

6-DoF Object Pose from Semantic Keypoints

Georgios Pavlakos¹, Xiaowei Zhou¹, Aaron Chan¹, Konstantinos G. Derpanis², and Kostas Daniilidis¹

Abstract—This paper presents a novel approach to estimating the continuous six degree of freedom (6-DoF) pose (3D translation and rotation) of an object from a single RGB image. The approach combines semantic keypoints predicted by a convolutional network (convnet) with a deformable shape model. Unlike prior work, we are agnostic to whether the object is textured or textureless, as the convnet learns the optimal representation from the available training image data. Furthermore, the same approach can be applied to instance- and class-based pose recovery. Empirically, we show that the proposed approach can accurately recover the 6-DoF object pose for both instance- and class-based scenarios with a cluttered background. For class-based object pose estimation, state-of-the-art accuracy is shown on the large-scale PASCAL3D+ dataset.

I. INTRODUCTION

This paper addresses the task of estimating the six degree of freedom (6-DoF) pose (3D translation and rotation) of an object from a single image. Despite its importance in a variety of applications, e.g., robotic manipulation, and its intense study, most solutions tend to treat objects on a case-by-case basis. For instance, approaches can be discerned by whether they address “sufficiently” textured objects with those that are textureless. Some approaches focus on instance-based object detection while others address object classes. In this work, we strive for an approach where the admissibility of objects considered is as wide as possible.

Our approach combines statistical models of appearance and the 3D shape layout of objects for pose estimation. It consists of two stages that first reasons about the 2D projected shape of an object captured by a set of 2D semantic keypoints and then estimates the 3D shape and camera viewpoint consistent with the keypoints. In the first stage, we use a high capacity convolutional network (convnet) to predict a set of semantic keypoints. Here, the network takes advantage of its ability to aggregate appearance information over a wide-field of view, as compared to localized part models, e.g., [1], to make reliable predictions of the semantic keypoints. In the second stage, the semantic keypoint predictions are used to explicitly reason about the 3D shape of an object and the camera pose modeled by a weak or full perspective camera model. Shape and pose estimates are realized by maximizing the geometric consistency between the parametrized shape model and the 2D semantic keypoints. In the instance-based object case where a corresponding point cloud is provided

with the image, our method can provide a robust way to initialize the iterative closest point (ICP) algorithm [2].

II. RELATED WORK

Estimating the 6-DoF pose of an object from a single image has attracted significant study. Given a rigid 3D object model and a set of 2D-to-3D point correspondences, various solutions have been explored, e.g., [3], [4]. This is commonly referred to as the Perspective-n-Point problem (PnP). To relax the assumption of known 2D landmarks, a number of approaches [5], [6], [7] have considered the detection of discriminative image keypoints, such as SIFT [8], with highly textured objects. A drawback with these approaches is that they are inadequate for addressing textureless objects and their performance is susceptible to scene clutter.

Holistic template-based approaches are one of the earliest approaches considered in the object detection literature. To accommodate appearance variation due to camera capture viewpoint, a set of template images of the object instance are captured about the view sphere and are compared to the input image at runtime. In recent years, template-based methods have received renewed interest due to the advent of accelerated matching schemes and their ability to detect textureless objects by way of focusing their model description on the object shape [9], [10], [11], [7], [12]. While impressive results in terms of accuracy and speed have been demonstrated, holistic template-based approaches are limited to instance-based object detection. To address class variability and viewpoint, various approaches have used a collection of 2D appearance-based part templates trained separately on discretized views [1], [13], [14], [15], [16].

Convolutional networks (convnets) [17], [18] have emerged as the method of choice for a variety of problems. Closest to the current work is their application in camera viewpoint and keypoint prediction. Convnets have been used to predict the camera’s viewpoint with respect to the object by way of direct regression or casting the problem as classification into a set discrete views [19], [20], [21]. While these approaches allow for object category pose estimation they do not provide fine-grained information about the 3D layout of the object. Convnet-based keypoint prediction for human pose estimation (e.g., [22], [23], [24], [25]) has attracted considerable study, while limited attention has been given to their application with generic object categories [26], [20]. Their success is due in part to the high discriminative capacity of the network. Furthermore, their ability to aggregate information over a wide field of view allows for the resolution of ambiguities (e.g., symmetry) and for localizing occluding joints.

¹G. Pavlakos, X. Zhou, A. Chan and K. Daniilidis are with the Department of Computer and Information Science, University of Pennsylvania, PA, USA, {pavlakos, xiaowz, aarchan, kostas}@seas.upenn.edu

²K. Derpanis is with the Department of Computer Science, Ryerson University, ON, Canada, kostas@scs.ryerson.ca

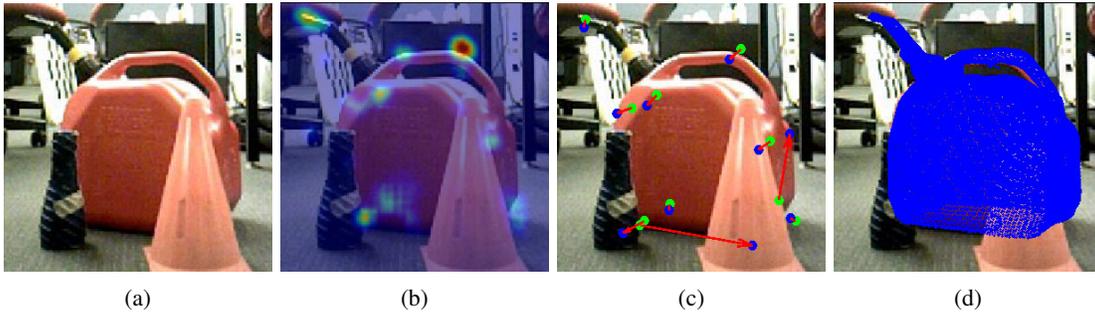


Fig. 1: Pipeline of our approach. Given a single RGB image of an object (a), we localize a set of class-specific keypoints using a convnet with the stacked hourglass design. The output of this step is a set of heatmaps for each keypoint (b), sometimes leading to false detections. In (c), green dots represent the detected keypoints and the corresponding blue dots (connected with an arrow) the groundtruth locations. For robustness against such localization errors, we solve a fitting problem to enforce global consistency of the keypoints, where the response of the heatmaps is used as a measure of certainty for each keypoint. The optimization recovers the full 6-DoF pose of the object (d).

