

Deep Fusion Transformer Network with Weighted Vector-Wise Keypoints Voting for Robust 6D Object Pose Estimation

Jun Zhou^{1,*}, Kai Chen^{2,*}, Linlin Xu³, Qi Dou², and Jing Qin¹

¹The Hong Kong Polytechnic University, ²The Chinese University of Hong Kong, ³University of Waterloo
juzhou@polyu.edu.hk, {kaichen, qdou}@cse.cuhk.edu.hk, l44xu@uwaterloo.ca, harry.qin@polyu.edu.hk

Abstract

One critical challenge in 6D object pose estimation from a single RGBD image is efficient integration of two different modalities, i.e., color and depth. In this work, we tackle this problem by a novel Deep Fusion Transformer (DFTr) block that can aggregate cross-modality features for improving pose estimation. Unlike existing fusion methods, the proposed DFTr can better model cross-modality semantic correlation by leveraging their semantic similarity, such that globally enhanced features from different modalities can be better integrated for improved information extraction. Moreover, to further improve robustness and efficiency, we introduce a novel weighted vector-wise voting algorithm that employs a non-iterative global optimization strategy for precise 3D keypoint localization while achieving near real-time inference. Extensive experiments show the effectiveness and strong generalization capability of our proposed 3D keypoint voting algorithm. Results on four widely used benchmarks also demonstrate that our method outperforms the state-of-the-art methods by large margins. Code is available at https://github.com/junzastar/DFTr_Voting.

1. Introduction

6D object pose estimation aims to recognize the 3D position and orientation of objects in the camera coordinate system. It is a widely studied task in both computer vision and robotics for its critical importance to many real-world applications, such as robotic grasping and manipulation [16, 37], augmented reality [1, 52] and autonomous navigation [14, 63]. Although significant progress has been made in recent years, critical challenges remain due to many factors such as varying illuminations, sensor noise, heavy occlusion, and the highly reflective surface of objects. Recently, along with the dramatic growth of RGB-D sensors, methods based on RGB-D data has attracted more attention. The * indicates equal contribution.




Figure 1. An illustration of typical challenges for object pose estimation, e.g. texture-less material, surface reflection, and inter-object occlusion.

tion, due to the fact that extra geometry information in the depth channel is complimentary to the color information for better alleviating difficulties in 6D pose estimation [24, 26, 59, 63].

However, so far, how to efficiently integrate these two modalities, i.e., color and depth, for better 6D pose estimation remains an open question. Prior works [9, 36, 61, 63] use the depth information as an additional clue to refine the final pose or concatenate the extracted RGB and geometric features directly for pose estimation. However, such methods do not make full use of the 3D geometry information and are sensitive to severe occlusions, and also ignore the global feature representation. Current methods are still limited, whether a simple point-wise fusion encoder [20, 58] or a k -Nearest Neighbor (k -NN) based feature query tactic [19], where inter-modality global semantic correlations are not considered. Such fusion strategies are more likely to be disturbed severely when the object has highly reflective surfaces also without any texture cues, e.g., metal objects, as shown in Fig. 1, because (1) the lack of one modality data will directly cause the failure of CNN or PCN (point cloud network) feature extraction, (2) the occlusion between objects results in data loss, and (3) the k -NN based feature clustering approach is sensitive to noise since the integrated features are likely absent on the query object.

Furthermore, given the fused representative RGB-D features, how to robustly estimate the object pose parameters

in challenging scenarios is another open problem. Works like [8, 42, 45, 61, 63] propose to regress the final pose parameters directly using the MLP-like prediction module. However, such methods often need a highly customized post-procedure for pose refinement. Conversely, the correspondence-based methods [3, 5, 7, 10] estimate pose by optimizing the pre-established correspondence, so they can achieve robust performance without the post-refinement procedure. Such methods can be subdivided into dense and sparse keypoint-based correspondence. Due to less computation cost and less hypothesis verification, sparse keypoint-based approaches have been widely used [19, 20]. However, most existing approaches, which output the point-wise translation offsets pointing to the keypoint directly without any scale constraint, is not conducive to network learning since the offset will change in scale due to the object’s size, thus severely degrading the keypoint localization accuracy. Besides, the keypoints voting method deployed in prior works, like MeanShift [15], is highly time-consuming for the iterative steps, which also limits performances in real-time applications.

In this work, we propose a *Deep Fusion Transformer* network for effective RGB-D fusion to estimate object 6D pose. The core of our network is to design a novel cross-modality fusion block named **Deep Fusion Transformer (DFTr)**. It implicitly aggregates distinguished features of two modality data by reasoning about the global semantic similarity between appearance and geometry information. Given two modality features from encoding or decoding layers of the network, our DFTr constructs a long-term dependence between them to extract cross-modality correlation for global semantic similarity modeling by using a transformer-based structure. We argue that the global semantic similarity modeling can alleviate perturbations in feature space caused by missing modality data and noises. Subsequently, with the learned fused RGB-D features, we adopt the keypoint-based workflow [19, 20] for 6D pose estimation, for their robustness to occlusion. Different from existing 3D keypoints voting methods, we propose to learn the 3D point-wise unit vector field and introduce an effective and non-iterative weighted vector-wise voting algorithm for 3D keypoints localization. In this way, the offsets with length constrained are easier for the network to learn and the inference speed is greatly improved while keeping comparable even superior location accuracy.

In summary, the main contributions of this work are:

- We propose an effective cross-modality feature aggregation network for 6DoF object pose estimation, in which a novel *Deep Fusion Transformer (DFTr)* block is designed and employed on a multi-scale level for robust representation learning.
- We propose an effective weighted vector-wise voting

algorithm, in which a global optimization scheme is deployed for 3D keypoint localization. We replace the original clustering method with the proposed algorithm in PVN3D [20] and FFB6D [19] framework, our approach is 1.7x faster than PVN3D and 2.7x faster than FFB6D when keeps a comparable even superior performance on the YCB dataset.

- We conduct extensive experiments on MP6D [11], YCB-Video [4], LineMOD [21], and Occlusion LineMOD [2] public benchmarks. Our method achieves dramatic performance improvements over other state-of-the-art methods without any post-refinement procedures.

