

HybridPose: 6D Object Pose Estimation under Hybrid Representations

Chen Song*, Jiaru Song*, Qixing Huang
The University of Texas at Austin

song@cs.utexas.edu, jiarus@cs.utexas.edu, huangqx@cs.utexas.edu

Abstract

We introduce *HybridPose*, a novel 6D object pose estimation approach. *HybridPose* utilizes a hybrid intermediate representation to express different geometric information in the input image, including keypoints, edge vectors, and symmetry correspondences. Compared to a unitary representation, our hybrid representation allows pose regression to exploit more and diverse features when one type of predicted representation is inaccurate (e.g., because of occlusion). Different intermediate representations used by *HybridPose* can all be predicted by the same simple neural network, and outliers in predicted intermediate representations are filtered by a robust regression module. Compared to state-of-the-art pose estimation approaches, *HybridPose* is comparable in running time and is significantly more accurate. For example, on *Occlusion Linemod* [3] dataset, our method achieves a prediction speed of 30 fps with a mean ADD(-S) accuracy of 79.2%, representing a 67.4% improvement from the current state-of-the-art approach. The implementation of *HybridPose* is available at <https://github.com/chensong1995/HybridPose>.

1. Introduction

Estimating the 6D pose of an object from an RGB image is a fundamental problem in 3D vision and has diverse applications in object recognition and robot-object interaction. Advances in deep learning have led to significant breakthroughs in this problem. While early works typically formulate pose estimation as end-to-end pose classification [39] or pose regression [16, 42], recent pose estimation methods usually leverage keypoints as an intermediate representation [38, 34], and align predicted 2D keypoints with ground-truth 3D keypoints. In addition to ground-truth pose labels, these methods incorporate keypoints as an intermediate supervision, facilitating smooth model training. Keypoint-based methods are built upon two assumptions: (1) a machine learning model can accurately predict 2D keypoint locations; and (2) these predictions provide suf-




Figure 1. HybridPose predicts keypoints, edge vectors, and symmetry correspondences. In (a), we show the input RGB image, in which the object of interest (driller) is partially occluded. In (b), red markers denote predicted 2D keypoints. In (c), edge vectors are defined by a fully-connected graph among all keypoints. In (d), symmetry correspondences connect each 2D pixel on the object to its symmetric counterpart. For illustrative purposes, we only draw symmetry correspondences of 50 random samples from 5755 detected object pixels in this example. The predicted pose (f) is obtained by jointly aligning all predictions with the 3D template, which involves solving a non-linear optimization problem.

ficient constraints to regress the underlying 6D pose. Both assumptions easily break in many real-world settings. Due to object occlusions and representational limitations of the prediction network, it is often impossible to accurately predict 2D keypoint coordinates from an RGB image alone.

In this paper, we introduce *HybridPose*, a novel 6D pose estimation approach that leverages multiple intermediate representations to express the geometric information in the input image. In addition to keypoints, *HybridPose* integrates a prediction network that outputs edge vectors between adjacent keypoints. As most objects possess a (partial) reflection symmetry, *HybridPose* also utilizes predicted dense pixel-wise correspondences that reflect the underlying symmetric relations between pixels. Compared to a unitary representation, this hybrid representation enjoys a multitude of advantages. First, *HybridPose* integrates more signals in the input image: edge vectors encode spa-

* Authors contributed equally

cial relations among object parts, and symmetry correspondences incorporate interior details. Second, HybridPose offers more constraints than using keypoints alone for pose regression, enabling accurate pose prediction even if a significant fraction of predicted elements are outliers (e.g., because of occlusion). Finally, it can be shown that symmetry correspondences stabilize the rotation component of pose prediction, especially along the normal direction of the reflection plane (details are provided in the supp. material).

Given the intermediate representation predicted by the first module, the second module of HybridPose performs pose regression. In particular, HybridPose employs trainable robust norms to prune outliers in predicted intermediate representation. We show how to combine pose initialization and pose refinement to maximize the quality of the resulting object pose. We also show how to train HybridPose effectively using a training set for the pose prediction module, and a validation set for the pose regression module.

We evaluate HybridPose on two popular benchmark datasets, Linemod [12] and Occlusion Linemod [3]. In terms of accuracy (under the ADD(-S) metric), HybridPose leads to considerable improvements from all state-of-the-art methods that merely utilize keypoints. On Occlusion Linemod [3], HybridPose achieves an accuracy of 79.2%, which represents a 67.4% improvement from DPOD [44], the current state-of-the-art method on this benchmark dataset.

Despite the gain in accuracy, our approach is efficient and runs at 30 frames per second on a commodity workstation. Compared to approaches which utilize sophisticated network architecture to predict one single intermediate representation (such as Pix2Pose [30]), HybridPose achieves considerably better performance by using a relative simple network to predict hybrid representations.

2. Related Works

Intermediate representation for pose. To express the geometric information in an RGB image, a prevalent intermediate representation is keypoints, which achieves state-of-the-art performance [34, 32, 36]. The corresponding pose estimation pipeline combines keypoint prediction and pose regression initialized by the PnP algorithm [18]. Keypoint predictions are usually generated by a neural network, and previous works use different types of tensor descriptors to express 2D keypoint coordinates. A common approach represents keypoints as peaks of heatmaps [28, 48], which becomes sub-optimal when keypoints are occluded, as the input image does not provide explicit visual cues for their locations. Alternative keypoint representations include vector-fields [34] and patches [14]. These representations allow better keypoint predictions under occlusion, and eventually lead to improvement in pose estimation accuracy. However, keypoints alone are a sparse representation of the

object pose, whose potential in improving estimation accuracy is limited.

Besides keypoints, another common intermediate representation is the coordinate of every image pixel in the 3D physical world, which provides dense 2D-3D correspondences for pose alignment, and is robust under occlusion [3, 4, 30, 20]. However, regressing dense object coordinates is much more costly than keypoint prediction. They are also less accurate than keypoints due to the lack of corresponding visual cues. In addition to keypoints and pixel-wise 2D-3D correspondences, depth is another alternative intermediate representation in visual odometry settings, which can be estimated together with pose in an unsupervised manner [47]. In practice, the accuracy of depth estimation is limited by the representational power of neural networks.

Unlike previous approaches, HybridPose combines multiple intermediate representations, and exhibits collaborative strength for pose estimation.

Multi-modal input. To address the challenges for pose estimation from a single RGB image, several works have considered inputs from multiple sensors. A popular approach is to leverage information from both RGB and depth images [47, 40, 42]. In the presence of depth information, pose regression can be reformulated as the 3D point alignment problem, which is then solved by the ICP algorithm [42]. Although HybridPose utilizes multiple intermediate representations, all intermediate representations are predicted from an RGB image alone. HybridPose handles situations in which depth information is absent.