Statistical shape-based models tackle recognition by aligning a shape subspace model to image features. While originally proposed in the context of 2D shape [27] they have proven useful for modelling the 3D shape of a host of object classes, e.g., faces [28], cars [29] and human pose [30]. In recent work [31], data-driven discriminative landmark hypotheses were combined with a 3D deformable shape model and a weak perspective camera model in a convex optimization framework to globally recover the shape and pose of an object in a single image. Here, we adapt this approach and extend it with a perspective camera model, in cases where the camera intrinsics are known.

Contributions In the light of previous work, the contributions of our work are as follows:

- We present an efficient approach that combines highly reliable (semantic) keypoints predicted by a convnet with a deformable shape model to estimate the continuous 6-DoF pose of an object. Unlike previous work, we are agnostic to whether the object is textured or textureless, as the convnet learns the optimal representation from the available training image data. Furthermore, the same approach can be applied to instance- and class-based pose recovery. For instance-based recovery, where the 3D object model and the point cloud are available, our approach can provide a robust initialization to ICP to further improve the pose estimate.
- Empirically, we demonstrate that the proposed approach yields accurate 6-DoF pose estimates in scenes with cluttered backgrounds without the requirement of any pose initialization. State-of-the-art performance is shown on the large-scale PASCAL3D+ dataset [15].

III. TECHNICAL APPROACH

The proposed pipeline includes object detection, keypoint localization and pose optimization. As object detection has been a well studied problem, we assume that a bounding box around the object has been provided by an off-the-shelf object detector, e.g., YOLO [32], and focus on the keypoint localization and pose optimization.

A. Keypoint localization

The keypoint localization step employs the “stacked hourglass” network architecture [24] that has been shown to be particularly effective for 2D human pose estimation. Motivated by this success, we use the same network design and train the network for object keypoint localization.

Network architecture A high level overview of the main network components is presented in Fig. 2. The network takes as input an RGB image, and outputs a set of heatmaps with the intensity at each location of the heatmap indicating the prediction uncertainty of the respective keypoint. The network consists of two hourglass components, where each component can be further subdivided into two main processing stages. In the first stage, a series of convolutional and max-pooling layers are applied to the input. After each max-pooling layer, the resolution of the feature maps decrease by a factor of two, allowing the next convolutional layer to process the features at a coarser scale. This sequence of processing continues until reaching the lowest resolution (4×4 feature maps), which is illustrated by the smallest layer in the middle of each module in Fig. 2. Following these downsampling layers, the processing continues with a series of convolutional upsampling layers. Each upsampling layer increases the resolution by a factor of two, until the feature maps reach the initial input resolution at the end of the module. A second hourglass component is stacked at the end of the first one to refine the output heatmaps. Following best practices [33], the groundtruth labels used to supervise the training are synthesized heatmaps with a 2D Gaussian with $\sigma = 1$ pixel, centered at each keypoint. The ℓ_2 loss is minimized during training. Optionally, intermediate supervision can be applied at the end of the first module, which provides a richer gradient signal to the network and guides the learning procedure towards a better optimum [34]. The heatmap responses of the last module are considered as the final output of the network and the peak in each heatmap indicates the most likely location for the corresponding keypoint.

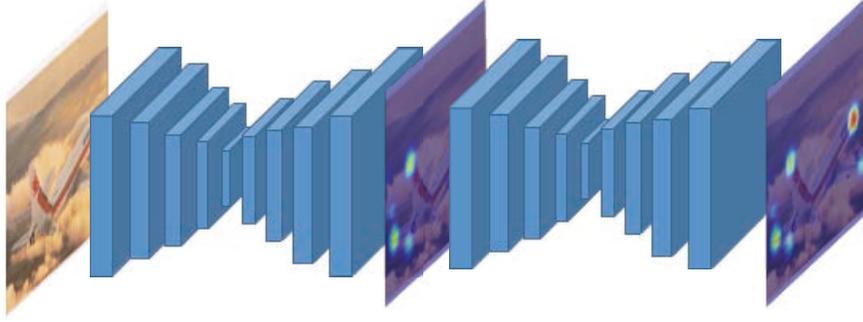


Fig. 2: Overview of the stacked hourglass architecture. Here, two hourglass modules are stacked together. The symmetric nature of the design allows for bottom-up processing in the first half of the module (feature map resolution decrease), and top-down processing in the second half (resolution increase). Intermediate supervision is applied after the first module. The heatmap responses of the second module represent the final output of the network that is used for keypoint localization.

Design benefits The most critical design element of the hourglass network is the symmetric combination of bottom-up and top-down processing that each hourglass module performs. Given the large appearance changes of objects due to in-class and viewpoint variation, both local and global cues are needed to effectively decide the locations of the keypoints in the image. The consolidation of features across different scales in the hourglass architecture allows the network to successfully integrate both local and global appearance information, and commit to a keypoint location only after this information has been made available to the network. Moreover, the stacking of the individual hourglass modules provides a form of iterative processing that has been shown to be effective with several other recent network designs [35], [25] and offers additional refinement of the network estimates. Additionally, the application of intermediate supervision at the end of each module has been validated as an effective training strategy, particularly ameliorating the practical issue of vanishing gradients when training a deep neural network [34]. Finally, the convolutional layers are implemented using residual modules which have achieved state-of-the-art results for many visual tasks, including object classification [36], instance segmentation [37], and 2D human pose estimation [24].

