

Multi-View Keypoints for Reliable 6D Object Pose Estimation

Alan Li and Angela P. Schoellig

Abstract—6D Object pose estimation is a fundamental component in robotics enabling efficient interaction with the environment. 6D pose estimation is particularly challenging in bin-picking applications, where many objects are low-feature and reflective, and self-occlusion between objects of the same type is common. We propose a novel multi-view approach leveraging known camera transformations from an eye-in-hand setup to combine heatmap and keypoint estimates into a probability density map over 3D space. The result is a robust approach that is scalable in the number of views. It relies on a confidence score composed of keypoint probabilities and point-cloud alignment error, which allows reliable rejection of false positives. We demonstrate an average pose estimation error of approximately 0.5 mm and 2 degrees across a variety of difficult low-feature and reflective objects in the ROBI dataset, while also surpassing the state-of-art correct detection rate, measured using the 10% object diameter threshold on ADD error.

I. INTRODUCTION

Pose estimation of surrounding objects is a crucial yet difficult task for robotic systems that are tasked to interact robustly with their environments. In bin-picking applications, for example, objects may be assorted randomly in a bin in various poses and with various levels of occlusion. In recent years, convolutional neural networks (CNNs) have demonstrated great success in the problem of object pose estimation, particularly on datasets such as LINEMOD and T-LESS [10], where objects are typically rich in texture and features, resulting in near-optimal color and depth images. However, single-view object pose estimation performs poorly when data is missing, sparse and noisy due to surfaces being shiny, glossy or transparent, or due to occlusion and ambiguity from other objects [3]. As a result, most methods detect large amounts of false positives, which are not easily separable from correct detections.

The goal of this work is to demonstrate accurate and reliable pose estimation in these sub-optimal conditions by leveraging estimates from multiple viewpoints to reduce estimation error and uncertainty. We leverage a known CAD model of the objects to generate training data for our estimation network, and to perform the ICP-based refinement step. There are three major technical contributions that contribute to higher overall accuracy and reliability compared to other work in the literature:

- 1) We use a scalable method of combining 2D heatmaps across multiple camera viewpoints, leveraging known

The authors are with the Dynamic Systems Lab, Institute for Aerospace Studies, University of Toronto, Canada, and affiliated with the Vector Institute for Artificial Intelligence. Angela P. Schoellig is also with the Technical University of Munich (TUM). E-mails: {firstname.lastname}@robotics.utias.utoronto.ca

This work was supported by resources from Epson Canada.




Fig. 1. The information from multiple viewpoints (with known relative camera transformation) is combined to estimate an object’s 6D pose. From single-view keypoint heatmaps, we calculate multi-view uncertainty estimates for filtering and ranking of candidate poses and demonstrate improved accuracy and reliability over current works.

camera transformations to project keypoint candidates and produce a probability density over 3D space. Compared to current approaches, this method allows for a more natural representation of pose uncertainty, while also enabling scaling to a higher number of viewpoints.

- 2) We use the predicted 2D object mask and keypoint uncertainties to filter the depth map for pose refinement. This allows for removal of outlier points from the background and other objects, resulting in more accurate point cloud alignment compared to the current state-of-the-art [11].
- 3) We use a novel and robust confidence score incorporating object rigidity constraints, keypoint probabilities, and point cloud alignment error to reliably estimate the accuracy of the pose, allowing for rejection of the majority of false detections.

II. RELATED WORKS

A. 6D Pose estimation

Object pose estimation from images is a core problem in computer vision, with numerous applications in robotic manipulation, augmented reality, and autonomous driving. Recent advances have leveraged the power of deep learning to surmount the shortcomings of traditional methods, which estimate pose based on local feature correspondences or template matching [1]. These methods have been revisited using the advantages of CNNs to automatically learn features and implicitly estimate pose [9] [7].