Edge features. Edges are known to capture important image features such as object contours [2], salient edges [23], and straight line segments [45]. Unlike these low-level image features, HybridPose leverages semantic edge vectors defined between adjacent keypoints. This representation, which captures correlations between keypoints and reveals underlying structure of object, is concise and easy to predict. Such edge vectors offer more constraints than keypoints alone for pose regressions and have clear advantages under occlusion. Our approach is similar to [5], which predicts directions between adjacent keypoints to link keypoints into a human skeleton. However, we predict both the direction and the magnitude of edge vectors, and use these vectors to estimate object poses.

Symmetry detection from images. Symmetry detection has received significant attention in computer vision. We refer readers to [22, 27] for general surveys, and [1, 41] for recent advances. Traditional applications of symmetry detection include face recognition [31], depth estimation [21], and 3D reconstruction [13, 43]. In the context of object pose estimation, people have studied symmetry from the perspective that it introduces ambiguities for pose estimation (c.f. [25, 36, 42]), since symmetric objects with different poses can have the same appearance in image. Several

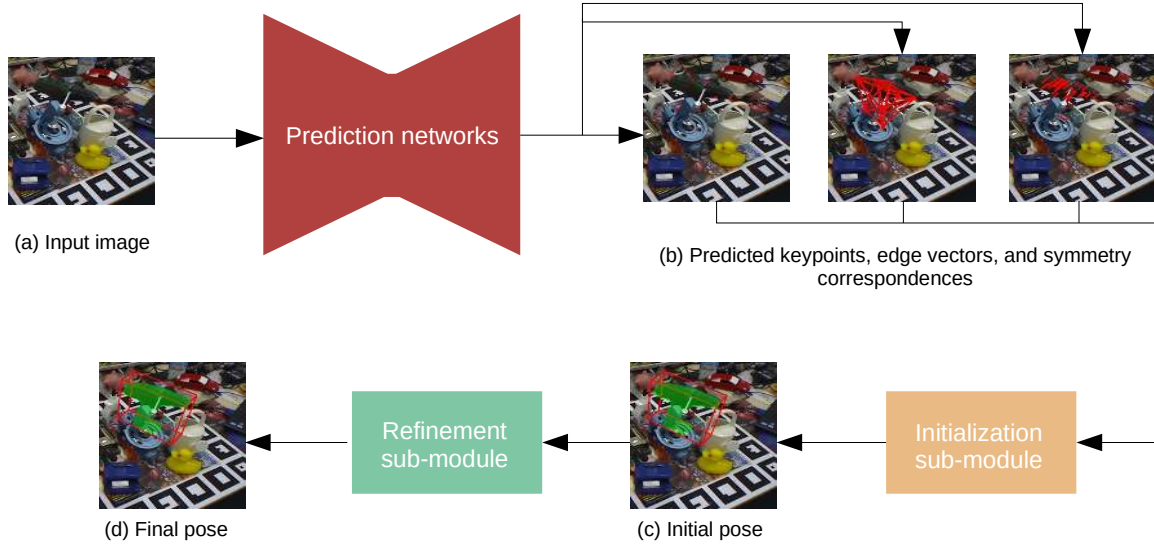


Figure 2. **Approach overview.** HybridPose consists of intermediate representation prediction networks and a pose regression module. The prediction networks take an image as input, and output predicted keypoints, edge vectors, and symmetry correspondences. The pose regression module consists of an initialization sub-module and a refinement sub-module. The initialization sub-module solves a linear system with predicted intermediate representations to obtain an initial pose. The refinement sub-module utilizes GM robust norm and optimizes (9) to obtain the final pose prediction.

works [36, 42, 6, 25, 30] have explored how to address such ambiguities, e.g., by designing loss functions that are invariant under symmetric transformations.

Robust regression. Pose estimation via intermediate representation is sensitive to outliers in predictions, which are introduced by occlusion and cluttered backgrounds [37, 32, 40]. To mitigate pose error, several works assign different weights to different predicted elements in the 2D-3D alignment stage [34, 32]. In contrast, our approach additionally leverages robust norms to automatically filter outliers in the predicted elements.

Besides the reweighting strategy, some recent works propose to use deep learning-based refiners to boost the pose estimation performance [19, 26, 44]. [44, 19] use point matching loss and achieve high accuracy. [26] predicts pose updates using contour information. Unlike these works, our approach considers the critical points and the loss surface of the robust objective function, and does not involve a fixed pre-determined iteration count used in recurrent network based approaches.

3. Approach

The input to HybridPose is an image I containing an object in a known class, taken by a pinhole camera with known intrinsic parameters. Assuming that the class of objects has a canonical coordinate system Σ (i.e. the 3D point cloud), HybridPose outputs the 6D camera pose $(R_I \in SO(3), t_I \in \mathbb{R}^3)$ of the image object under Σ , where R_I is the rotation and t_I is the translation component.

3.1. Approach Overview

As illustrated in Figure 2, HybridPose consists of a prediction module and a pose regression module.

Prediction module (Section 3.2). HybridPose utilizes three prediction networks f_θ^K , f_ϕ^E , and f_γ^S to estimate a set of keypoints $\mathcal{K} = \{p_k\}$, a set of edges between keypoints $\mathcal{E} = \{v_e\}$, and a set of symmetry correspondences between image pixels $\mathcal{S} = \{(q_{s,1}, q_{s,2})\}$. \mathcal{K} , \mathcal{E} , and \mathcal{S} are all expressed in 2D. θ , ϕ , and γ are trainable parameters.

The keypoint network f_θ^K employs an off-the-shelf prediction network [34]. The other two prediction networks, f_ϕ^E , and f_γ^S , are introduced to stabilize pose regression when keypoint predictions are inaccurate. Specifically, f_ϕ^E predicts edge vectors along a pre-defined graph of keypoints, which stabilizes pose regression when keypoints are cluttered in the input image. f_γ^S predicts symmetry correspondences that reflect the underlying (partial) reflection symmetry. A key advantage of this symmetry representation is that the number of symmetry correspondences is large: every image pixel on the object has a symmetry correspondence. As a result, even with a large outlier ratio, symmetry correspondences still provide sufficient constraints for estimating the plane of reflection symmetry for regularizing the underlying pose. Moreover, symmetry correspondences incorporate more features within the interior of the underlying object than keypoints and edge vectors.

Pose regression module (Section 3.3). The second module of HybridPose optimizes the object pose (R_I, t_I) to fit the output of the three prediction networks. This module combines a trainable initialization sub-module and a train-

able refinement sub-module. In particular, the initialization sub-module performs SVD to solve for an initial pose in the global affine pose space. The refinement sub-module utilizes robust norms to filter out outliers in the predicted elements for accurate object pose estimation.

Training HybridPose (Section 3.4). We train HybridPose by splitting the dataset into a training set and a validation set. We use the training set to learn the prediction module, and the validation set to learn the hyper-parameters of the pose regression module. We have tried training HybridPose end-to-end using one training set. However, the difference between the prediction distributions on the training set and testing set leads to sub-optimal generalization performance.

3.2. Hybrid Representation

This section describes three intermediate representations used in HybridPose.