B. Pose optimization

Given the keypoint locations on the 3D model as well as their correspondences in the 2D image, one naive approach is to simply apply an existing PnP algorithm to solve for the 6-DoF pose. This approach is problematic because the keypoint predictions by the convnet can be rendered imprecise due to occlusions and false detections in the background. Moreover, the exact 3D model of the object instance in the testing image is often unavailable. To address these difficulties, we propose to fit a deformable shape model to the 2D detections while considering the uncertainty in keypoint predictions.

A deformable shape model is built for each object category using 3D CAD models with annotated keypoints. More

specifically, the p keypoint locations on a 3D object model are denoted by $\mathbf{S} \in \mathbb{R}^{3 \times p}$ and

$$\mathbf{S} = \mathbf{B}_0 + \sum_{i=1}^k c_i \mathbf{B}_i, \quad (1)$$

where \mathbf{B}_0 is the mean shape of the given 3D model and $\mathbf{B}_1, \dots, \mathbf{B}_k$ are several modes of possible shape variability computed by Principal Component Analysis (PCA).

Given detected keypoints in an image, which are denoted by $\mathbf{W} \in \mathbb{R}^{2 \times p}$, the goal is to estimate the rotation $\mathbf{R} \in \mathbb{R}^{3 \times 3}$ and translation $\mathbf{T} \in \mathbb{R}^{3 \times 1}$ between the object and camera frames as well as the coefficients of the shape deformation $\mathbf{c} = [c_1, \dots, c_k]^\top$.

The inference is formulated as the following optimization problem:

$$\min_{\theta} \frac{1}{2} \left\| \xi(\theta) \mathbf{D}^{\frac{1}{2}} \right\|_F^2 + \frac{\lambda}{2} \|\mathbf{c}\|_2^2, \quad (2)$$

where θ is the set of unknowns, $\xi(\theta)$ denotes the fitting residuals dependent on θ , and the Tikhonov regularizer $\|\mathbf{c}\|_2^2$ is introduced to penalize large deviations from the mean shape.

To incorporate the uncertainty in 2D keypoint predictions, a diagonal weighting matrix $\mathbf{D} \in \mathbb{R}^{p \times p}$ is introduced:

$$\mathbf{D} = \begin{bmatrix} d_1 & 0 & \cdots & 0 \\ 0 & d_2 & \cdots & 0 \\ \vdots & \vdots & \ddots & \vdots \\ 0 & 0 & \cdots & d_p \end{bmatrix}, \quad (3)$$

where d_i indicates the localization confidence of the i th keypoint in the image. In our implementation, d_i is assigned the peak value in the heatmap corresponding to the i th keypoint. As shown previously [24], the peak intensity of the heatmap provides a good indicator for the visibility of a keypoint in the image.

The fitting residuals, $\xi(\theta)$, measure the differences between the given 2D keypoints, provided by the previous

processing stage, and the projections of 3D keypoints. Two camera models are next considered.

1) *Weak perspective model*: If the camera intrinsic parameters are unknown, the weak perspective camera model is adopted, which is usually a good approximation to the full perspective case when the camera is relatively far away from the object. In this case, the reprojection error is written as

$$\xi(\theta) = \mathbf{W} - s\bar{\mathbf{R}} \left(\mathbf{B}_0 + \sum_{i=1}^k c_i \mathbf{B}_i \right) - \bar{\mathbf{T}} \mathbf{1}^\top, \quad (4)$$

where s is a scalar, $\bar{\mathbf{R}} \in \mathbb{R}^{2 \times 3}$ and $\bar{\mathbf{T}} \in \mathbb{R}^2$ denote the first two rows of \mathbf{R} and \mathbf{T} , respectively, and $\theta = \{s, \mathbf{c}, \bar{\mathbf{R}}, \bar{\mathbf{T}}\}$.

The problem in (2) is continuous and in principal can be locally solved by any gradient-based method. We solve it with a block coordinate descent scheme because of its fast convergence and the simplicity in implementation. We alternately update each of the variables while fixing the others. The updates of s , \mathbf{c} and $\bar{\mathbf{T}}$ are simply solved using closed-form least squares solutions. The update of $\bar{\mathbf{R}}$ should consider the $SO(3)$ constraint. Here, the Manopt toolbox [38] is used to optimize $\bar{\mathbf{R}}$ over the Stiefel manifold. As the problem in (2) is non-convex, we further adopt a convex relaxation approach [39] to initialize the optimization. More specifically, we only estimate the pose parameters while fixing the 3D model as the mean shape in the initialization stage. By setting $\mathbf{c} = \mathbf{0}$ and replacing the orthogonality constraint on $\bar{\mathbf{R}}$ by the spectral norm regularizer, the problem in (2) can be converted to a convex program and solved with global optimality [39].

2) *Full perspective model*: If the camera intrinsic parameters are known, the full perspective camera model is used, and the residuals are defined as

$$\xi(\theta) = \tilde{\mathbf{W}} \mathbf{Z} - \mathbf{R} \left(\mathbf{B}_0 + \sum_{i=1}^k c_i \mathbf{B}_i \right) - \mathbf{T} \mathbf{1}^\top, \quad (5)$$

where $\tilde{\mathbf{W}} \in \mathbb{R}^{3 \times p}$ represents the normalized coordinates of the 2D keypoints and \mathbf{Z} is a diagonal matrix:

$$\mathbf{Z} = \begin{bmatrix} z_1 & 0 & \cdots & 0 \\ 0 & z_2 & \cdots & 0 \\ \vdots & \vdots & \ddots & \vdots \\ 0 & 0 & \cdots & z_p \end{bmatrix}, \quad (6)$$

where z_i is the depth for the i th keypoint in 3D. Intuitively, the distances from the 3D points to the rays crossing the corresponding 2D points are minimized. In this case, the unknown parameter set θ is given by $\{\mathbf{Z}, \mathbf{c}, \mathbf{R}, \mathbf{T}\}$.