2. Related Works

2.1. RGB-D Fusion-based Pose Estimation

Traditional approaches [33, 36, 61] employ the coarse-to-fine strategy, which computes the initial poses from RGB images and regards the depth map as compensatory cues used in subsequent pose-refinement procedures. Others [35, 36, 40] treat the depth information as an extra channel of RGB images or convert it into a bird-eye-view (BEV) image and are fed to a CNN-based network. However, these methods are time-consuming because of the expensive pose-processing step, and also the spatial geometric structure information is not fully explored. Instead, works like [57, 58, 63, 67] leverage two separate branch networks to extract appearance and geometric features and then deploy ‘later fusion’ tactics for pose computing. These methods are more effective but sensitive to modality data loss. Recently, FFB6D [19] propose a novel ‘early fusion’ module to enhance the communication between the two feature extraction branches but sensitive to local noises. To this end, we introduce a novel cross-modality fusion transformer block for deep implicit local-global RGB-D features aggregation.

2.2. Keypoints-based Pose Estimation

The classical way of this kind of method [23, 32, 43] establishes 2D-3D or 3D-3D correspondences in feature space and then recovers pose parameters by utilizing the PnP or Least-Squares Fitting algorithm. These methods require handcrafted feature descriptors which are designed on the surface of the object. However, they cannot handle textureless objects and are not robust to complex scenarios. Another category is regression-based methods [30, 50, 54], which directly regresses the coordinate of the keypoints by using neural networks. For better robustness to highly occluded scenes, furthermore, the pixel/point-wise voting methods [20, 47, 49, 62] are proposed to vote for the keypoints position. Keypoints-based methods can achieve satisfactory performance even in complex scenes, i.e. inter-class occlusion caused by object stacking and self-modality

data loss caused by the surface reflection. However, the accuracy and inference efficiency of keypoints localization is severely limited due to the current iteration and scale-free based keypoints voting scheme. In this paper, we present a stronger keypoints voting method by utilizing a global non-iterative optimization strategy with scale constraints.

3. Methodology

Given an RGB-D image of a test scene, the objective of this task is to estimate a transformation matrix between the target object coordinate system and the camera coordinate system. This transformation consists of an orientation component $R \in SO(3)$ and a translation component $t \in \mathbb{R}^3$. To this end, what we need to do is to design an effective deep model that can fully integrate cross-modality cues to satisfy the mapping ψ , which is formulated as follows:

$$[\psi(\Theta): I(I_o, D_o) \mapsto F_{rgb}] \rightsquigarrow O(R, t), \quad (1)$$

where Θ , $I(\cdot)$, F , and $O(\cdot)$ denote parameters of the deep model, the input set, the aggregated features, and the output set respectively.

3.1. Overview

As illustrated in Fig. 2, we propose the DFTr network for 6D object pose estimation. The network first extracts appearance features and geometric features from $I_o \in \mathbb{R}^{H \times W \times 3}$ and $P_o \in \mathbb{R}^{N \times 3}$ by using CNN and point cloud network respectively, where N and (H, W) denotes the number of scene points and the size of RGB image. And a DenseFuion [58] module is added for inter-modality pointwise local features aggregation. Before obtaining the final learned feature from these two branches, deep fusion transformer blocks are employed in each layer for cross-modality communication. It models the global semantic similarity between appearance and geometric features, and the highlighting features from the corresponding modality are integrated into its own branch to enhance their representation learning. Moreover, instead of the explicit utilization of RGBD correspondence for feature aggregation [19], we employ an implicit strategy by adding positional embedding on the whole feature elements sequence. With the extracted pointwise fused features, an instance segmentation module is utilized to obtain the mask of each object, and a weighted vector-wise voting module is adopted to localize the per-object 3D keypoints in the scenario with the predicted vector field. Finally, a 3D-3D correspondence-based algorithm is employed to recover the pose parameters.

3.2. Deep Cross-Modality Feature Embedding

Taking an RGB-D image as input, we first convert the depth image into a point cloud $P_o \in \mathbb{R}^{N \times 3}$ by using the camera intrinsic matrix. And then two branch networks are

applied to extract features from I_o and P_o respectively. For each layer, the learned color and geometric features would be fed into a DFTr block for cross-modality information aggregation. To this extent, multi-scale local and global information from these two modality data can be integrated into the feature extraction flow of the two branches even in the early stage of networks.

Deep fusion transformer block. The DFTr block aims to explore the inherent correspondence between these two modality features. Prior works, such as [58] and [20], first fuse RGB and geometric features in a pixel-wise concatenation manner, and generate a global feature vector which is then stacked to each concatenated feature to facilitate the global representative. Another fusion way [19] expands the searching scope of the neighborhood for a query pixel (point), and gathers these K nearest pixels (points) to generate a non-local feature by passing it to a max pooling layer. The generated features are then concatenated to their corresponding point (pixel) feature one by one. However, this kind of fusion strategy does not make full use of the global feature representation and is highly dependent on the quality of input data. In other words, the performance of the algorithm tends to be dramatically disturbed when uncertainty occurs in one of the input modality data, i.e., missing data due to the object surface reflection. Moreover, RGBD image misalignment caused by sensor calibration error is also one of the potential peril of this explicit fusion strategy. Instead, we design a novel bidirectional cross-modality fusion block, in which we treat the elements of RGB and geometric features as long sequence tokens and model them globally to implicitly integrate these features.

Specifically, we propose to implement the DFTr block with a transformer-based architecture, since it showed promising results in various vision tasks, and is proven to enjoy the ability to capture long-range dependencies in the data. We utilize this property to enhance the global representative of the model. As shown in Fig. 3, given the RGB feature map $F_I \in \mathbb{R}^{\hat{H} \times \hat{W} \times C}$ and geometric feature map $F_p \in \mathbb{R}^{\hat{N} \times C}$ of the i -th layer from the two feature extraction branches respectively, we first employ a dual-branch inter-modality interaction operation by cross-attention module, and then a transformer-based module is utilized to model the sequence of the gross elements of these two modality data. In detail, inspired by [6], they introduce an effective image patch-wise cross-information aggregation strategy for the classification task. We deploy their insight to cross-modality feature integration. We first flatten the feature map F_I and F_p to form a sequence and then use max pooling to generate global features $F_I^{max} \in \mathbb{R}^{1 \times C}$ and $F_p^{max} \in \mathbb{R}^{1 \times C}$. Then following the terminology in [56], we compute the *query* Q , *key* K , and *value* V matrices for

