2}$ and $\hat{d}_{c_1 c_2}$ are independent. We would like to fit the model l_1, l_2 that maximizes score s based

on the aforementioned objectives. We compute the score s of a hypothesis l_1, l_2 as the likelihood function given by

$$P(\hat{L}, \hat{d}_{l_1 l_2}, \hat{d}_{c_1 c_2} | l_1, l_2) = P(\hat{L} | l_1, l_2) P(\hat{d}_{l_1 l_2} | l_1, l_2) P(\hat{d}_{c_1 c_2} | l_1, l_2). \quad (4)$$

The likelihoods $P(\hat{L} | l_1, l_2)$, $P(\hat{d}_{l_1 l_2} | l_1, l_2)$ and $P(\hat{d}_{c_1 c_2} | l_1, l_2)$ of the three individual observations are given by the normal distributions. Therefore, the score of the fitting model at each iteration is given by

$$s = \mathcal{N}(\hat{L}; \mu_{\hat{L}}, \sigma_{\hat{L}}^2) \mathcal{N}(\hat{d}_{l_1 l_2}; \mu_{\hat{d}_{l_1 l_2}}, \sigma_{\hat{d}_{l_1 l_2}}^2) \mathcal{N}(\hat{d}_{c_1 c_2}; \mu_{\hat{d}_{c_1 c_2}}, \sigma_{\hat{d}_{c_1 c_2}}^2). \quad (5)$$

The mean values of these distributions are set as follows, $\mu_{\hat{L}}$ is set to 1 to advocate the inherent notion of RANSAC that the consensus set of the fitting model consists of the most number of inliers. And the geometry of the screw allows us to set $\mu_{\hat{d}_{l_1 l_2}} > 0$, in our experiments 0.5 and $\mu_{\hat{d}_{c_1 c_2}}$ to 0 which attributes to the width of the screw and symmetry about its principal axis, respectively. The variances of these normal distributions are easily set empirically; alternatively, they can be estimated from training data. The best fitting model to the observations is associated with the best score s_{best} which is updated iteratively within RANSAC. Our RANSAC parallel-line fitting method is described in Alg. 1.

Algorithm 1 Proposed RANSAC Parallel-Line (RPL) model

1. Select three random points p_1, p_2 and p_3 from n data points.
 2. Define parallel lines l_1 and l_2 passing through p_1, p_2 and p_3 respectively.
 3. If $l_1 = l_2$ discard p_1, p_2, p_3 and goto step 1.
 4. Find all the inliers \mathcal{L} given a distance threshold t .
 5. Compute score s as given by eq. 5.
 6. If $s_{\text{best}} < s$ discard p_1, p_2, p_3 and goto step 1, else update $s_{\text{best}} = s$.
 7. Repeat these steps until a good model has been found with high confidence.
-

Some hypotheses are shown in Fig. 3. Evidently, the hypothesis with the highest score is the best fit to the observed data points, as shown in Fig. 3e. The location of the screw detected by our RPL model can be seen in Fig. 4c. It can be seen that despite the presence of spurious corner features the proposed method performs robustly in practice.

3 Tip Detection and Pose Estimation using Image Moments

The outcome of both aforementioned object detection strategies are satisfying from an object-detection point of view. Given accurate, object-specific models, these would be sufficient for accurate tip pose estimation. However, even a slight perturbation in the estimated object pose can lead to substantial errors in the estimated tip location. If we fit the learned model to the test images as shown in Fig. 2a, we can see that the




Fig. 3: RANSAC Parallel-line (RPL) model: (a)–(e) Some of the parallel-line hypotheses, (e) Best fitting model with the highest score.




Fig. 4: (a)–(c) Screw detection using our RANSAC Parallel-Line (RPL) method. (d)–(e) Screw tip detection using image moments (section 3); $\theta = 91.98^\circ$.

estimated tip is off the actual tip. Though the probabilistic framework can detect the objects reliably, precise estimation of the tip location is a challenge. To overcome this problem, we adopt image moments as described by Hu [6]. A shape can be summarized by a function of few lower-order moments. We can characterize the image properties like area, centroid, orientation using these moments. For example, to detect object pose we can compute image orientation using second-order central moments.

Image moments have been extensively used to compute geometric features of a shape in a binary image. We adopt image moments to detect the tip of the object. Let us consider an example of the detection of the bit of the screw driver and the screw as shown in Fig. 5a and Fig. 4c, respectively. The location of the object is given by the bounding box (bb). First, the image features like edges (x_{bb}) within the bounding box I_{bb} are retained as shown in Fig. 5b. Then a binary mask is created from those edge points. In the case of screws, the extreme corner features are connected to create the binary mask as shown in Fig. 4c. Morphological operations are then carried out to eliminate islands or unwanted features. Finally, the major axis and orientation θ of the shape is computed using second-order central moments, which is equivalent to fitting an ellipse. It can be seen in Figs. 5c and 4d that one of the end points of the major axis is located closest to the tip. Now the problem is reduced to finding a nearest neighbour. Since we know the pose of the robot hand, we also know whether the tip is facing vertically *up* or *down*. Based on this the appropriate end point ep_j of the major axis, either ep_1 or ep_2 , is selected. The estimated tip of the object is the nearest edge point to ep_j . It can be seen in Fig. 5d and Fig. 4e that the outcome of the proposed tip detection method results in a very precise estimation.