The optimization here is similar to the alternating scheme in the weak perspective case. The update of \mathbf{Z} also admits a closed-form solution and the update of \mathbf{R} can be analytically solved by the orthogonal Procrustes analysis. To avoid local minima, the optimization is initialized by the weak perspective solution.

IV. EXPERIMENTS

A. Instance-based pose recovery: gas canister

This section considers the recovery of pose for a specific object instance. This case fits well with many robotics applications where the objects in the environment are known. Moreover, it allows us to establish the accuracy of our approach in a relatively simple setting before dealing with the more challenging object class scenario.

We collected a dataset of roughly 175 RGB-D images of a textureless gas canister. The depth data was only used to generate the groundtruth. More specifically, a complete 3D model of the gas canister was reconstructed using KinectFusion [40] and the groundtruth object pose for each image was calculated by ICP with careful manual initialization. Then, 10 keypoints were manually defined on the 3D model and projected to the images yielding groundtruth keypoint locations in 2D for training the convnet.

A random 85%/15% split was used for training/test data. A stacked hourglass network with two hourglass modules was trained. The output heatmaps for the testing images are visualized in the second column of Fig. 3. As can be seen, the hourglass network is able to locate the keypoints reliably in the presence of viewpoint variety and occlusions. The non-visible keypoints are also well localized thanks to the network's ability to take global context into account. The estimated object poses are shown in the last two columns of Fig. 3. The projected 3D models align accurately with the image; the full-perspective solution is more precise than the weak-perspective one. It is worth noting that only 150 images were used to train the network from scratch. Overfitting might be an issue with such a small training set, but the empirical results suggest that the hourglass model captures the object appearance very well in this single instance case. More challenging examples with large intra-class variability are considered in Section IV-B.

The 6-DoF pose was estimated with the known 3D model and camera intrinsic parameters using the optimization in Section III-B. The following geodesic distance was used to measure the rotation error between a pose estimate, R_1 , and the groundtruth, R_2 :

$$\Delta(R_1, R_2) = \frac{\|\log(R_1^T R_2)\|_F}{\sqrt{2}}. \quad (7)$$

As a simple baseline, the following greedy approach was implemented: the maximum response locations in the heatmaps were selected as 2D keypoint locations and the standard PnP problem was solved by the EPnP algorithm [4] to estimate the object pose. The results are presented in Table I. While the weak-perspective solutions (Proposed WP) are on average worse than EPnP due to the inaccurate camera model, the full-perspective solutions (Proposed FP) are much more precise than those of EPnP. The remarkably small pose errors returned by the proposed approach based on a single RGB image are in the range suitable for a general grasping system.

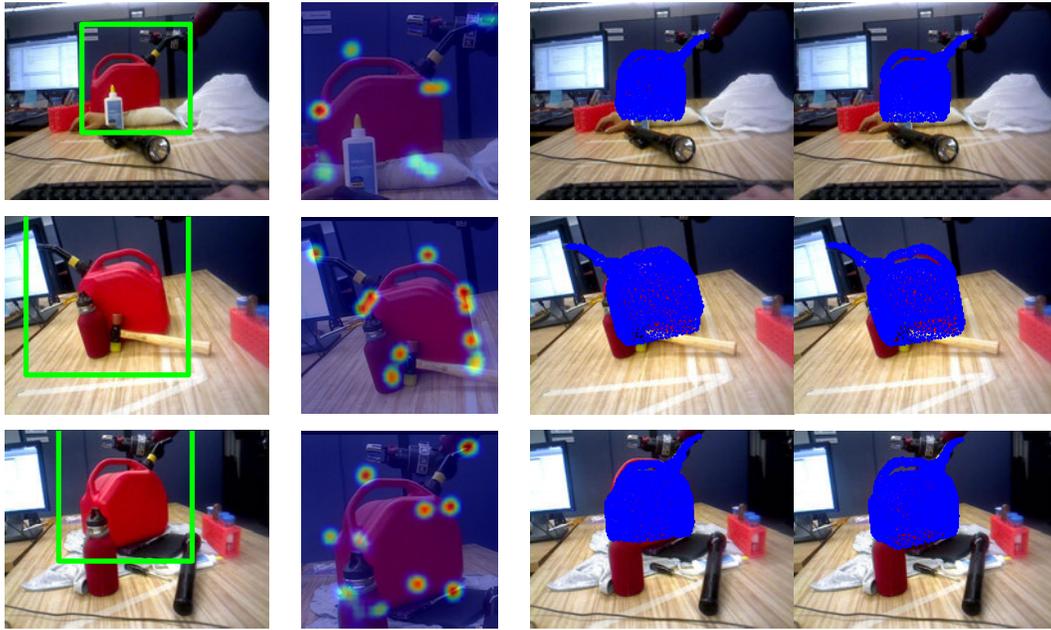


Fig. 3: Qualitative results on the gas canister dataset. From left to right: RGB images with bounding boxes, heatmaps from the convnet, projections of the 3D model with estimated poses using the weak-perspective model and full-perspective model, respectively. Note the better alignment near the handle with the full-perspective model.

TABLE I: Pose estimation errors on the gas canister dataset.