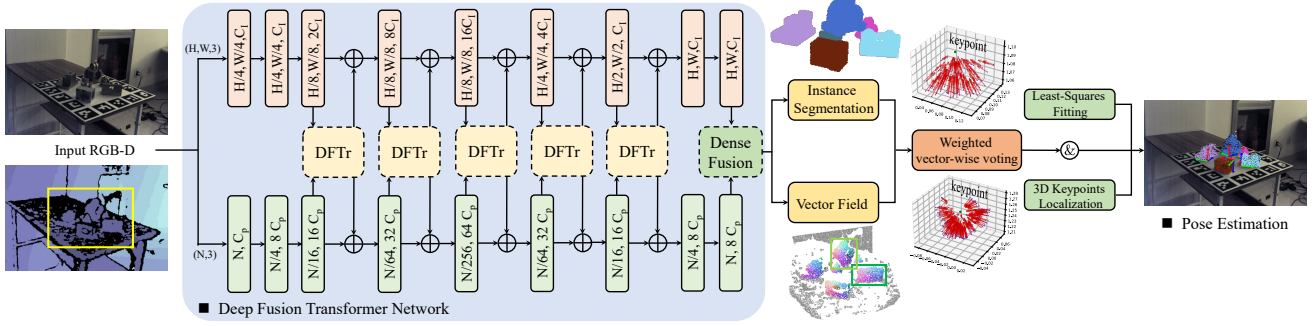


Figure 2. An overview of our proposed method. Taking the RGB-D images as input, two branch networks are utilized to extract color and geometric features respectively. Between them, the deep fusion transformer blocks are integrated for cross-modal features aggregation. With the fused features, two prediction heads are utilized to obtain segmented masks and vector fields. By using the weighted vector-wise voting algorithm, the accurate keypoints are obtained and the final poses are computed by a least-squares fitting algorithm.




Figure 3. An illustration of the deep fusion transformer block.

the input features by linear transformations as follows:

$$(Q^{max}, K, V) = F_{in}^{(max)} \cdot (W_q^{max}, W_k, W_v), \quad (2)$$

where W_q^{max}, W_k and W_v all $\in \mathbb{R}^{C \times C}$ are learnable projection matrices. $F_{in} \in \mathbb{R}^{N_{in} \times C}$ denotes the flattened feature of F_I (F_P), where $N_{in} = \hat{H}\hat{W}(\hat{N})$. $F_{in}^{max} \in \mathbb{R}^{1 \times C}$ denotes the pooling feature F_p^{max} (F_I^{max}). We then perform the multi-head cross-attention operation to infer the refined feature F_{out} as follows:

$$CA_i = \text{C-Attention}(Q_i^{max}, K_i, V_i) = \sigma(Q_i^{max}(K_i)^T / \sqrt{d_k})V_i, \quad (3)$$

$$F_{out} = \text{MultiHead}(Q, K, V) = \text{Concat}(CA_1, \dots, CA_h)W^O, \quad (4)$$

where $1/\sqrt{d_k}$ is a scaling factor and $d_k = C/h$. $\sigma(\cdot)$ is the standard *softmax* normalization function. $W^O \in \mathbb{R}^{C \times C}$ is the projection matrix, h denotes number of heads. Then we add the output feature F_{out} from the cross-attention module to each element of the origin input feature F_I (F_P)

to obtain feature $F_{P \rightarrow I}$ ($F_{I \rightarrow P}$), which are stacked into one sequence in spatial dimension as the input of following transformer-based module:

$$F_{rgb,d} = F_{P \rightarrow I} \oplus F_{I \rightarrow P}, \quad (5)$$

where $F_{P \rightarrow I} = F_I + \mathcal{R}(F_{out}^{P \rightarrow I})$ and $F_{I \rightarrow P} = F_P + \mathcal{R}(F_{out}^{I \rightarrow P})$, \mathcal{R} denotes the repeat operation, $+$ is element-wise addition, and \oplus is the concatenate operation.

After the bidirectional cross-attention module, we obtain the feature sequence $F_{rgb,d} \in \mathbb{R}^{L \times C}$, where $L = \hat{H}\hat{W} + \hat{N}$. We then feed $F_{rgb,d}$ into a transformer-based module $TrM(\cdot)$, as shown in Fig. 3. As mentioned above, in order to implicitly encode the spatial information between different feature elements of the two feature sequences, we insert a learnable positional embedding into the DFT block following [17, 56]. Given the input feature sequence $F_{rgb,d}$, for each layer $TrM_\ell(\cdot)$, the output $\mathcal{F}_\ell = TrM_\ell(F_{rgb,d})$ is formulated as follows:

$$\mathcal{F}'_\ell = \text{MSA}(\text{LN}(\mathcal{F}'_0)) + \mathcal{F}'_0, \quad (6)$$

$$\mathcal{F}_\ell = \text{MLP}(\text{LN}(\mathcal{F}'_\ell)) + \mathcal{F}'_\ell, \ell = 1 \dots L \quad (7)$$

where $\mathcal{F}'_0 = F_{rgb,d} + \sigma_{pos}$, $\sigma_{pos} \in \mathbb{R}^{L \times C}$ is the position embeddings. MSA, LN, and MLP denote multiheaded self-attention, layernorm, and Multi-Layer Perceptrons respectively. The final feature sequence \mathcal{F}_ℓ is then split into two feature sequences F'_I and F'_P by following the initial permutation order, which is integrated into the original modality branch as complementary features.

Dense RGB-D feature fusion. Through the proposed DFT block and two feature extraction branches, we get dense features from the two modality inputs. We follow [19] to obtain the pair-wise RGBD feature, which is then fed into the Densefusion [58] module to generate dense fused features. This early and later fusion strategy has shown a remarkable performance in our experiments. The generated features are utilized to predict instance masks and keypoints vector field for subsequent pose estimation.

3.3. Weighted Vector-Wise Voting for 3D Keypoints Localization and Pose Estimation

In most pose estimation methods, for better performance, the keypoints of the object are predicted first rather than the pose parameters directly [19, 20]. We also follow this workflow but further introduce a novel 3D keypoint localization algorithm to improve the accuracy and efficiency of the model. Concretely, we first utilize the predicted mask and keypoints vector field to locate 3D keypoints and then employ a correspondence-based approach to estimate object pose parameters.