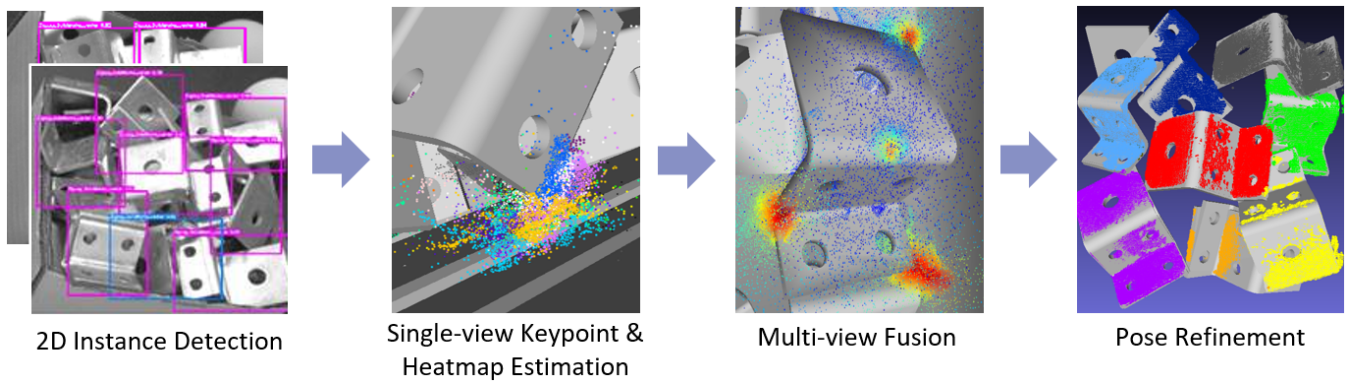


Fig. 2. An overview of the multi-view pipeline, showing object instance detection, followed by keypoint/heatmap computation, multi-view fusion, and finally the pose refinement step.

B. Keypoint Representation

The representation of 6D object pose through object keypoints has also regained popularity with the introduction of CNNs. Networks are trained to estimate the 2D or 3D location of several predefined points on an object model, which can then be used to retrieve the 6D pose of the object in the image using the correspondences between the locations of the predicted keypoints and the ones defined on the object model [2]. This not only simplifies the pose estimation problem into more interpretable steps, but also allows for techniques to refine and combine estimates from different viewpoints at a more fundamental level. We leverage this keypoint representation for our multi-view pipeline, allowing for more robust fusion of occluded or even partially incorrect estimates, compared to the typical 6D clustering of pose candidates of other multi-view approaches.

C. Multi-view Detection

Several recent works have explored multi-view detection using keypoints, although typically using only two or three viewpoints, often with unknown transformations between each viewpoint. A recent work combines Stereo 2D keypoint detections within a trained network to estimate the 3D positions of each keypoint for transparent and translucent objects [3]. The results show an improvement over the state of the art by a factor of 1.5-3.5, using only two known views from a stereo camera. The objects are unoccluded as opposed to placed in bins, however, and only one instance of the object can be present in the scene. While this approach is faster, the network architecture is fixed to take in exactly two stereo images as input, and to output exactly one keypoint prediction. Multi-object detection will likely prove difficult as repeated features and occlusions from different instances of the same object may confuse the depth estimation within the network.

Another method proposes the use of multiple monocular 2D keypoint estimates from spread out viewpoints to estimate the poses of vehicles [4]. The transformations between these viewpoints are unknown; instead the approach adds constraints based on object rigidity and the relative positions of different keypoints in the image to solve for the vehicle

poses. We borrow the object rigidity constraint from this paper, but also make use of known transformations between viewpoints from a robot arm, to allow more precise fusion of estimates between views.

CosyPose is able to achieve higher performance in more cluttered scenes by utilizing a larger number of views [14]. The scene is assumed to be static across the different viewpoints, and the poses of all objects within the scene are estimated. This is then used to estimate the camera pose across different viewpoints, followed by bundle adjustment to refine the estimates and generate a globally consistent scene across all views.

In our proposed approach, the known transformation between camera viewpoints is leveraged to transform all estimates into the world frame, and fusion of estimates from different viewpoints occurs at the keypoint level, as opposed to the pose level, improving the robustness of the resulting pose estimates in the presence of occlusion.

III. METHOD

The proposed approach involves a pipeline of several steps, with the main contributions found in the multi-view fusion and pose refinement steps, shown in Fig 2. First, the 2D bounding boxes for each object instance in each viewpoint are detected; the cropped bounding boxes are then used to detect the keypoint locations and heatmaps. The detections are then projected into the same frame of reference and used together using the known camera poses, and finally the extracted object poses are refined using the filtered point cloud of the scene.