Keypoints. The first intermediate representation consists of keypoints, which have been widely used for pose estimation. Given the input image I , we train a neural network $f_{\theta}^{\mathcal{K}}(I) \in \mathbb{R}^{2 \times |\mathcal{K}|}$ to predict 2D coordinates of a pre-defined set of $|\mathcal{K}|$ keypoints. In our experiments, HybridPose incorporates an off-the-shelf architecture called PVNet [34], which is the state-of-the-art keypoint-based pose estimator that employs a voting scheme to predict both visible and invisible keypoints.

Besides outliers in predicted keypoints, another limitation of keypoint-based techniques is that when the difference (direction and distance) between adjacent keypoints characterizes important information of the object pose, inexact keypoint predictions incur large pose error.

Edges. The second intermediate representation, which consists of edge vectors along a pre-defined graph, explicitly models the displacement between every pair of keypoints. As illustrated in Figure 2, HybridPose utilizes a simple network $f_{\phi}^{\mathcal{E}}(I) \in \mathbb{R}^{2 \times |\mathcal{E}|}$ to predict edge vectors in the 2D image plane, where $|\mathcal{E}|$ denotes the number of edges in the pre-defined graph. In our experiments, \mathcal{E} is a fully-connected graph, i.e., $|\mathcal{E}| = \frac{|\mathcal{K}| \cdot (|\mathcal{K}| - 1)}{2}$.

Symmetry correspondences. The third intermediate representation consists of predicted pixel-wise symmetry correspondences that reflect the underlying reflection symmetry. In our experiments, HybridPose extends the network architecture of FlowNet 2.0 [15] that combines a dense pixel-wise flow and the semantic mask predicted by PVNet. The resulting symmetry correspondences are given by predicted pixel-wise flow within the mask region. Compared to the first two representations, the number of symmetry correspondences is significantly larger, which provides rich constraints even for occluded objects. However, symmetry correspondences only constrain two degrees of freedom in the rotation component of the object pose (c.f. [24]). It is necessary to combine symmetry correspondences with other intermediate representations.

A 3D model may possess multiple reflection symmetry planes. For these models, we train HybridPose to predict symmetry correspondences with respect to the most salient reflection symmetry plane, i.e., one with the largest number of symmetry correspondences on the original 3D model.

Summary of network design. In our experiments, $f_{\theta}^{\mathcal{K}}(I)$, $f_{\phi}^{\mathcal{E}}(I)$, and $f_{\gamma}^{\mathcal{S}}$ are all based on ResNet [11], and the implementation details are discussed in Section 4.1. Trainable parameters are shared across all except the last convolutional layer. Therefore, the overhead of introducing the edge prediction network $f_{\phi}^{\mathcal{E}}(I)$ and the symmetry prediction network $f_{\gamma}^{\mathcal{S}}$ is insignificant.

3.3. Pose Regression

The second module of HybridPose takes predicted intermediate representations $\{\mathcal{K}, \mathcal{E}, \mathcal{S}\}$ as input and outputs a 6D object pose ($R_I \in SO(3), t_I \in \mathbb{R}^3$) for the input image I . Similar to state-of-the-art pose regression approaches [35], HybridPose combines an initialization sub-module and a refinement sub-module. Both sub-modules leverage all predicted elements. The refinement sub-module additionally leverages a robust function to model outliers in the predicted elements.

In the following, we denote 3D keypoint coordinates in the canonical coordinate system as $\bar{\mathbf{p}}_k, 1 \leq k \leq |\mathcal{K}|$. To make notations uncluttered, we denote output of the first module, i.e., predicted keypoints, edge vectors, and symmetry correspondences as $\mathbf{p}_k \in \mathbb{R}^2, 1 \leq k \leq |\mathcal{K}|$, $\mathbf{v}_e \in \mathbb{R}^2, 1 \leq e \leq |\mathcal{E}|$, and $(\mathbf{q}_{s,1} \in \mathbb{R}^2, \mathbf{q}_{s,2} \in \mathbb{R}^2), 1 \leq s \leq |\mathcal{S}|$, respectively. Our formulation also uses the homogeneous coordinates $\hat{\mathbf{p}}_k \in \mathbb{R}^3, \hat{\mathbf{v}}_e \in \mathbb{R}^3, \hat{\mathbf{q}}_{s,1} \in \mathbb{R}^3$ and $\hat{\mathbf{q}}_{s,2} \in \mathbb{R}^3$ of $\mathbf{p}_k, \mathbf{v}_e, \mathbf{q}_{s,1}$ and $\mathbf{q}_{s,2}$ respectively. The homogeneous coordinates are normalized by the camera intrinsic matrix.

Initialization sub-module. This sub-module leverages constraints between (R_I, t_I) and predicted elements and solves (R_i, t_i) in the affine space, which are then projected to $SE(3)$ in an alternating optimization manner. To this end, we introduce the following difference vectors for each type of predicted elements:

$$\bar{\mathbf{r}}_{R,t}^{\mathcal{K}}(\mathbf{p}_k) := \hat{\mathbf{p}}_k \times (R\bar{\mathbf{p}}_k + \mathbf{t}), \quad (1)$$

$$\bar{\mathbf{r}}_{R,t}^{\mathcal{E}}(\mathbf{v}_e, \mathbf{p}_{e_s}) := \hat{\mathbf{v}}_e \times (R\bar{\mathbf{p}}_{e_t} + \mathbf{t}) + \hat{\mathbf{p}}_{e_s} \times (R\bar{\mathbf{v}}_e) \quad (2)$$

$$\mathbf{r}_{R,t}^{\mathcal{S}}(\mathbf{q}_{s,1}, \mathbf{q}_{s,2}) := (\hat{\mathbf{q}}_{s,1} \times \hat{\mathbf{q}}_{s,2})^T R\bar{\mathbf{n}}_r. \quad (3)$$

where e_s and e_t are end vertices of edge e , $\bar{\mathbf{v}}_e = \bar{\mathbf{p}}_{e_t} - \bar{\mathbf{p}}_{e_s} \in \mathbb{R}^3$, and $\bar{\mathbf{n}}_r \in \mathbb{R}^3$ is the normal of the reflection symmetry plane in the canonical system.

HybridPose modifies the framework of EPnP [18] to generate the initial poses. By combining these three constraints from predicted elements, we generate a linear system of the form $A\mathbf{x} = \mathbf{0}$, where A is matrix and its dimension is $(3|\mathcal{K}| + 3|\mathcal{E}| + |\mathcal{S}|) \times 12$. $\mathbf{x} = [\mathbf{r}_1^T, \mathbf{r}_2^T, \mathbf{r}_3^T, \mathbf{t}^T]_{12 \times 1}^T$ is a vector that contains rotation and translation parameters in

affine space. To model the relative importance among keypoints, edge vectors, and symmetry correspondences, we rescale (2) and (3) by hyper-parameters α_E and α_S , respectively, to generate A .