Instance-level 3D keypoint localization. Given the dense fused features, we deploy two head networks for instances segmentation and keypoints vector field prediction. As a powerful guidance signal in the training phase, variables with clear boundaries or constraints are more conducive to network learning. In this way, instead of regressing point-wise offsets to the predefined keypoints directly, we propose to predict the unit vector that represents the direction from the point p_i to a 3D keypoint k_j of the object, like [47] in 2D. More specifically, given the segmented object 3D points $P_o = \{p_i | i = 1 \dots M\} \in \mathbb{R}^{M \times 3}$ and its corresponding vector field $V_o = \{v_j | j = 1 \dots K\} \in \mathbb{R}^{K \times M \times 3}$ from the prediction head, where K is the number of keypoints. For a keypoint k_j , we formulate this problem as follows:

$$\begin{aligned} D(k_j; P_o, V_{o-j}, c) &= \sum_{i=1}^M D(k_j; P_o^i, V_{o-j}^i, c_i) \\ &= \sum_{i=1}^M c_i (P_o^i - k_j)^T (I - V_{o-j}^i (V_{o-j}^i)^T) (P_o^i - k_j), \end{aligned} \quad (8)$$

where $V_{o-j} \in \mathbb{R}^{M \times 3}$ denotes the j -th predicted keypoint vector field, and c_i is the weight of each vector learned by the proposed network. We solve this optimization problem by minimizing the sum of squared distances $D(\cdot)$. Thus, the objective is:

$$\hat{k}_j = \arg \min_{k_j} D(k_j; P_o, V_{o-j}, c) \quad (9)$$

Taking derivatives with respect to k_j , we have $\partial D / \partial k_j = \sum_{i=1}^M -2(I - V_{o-j}^i (V_{o-j}^i)^T) (P_o^i - k_j) = 0$. Finally, we can get a linear system of equations:

$$A k_j = b, \quad (10)$$

$$A = \sum_{i=1}^M (I - V_{o-j}^i (V_{o-j}^i)^T), \quad b = \sum_{i=1}^M (I - V_{o-j}^i (V_{o-j}^i)^T) P_o^i, \quad (11)$$

We then obtain k_j by applying the Moore-Penrose pseudoinverse for Eq. 10: $k_j = \hat{k}_j = A^\dagger b$. Compare to the

MeanShift [15] clustering algorithm employed by [19, 20], our method can obtain the weighted least squares solution without any iterations. Experiments show that our method can achieve superior efficiency and comparable, or even better accuracy in the inference phase.

Keypoint-based object pose estimation. Given the predicted 3D keypoints $\{k_j\}_{j=1}^K$ and the corresponding 3D keypoints $\{k_j^*\}_{j=1}^K$ in the object coordinate system, a correspondence based method [55] is adopted to compute the pose parameters.

Overall multi-task loss function. We supervise our network with the following loss function:

$$\mathcal{L} = \lambda_1 \mathcal{L}_{seg} + \lambda_2 \mathcal{L}_{vecf}, \quad (12)$$

where \mathcal{L}_{seg} and \mathcal{L}_{vecf} are the instance segmentation loss and keypoints vector field prediction loss respectively, where $\mathcal{L}_{vecf} = \frac{1}{M} \sum_i (\mathcal{L}_{kps} c_i - w \cdot \log(c_i))$, $i = 1 \dots M$, c_i denotes the weight of each vector, w is a balancing hyperparameter, which is set to 0.015 in our experiments. For \mathcal{L}_{seg} and \mathcal{L}_{kps} , we use Focal Loss [41] and L1 Loss as in [20] respectively. We set $\lambda_1 = \lambda_2 = 1.0$ in our experiments.

4. Experiments and Results

4.1. Experiments Settings

Datasets. We evaluate our method on four benchmark datasets. **MP6D** [11] contains 77 RGBD video segments (20100 frames in total) which capture scenes with high occlusion and illumination changes of 20 metal parts. The selected metal parts are collected from natural industrial environments, and all objects are *texture-less*, *symmetric*, *of complex shape*, *high reflectivity*, and *uniform color*, which make this dataset challenge. We follow [11] to split the training and testing set, and the hole completion algorithm [34] is deployed following [20]. **YCB-Video** [4] has 92 RGBD videos. Each video captures a subset of the 21 objects in varying scenes. We follow prior works [19, 20, 58, 61] to prepare our training and testing set, including data processing. **LineMOD** [21] consists of 13 low-textured objects in 13 videos with annotated 6D pose and instance mask. We use synthesis images in the training phase following [19, 20, 47] and follow previous works [47, 61] to split the training and testing set. **Occlusion LINEMOD** [2] is a subset of the LINEMOD datasets created by additionally annotating. Each scene has multi-labeled instances with heavy occlusion, making pose estimation a great challenge.

Evaluation Metrics. We use three metrics for our method evaluation (i.e., the Average Distance of Model Points (ADD) [22], the Average Closest Point Distance (ADD-S) [61] and the Visible Surface Discrepancy (VSD) [25]. For asymmetric objects, the ADD metric is utilized to compute the mean distance between the two object point

Table 1. Quantitative comparison results (ADD-S [61] AUC, VSD [25]) on the MP6D Dataset with the state-of-the-art frameworks. DF (per-pixel) means DenseFusion (per-pixel).