A. 2D Object Instance Detection

In the first stage, object instances are detected in the form of 2D bounding boxes. The network is derived from YOLO object detection [15], specifically trained using multi-view RGB-D data that is rendered synthetically using domain randomization [7]. The network is trained to output a tight 2D bounding box centered on each object instance; the detections are then cropped and resized to 128x128 before being passed into the keypoint and heatmap detection network.




Fig. 3. Reconstruction of the object instance point cloud by applying the predicted mask to the depth images and projecting them to the world frame. The process results in noise and outliers around the edges due to imperfect mask prediction.




Fig. 4. Visualization of point cloud filtering by keypoint correspondence. By selecting only the highest probability keypoints, we can filter out outliers and noisy points in the object point cloud.

B. Single-view Keypoint and Heatmap Estimation

The keypoint detection network is based off PVNet, a pixel-wise voting network for 6D pose estimation [7]. The network takes as input an image patch containing an object of interest, and predicts a keypoint location from each pixel input, in the form of an offset to the pixel location itself. The result is a large number of votes, one from each input pixel, of the predicted keypoint location; these votes form a cluster of keypoint candidates in 3D space. The network also predicts an object mask, which is used to filter out pixels that are not relevant to the object. Although only three keypoints are needed to determine 6D pose, we choose to instead estimate 10 to add redundancies and improve the robustness in the presence of occlusions. The main modification to the network is addition of a heatmap output, where the probability distribution of the keypoint location over the input image is estimated. This heatmap estimate is used in the next step to estimate the multi-view probabilities of keypoint candidates. The training data is generated synthetically by simulating the random arrangement of parts within a bin, followed by realistic rendering of the scene to RGB in Blender. The depth images are rendered using NxDView, a depth camera simulator supplied with the Ensenso camera A total of 500 different scenes are rendered for use as training data, which are cropped around individual object instances and augmented using Domain Randomization [21].

C. Multiview Fusion

After obtaining keypoint estimates from each viewpoint, the keypoints are projected into the world frame using the

known camera poses of each view. This allows for multi-view clustering and object association between different views. Objects are clustered based on their center keypoint, and all keypoints within 2mm of the cluster center are determined to be the same object. Since each viewpoint also has a corresponding 2D heatmap of keypoint probabilities, we can calculate the estimated probability over all views of a given 3D keypoint candidate in world by projecting the 3D point into each of the heatmaps and retrieving the probability at the resulting pixel, as shown in Fig. 1.

The probability of a keypoint being at a point \mathbf{p}_k in the world frame is as following:

$$Pr(\mathbf{p}_k) = \prod_i^n h_i(u_k^i, v_k^i), \quad (1)$$

where $h_i(u, v)$ represents the heatmap probability at pixel (u, v) in viewpoint i , and n represents the total number of viewpoints. The pixel coordinates u_k^i and v_k^i of keypoint candidate \mathbf{p}_k in viewpoint i are found by projecting the keypoint to the i^{th} camera image coordinates using intrinsic matrix \mathbf{K} and transformation from the world frame to the i^{th} camera pose $\mathbf{T}_w^{c_i}$:

$$\begin{bmatrix} u_k^i \\ v_k^i \\ 1 \end{bmatrix} = \mathbf{K} \frac{1}{z_k^{c_i}} \mathbf{T}_w^{c_i} \mathbf{p}_k, \quad (2)$$

where $z_k^{c_i}$ is the z-coordinate of the keypoint candidate projected to the camera frame. Points which end up being projected outside of the image patch receive a probability of zero. Taking the negative log likelihood, the uncertainty of the k^{th} keypoint $U(\mathbf{p}_k)$ can be represented as:

$$U(\mathbf{p}_k) = \sum_i^n -\log(h_i(u_k^i, v_k^i)). \quad (3)$$

The result after calculating the uncertainty for all candidates over all viewpoints is a probability density in 3D space, representing the likelihood of a given keypoint position.

The next step is the retrieval of the 6D object pose from the keypoint correspondences between the scene and the original CAD model. We use a RANSAC-based approach to sample and estimate candidate poses based on the previously calculated probability densities [5]. The goal is to find the object pose which minimizes the overall keypoint uncertainty subject to the object rigidity constraint, which constrains the relative positions between the 10 keypoints. We define the rigidity constraint as a 6D rigid transformation matrix from the original keypoints defined on the object to the candidate object pose.