Following EPnP [18], we compute \mathbf{x} as

$$\mathbf{x} = \sum_{i=1}^N \gamma_i \mathbf{v}_i \quad (4)$$

where \mathbf{v}_i is the i_{th} smallest right singular vector of A . Ideally, when predicted elements are noise-free, $N = 1$ with $\mathbf{x} = \mathbf{v}_1$ is an optimal solution. However, this strategy performs poorly given noisy predictions. Same as EPnP [18], we choose $N = 4$. To compute the optimal \mathbf{x} , we optimize latent variables γ_i and the rotation matrix R in an alternating optimization procedure with following objective function:

$$\min_{R \in \mathbb{R}^{3 \times 3}, \gamma_i} \left\| \sum_{i=1}^4 \gamma_i R_i - R \right\|_{\mathcal{F}}^2 \quad (5)$$

where $R_i \in \mathbb{R}^{3 \times 3}$ is reshaped from the first 9 elements of \mathbf{v}_i . After obtaining optimal γ_i , we project the resulting affine transformation $\sum_{i=1}^4 \gamma_i R_i$ into a rigid transformation. Due to space constraint, we defer details to the supp. material.

Refinement sub-module. Although (5) combines hybrid intermediate representations and admits good initialization, it does not directly model outliers in predicted elements. Another limitation comes from (1) and (2), which do not minimize the projection errors (i.e., with respect to keypoints and edges), which are known to be effective in keypoint-based pose estimation (c.f. [35]).

Benefited from having an initial object pose ($R^{init}, \mathbf{t}^{init}$), the refinement sub-module performs local optimization to refine the object pose. We introduce two difference vectors that involve projection errors: $\forall k, e, s$,

$$\mathbf{r}_{R,t}^{\mathcal{K}}(\mathbf{p}_k) := \mathcal{P}_{R,t}(\bar{\mathbf{p}}_k) - \mathbf{p}_k, \quad (6)$$

$$\mathbf{r}_{R,t}^{\mathcal{E}}(\mathbf{v}_e) := \mathcal{P}_{R,t}(\bar{\mathbf{p}}_{e_t}) - \mathcal{P}_{R,t}(\bar{\mathbf{p}}_{e_s}) - \mathbf{v}_e, \quad (7)$$

where $\mathcal{P}_{R,t} : \mathbb{R}^3 \rightarrow \mathbb{R}^2$ is the projection operator induced from the current pose (R, \mathbf{t}) .

To prune outliers in the predicted elements, we consider a generalized German-McClure (or GM) robust function

$$\rho(x, \beta) := \beta_1^2 / (\beta_2^2 + x^2). \quad (8)$$

With this setup, HybridPose solves the following non-linear

optimization problem for pose refinement:

$$\begin{aligned} \min_{R,t} & \sum_{k=1}^{|\mathcal{K}|} \rho(\|\mathbf{r}_{R,t}^{\mathcal{K}}(\mathbf{p}_k)\|, \beta_{\mathcal{K}}) \|\mathbf{r}_{R,t}^{\mathcal{K}}(\mathbf{p}_k)\|_{\Sigma_k}^2 \\ & + \frac{|\mathcal{K}|}{|\mathcal{E}|} \sum_{e=1}^{|\mathcal{E}|} \rho(\|\mathbf{r}_{R,t}^{\mathcal{E}}(\mathbf{v}_e)\|, \beta_{\mathcal{E}}) \|\mathbf{r}_{R,t}^{\mathcal{E}}(\mathbf{v}_e)\|_{\Sigma_e}^2 \\ & + \frac{|\mathcal{K}|}{|\mathcal{S}|} \sum_{s=1}^{|\mathcal{S}|} \rho(r_{R,t}^{\mathcal{S}}(\mathbf{q}_{s,1}, \mathbf{q}_{s,2}), \beta_{\mathcal{S}}) \end{aligned} \quad (9)$$

where $\beta_{\mathcal{K}}$, $\beta_{\mathcal{E}}$, and $\beta_{\mathcal{S}}$ are separate hyper-parameters for keypoints, edges, and symmetry correspondences. Σ_k and Σ_e denote the covariance information attached to the keypoint and edge predictions. $\|\mathbf{x}\|_A = (\mathbf{x}^T A \mathbf{x})^{\frac{1}{2}}$. When covariances of predictions are unavailable, we simply set $\Sigma_k = \Sigma_e = I_2$. The above optimization problem is solved by Gauss-Newton method starting from R^{init} and \mathbf{t}^{init} .

In the supp. material, we provide a stability analysis of (9), and show how the optimal solution of (9) changes with respect to noise in predicted representations. We also show collaborative strength among all three intermediate representations. While keypoints significantly contribute to the accuracy of \mathbf{t} , edge vectors and symmetry correspondences can stabilize the regression of R .

3.4. HybridPose Training

This section describes how to train the prediction networks and hyper-parameters of HybridPose using a labeled dataset $\mathcal{T} = \{I, (\mathcal{K}_I^{gt}, \mathcal{E}_I^{gt}, \mathcal{S}_I^{gt}, (R_I^{gt}, \mathbf{t}_I^{gt}))\}$. With $I, \mathcal{K}_I^{gt}, \mathcal{E}_I^{gt}, \mathcal{S}_I^{gt}$, and $(R_I^{gt}, \mathbf{t}_I^{gt})$, we denote the RGB image, labeled keypoints, edges, symmetry correspondences, and ground-truth object pose, respectively. A popular strategy is to train the entire model end-to-end, e.g., using recurrent networks to model the optimization procedure and introducing loss terms on the object pose output as well as the intermediate representations. However, we found this strategy sub-optimal. The distribution of predicted elements on the training set differs from that on the testing set. Even by carefully tuning the trade-off between supervisions on predicted elements and the final object pose, the pose regression model, which fits the training data, generalizes poorly on the testing data.

Our approach randomly divides the labeled set $\mathcal{T} = \mathcal{T}_{train} \cup \mathcal{T}_{val}$ into a training set and a validation set. \mathcal{T}_{train} is used to train the prediction networks, and \mathcal{T}_{val} trains the hyper-parameters of the pose regression model. Implementation and training details of the prediction networks are presented in Section 4.1. In the following, we focus on training the hyper-parameters using \mathcal{T}_{val} .

Initialization sub-module. Let R_I^{init} and \mathbf{t}_I^{init} be the output of the initialization sub-module. We obtain the optimal hyper-parameters α_E and α_S by solving the following opti-

mization problem:

$$\min_{\alpha_E, \alpha_S} \sum_{I \in \mathcal{T}_{\text{pose}}} (\|R_I^{\text{init}} - R_I^{\text{gt}}\|_{\mathcal{F}}^2 + \|\mathbf{t}_I^{\text{init}} - \mathbf{t}_I^{\text{gt}}\|^2). \quad (10)$$

Since the number of hyper-parameters is rather small, and the pose initialization step does not admit an explicit expression, we use the finite-difference method to compute numerical gradient, i.e., by fitting the gradient to samples of the hyper-parameters around the current solution. We then apply back-track line search for optimization.