Fig. 5: (a) Screw driver bit detection using probabilistic method [12]. (b)–(d) Tip detection (section 3); $\theta = 123.23^\circ$.

4 Object Insertion

A block diagram of the overall object insertion process can be seen in the Fig. 6. At the start, the pose \mathbf{p}_h of the hole is assumed to be known and the estimation for the pose $\hat{\mathbf{p}}_o$ of the object tip is unknown. All the poses in the diagram include the 6-DoF variables ($\mathbf{p} \in \mathbb{R}^3 \times \mathcal{S}^3$). The difference between \mathbf{p}_h and $\hat{\mathbf{p}}_o$ is given as the error \mathbf{e} to the control system. The decision maker in Fig. 6 plans the motion of the robot such that the camera gathers maximum possible information to estimate the pose of the object tip. This is achieved by taking orthogonal views in this study. As mentioned earlier we integrate our system with the KOMO motion planner provided by Toussaint [13]. It should be noted that due to inaccuracies in the encoder readings of the robot ($\tilde{\mathbf{J}}_R$), there will be a steady-state error between the desired input \mathbf{u}_R and the actual output pose of the robot. Since these errors are not systematic, in our evaluation we measure the final error between the desired and the actual object poses.

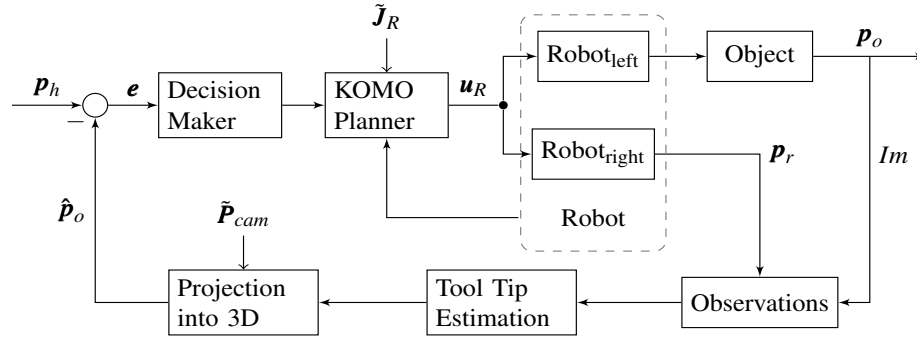


Fig. 6: Block diagram of the overall system.

The experimental setup is situated in a world frame (XYZ) as shown in Fig. 1c. It consists of an eye-in-hand camera setup for the right arm while a hand mounted on the left arm grasps the object. The camera initially looks towards the left arm from a

distance d . The robot first moves to a known 3D location in the world frame. To estimate the pose \mathbf{p}_o of the object after robot motion, an observation is made by acquiring the image Im and estimating the current pose \mathbf{p}_r of the camera from the encoder readings of the robot. Here, the first pair of orthogonal views u, u' of the object in hand is captured by rotating the end effector of the left arm. The tip pose estimation of the object aligns its axes with the hole. The tip detection results of the screw in the first pair of orthogonal views can be seen in Fig. 7. The tip location $(t_x, t_y)_u, (t_x, t_y)_{u'}$ and orientation $\theta_u, \theta_{u'}$ is computed in both views. Then we can estimate the object tip in the world frame by using the camera projection matrix. We observe that mere alignment of the object tip is not enough for the insertion. There will be a small error in the projection due to inaccuracies in the calibration of the camera parameters ($\tilde{\mathbf{P}}_{cam}$), which is inevitable. These errors affect the position of the object tip mainly along the X and Y axes.




Fig. 7: *L-R*: Tip pose estimation in two pairs of orthogonal views $(u, u')_1, (u, u')_2$. Views $(u, u')_1$ are chosen to align the axes of the object and the hole. Views $(u, u')_2$ are chosen to compute translation error of the object tip with respect to the hole in X and Y axes, respectively, in the world frame. $\theta_{u_1} = 91.69^\circ$, $\theta_{u'_1} = 75.00^\circ$, $\theta_{u_2} = 93.90^\circ$, $\theta_{u'_2} = 90.38^\circ$.