Thus we can represent each keypoint candidate \mathbf{p}_k as the original keypoint o_k defined on the object transformed by the candidate object pose, represented as rotation matrix \mathbf{R} and translation \mathbf{t} :

$$\mathbf{p}_k = \mathbf{R}o_k + \mathbf{t}. \quad (4)$$




Fig. 5. An overview of the ROBI test dataset, which includes seven objects and two bin fill levels [16].

The maximum likelihood estimate would then be the transformation which minimizes the keypoint uncertainties across each of the 10 keypoints k in K :

$$\min_{\mathbf{R}, \mathbf{t}} \sum_{k \in K} U(\mathbf{p}_k) = \min_{\mathbf{R}, \mathbf{t}} \sum_{k \in K} U(\mathbf{R}o_k + \mathbf{t}). \quad (5)$$

We approach this optimization with a heuristic method, using RANSAC to robustify against outliers. We produce candidate poses \mathbf{R} and \mathbf{t} by sampling three keypoints out of the ten keypoint groups according to their probability density maps, and estimating the transformation to the corresponding origin keypoints on the object model. The transformation is then used to reproject all ten keypoints on the model into the probability density maps, allowing us to calculate the overall keypoint uncertainty over all ten keypoints. The top 5 poses which result in the lowest uncertainty are selected for the next refinement step.

D. Pose Refinement using Filtered ICP

Following the poses estimated from the previous step, a refinement step is used to further align the pose estimates with the depth images obtained from the scene. Using the object mask predicted from the keypoint network, the depth images are masked such that only the object of interest remains.

These filtered depth images are then projected to into the world frame and combined using the known camera poses of each viewpoint, producing a rough 3D point cloud of each object. These point clouds contain noise and outlier points from surrounding objects and surfaces which may affect alignment quality, however, and thus a filtering step is performed first.

Recall that for keypoint prediction, each pixel within the predicted object mask is used to estimate an offset towards the keypoint position. Additionally, each keypoint candidate is assigned an uncertainty score, based on the heatmap probability after projecting it to each viewpoint. This probability can also be related back to the original pixel within the mask which produced said keypoint prediction, giving a quality score for each point in the object point cloud. This enables the filtering of the point cloud by keypoint

uncertainty, removing outliers and noisy points, as points which do not lay on the surface of the object have low probability of producing an accurate keypoint estimate.

These reconstructed, filtered 3D point clouds are then used in the Iterative Closest Point algorithm (ICP) [11] to align the candidate poses with the observed depth data, producing the final refined outputs.

E. Confidence/Uncertainty Score Calculation.

The final confidence score given to each candidate pose is a combination of the ICP alignment error, along with the overall keypoint uncertainty. Since the camera viewpoints all tend to be within a cone above the object, the keypoint uncertainty tends to have higher variance in the Z-direction, parallel to the camera view, and lower variance, and therefore higher sensitivity in the X and Y directions. In contrast, the ICP alignment error measures mainly the Z-error, as the point cloud defines a surface of depth values spread across the X and Y axes. By combining these two metrics, we are able to create a reliable uncertainty score which can be used to safely reject the majority of false detections, while minimizing the number of correct detections being rejected.

IV. EXPERIMENTS

The effectiveness of the multi-view approach is shown using the newly released ROBI dataset. This dataset includes a total of 63 bin-picking scenes captured with an active stereo depth Ensenso N35 sensor [16]. For each scene, a view sphere totalling 88 RGB images and depth maps are captured from an Ensenso N35 stereo depth camera, and are annotated with accurate 6D poses of visible objects and an associated visibility score. A total of seven different objects are found in the dataset, arranged randomly within bins. The objects are metallic and highly reflective, with varying levels of symmetry. This results in a dataset representing the most difficult scenes typically found in bin picking applications, with self-occlusions between multiple instances of the same object, and reflections causing missing depth data and false RGB edges.