Refinement sub-module. Let $\beta = \{\beta_{\mathcal{K}}, \beta_{\mathcal{E}}, \beta_S\}$ be the hyper-parameters of this sub-module. For each instance $(I, (\mathcal{K}_I^{\text{gt}}, \mathcal{E}_I^{\text{gt}}, \mathcal{S}_I^{\text{gt}}, (R_I^{\text{gt}}, \mathbf{t}_I^{\text{gt}}))) \in \mathcal{T}_{\text{val}}$, denote the objective function in (9) as $f_I(\mathbf{c}, \beta)$, where $\mathbf{c} = (\mathbf{c}^T, \bar{\mathbf{c}}^T)^T \in \mathbb{R}^6$ is a local parameterization of R_I and \mathbf{t}_I , i.e., $R_I = \exp(\mathbf{c} \times) R_I^{\text{gt}}$, $\mathbf{t}_I = \mathbf{t}_I^{\text{gt}} + \bar{\mathbf{c}}$. \mathbf{c} encodes the different the current estimated pose and the ground-truth pose in $SE(3)$.

The refinement module solves an unconstrained optimization problem, whose optimal solution is dictated by its critical points and the loss surface around the critical points. We consider two simple objectives. The first objective forces $\frac{\partial f_I}{\partial \mathbf{c}}(\mathbf{0}, \beta) \approx 0$, or in other words, the ground-truth is approximately a critical point. The second objective minimizes the condition number $\kappa(\frac{\partial^2 f_I}{\partial^2 \mathbf{c}}(\mathbf{0}, \beta)) = \lambda_{\max}(\frac{\partial^2 f_I}{\partial^2 \mathbf{c}}(\mathbf{0}, \beta)) / \lambda_{\min}(\frac{\partial^2 f_I}{\partial^2 \mathbf{c}}(\mathbf{0}, \beta))$. This objective regularizes the loss surface around each optimal solution, promoting a large converge radius for $f_I(\mathbf{c}, \beta)$. With this setup, we formulate the following objective function to optimize β :

$$\min_{\beta} \sum_{I \in \mathcal{T}_{\text{val}}} \left\| \frac{\partial f_I}{\partial \mathbf{c}}(\mathbf{0}, \beta) \right\|^2 + \gamma \kappa\left(\frac{\partial^2 f_I}{\partial^2 \mathbf{c}}(\mathbf{0}, \beta)\right) \quad (11)$$

where γ is set to 10^{-4} . The same strategy used in (10) is then applied to optimize (11).

4. Experimental Evaluation

This section presents an experimental evaluation of the proposed approach. Section 4.1 describes the experimental setup. Section 4.2 quantitatively and qualitatively compares HybridPose with other 6D pose estimation methods. Section 4.3 presents an ablation study to investigate the effectiveness of symmetry correspondences, edge vectors, and the refinement sub-module.

4.1. Experimental Setup

Datasets. We consider two popular benchmark datasets that are widely used in the 6D pose estimation problem, Linemod [12] and Occlusion Linemod [3]. In comparison to Linemod, Occlusion Linemod contains more examples where the objects are under occlusion. Our keypoint annotation strategy follows that of [34], i.e., we choose $|\mathcal{K}| = 8$

keypoints via the farthest point sampling algorithm. Edge vectors are defined as vectors connecting each pair of keypoints. In total, each object has $|\mathcal{E}| = \frac{|\mathcal{K}| \cdot (|\mathcal{K}| - 1)}{2} = 28$ edges. We further use the algorithm proposed in [8] to annotate Linemod and Occlusion Linemod with reflection symmetry labels. On each dataset, we randomly select 80% of the examples for training, 20 instances for validation, and the rest for testing.

Implementation details. We use ResNet [11] with pre-trained weights on ImageNet [7] to build the prediction networks $f_{\theta}^{\mathcal{K}}$, $f_{\phi}^{\mathcal{E}}$, and $f_{\gamma}^{\mathcal{S}}$. The prediction networks take an RGB image I of size $(3, H, W)$ as input, and output a tensor of size (C, H, W) , where (H, W) is the image resolution, and $C = 1 + 2|\mathcal{K}| + 2|\mathcal{E}| + 2$ is the number of channels in the output tensor.

The first channel in the output tensor is a binary segmentation mask M . If $M(x, y) = 1$, then (x, y) corresponds to a pixel on the object of interest in the input image I . The segmentation mask is trained using the cross-entropy loss.

The $2|\mathcal{K}|$ channels afterwards in the output tensor give x and y components of all $|\mathcal{K}|$ keypoints. A voting-based keypoint localization scheme [34] is applied to extract the coordinates of 2D keypoints from this $2|\mathcal{K}|$ -channel tensor and the segmentation mask M .

The next $2|\mathcal{E}|$ channels in the output tensor give the x and y components of all $|\mathcal{E}|$ edges, which we denote as $Edge$. Let i ($0 \leq i < |\mathcal{E}|$) be the index of an edge. Then

$$Edge_i = \{(Edge(2i, x, y), Edge(2i + 1, x, y)) | M(x, y) = 1\}$$

is a set of 2-tuples containing pixel-wise predictions of the i^{th} edge vector in $Edge$. The mean of $Edge_i$ is extracted as the predicted edge.

The final 2 channels in the output tensor define the x and y components of symmetry correspondences. We denote this 2-channel ‘‘map’’ of symmetry correspondences as Sym . Let (x, y) be a pixel on the object of interest in the input image, i.e. $M(x, y) = 1$. Assuming $\Delta x = Sym(0, x, y)$ and $\Delta y = Sym(1, x, y)$, we consider (x, y) and $(x + \Delta x, y + \Delta y)$ to be symmetric with respect to the reflection symmetry plane.

We train all three intermediate representations using the smooth ℓ_1 loss described in [9]. Network training employs the Adam [17] optimizer for 500 epochs. The learning rate is 0.02 before the 200th epoch, and 0.002 after the 200th epoch. Training weights of the segmentation mask, keypoints, edge vectors, and symmetry correspondences are 1.0, 10.0, 0.1, and 0.1, respectively.

The architecture described above achieves good performance in terms of detection accuracy. Nevertheless, it should be emphasized that the framework of HybridPose can incorporate future improvements in keypoint, edge vector, and symmetry correspondence detection techniques. Besides, Hybridpose can be extended to handling multiple objects within an image. One approach is to pre-




Figure 3. **Pose regression results.** HybridPose is able to accurately predict 6D poses from RGB images. HybridPose handles situations where the object has no occlusion (a, d, f, h), light occlusion (b, c), and severe occlusion (e, g). For illustrative purposes, we only draw 50 randomly selected symmetry correspondences in each example.

dict instance-level rather than semantic-level segmentation masks by methods such as Mask R-CNN [10]. Intermediate representations are then extracted from each instance, and fed to the pose regression module in 3.3.

Evaluation protocols. We use two metrics to evaluate the performance of HybridPose:

1. ADD(-S) [12, 42] first calculates the distance between two point sets transformed by predicted pose and ground-truth pose respectively, and then extracts the mean distance. When the object possesses symmetric pose ambiguity, the mean distance is computed from the closest points between two transformed sets. ADD(-S) accuracy is defined as the percentage of examples whose calculated mean distance is less than 10% of the model diameter.