The tip estimation is mainly affected by the accumulation of different errors in the robot system due to transformation \mathbf{e}_r , calibration \mathbf{e}_c and tolerances of the motion planner. To limit these errors we repeat the procedure to capture a second pair of the orthogonal views. We re-capture view u and estimate the tip location to achieve position correction along the X axis. Since the object is aligned with the hole, at this step, we move camera to an orthogonal location with respect to view u to capture view u' . The tip estimation in view u' reduces the error along the Y axis. Furthermore, this procedure can be repeated multiple times until position corrections along X and Y axes converge to zero. In practice we achieve tip insertion with only two iterations within admissible error values. In the evaluation we demonstrate that the combined position corrections along X and Y axes increase the number of successful insertions. Tip detection of the screw in the second pair of orthogonal views can be seen in Fig. 7.

5 Experiments and Results

The experiments were carried out with the wooden peg, different types of screws and screw-driver bits in the experimental setup as shown in Fig. 1c to verify the proposed tip detection and insertion method. The setup consists of two 6-DoF KUKA-LWR (Light-Weight Robots), a Schunk SDH hand, and a BlackFly camera (BFLY-PGE-13E4C-CS).

The parameters of interest for the KOMO motion planner are position precision $P_p = 10^{-4}\text{m}$, alignment precision $A_p = 10^{-6}\text{m}$, position tolerance $P_t = 0.005\text{m}$ and angular tolerance $A_t = 0.1\text{rad}$. These parameters were experimentally determined.

The effectiveness of the proposed tip detection algorithm was verified by several trials of insertion. We divide the objects in two groups based on the width of the tip: (1) *thick* objects with a tip of width $\geq 5\text{mm}$ and (2) *thin* objects with a tip of width $< 5\text{mm}$ as shown in the top and bottom rows of Fig. 8a, respectively. Six insertion trials are made with thick objects and five insertion trials are made with thin objects. The results in Fig. 8b show that the proposed method can work under different lighting conditions and background. On visual inspection it can be seen that the proposed method can successfully estimate the tip location in the image for a variety of objects.




Fig. 8: (a) Tip detection results on a variety of screws and a peg; $L-R$: $\theta_1 = 83.79^\circ$, $\theta_2 = 90.46^\circ$, $\theta_3 = 94.79^\circ$, $\theta_4 = 89.34^\circ$, $\theta_5 = 76.29^\circ$, $\theta_6 = 65.85^\circ$. (b) Tip detection in distinct backgrounds; $T-B$: $\theta = 91.66^\circ$, $\theta = 95.68^\circ$

We consider tip pose estimation to be successful if the object tip can be inserted into the hole. Quantitative results can be seen in Fig. 9a and Fig. 9b, where independent trials are coded with distinct colors. Trials that fall inside the dotted circles have errors less than 3.5mm and 3mm for thick and thin object tips, respectively. Three insertion attempts are made for each trial discussed as follows. The points marked \star are the attempts after aligning the object tip with the hole as described in section 4. It can be seen from the plot that the insertion attempts with mere object tip alignment are often quite far from the hole. As mentioned earlier, to overcome this the position corrections along X and Y axes are performed by capturing a second pair of orthogonal views $(u, u')_2$, respectively. The insertion attempts after position correction only along the X axis are marked with \star , and the insertion attempts with position corrections along X and Y axes combined together are marked with \bullet . These position corrections bring the object tip on top of the hole.

We include these position corrections to overcome the inevitable accumulation of camera calibration and steady-state errors along with tolerances of the motion planner to some extent. Moreover, tip insertion significantly depends on estimation of the tip in the world frame. Despite accurate tip detection in the camera frame, the estimation in

the world frame also depends on the width of the object tip. It was observed that the estimates were more accurate for the thin objects as compared to the thick objects, as expected. For six insertion trials of thick objects, the number of successful insertions are zero, three and five after no corrections, correction along the X axis and correction along the X and Y axes, respectively. For five insertion trials of thin objects, the number of successful insertions are two, one and five after no corrections, correction along the X axis and correction in X and Y axes, respectively. It can be seen that the errors are reduced dramatically after corrections along the X and Y axes.

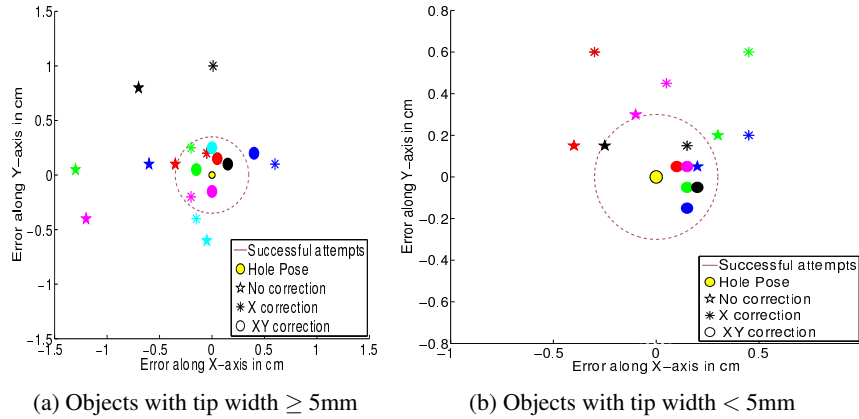


Fig. 9: Quantitative results of tip insertion. Trials that fall inside the dotted circle are successful attempts of insertion.