Object	Hodan [27]		PointFusion [63]		DCF [40]		DF (per-pixel) [58]		MaskedFusion [48]		G2L-Net [12]		PVN3D [20]		FFB6D [19]		Ours	
	ADDS	VSD	ADDS	VSD	ADDS	VSD	ADDS	VSD	ADDS	VSD	ADDS	VSD	ADDS	VSD	ADDS	VSD	ADDS	VSD
Obj_01	83.42	73.14	84.33	73.45	86.06	74.09	89.35	75.35	88.95	76.01	89.51	78.39	90.28	85.06	93.28	80.35	95.44	94.93
Obj_02	80.23	70.35	81.01	72.36	85.36	73.42	87.78	76.84	89.19	75.98	89.03	80.04	91.88	88.43	92.83	81.47	96.51	94.12
Obj_03	65.78	35.69	64.74	35.95	65.33	36.08	72.45	39.51	70.03	37.55	74.93	38.42	76.67	35.68	79.51	43.50	84.93	57.36
Obj_04	70.56	57.52	72.50	55.01	73.95	55.90	77.98	60.96	74.68	58.13	85.39	60.55	88.13	68.34	84.98	64.93	92.02	79.95
Obj_05	69.78	51.35	68.96	51.39	67.19	52.37	71.23	54.62	75.69	55.92	72.13	56.82	73.46	58.96	76.33	60.20	86.24	71.23
Obj_06	72.36	55.84	70.66	55.82	71.65	54.21	75.34	57.15	78.31	59.01	85.08	62.95	87.16	76.39	83.98	63.70	96.10	88.96
Obj_07	80.79	74.95	81.12	76.31	82.07	75.58	88.63	83.58	85.25	81.32	89.09	89.37	94.81	94.63	94.94	90.29	97.51	97.94
Obj_08	80.71	67.98	81.37	69.80	82.39	68.29	84.78	70.12	85.38	69.71	90.10	72.91	93.76	73.21	89.76	74.73	96.75	87.12
Obj_09	69.8	38.27	65.98	34.32	68.27	35.22	73.67	40.52	75.46	38.44	79.91	41.59	82.71	40.92	81.25	46.63	91.23	70.15
Obj_10	75.32	65.69	77.19	66.08	79.10	67.92	80.54	68.65	77.62	69.98	86.03	71.32	86.16	76.21	88.92	71.07	94.98	90.92
Obj_11	72.56	43.88	71.98	41.99	70.96	42.35	79.65	44.71	75.91	45.36	82.01	46.09	81.21	56.09	84.87	47.12	92.36	78.46
Obj_12	74.13	44.57	76.32	45.23	77.03	46.07	78.88	46.78	76.98	45.17	77.93	47.91	79.00	45.97	84.82	54.86	89.99	66.21
Obj_13	78.63	48.41	77.05	49.02	75.15	48.31	80.12	50.26	80.58	49.33	85.38	51.98	86.69	52.22	85.42	52.57	95.04	85.34
Obj_14	76.89	49.68	77.90	52.39	76.98	50.23	80.89	51.28	81.15	48.92	84.54	48.39	87.06	49.09	87.99	56.16	94.13	65.08
Obj_15	64.53	8.68	67.36	6.08	66.23	8.19	68.45	10.81	66.30	8.98	72.92	11.94	74.17	17.96	75.01	13.08	86.97	40.29
Obj_16	69.88	35.88	72.28	38.69	73.08	39.15	75.81	40.78	73.86	38.71	79.38	41.49	81.35	40.80	83.95	41.01	92.14	71.13
Obj_17	77.42	65.11	85.93	67.00	84.68	68.91	89.16	69.02	88.11	71.25	92.08	74.35	93.47	85.86	93.19	71.22	94.25	94.58
Obj_18	75.63	73.85	81.46	76.19	80.91	77.49	83.23	78.65	85.94	77.93	88.13	76.92	87.57	74.82	91.73	81.35	94.69	92.47
Obj_19	72.89	50.54	76.82	50.98	78.07	52.04	81.98	52.36	79.37	54.98	85.31	60.18	88.82	63.76	87.28	58.01	95.03	87.24
Obj_20	72.65	50.18	75.91	55.39	74.20	54.39	76.59	53.58	78.93	53.29	81.41	59.17	88.10	59.26	85.75	58.53	93.92	92.08
ALL	74.20	53.08	75.54	53.67	75.93	54.01	79.84	56.28	79.38	55.80	83.51	58.54	85.42	62.18	86.29	60.54	93.01	80.23

Table 2. Quantitative evaluation results (ADD-S [61] and ADD(S) [22] AUC) on the YCB-Video Dataset.

Refinement?	Method	ADDS	ADD(S)
w/o	PoseCNN [61]	75.8	59.9
	PointFusion [63]	83.9	-
	DCF [40]	85.7	77.9
	DF (per-pixel) [58]	91.2	82.9
	PVN3D [20]	95.5	91.8
	PR-GCN [66]	95.8	-
	FFB6D [19]	96.6	92.7
	RCVPose [60]	96.6	95.2
	Ours	96.7	94.4
w/	PoseCNN+ICP [61]	93.0	85.4
	DF (iterative) [58]	93.2	86.1
	MoreFusion [57]	95.7	91.0
	PVN3D [20]+ICP	96.1	92.3
	FFB6D [19]+ICP	97.0	93.1
	RCVPose [60]+ICP	97.2	95.9
	Ours+ICP	97.3	94.8

sets transformed by the estimated pose $[R, t]$ and the ground truth pose $[R^*, t^*]$, formulated as follows:

$$ADD = \frac{1}{N} \sum_{p \in \mathcal{O}} \|(Rp + t) - (R^*p + t^*)\|. \quad (13)$$

where p is a point of totally N points in the object \mathcal{O} . The ADD-S metric is designed for symmetric objects based on the closest point distance:

$$ADD-S = \frac{1}{N} \sum_{p_1 \in \mathcal{O}} \min_{p_2 \in \mathcal{O}} \|(Rp_1 + t) - (R^*p_2 + t^*)\|. \quad (14)$$

The VSD metric is ambiguity-invariant to object symmetries, determined by the distance between the estimated and ground truth visible object depth surfaces. For the MP6D




Figure 4. Performance of different methods at increasing occlusion levels on YCB-Video dataset.

datasets, as in [11] and the BOP challenge [26], we report the recall of correct poses at $e_{vsd} < 0.3$ with the tolerance $\tau = 5cm$ and $\delta = 1.5cm$. We also report the area under the accuracy-threshold curve computed by varying the distance threshold (ADD-S AUC) following [11, 20, 61]. For the YCB-Video datasets, we report the ADD-S AUC and the ADD(S) AUC following [19]. For the LineMOD and Occlusion LineMOD datasets, we report the accuracy of distance less than 10% of the object’s diameter (ADD-0.1d) as in [22, 47].

Implementation Details. For the image branch embedding network, we employ a pre-trained ResNet34 [18] encoder followed by a four-level PSPNet [65] as the decoder. The RandLA-Net [28] is used to extract geometric features with randomly sampled 12800 points from the depth image as input. Between each encoding and decoding layer, the DFTr block is inserted to integrate two modality features.

Table 3. Quantitative evaluation using the ADD-0.1d [22] metric on the LineMOD Dataset. Symmetric objects are in bold.