2. In the ablation study, we compute and report the the angular rotation error $\left\| \frac{\log(R_{gt}^T R_I)}{2} \right\|$ and the relative translation error $\frac{\|t_I - t_{gt}\|}{d}$ between the predicted pose (R_I, t_I) and the ground-truth pose (R_{gt}, t_{gt}) , where d is object diameter.

4.2. Analysis of Results

As shown in Table 1, Table 2, and Figure 3, HybridPose leads to accurate pose estimation. On Linemod and Occlusion Linemod, HybridPose has an average ADD(-S) accuracy of 94.5 and 79.2, respectively. The result on Linemod outperforms all except one state-of-the-art approaches that

object	Tekin	BB8	Pix2Pose	PVNet	CDPN	DPOD	Ours
ape	21.6	40.4	58.1	43.6	64.4	87.7	77.6
benchvise	81.8	91.8	91.0	99.9	97.8	98.5	99.6
cam	36.6	55.7	60.9	86.9	91.7	96.1	95.9
can	68.8	64.1	84.4	95.5	95.9	99.7	93.6
cat	41.8	62.6	65.0	79.3	83.8	94.7	93.5
driller	63.5	74.4	76.3	96.4	96.2	98.8	97.2
duck	27.2	44.3	43.8	52.6	66.8	86.3	87.0
eggbox [†]	69.6	57.8	96.8	99.2	99.7	99.9	99.6
glue [†]	80.0	41.2	79.4	95.7	99.6	96.8	98.7
holepuncher	42.6	67.2	74.8	81.9	85.8	86.9	92.5
iron	75.0	84.7	83.4	98.9	97.9	100.0	98.1
lamp	71.1	76.5	82.0	99.3	97.9	96.8	96.9
phone	47.7	54.0	45.0	92.4	90.8	94.7	98.3
average	56.0	62.7	72.4	86.3	89.9	95.2	94.5

Table 1. **Quantitative evaluation: ADD(-S) accuracy on Linemod.** Baseline approaches: Tekin et al. [38], BB8 [36], Pix2Pose [30], PVNet [34], CDPN [20], and DPOD [44]. Objects annotated with (†) possess symmetric pose ambiguity.

regress poses from intermediate representations. The result on Occlusion-Linemod outperforms all state-of-the-art approaches.

Baseline comparison on Linemod. HybridPose outperforms PVNet [34], the backbone model we use to predict keypoints. The improvement is consistent across all except one object classes, which demonstrates clear advantage of using a hybrid as opposed to unitary intermediate representation. HybridPose shows competitive results against

object	PoseCNN	Oberweger	Hu	Pix2Pose	PVNet	DPOD	Ours
ape	9.6	12.1	17.6	22.0	15.8	-	53.3
can	45.2	39.9	53.9	44.7	63.3	-	86.5
cat	0.93	8.2	3.3	22.7	16.7	-	73.4
driller	41.4	45.2	62.4	44.7	65.7	-	92.8
duck	19.6	17.2	19.2	15.0	25.2	-	62.8
eggbox [†]	22.0	22.1	25.9	25.2	50.2	-	95.3
glue [†]	38.5	35.8	39.6	32.4	49.6	-	92.5
holepuncher	22.1	36.0	21.3	49.5	39.7	-	76.7
average	24.9	27.0	27.0	32.0	40.8	47.3	79.2

Table 2. **Quantitative evaluation: ADD(-S) accuracy on Occlusion Linemod.** Baseline approaches: PoseCNN [42], Oberweger et al. [29], Hu et al. [14], PVNet [34], and DPOD [44]. Objects annotated with (†) possess symmetric pose ambiguity.

DPOD [44], winning on six object classes. The advantage of DPOD comes from data augmentation and explicit modeling of dense correspondences between input and projected images, both of which cater to situations without object occlusion. A detailed analysis reveals that the classes of objects on which HybridPose exhibits sub-optimal performance are among the smallest objects in Linemod. It suggests that pixel-based descriptors used in our pipeline are limited by image resolution.

Baseline comparison on Occlusion Linemod. HybridPose outperforms all baselines by a considerable margin. In terms of ADD(-S), our approach improves PVNet [34] from 40.8 to 79.2, representing a 94.1% enhancement, which clearly shows the advantage of HybridPose on occluded objects, where predictions of invisible keypoints can be noisy, and visible keypoints may not provide sufficient constraints for pose regression alone. HybridPose also outperforms DPOD, the state-of-the-art model on this dataset, by 67.4%. One explanation is that rendering-based approaches like DPOD work less well on occluded objects, due to the difficulty in modeling occlusions in data augmentation and correspondence computation.

Running time. On a desktop with 16-core Intel(R) Xeon(R) E5-2637 CPU and GeForce GTX 1080 GPU, HybridPose takes 0.6 second to predict the intermediate representations, 0.4 second to regress the pose. Assuming a batch size of 30, this gives an average processing speed of around 30 fps, enabling real-time analysis.

4.3. Ablation Study

Table 3 summarizes the performance of HybridPose using different predicted intermediate representations on the Occlusion Linemod dataset. Ablation study on Linemod is deferred to the supp. material.

With keypoints. As a baseline approach, we estimate object poses by only utilizing keypoint information. As shown in Table 3, this gives a mean absolute rotation error of 1.648°, and a mean relative translation error of 0.100.

With keypoints and symmetry. Adding symmetry correspondences to keypoints leads to noticeable performance gains in rotation. The relative performance gain is 3.52%,

	keypoints		keypoints + symmetries		full model	
	Rot.	Trans.	Rot.	Trans.	Rot.	Trans.
ape	1.914°	0.107	1.809°	0.113	1.543°	0.092
can	1.472°	0.059	1.710°	0.073	0.912°	0.041
cat	1.039°	0.119	0.888°	0.117	0.751°	0.055
driller	1.180°	0.057	1.180°	0.057	0.803°	0.027
duck	1.773°	0.116	1.679°	0.115	1.439°	0.068
eggbox	1.675°	0.107	1.587°	0.105	1.052°	0.052
glue	1.796°	0.097	1.681°	0.099	1.224°	0.066
holepuncher	2.319°	0.141	2.192°	0.140	1.704°	0.051
mean	1.648°	0.100	1.590°	0.102	1.179°	0.057

Table 3. **Quantitative evaluation with different intermediate representations.** We report errors using two metrics: the median of absolute angular error in rotation, and the median of relative error in translation with respect to object diameter.

and such improvement is almost consistent across all object categories, which demonstrates the effectiveness of symmetry correspondences. On the other hand, the translation error of using keypoints and using keypoints + symmetry remains almost the same. One explanation is that symmetry correspondences only constrain two degrees of freedom in a total of three rotation parameters, and provide no constraint on translation parameters (see (3)).