Approach	Rotation (degree)		Translation (mm)	
	Mean	Median	Mean	Median
Proposed WP	7.99	7.61	N/A	N/A
Proposed FP	3.57	3.11	12.05	8.82
EPnP [4]	7.17	5.21	43.45	21.51

B. Class-based pose recovery: PASCAL3D+

Moving to a more challenging scenario, we demonstrate the full strength of our approach using the large-scale PASCAL3D+ dataset [15]. The stacked hourglass network was trained from scratch with the training set of PASCAL3D+. Instead of training separate models for different object classes, a single network was trained to output heatmap predictions for all of the 124 keypoints from all classes. Using a single network for all keypoints allows us to share features across the available classes and significantly decreases the number of parameters needed for the network. At test time, given the class of the test object, the heatmaps corresponding to the keypoints belonging to this class were extracted. For pose optimization, two cases were tested: (i) the CAD model for the test image was known; and (ii) the CAD model was unknown and the pose was estimated with a deformable model whose basis was learned by PCA on all CAD models for each class in the dataset. Two principal components were used ($k = 2$) for each class, which was sufficient to explain greater than 95% of the shape variation. The 3D model was fit to the 2D keypoints with a weak-perspective model as the camera intrinsic parameters were not available.

Semantic correspondences A crucial component of our approach is the powerful learning procedure that is particularly successful at establishing correspondences across the semantically related keypoints of each class. To demonstrate this network property, in Fig. 4 we present a subset of the keypoints for each class along with the localizations of these keypoints in a randomly selected set of images among the ones with the top 50 responses. It is interesting to note that despite the large appearance differences due to extreme viewpoint and intra-class variability, the predictions are very consistent and preserve the semantic relation across various class instances.

Pose estimation The quantitative evaluation for pose estimation on PASCAL3D+ is presented in Table II. Only the errors for rotations are reported as the 3D translation cannot be determined in the weak perspective case and the ground truth is not available as well. The rotational error is calculated using the geodesic distance, (7). The proposed method shows improvement across most categories with respect to the state-of-the-art. The best results are achieved in the case where the fine subclass for the object is known and there exists an accurate CAD model correspondence. A subset of results are visualized in Fig. 5.

Failure cases In Table II we observed much higher errors than the state-of-the-art for two classes, namely boat and TV monitor. For most images of TV monitor, there are only four coplanar keypoints annotated. This makes the pose estimation an ill-posed problem for the weak perspective case. Figure 6 illustrates some failure cases because of this ambiguity. For boat we observed that in many cases the

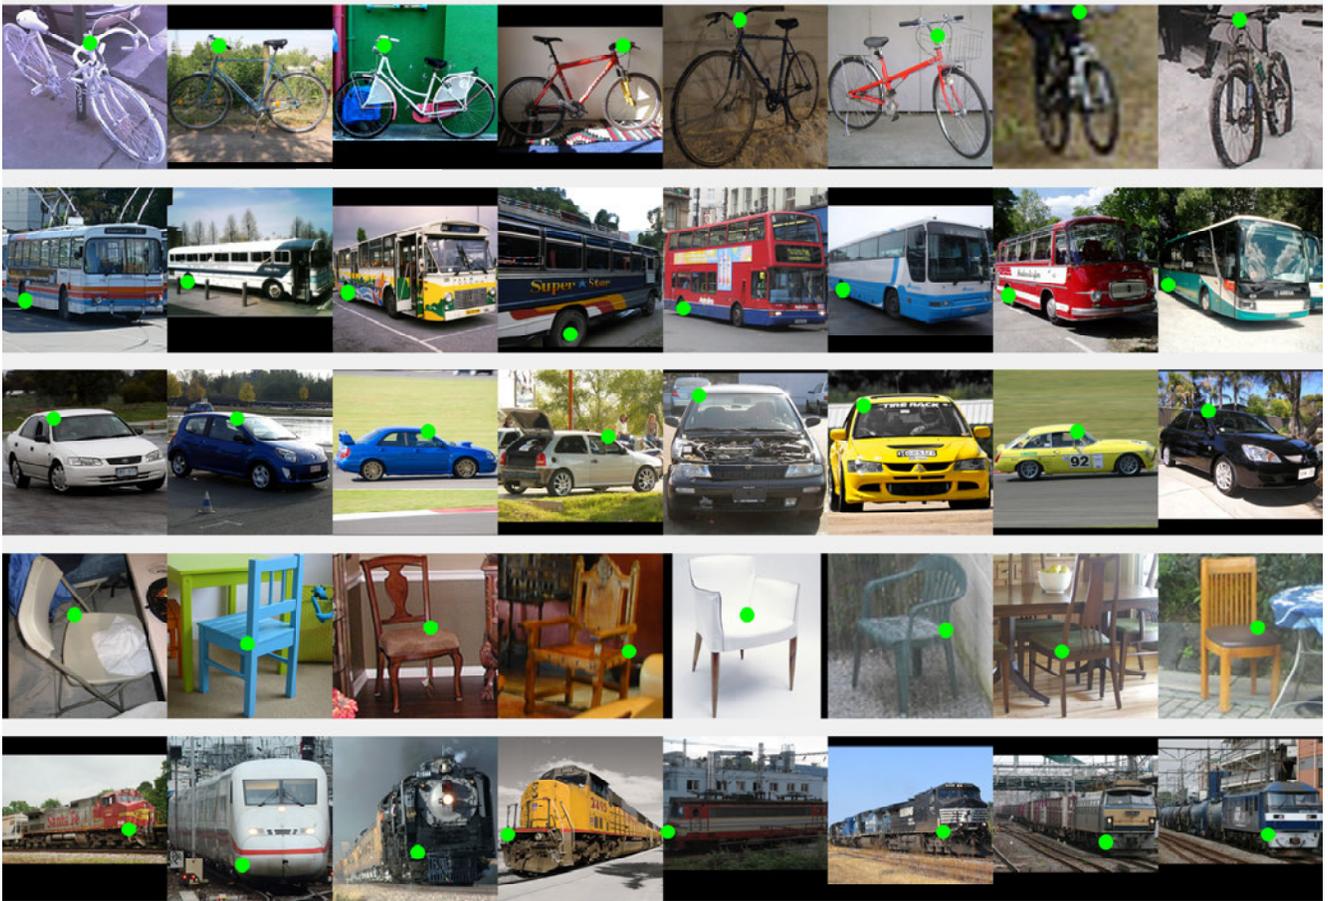


Fig. 4: Localization results for diverse keypoint categories. We visualize eight images selected randomly from the top 50 responses for each keypoint. The keypoint localization network is particularly successful at establishing semantic correspondences across the instances of a class, despite the significant intra-class variation and wide ranging camera viewpoints.