	RGB				RGB-D					
	PoseCNN DeepIM [38, 61]	PVNet [47]	CDPN [39]	DPOD [64]	PointFusion [63]	DenseFusion [58]	G2L-Net [13]	PVN3D [20]	FFB6D [19]	Ours
ape	77.0	43.6	64.4	87.7	70.4	92.3	96.8	97.3	98.4	98.6
benchvise	97.5	99.9	97.8	98.5	80.7	93.2	96.1	99.7	100.0	100.0
camera	93.5	86.9	91.7	96.1	60.8	94.4	98.2	99.6	99.9	100.0
can	96.5	95.5	95.9	99.7	61.1	93.1	98.0	99.5	99.8	100.0
cat	82.1	79.3	83.8	94.7	79.1	96.5	99.2	99.8	99.9	100.0
driller	95.0	96.4	96.2	98.8	47.3	87.0	99.8	99.3	100.0	100.0
duck	77.7	52.6	66.8	86.3	63.0	92.3	97.7	98.2	98.4	99.1
eggbox	97.1	99.2	99.7	99.9	99.9	99.8	100.0	99.8	100.0	100.0
glue	99.4	95.7	99.6	96.8	99.3	100.0	100.0	100.0	100.0	100.0
holepuncher	52.8	82.0	85.8	86.9	71.8	92.1	99.0	99.9	99.8	100.0
iron	98.3	98.9	97.9	100.0	83.2	97.0	99.3	99.7	99.9	99.9
lamp	97.5	99.3	97.9	96.8	62.3	95.3	99.5	99.8	99.9	100.0
phone	87.7	92.4	90.8	94.7	78.8	92.8	98.9	99.5	99.7	99.6
MEAN	88.6	86.3	89.9	95.2	73.7	94.3	98.7	99.4	99.7	99.8

Table 4. Quantitative evaluation using the ADD-0.1d [22] metric on the Occlusion-LineMOD Dataset. Symmetric objects are in bold.

Method	PoseCNN [61]	Pix2Pose [46]	PVNet [47]	Hu et al. [29]	HybridPose [51]	PVN3D [20]	PR-GCN [66]	FFB6D [19]	RCVPose [60]	Nguyen et al. [44]	ZebraPose [53]	Ours
ape	9.6	22.0	15.8	19.2	20.9	33.9	40.2	47.2	-	53.8	57.9	64.1
can	45.2	44.7	63.3	65.1	75.3	88.6	76.2	85.2	-	89.7	95.0	96.1
cat	0.9	22.7	16.7	18.9	24.9	39.1	57.0	45.7	-	45.1	60.6	52.2
driller	41.4	44.7	65.7	69.0	70.2	78.4	82.3	81.4	-	84.4	94.8	95.8
duck	19.6	15.0	25.2	25.3	27.9	41.9	30.0	53.9	-	87.2	64.5	72.3
eggbox	22.0	25.2	50.2	52.0	52.4	80.9	68.2	70.2	-	76.9	70.9	75.3
glue	38.5	32.4	49.6	51.4	53.8	68.1	67.0	60.1	-	89.9	88.7	79.3
holepuncher	22.1	49.5	39.7	45.6	54.2	74.7	97.2	85.9	-	83.3	83.0	86.8
MEAN	24.9	32.0	40.8	43.3	47.5	63.2	65.0	66.2	70.2	76.3	76.9	77.7

The final fused RGBD features are then fed into the head network consisting of shared MLPs for instance segmentation and keypoints vector field prediction. The input of the RGB embedding branch is a scene image with the size of $480 \times 640 \times 3$, and a point set with a size of $12800 \times C_{in}$ as the input for geometric representation learning, where C_{in} denotes the input coordinate, color and normal information of each point.

4.2. Comparison With State-of-The-Arts.

Evaluation on the MP6D dataset. Tab. 1 shows the quantitative evaluation results for all 20 objects in the MP6D dataset. We compare our method with other single-view RGB-D fusion-based methods without iterative refinement. As shown in the table, our method significantly outperforms other approaches by a large margin. In particular, our model advances FFB6D [19] and PVN3D [20] by 6.72% and 7.59% respectively on the ADDS metric, and achieves 19.69% and 18.05% improvement on the VSD metric. These experimental results reveal the effectiveness of our method. We further present qualitative comparison results of FFB6D and our model in Fig. 5. Compared to FFB6D, our method is more robust towards objects with texture-less or heavy reflective surfaces as well as severe occlusions.

Evaluation on the YCB-Video dataset. We then evaluate our approach on the YCB-Video dataset, as illustrated in Tab. 2. Our model surpasses FFB6D by 0.1% and 1.7%

Table 5. Effect of different model types on the MP6D Dataset.

	β -w/DF	M_s -w/o DF	M_s	M_{mL}	M_{mu}	M_L
ADD-S	86.94	89.99	93.01	90.98	89.74	88.26
VSD	64.92	70.49	80.23	74.48	71.05	68.31
Params	40.6M	116.7M	138.6M	162.1M	172.4M	175M

on the ADDS and ADD(S) metrics respectively in terms of average accuracy. With the extra iterative refinement (e.g. ICP), our method can achieve the best performance on ADD-S metric compared with RCVPose [60]. In Fig. 4, we evaluate the robustness of our algorithm to occlusion. Our approach is capable of maintain robust performance with varying occlusion levels, compared with FFB6D [19] and Uni6D [31]. We think that is because the global modeling of semantic similarity between two modalities helps the network make full use of cues from visible parts of objects, which is beneficial to locate 3D keypoints precisely from the vector field. Fig. 5 also provides qualitative comparison results, in which our approach performs more robustly with better performances.

Evaluation on the LineMOD dataset & Occlusion

LineMOD dataset. Tab. 3 and Tab. 4 show the quantitative comparison results of the proposed method on the LineMOD and Occlusion LineMOD datasets respectively. Our approach achieves state-of-the-art performance on both datasets, demonstrating its effectiveness and robustness towards severe occlusion.