Full model. Adding edge vectors to keypoints and symmetry correspondences leads to salient performance gain in both rotation and translation estimations. The relative performance gains in rotation and translation are 25.85% and 44.12%, respectively. One explanation is that edge vectors provide more constraints on both translation and rotation (see (2)). Edge vectors provide more constraints on translation than keypoints as they represent adjacent keypoints displacement and provide gradient information for regression. Unlike symmetry correspondences, edge vectors constrain 3 degrees of freedom on rotation parameters which further boosts the performance of rotation estimation.

5. Conclusions and Future Work

In this paper, we introduce HybridPose, a 6D pose estimation approach that utilizes keypoints, edge vectors, and symmetry correspondences. Experiments show that HybridPose enjoys real-time prediction and outperforms current state-of-the-art pose estimation approaches in accuracy. HybridPose is robust to occlusion and extreme poses. In the future, we plan to extend HybridPose to include more intermediate representations such as shape primitives, normals, and planar faces. Another possible direction is to enforce consistency across different representations in a similar way to [46] as a self-supervision loss in network training.

6. Acknowledgement

We would like to acknowledge the support of this research from NSF DMS-1700234, a Gift from Snap Research, and a hardware donation from NVIDIA.

References

- [1] Ibragim R. Atadjanov and Seungkyu Lee. Reflection symmetry detection via appearance of structure descriptor. In *Computer Vision - ECCV 2016 - 14th European Conference, Amsterdam, The Netherlands, October 11-14, 2016, Proceedings, Part III*, pages 3–18, 2016. [2](#)
- [2] Gedas Bertasius, Jianbo Shi, and Lorenzo Torresani. Deepedge: A multi-scale bifurcated deep network for top-down contour detection. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 4380–4389, 2015. [2](#)
- [3] Eric Brachmann, Alexander Krull, Frank Michel, Stefan Gumhold, Jamie Shotton, and Carsten Rother. Learning 6d object pose estimation using 3d object coordinates. In *European conference on computer vision*, pages 536–551. Springer, 2014. [1](#), [2](#), [6](#)
- [4] Eric Brachmann, Frank Michel, Alexander Krull, Michael Ying Yang, Stefan Gumhold, et al. Uncertainty-driven 6d pose estimation of objects and scenes from a single rgb image. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 3364–3372, 2016. [2](#)
- [5] Zhe Cao, Tomas Simon, Shih-En Wei, and Yaser Sheikh. Realtime multi-person 2d pose estimation using part affinity fields. In *2017 IEEE Conference on Computer Vision and Pattern Recognition, CVPR 2017, Honolulu, HI, USA, July 21-26, 2017*, pages 1302–1310, 2017. [2](#)
- [6] Enric Corona, Kaustav Kundu, and Sanja Fidler. Pose estimation for objects with rotational symmetry. In *2018 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, pages 7215–7222. IEEE, 2018. [3](#)
- [7] J. Deng, W. Dong, R. Socher, L. Li, Kai Li, and Li Fei-Fei. Imagenet: A large-scale hierarchical image database. In *2009 IEEE Conference on Computer Vision and Pattern Recognition*, pages 248–255, June 2009. [6](#)
- [8] Aleksandrs Ecins, Cornelia Fermüller, and Yiannis Aloimonos. Seeing behind the scene: Using symmetry to reason about objects in cluttered environments. In *2018 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, pages 7193–7200. IEEE, 2018. [6](#)
- [9] Ross Girshick. Fast r-cnn. *2015 IEEE International Conference on Computer Vision (ICCV)*, Dec 2015. [6](#)
- [10] Kaiming He, Georgia Gkioxari, Piotr Dollar, and Ross Girshick. Mask r-cnn. *2017 IEEE International Conference on Computer Vision (ICCV)*, Oct 2017. [7](#)
- [11] Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. Deep residual learning for image recognition. *arXiv preprint arXiv:1512.03385*, 2015. [4](#), [6](#)
- [12] Stefan Hinterstoisser, Vincent Lepetit, Slobodan Ilic, Stefan Holzer, Gary Bradski, Kurt Konolige, and Nassir Navab. Model based training, detection and pose estimation of texture-less 3d objects in heavily cluttered scenes. In *Proceedings of the 11th Asian Conference on Computer Vision - Volume Part I, ACCV'12*, pages 548–562, Berlin, Heidelberg, 2013. Springer-Verlag. [2](#), [6](#), [7](#)
- [13] Wei Hong, Allen Yang Yang, Kun Huang, and Yi Ma. On symmetry and multiple-view geometry: Structure, pose, and calibration from a single image. *International Journal of Computer Vision*, 60(3):241–265, 2004. [2](#)
- [14] Yinlin Hu, Joachim Hugonot, Pascal Fua, and Mathieu Salzmann. Segmentation-driven 6d object pose estimation. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 3385–3394, 2019. [2](#), [8](#)
- [15] Eddy Ilg, Nikolaus Mayer, Tommo Saikia, Margret Keuper, Alexey Dosovitskiy, and Thomas Brox. Flownet 2.0: Evolution of optical flow estimation with deep networks. In *2017 IEEE Conference on Computer Vision and Pattern Recognition, CVPR 2017, Honolulu, HI, USA, July 21-26, 2017*, pages 1647–1655, 2017. [4](#)
- [16] Alex Kendall, Matthew Grimes, and Roberto Cipolla. Posenet: A convolutional network for real-time 6-dof camera relocalization. In *Proceedings of the IEEE international conference on computer vision*, pages 2938–2946, 2015. [1](#)
- [17] Diederik P. Kingma and Jimmy Ba. Adam: A method for stochastic optimization, 2014. [6](#)
- [18] Vincent Lepetit, Francesc Moreno-Noguer, and Pascal Fua. Epnnp: An accurate o(n) solution to the pnp problem. *International journal of computer vision*, 81(2):155, 2009. [2](#), [4](#), [5](#)
- [19] Yi Li, Gu Wang, Xiangyang Ji, Yu Xiang, and Dieter Fox. Deepim: Deep iterative matching for 6d pose estimation. In *Proceedings of the European Conference on Computer Vision (ECCV)*, pages 683–698, 2018. [3](#)
- [20] Zhigang Li, Gu Wang, and Xiangyang Ji. Cdpn: Coordinates-based disentangled pose network for real-time rgb-based 6-dof object pose estimation. In *Proceedings of the IEEE International Conference on Computer Vision*, pages 7678–7687, 2019. [2](#), [7](#)
- [21] Guilin Liu, Chao Yang, Zimo Li, Duygu Ceylan, and Qixing Huang. Symmetry-aware depth estimation using deep neural networks. *arXiv preprint arXiv:1604.06079*, 2016. [2](#)
- [22] Yanxi Liu, Hagit Hel-Or, Craig S. Kaplan, and Luc Van Gool. Computational symmetry in computer vision and computer graphics. *Foundations and Trends in Computer Graphics and Vision*, 5(1-2):1–195, 2010. [2](#)
- [23] Yu Liu and Michael S Lew. Learning relaxed deep supervision for better edge detection. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 231–240, 2016. [2](#)
- [24] Yi Ma, Stefano Soatto, Jana Kosecka, and S. Shankar Sastry. *An Invitation to 3-D Vision: From Images to Geometric Models*. SpringerVerlag, 2003. [4](#)
- [25] Fabian Manhardt, Diego Martin Arroyo, Christian Rupprecht, Benjamin Busam, Tolga Birdal, Nassir Navab, and Federico Tombari. Explaining the ambiguity of object detection and 6d pose from visual data. In *Proceedings of the IEEE International Conference on Computer Vision*, pages 6841–6850, 2019. [2](#), [3](#)
- [26] Fabian Manhardt, Wadim Kehl, Nassir Navab, and Federico Tombari. Deep model-based 6d pose refinement in rgb. In *Proceedings of the European Conference on Computer Vision (ECCV)*, pages 800–815, 2018. [3](#)
- [27] Niloy J. Mitra, Mark Pauly, Michael Wand, and Duygu Ceylan. Symmetry in 3d geometry: Extraction and applications. *Comput. Graph. Forum*, 32(6):1–23, Sept. 2013. [2](#)
- [28] Alejandro Newell, Kaiyu Yang, and Jia Deng. Stacked hourglass networks for human pose estimation. In *Computer Vi-*