TABLE II: Viewpoint Estimation Median Error on PASCAL3D+ dataset.

Approach	aero	bike	bottle	bus	car	chair	sofa	train	TV monitor	boat
[20]	13.8	17.7	12.9	5.8	9.1	14.8	15.2	8.7	15.4	21.3
ours - PCA basis	11.4	14.8	12.1	4.4	7.5	15.5	20.0	7.0	36.1	44.9
ours - CAD basis	8.2	12.0	11.4	2.2	6.5	13.3	10.0	7.2	37.7	45.3

objects are very small and there are not sufficient cues to discriminate between the front and the back, which makes the keypoint localization extremely hard. In these extreme cases, holistic and discrete viewpoint prediction might be more robust, which could in practice provide a prior to regularize our continuous prediction. In Table II, we exclude the results for two classes from PASCAL3D+ dataset, namely dining table and motorbike, as we observed some inconsistency of left-right definition in groundtruth annotations, which lead to erroneous training data. Since other competing approaches like [20] rely on discrete viewpoint annotations only, this is not a reported issue.

C. Processing time

On a desktop with an Intel i7 3.4GHz CPU, 8G RAM and a GeForce GTX Titan X 6GB GPU, our pipeline needs around 0.2 seconds for the keypoint localization step and less than 0.1 seconds for the shape fitting step, for a total running time under 0.3 seconds. This makes our approach particularly suitable for applications where near real-time is desired. Moreover, further improvements in the running time are anticipated due to improvements in hardware, particularly with GPUs.

V. SUMMARY

In this paper, we proposed an efficient approach to estimate the 6-DoF pose of an object from a single RGB image. Capitalizing on the robust semantic keypoint predictions