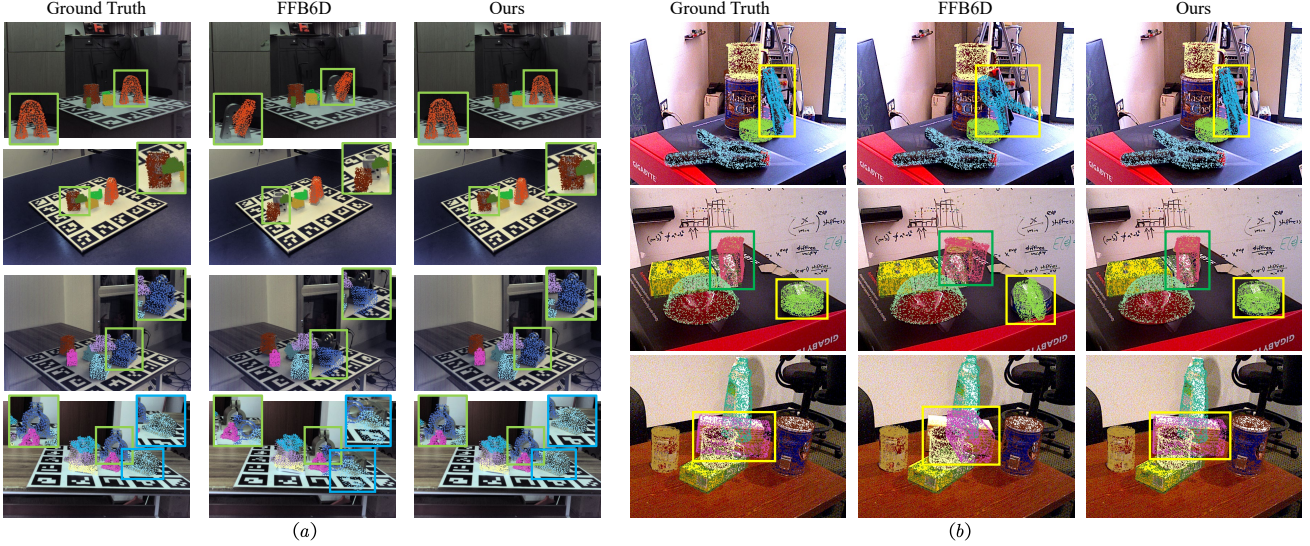


Figure 5. Qualitative comparison results on (a) MP6D dataset and (b) YCB-Video dataset.

Table 6. Effect of components of DFTr block on the MP6D dataset. CMA: cross-modality attention, PE: positional embedding.

		Pose Result		Query Source		Pose Result	
	CMA	ADD(S)	VSD	RGB	D	ADD(S)	VSD
	PE	87.73	63.75	✓		91.57	76.94
✓		91.46	77.48		✓	90.13	73.89
	✓	89.97	71.59	✓	✓	93.01	80.23
✓	✓	93.01	80.23	✓	✓	93.01	80.23

Table 7. Effect of the weighted vector-wise voting (WVWV) algorithm. ms/f: ms per frame. MS: MeanShift algorithm.

	ADD-S	ADD(S)	Run-time	MS	WVWV
PVN3D	95.5	91.8	367 ms/f	KP err. (cm)	0.0379
PVN3D+WVWV	95.6	92.0	209 ms/f	Run-time (ms/f)	18
FFB6D	96.6	92.7	295 ms/f	ADD-S	92.98
FFB6D+WVWV	96.4	92.9	109 ms/f	VSD	80.12
					93.01
					80.23

4.3. Ablation Studies.

We comprehensively conduct the following ablation studies on our design choice and explore the effect of individual components.

How many DFTr blocks do you need? In order to verify the fusion power of DFTr block, we select a dozen various models, as shown in Tab. 5. Concretely, we define $ds-n$ and $up-n$ as the n -th layer in the downsampling and upsampling process respectively. \mathcal{B} -w/ DF: baseline model only with the DenseFusion module. M_S -w/o DF and M_S : without/with DenseFusion module, add DFTr module to $[ds-5; up-1]$. Similarly, M_{mL} , M_{mu} and M_L : add DFTr module to $[ds-4, 5; up-1]$, $[ds-4, 5; up-1, 2]$ and $[ds-3, 4, 5; up-1, 2]$ on baseline model respectively. Compared with the baseline model, networks with DFTr block can all show performance gains. Among them, model M_S exhibits the highest improvement. We think that the deployment of the DFTr in the deeper and high-dimensional feature maps can maximize the network to perceive the global information of the scene and share it with every neuron. With the increase of DFTr blocks, the performance is no longer improved because RGBD features have been fully integrated, resulting in overfitting of the network.

Effect of components of the DFTr block. We ablate the bidirectional cross-modality attention (CMA) and positional embedding (PE) components in the DFTr block to validate their impact, as shown in Tab. 6 (left). Compared

with the base model, network with the CMA achieves significant improvements. Combining with PE obtains the best results. We believe that CMA can enhance the feature representation of the DFTr block by filtering out non-salient features and establishing connections between salient global features across modalities. We also ablate the query source in the CMA, as shown in Tab. 6 (right). Compared with no CMA (only PE), querying from the RGB or Depth can all boost performance, demonstrating its effectiveness in integrating cross-modal features. The combination of them achieves the best results.

Effect of the weighted vector-wise voting algorithm. We integrate the weighted vector-wise voting algorithm (WVWV) into two state-of-the-art frameworks and evaluate their generalization performance on YCB-Video, as in Tab. 7 (left). Compared with the original methods, our keypoints voting algorithm is 1.7x faster than PVN3D and 2.7x faster than FFB6D, with better performance on the ADD(S) metric. The results reveal that WVWV benefits other frameworks for pose estimation. We also compared the performance of WVWV and MeanShift algorithm on MP6D, as in Tab. 7 (right). Our method has higher accuracy and inference speed.

Time efficiency. In Fig. 6, we present comparison results in terms of inference time and performance on YCB-Video. Our method achieves better performance with on-par efficiency (48 ms/frame) compared with Uni6D [31].




Figure 6. Balance the speed and accuracy.

5. Conclusion

In this paper, we propose the DFTr network, a novel 6D object pose estimator with a powerful deep fusion transformer block for cross-modalities feature aggregation. We further introduce a new efficient weighted vector-wise voting algorithm for 3D keypoints detection, which employs a non-iterative global optimization strategy to ensure location accuracy and greatly reduce computing costs. Extensive experiments on four benchmark datasets demonstrate that our method achieves significant performance increase over other SOTA approaches. This work can potentially generalize to other RGBD-based applications, such as object perception and robot manipulation.

Acknowledgment. This work is partly supported by Hong Kong RGC Theme-based Research Scheme (No.T45-401/22-N), Hong Kong RGC General Research Fund (No.15218521), Hong Kong Research Grants Council Project (No.24209223), and Science, Technology and Innovation Commission of Shenzhen Municipality Project (No.SGDx20220530111201008).