Table 2 shows the pose estimation results on several objects in the ROBI dataset. A total of 8 different viewpoints

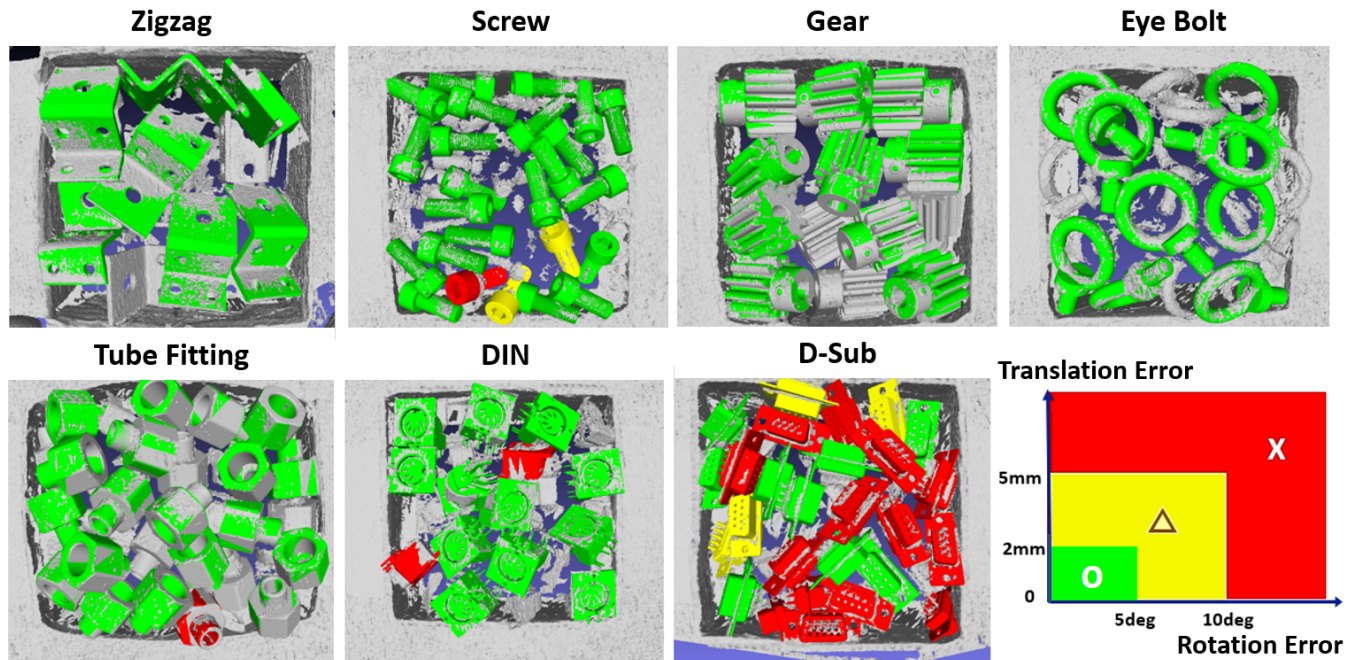


Fig. 6. Sample detections for bins of each object. Estimation performance is shown by color, with the color scheme being defined at the bottom right. Higher error cases can be seen when entire faces of the object are occluded by the bin walls or other objects. Correct detections (green) show an average accuracy of approximately 0.5 mm and 2 degrees on average across all objects.

TABLE I
POSE ESTIMATION ACCURACY AVERAGED OVER 5 BINS

Type	Percentage of Detections and Average Error (mm, deg)						
	Zigzag	Screw	Gear	Eye Bolt	Tube Fitting	DIN	D-Sub
○	98.2 (0.6, 1.2)	86.3 (0.6, 2.2)	79.6 (0.2, 1.4)	84.8 (0.5, 1.9)	86.1 (0.3, 2.07)	69.7 (0.5, 2.1)	35.3 (0.6, 2.2)
△	1.7 (2.6, 5.5)	7.7 (1.1, 7.9)	7.4 (1.6, 7.5)	11.1 (0.8, 7.7)	8.9 (0.4, 7.5)	6.6 (1.3, 8.1)	12.0 (1.2, 8.4)
×	0 (N/A)	6.0 (3.0, 16.8)	12.9 (3.4, 14.6)	4.2 (2.9, 20.0)	1 (1.9, 17.1)	23.7 (3.3, 72.4)	52.7 (4.0, 129)