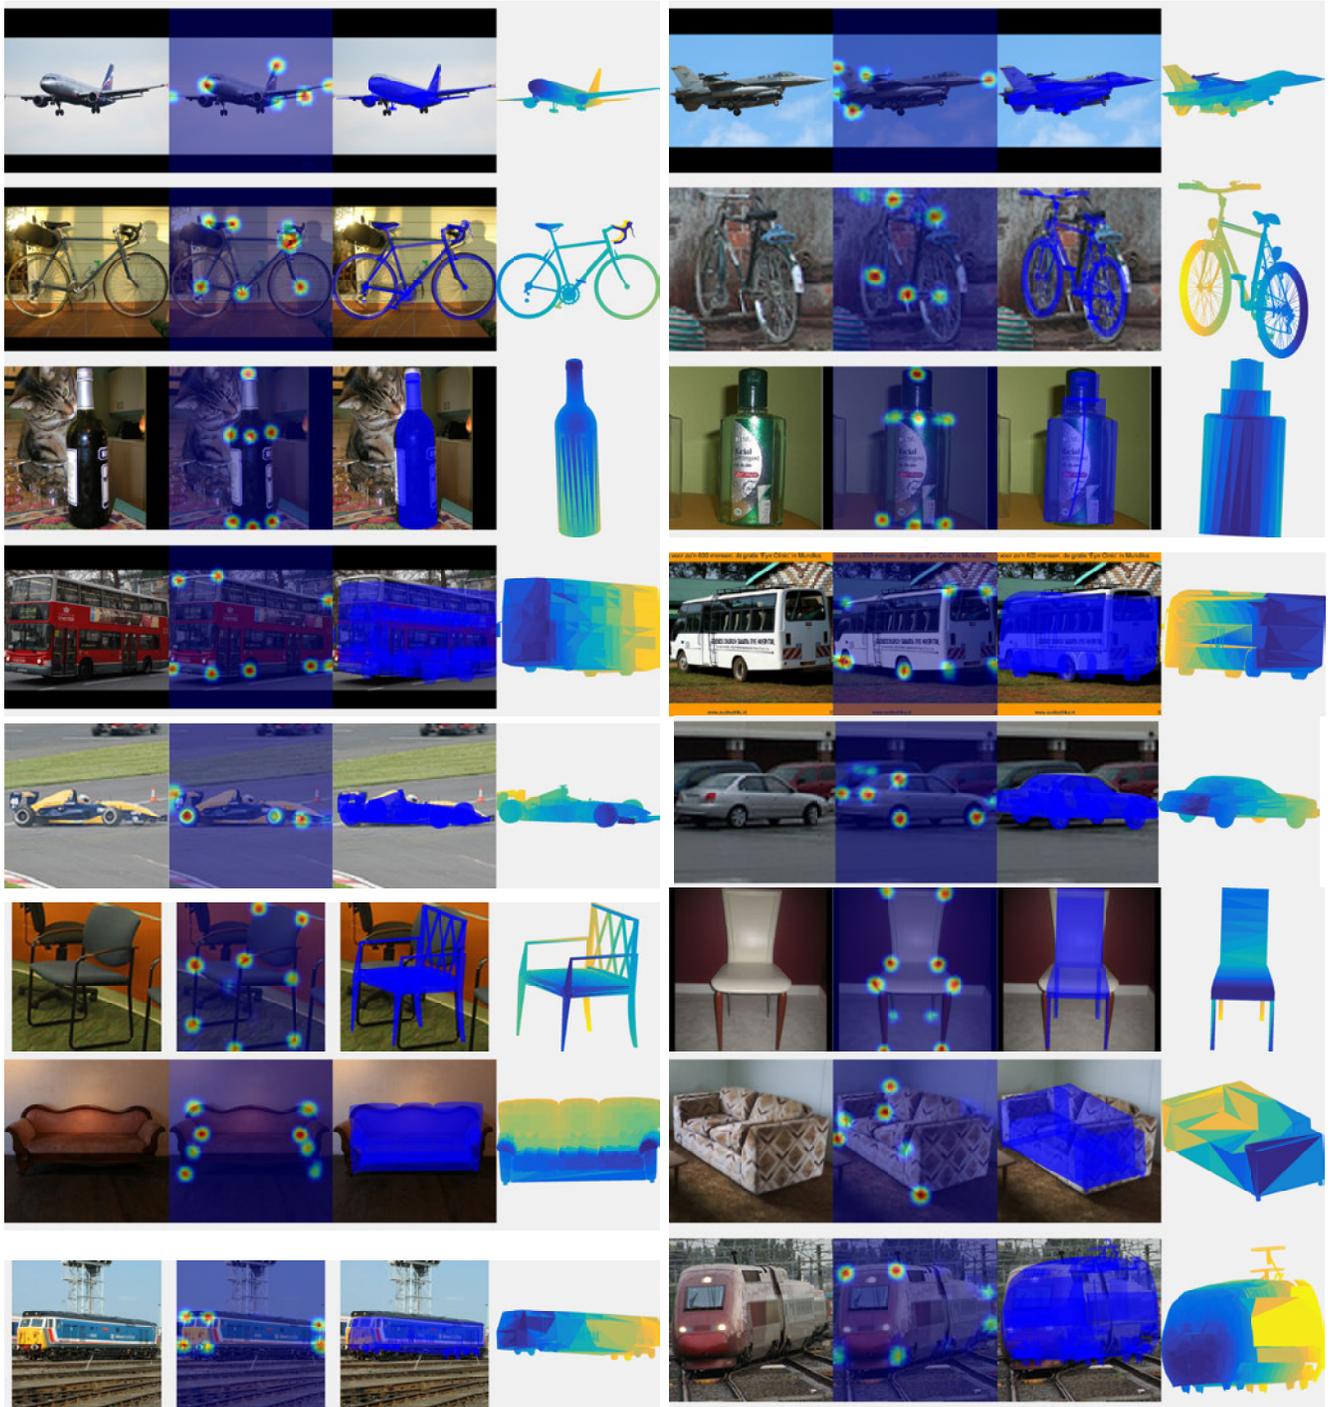


Fig. 5: Example results of our approach for PASCAL3D+ dataset. For each example from left to right: the RGB image of the object we estimate the pose for, heatmap responses for the keypoints of the specific class, the CAD model projected to 2D after pose estimation, and the CAD model visualized in 3D. The CAD model associated with each instance in the dataset is used for visualization.



Fig. 6: Failure cases for TV monitor class of PASCAL3D+. Although the visible keypoints are localized successfully the pose optimization fails because the 4 points are coplanar and the problem becomes ill-posed.

provided by a state-of-the-art convnet, we proposed a pose optimization scheme that fits a deformable shape model to the 2D keypoints and recovers the 6-DoF pose of the object. To ameliorate the effect of false detections, our pose optimization integrates the heatmap response values in the optimization scheme to model the certainty of each detection. Both the weak perspective and the full perspective case were investigated. The experimental validation included an instance-based scenario as well as full-scale evaluation on the PASCAL3D+ dataset, where we demonstrated state-of-the-art results for viewpoint estimation. Additionally, our method is accompanied by an efficient implementation with a running time under 0.3 seconds, making it a good fit for near real-time robotics applications.