- sion - ECCV 2016 - 14th European Conference, Amsterdam, The Netherlands, October 11-14, 2016, Proceedings, Part VIII, pages 483–499, 2016. [2](#)
- [29] Markus Oberweger, Mahdi Rad, and Vincent Lepetit. Making deep heatmaps robust to partial occlusions for 3d object pose estimation. *Lecture Notes in Computer Science*, page 125141, 2018. [8](#)
- [30] Kiru Park, Timothy Patten, and Markus Vincze. Pix2pose: Pixel-wise coordinate regression of objects for 6d pose estimation. In *Proceedings of the IEEE International Conference on Computer Vision*, pages 7668–7677, 2019. [2](#), [3](#), [7](#)
- [31] Georgios Passalis, Panagiotis Perakis, Theoharis Theoharis, and Ioannis A Kakadiaris. Using facial symmetry to handle pose variations in real-world 3d face recognition. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 33(10):1938–1951, 2011. [2](#)
- [32] Georgios Pavlakos, Xiaowei Zhou, Aaron Chan, Konstantinos G Derpanis, and Kostas Daniilidis. 6-dof object pose from semantic keypoints. In *2017 IEEE International Conference on Robotics and Automation (ICRA)*, pages 2011–2018. IEEE, 2017. [2](#), [3](#)
- [33] F. Pedregosa, G. Varoquaux, A. Gramfort, V. Michel, B. Thirion, O. Grisel, M. Blondel, P. Prettenhofer, R. Weiss, V. Dubourg, J. Vanderplas, A. Passos, D. Cournapeau, M. Brucher, M. Perrot, and E. Duchesnay. Scikit-learn: Machine learning in Python. *Journal of Machine Learning Research*, 12:2825–2830, 2011.
- [34] Sida Peng, Yuan Liu, Qixing Huang, Hujun Bao, and Xiaowei Zhou. Pynet: Pixel-wise voting network for 6dof pose estimation. *CoRR*, abs/1812.11788, 2018. [1](#), [2](#), [3](#), [4](#), [6](#), [7](#), [8](#)
- [35] Mikael Persson and Klas Nordberg. Lambda twist: An accurate fast robust perspective three point (p3p) solver. In *The European Conference on Computer Vision (ECCV)*, September 2018. [4](#), [5](#)
- [36] Mahdi Rad and Vincent Lepetit. Bb8: A scalable, accurate, robust to partial occlusion method for predicting the 3d poses of challenging objects without using depth. In *Proceedings of the IEEE International Conference on Computer Vision*, pages 3828–3836, 2017. [2](#), [3](#), [7](#)
- [37] Martin Sundermeyer, Zoltan-Csaba Marton, Maximilian Durner, Manuel Brucker, and Rudolph Triebel. Implicit 3d orientation learning for 6d object detection from rgb images. In *Proceedings of the European Conference on Computer Vision (ECCV)*, pages 699–715, 2018. [3](#)
- [38] Bugra Tekin, Sudipta N. Sinha, and Pascal Fua. Real-time seamless single shot 6d object pose prediction. In *2018 IEEE Conference on Computer Vision and Pattern Recognition, CVPR 2018, Salt Lake City, UT, USA, June 18-22, 2018*, pages 292–301, 2018. [1](#), [7](#)
- [39] Shubham Tulsiani and Jitendra Malik. Viewpoints and keypoints. In *IEEE Conference on Computer Vision and Pattern Recognition, CVPR 2015, Boston, MA, USA, June 7-12, 2015*, pages 1510–1519, 2015. [1](#)
- [40] Chen Wang, Danfei Xu, Yuke Zhu, Roberto Martín-Martín, Cewu Lu, Li Fei-Fei, and Silvio Savarese. Densefusion: 6d object pose estimation by iterative dense fusion. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 3343–3352, 2019. [2](#), [3](#)
- [41] Zhaozhong Wang, Zesheng Tang, and Xiao Zhang. Reflection symmetry detection using locally affine invariant edge correspondence. *IEEE Trans. Image Processing*, 24(4):1297–1301, 2015. [2](#)
- [42] Yu Xiang, Tanner Schmidt, Venkatraman Narayanan, and Dieter Fox. Posecnn: A convolutional neural network for 6d object pose estimation in cluttered scenes. In *Robotics: Science and Systems XIV, Carnegie Mellon University, Pittsburgh, Pennsylvania, USA, June 26-30, 2018*, 2018. [1](#), [2](#), [3](#), [7](#), [8](#)
- [43] Tianfan Xue, Jianzhuang Liu, and Xiaoou Tang. Symmetric piecewise planar object reconstruction from a single image. In *CVPR 2011*, pages 2577–2584. IEEE, 2011. [2](#)
- [44] Sergey Zakharov, Ivan Shugurov, and Slobodan Ilic. Dpod: 6d pose object detector and refiner, 2019. [2](#), [3](#), [7](#), [8](#)
- [45] Ziheng Zhang, Zhengxin Li, Ning Bi, Jia Zheng, Jinlei Wang, Kun Huang, Weixin Luo, Yanyu Xu, and Shenghua Gao. Ppnet: Learning point-pair graph for line segment detection. *arXiv preprint arXiv:1905.03415*, 2019. [2](#)
- [46] Zaiwei Zhang, Zhenxiao Liang, Lemeng Wu, Xiaowei Zhou, and Qixing Huang. Path-invariant map networks. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 11084–11094, 2019. [8](#)
- [47] Tinghui Zhou, Matthew Brown, Noah Snavely, and David G Lowe. Unsupervised learning of depth and ego-motion from video. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 1851–1858, 2017. [2](#)
- [48] Xingyi Zhou, Arjun Karpur, Linjie Luo, and Qixing Huang. Starmap for category-agnostic keypoint and viewpoint estimation. In *Proceedings of the European Conference on Computer Vision (ECCV)*, pages 318–334, 2018. [2](#)