are used as a balance between speed and accuracy. Sensitivity analysis of the keypoint accuracy to the number of views showed the majority of the performance improvements leveled off past 10 viewpoints. The ROBI evaluation consists of the top layer of objects which are less than 40% occluded in the majority of the chosen views. The baseline performance is measured using a method proposed by Drost et al [17], called point-pair features, which relies on 3D point-cloud data as input. It was chosen as it ranks amongst the top methods in the BOP dataset [12], while still benefiting from multi-view data through the fusion of depth maps. Also tested was a multi-view estimation pipeline called CosyPose [14], ranking number one overall in the BOP dataset for correct detection rate. Its approach relies on single-view pose estimation in each viewpoint, followed by a pose clustering step to combine estimates into a multi-view estimate. The implementation was taken directly from the publicly available source code online. Additionally, we use another state-of-the-art object pose estimation method as a single-view baseline: DC-Net [19]. The correct detection rate is defined as the percentage of visible objects in the bin associated

with an output pose less than 10% ADD error, as proposed by Hinterstoisser et al [20]. The average distance between vertices on the ground truth pose of the model and the corresponding vertices on the predicted pose is calculated, and predictions with less than 10% average error w.r.t. the object diameter are considered correct detections.

TABLE II
EVALUATION AND COMPARISON RESULTS ON THE ROBI DATASET

Object	Correct Detection Rate (<10% ADD Error)			
	DC-Net	CosyPose	PPF	Ours
Zigzag	30.9	92.4	86.7	98.5
Eye Bolt	53.5	75.0	88.7	96.8
DIN	18.7	51.9	61.9	70.7
Gear	77.6	82.1	76.8	86.0
Tube Fit.	74.8	85.8	89.5	87.0
Screw	67.7	90.3	78.1	96.2
D-Sub	10.6	4.1	22.9	31.3

V. DISCUSSION AND LIMITATIONS

The correct detection rate of our method outperforms the others in all but one object, the Tube Fitting, where PPF




Fig. 7. Histogram of detection quality versus the algorithm log-uncertainty for three objects. For all objects other than D-Sub, the green detections can be separated from the majority of the yellow and red detections by thresholding based on an uncertainty score of around 1.0-1.5.




Fig. 8. Plots of the precision-recall curve for Zigzag and DIN objects, demonstrating the effectiveness of our confidence score in maintaining maximum precision while also achieving high detection rate.

performs slightly better. However, looking at the precision of PPF reveals that it achieves this higher detection rate at the cost of many false positives; its confidence score is unable to sufficiently distinguish between correct and incorrect detections due to the lack of distinctive features in its depth-only approach. While CosyPose achieves higher overall precision in two objects, our method attains a larger area under the precision-recall curve, meaning the lower accuracy detections are given lower confidences and can be filtered out while minimizing the effect on correct detections, as seen in Fig. 8. The single view DC-Net suffers particularly for Zigzag and Eye Bolt objects, which are reflective and tend to have large areas of missing depth—this is where the multi-view approaches show the largest gains. We found that the majority of false detections of our approach occur when the object is occluded on multiple sides, particularly by the bin walls. This results in lower accuracy, as the bin effectively blocks visibility of one side of the object, regardless of the viewpoint used. This is reflected in the uncertainty score however, which can be used to filter out the majority of these cases. The D-Sub connector proves difficult for all approaches, as the near symmetry in two axes gives four poses which all appear nearly identical from

many viewpoints.

Fig. 7 visualizes the separability of the three detection types, in the form of a histogram sorted by the algorithm uncertainty score. By thresholding detections based on the uncertainty, we can retain perfect precision for approximately 50% of all correct detections in green before the higher error yellow and eventually red detections are observed. This allows us to select the confidence threshold depending on the precision required for the use case, with more detections coming at the expense of a slightly higher chance of error.

VI. CONCLUSION

The main contribution of this paper is a novel method of joining keypoint estimates from different views in a natural, probabilistic way by leveraging 2D heatmaps and known camera transformations from an eye-in-hand camera, allowing for a highly-multiview approach to increase accuracy and produce reliable certainty estimates for each pose. The key advantages over existing methods include a scalable approach allowing for higher accuracies with an increasing number of views, as well as a large improvement in reliability stemming from the novel confidence score that rejects poses with higher error efficiently.