REFERENCES

- [1] C. Gu and X. Ren, "Discriminative mixture-of-templates for viewpoint classification," in *ECCV*, 2010, pp. 408–421.
- [2] P. J. Besl and N. D. McKay, "A method for registration of 3-D shapes," *PAMI*, vol. 14, no. 2, pp. 239–256, 1992.
- [3] M. A. Fischler and R. C. Bolles, "Random sample consensus: A paradigm for model fitting with applications to image analysis and automated cartography," *Commun. ACM*, vol. 24, no. 6, pp. 381–395, 1981.
- [4] V. Lepetit, F. Moreno-Noguer, and P. Fua, "Eppm: An accurate $O(n)$ solution to the pmp problem," *IJCV*, vol. 81, no. 2, pp. 155–166, 2009.
- [5] A. Collet, D. Berenson, S. S. Srinivasa, and D. Ferguson, "Object recognition and full pose registration from a single image for robotic manipulation," in *ICRA*, 2009, pp. 48–55.
- [6] A. Collet, M. Martinez, and S. S. Srinivasa, "The MOPED framework: Object recognition and pose estimation for manipulation," *IJRR*, vol. 30, no. 10, pp. 1284–1306, 2011.
- [7] Z. Xie, A. Singh, J. Uang, K. S. Narayan, and P. Abbeel, "Multimodal blending for high-accuracy instance recognition," in *IROS*, 2013, pp. 2214–2221.
- [8] D. G. Lowe, "Distinctive image features from scale-invariant keypoints," *IJCV*, vol. 60, no. 2, pp. 91–110, 2004.
- [9] M. Muja, R. B. Rusu, G. R. Bradski, and D. G. Lowe, "REIN - A fast, robust, scalable recognition infrastructure," in *ICRA*, 2011, pp. 2939–2946.
- [10] S. Hinterstoisser, C. Cagniart, S. Ilic, P. F. Sturm, N. Navab, P. Fua, and V. Lepetit, "Gradient response maps for real-time detection of textureless objects," *PAMI*, vol. 34, no. 5, pp. 876–888, 2012.
- [11] R. Rios-Cabrera and T. Tuytelaars, "Discriminatively trained templates for 3D object detection: A real time scalable approach," in *ICCV*, 2013, pp. 2048–2055.
- [12] Z. Cao, Y. Sheikh, and N. Banerjee, "Real-time scalable 6DOF pose estimation for textureless objects," in *ICRA*, 2016.
- [13] S. Fidler, S. J. Dickinson, and R. Urtasun, "3D object detection and viewpoint estimation with a deformable 3D cuboid model," in *NIPS*, 2012, pp. 620–628.
- [14] B. Pepik, M. Stark, P. V. Gehler, and B. Schiele, "Teaching 3D geometry to deformable part models," in *CVPR*, 2012, pp. 3362–3369.
- [15] Y. Xiang, R. Mottaghi, and S. Savarese, "Beyond PASCAL: A benchmark for 3D object detection in the wild," in *WACV*, 2014, pp. 75–82.
- [16] M. Zhu, K. G. Derpanis, Y. Yang, S. Brahmabhatt, M. Zhang, C. Phillips, M. Lecce, and K. Daniilidis, "Single image 3D object detection and pose estimation for grasping," in *ICRA*, 2014, pp. 3936–3943.
- [17] Y. LeCun, B. E. Boser, J. S. Denker, D. Henderson, R. E. Howard, W. E. Hubbard, and L. D. Jackel, "Backpropagation applied to handwritten zip code recognition," *Neural Computation*, vol. 1, no. 4, pp. 541–551, 1989.
- [18] A. Krizhevsky, I. Sutskever, and G. E. Hinton, "ImageNet classification with deep convolutional neural networks," in *NIPS*, 2012, pp. 1106–1114.
- [19] F. Massa, M. Aubry, and R. Marlet, "Convolutional neural networks for joint object detection and pose estimation: A comparative study," *CoRR*, vol. abs/1412.7190, 2014.
- [20] S. Tulsiani and J. Malik, "Viewpoints and keypoints," in *CVPR*, 2015, pp. 1510–1519.
- [21] H. Su, C. R. Qi, Y. Li, and L. J. Guibas, "Render for CNN: viewpoint estimation in images using cnns trained with rendered 3D model views," in *ICCV*, 2015, pp. 2686–2694.
- [22] A. Toshev and C. Szegedy, "DeepPose: Human pose estimation via deep neural networks," in *CVPR*, 2014, pp. 1653–1660.
- [23] X. Zhou, M. Zhu, S. Leonardos, K. Derpanis, and K. Daniilidis, "Sparseness meets deepness: 3D human pose estimation from monocular video," in *CVPR*, 2015.
- [24] A. Newell, K. Yang, and J. Deng, "Stacked hourglass networks for human pose estimation," in *ECCV*, 2016.
- [25] S.-E. Wei, V. Ramakrishna, T. Kanade, and Y. Sheikh, "Convolutional pose machines," in *CVPR*, 2016.
- [26] J. Long, N. Zhang, and T. Darrell, "Do convnets learn correspondence?" in *NIPS*, 2014, pp. 1601–1609.
- [27] T. F. Cootes, C. J. Taylor, D. H. Cooper, and J. Graham, "Active shape models-Their training and application," *CVIU*, vol. 61, no. 1, pp. 38–59, 1995.
- [28] C. Cao, Y. Weng, S. Lin, and K. Zhou, "3D shape regression for real-time facial animation," *TOG*, vol. 32, no. 4, p. 41, 2013.
- [29] M. Z. Zia, M. Stark, B. Schiele, and K. Schindler, "Detailed 3D representations for object recognition and modeling," *PAMI*, vol. 35, no. 11, pp. 2608–2623, 2013.
- [30] V. Ramakrishna, T. Kanade, and Y. Sheikh, "Reconstructing 3D human pose from 2D image landmarks," in *ECCV*, 2012, pp. 573–586.
- [31] M. Zhu, X. Zhou, and K. Daniilidis, "Single image pop-up from discriminatively learned parts," in *ICCV*, 2015, pp. 927–935.
- [32] J. Redmon, S. K. Divvala, R. B. Girshick, and A. Farhadi, "You only look once: Unified, real-time object detection," in *CVPR*, 2016.
- [33] J. Tompson, R. Goroshin, A. Jain, Y. LeCun, and C. Bregler, "Efficient object localization using convolutional networks," in *CVPR*, 2015, pp. 648–656.
- [34] C.-Y. Lee, S. Xie, P. Gallagher, Z. Zhang, and Z. Tu, "Deeply-supervised nets," in *AISTATS*, vol. 2, no. 3, 2015, p. 6.
- [35] J. Carreira, P. Agrawal, K. Fragkiadaki, and J. Malik, "Human pose estimation with iterative error feedback," in *CVPR*, 2016.
- [36] K. He, X. Zhang, S. Ren, and J. Sun, "Deep residual learning for image recognition," in *CVPR*, 2016.
- [37] J. Dai, K. He, and J. Sun, "Instance-aware semantic segmentation via multi-task network cascades," in *CVPR*, 2016.
- [38] N. Boumal, B. Mishra, P.-A. Absil, and R. Sepulchre, "Manopt, a matlab toolbox for optimization on manifolds," *Journal of Machine Learning Research*, vol. 15, pp. 1455–1459, 2014.
- [39] X. Zhou, S. Leonardos, X. Hu, and K. Daniilidis, "3D shape estimation from 2D landmarks: A convex relaxation approach," in *CVPR*, 2015, pp. 4447–4455.
- [40] R. A. Newcombe, S. Izadi, O. Hilliges, D. Molyneaux, D. Kim, A. J. Davison, P. Kohli, J. Shotton, S. Hodges, and A. W. Fitzgibbon, "Kinectfusion: Real-time dense surface mapping and tracking," in *ISMAR*, 2011, pp. 127–136.