6 Conclusion

In this work, we propose a novel object tip detection and pose estimation method for elongated objects. The proposed method was tested with two 6-DOF KUKA Light Weight Robots to perform peg-in-hole style tasks. The principal contribution is that this method is able to extract object tip poses with high accuracy, without using object-specific models. This allows our method to be used flexibly in a wide range of application settings. The only key assumption is that the tip to be estimated lies on the major axis of the elongated object.

The following conclusions were drawn from this study:

- We achieve object insertion task with only two pairs of orthogonal views.
- A success rate of 100% can be achieved by the system for objects with a sharp tip. However, the errors increase as the width of the object tip increases.
- The proposed method can estimate the tip location even with high occlusion levels.

- Since this study focuses on visual object tip localization, we validated accuracy using stiff robot position control. Using impedance control instead, practical insertion performance can be significantly increased, especially for wider objects.

The research leading to these results has received funding from the European Community's Seventh Framework Programme FP7/2007-2013 (Specific Programme Cooperation, Theme 3, Information and Communication Technologies) under grant agreement no. 610878, 3rd HAND.

References

1. Chan, T.F., Vese, L.A.: Active contours without edges. *Image processing, IEEE transactions on* 10(2), 266–277 (2001)
2. Chaumette, F.: Image moments: a general and useful set of features for visual servoing. *Robotics, IEEE Transactions on* 20(4), 713–723 (2004)
3. Dantam, N.T., Amor, H.B., Christensen, H.I., Stilman, M.: Online Multi-Camera Registration for Bimanual Workspace Trajectories. In: 14th IEEE-RAS International Conference on Humanoid Robots (Humanoids). pp. 588–593. IEEE (2014)
4. Fischler, M.A., Bolles, R.C.: Random sample consensus: a paradigm for model fitting with applications to image analysis and automated cartography. *Communications of the ACM* 24(6), 381–395 (1981)
5. Hoffmann, H., Chen, Z., Earl, D., Mitchell, D., Salemi, B., Sinapov, J.: Adaptive robotic tool use under variable grasps. *Robotics and Autonomous Systems* 62(6), 833–846 (2014)
6. Hu, M.K.: Visual pattern recognition by moment invariants. *Information Theory, IRE Transactions on* 8(2), 179–187 (1962)
7. Huang, S., Yamakawa, Y., Senoo, T., Ishikawa, M.: Dynamic compensation by fusing a high-speed actuator and high-speed visual feedback with its application to fast peg-and-hole alignment. *Advanced Robotics* 28(9), 613–624 (2014)
8. Liu, S., Xie, W.F., Su, C.Y.: Image-based visual servoing using improved image moments. In: *Information and Automation, 2009. ICIA'09. International Conference on*. pp. 577–582. IEEE (2009)
9. Song, H.C., Kim, Y.L., Song, J.B.: Automated guidance of peg-in-hole assembly tasks for complex-shaped parts. In: *Intelligent Robots and Systems (IROS 2014), 2014 IEEE/RSJ International Conference on*. pp. 4517–4522. IEEE (2014)
10. Stückler, J., Behnke, S.: Adaptive Tool-Use Strategies for Anthropomorphic Service Robots. In: 14th IEEE-RAS International Conference on Humanoid Robots. IEEE (2014)
11. Teney, D., Piater, J.: Continuous Pose Estimation in 2D Images at Instance and Category Levels. In: *Tenth Conference on Computer and Robot Vision*. pp. 121–127. IEEE (5 2013), <https://iis.uibk.ac.at/public/papers/Teney-2013-CRV.pdf>
12. Teney, D., Piater, J.: Multiview feature distributions for object detection and continuous pose estimation. *Computer Vision and Image Understanding* 125, 265–282 (8 2014), <https://iis.uibk.ac.at/public/papers/Teney-2014-CVIU.pdf>
13. Toussaint, M.: Newton methods for k-order markov constrained motion problems. *CoRR abs/1407.0414* (2014), <http://arxiv.org/abs/1407.0414>
14. Wang, J., Cho, H.: Micropeg and hole alignment using image moments based visual servoing method. *Industrial Electronics, IEEE Transactions on* 55(3), 1286–1294 (2008)