REFERENCES

- [1] Lowe, D.G. Object recognition from local scale-invariant features. In *Proceedings of the Seventh IEEE International Conference on Computer Vision*. Volume 2, 1150–1157 (1999)
- [2] S. T. a. J. Malik. Viewpoints and keypoint. In *IEEE Conference on Computer Vision and Pattern Recognition, CVPR*, [Online]. (2015)
- [3] R. J. Xingyu Liu KeyPose: Multi-View 3D Labeling and Keypoint Estimation for Transparent Objects, Robotics at Google, [Online]. Available: <https://arxiv.org/pdf/1912.02805.pdf>. (2019)
- [4] S. L. Wenhao Ding, Vehicle Pose and Shape Estimation through Multiple Monocular Vision, [Online]. Available: <https://arxiv.org/pdf/1802.03515.pdf>. (2019)
- [5] M. Fischler and R. Bolles Random Sample Consensus: A Paradigm for Model Fitting with Applications to Image Analysis and Automated Cartography In *Communications ACM*, 24, 381–395, (1981)
- [6] K. He, X. Zhang, S. Ren, and J. Sun. Deep Residual Learning for Image Recognition. In *Conference on Computer Vision and Pattern Recognition*, 770–778, (2016)
- [7] Peng, S., Liu, Y., Huang, Q., Zhou, X., Bao, H. Pvnnet: Pixel-wise voting network for 6dof pose estimation. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*. (2019)
- [8] Park, K., Patten, T., Vincze, M. Pix2pose: Pixel-wise coordinate regression of objects for 6d pose estimation. In *Proceedings of the IEEE International Conference on Computer Vision*. (2019)
- [9] Xiang, Y., Schmidt, T., Narayanan, V., Fox, D. PoseCNN: A convolutional neural network for 6D object pose estimation in cluttered scenes. In *Robotics: Science and Systems XIV*. (2018)
- [10] Hodan, T., Haluza, P., Obdr̃z'alek, S., Matas, J., Lourakis, M., Zabulis, X. T-LESS: ~ An RGB-D dataset for 6D pose estimation of Texture-Less objects. In *IEEE Winter Conference on Applications of Computer Vision (WACV)*. (2017)
- [11] Zhang, Z. Iterative point matching for registration of free-form curves and surfaces. In *International journal of computer vision* 13(2), 119–152, (1994)
- [12] Hodan, T., Michel, F., Brachmann, E., Kehl, W., GlentBuch, A., Kraft, D., Drost, B., Vidal, J., Ihrke, S., Zabulis, X., et al. Bop: Benchmark for 6d object pose estimation. In *Proceedings of the European Conference on Computer Vision (ECCV)*. 19–34, (2018)
- [13] Tobin, J., Fong, R., Ray, A., Schneider, J., Zaremba, W., Abbeel, P. Domain randomization for transferring deep neural networks from simulation to the real world. In *2017 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*. 23–30, (2017)
- [14] Yann Labbe, CosyPose: Consistent multi-view multi-object 6D pose estimation. [Online]. Available: <https://arxiv.org/pdf/2008.08465.pdf>. (2020)
- [15] S. D. Joseph Redmon, You Only Look Once: Unified, Real-Time Object Detection. [Online]. Available: <https://arxiv.org/pdf/1506.02640.pdf>. (2016)
- [16] Jun Yang ROBI: A Multi-View Dataset for Reflective Objects in Robotic Bin-Picking, [Online]. Available: <https://arxiv.org/pdf/2105.04112.pdf>. (2021)
- [17] B. Drost, M. Ulrich, N. Navab, and S. Ilic, Model globally, match locally: Efficient and robust 3d object recognition, in 2010 IEEE computer society conference on computer vision and pattern recognition, 998–1005, (2010)
- [18] Sundermeyer, M., Marton, Z.C., Durner, M., Triebel, R., Augmented autoencoders: Implicit 3d orientation learning for 6d object detection, *International Journal of Computer Vision* 128(3), 714–729, (2020)
- [19] Tian, M., Pan, L., Ang, M.H., Lee, G.H., Robust 6d object pose estimation by learning rgb-d features, 2020 IEEE International Conference on Robotics and Automation (ICRA). pp. 6218–6224, (2020)
- [20] Hinterstoisser, S., Lepetit, V., Ilic, S., Holzer, S., Bradski, G., Konolige, K., Navab, N. Model based training, detection and pose estimation of texture-less 3D objects in heavily cluttered scenes, *ACCV*, (2012)
- [21] Tobin,J., Fong, R., Ray, A., Schneider,J., Zaremba, W., Abbeel, P. Domain Randomization for Transferring Deep Neural Networks from Simulation to the Real World, *IROS*, (2017)