References

- [1] Hayet Belghit, Abdelkader Bellarbi, Nadia Zenati, and Samir Otmane. Vision-based pose estimation for augmented reality: Comparison study. In *3rd IEEE International Conference on Pattern Analysis and Intelligent Systems (PAIS 2018)*. IEEE, 2018. 1
- [2] 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. 2, 5
- [3] Ming Cai and Ian Reid. Reconstruct locally, localize globally: A model free method for object pose estimation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 3153–3163, 2020. 2
- [4] Berk Calli, Arjun Singh, Aaron Walsman, Siddhartha Srinivasa, Pieter Abbeel, and Aaron M Dollar. The ycb object and model set: Towards common benchmarks for manipulation research. In *2015 international conference on advanced robotics (ICAR)*, pages 510–517. IEEE, 2015. 2, 5
- [5] Tuo Cao, Fei Luo, Yanping Fu, Wenxiao Zhang, Shengjie

- Zheng, and Chunxia Xiao. Dgecn: A depth-guided edge convolutional network for end-to-end 6d pose estimation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 3783–3792, 2022. 2
- [6] Chun-Fu (Richard) Chen, Quanfu Fan, and Rameswar Panda. CrossViT: Cross-Attention Multi-Scale Vision Transformer for Image Classification. In *International Conference on Computer Vision (ICCV)*, 2021. 3
- [7] Dengsheng Chen, Jun Li, Zheng Wang, and Kai Xu. Learning canonical shape space for category-level 6d object pose and size estimation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 11973–11982, 2020. 2
- [8] Kai Chen, Rui Cao, Stephen James, Yichuan Li, Yun-Hui Liu, Pieter Abbeel, and Qi Dou. Sim-to-real 6d object pose estimation via iterative self-training for robotic bin picking. In *European Conference on Computer Vision*, pages 533–550. Springer, 2022. 2
- [9] Kai Chen and Qi Dou. Sgpa: Structure-guided prior adaptation for category-level 6d object pose estimation. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 2773–2782, 2021. 1
- [10] Kai Chen, Stephen James, Congying Sui, Yun-Hui Liu, Pieter Abbeel, and Qi Dou. Stereopose: Category-level 6d transparent object pose estimation from stereo images via back-view nocs. In *2023 IEEE International Conference on Robotics and Automation (ICRA)*, pages 2855–2861. IEEE, 2023. 2
- [11] Long Chen, Han Yang, Chenrui Wu, and Shiqing Wu. Mp6d: An rgb-d dataset for metal parts’ 6d pose estimation. *IEEE Robotics and Automation Letters*, 7(3):5912–5919, 2022. 2, 5, 6
- [12] Wei Chen, Xi Jia, Hyung Jin Chang, Jinming Duan, and Ales Leonardis. G2l-net: Global to local network for real-time 6d pose estimation with embedding vector features. In *IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, June 2020. 6
- [13] Wei Chen, Xi Jia, Hyung Jin Chang, Jinming Duan, and Ales Leonardis. G2l-net: Global to local network for real-time 6d pose estimation with embedding vector features. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 4233–4242, 2020. 7
- [14] Xiaozhi Chen, Huimin Ma, Ji Wan, Bo Li, and Tian Xia. Multi-view 3d object detection network for autonomous driving. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 1907–1915, 2017. 1
- [15] Dorin Comaniciu and Peter Meer. Mean shift: A robust approach toward feature space analysis. *IEEE Transactions on Pattern Analysis & Machine Intelligence*, (5):603–619, 2002. 2, 5
- [16] Xinke Deng, Yu Xiang, Arsalan Mousavian, Clemens Eppner, Timothy Bretl, and Dieter Fox. Self-supervised 6d object pose estimation for robot manipulation. In *2020 IEEE International Conference on Robotics and Automation (ICRA)*, pages 3665–3671. IEEE, 2020. 1
- [17] Alexey Dosovitskiy, Lucas Beyer, Alexander Kolesnikov, Dirk Weissenborn, Xiaohua Zhai, Thomas Unterthiner, Mostafa Dehghani, Matthias Minderer, Georg Heigold, Syl-

- vain Gelly, Jakob Uszkoreit, and Neil Houlsby. An image is worth 16x16 words: Transformers for image recognition at scale. *ICLR*, 2021. 4
- [18] Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. Deep residual learning for image recognition. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 770–778, 2016. 6
- [19] Yisheng He, Haibin Huang, Haoqiang Fan, Qifeng Chen, and Jian Sun. Ffb6d: A full flow bidirectional fusion network for 6d pose estimation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 3003–3013, 2021. 1, 2, 3, 4, 5, 6, 7
- [20] Yisheng He, Wei Sun, Haibin Huang, Jianran Liu, Haoqiang Fan, and Jian Sun. Pvn3d: A deep point-wise 3d keypoints voting network for 6dof pose estimation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 11632–11641, 2020. 1, 2, 3, 5, 6, 7
- [21] Stefan Hinterstoisser, Stefan Holzer, Cedric Cagniard, Slobodan Ilic, Kurt Konolige, Nassir Navab, and Vincent Lepetit. Multimodal templates for real-time detection of texture-less objects in heavily cluttered scenes. In *2011 international conference on computer vision*, pages 858–865. IEEE, 2011. 2, 5
- [22] 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 *Asian conference on computer vision*, pages 548–562. Springer, 2012. 5, 6, 7
- [23] Stefan Hinterstoisser, Vincent Lepetit, Naresh Rajkumar, and Kurt Konolige. Going further with point pair features. In *Computer Vision—ECCV 2016: 14th European Conference, Amsterdam, The Netherlands, October 11–14, 2016, Proceedings, Part III 14*, pages 834–848. Springer, 2016. 2
- [24] Tomáš Hodan, Pavel Haluza, Štěpán Obdržálek, Jiri Matas, Manolis Lourakis, and Xenophon Zabulis. T-less: An rgb-d dataset for 6d pose estimation of texture-less objects. In *2017 IEEE Winter Conference on Applications of Computer Vision (WACV)*, pages 880–888. IEEE, 2017. 1
- [25] Tomáš Hodaň, Jiří Matas, and Štěpán Obdržálek. On evaluation of 6d object pose estimation. In *European Conference on Computer Vision*, pages 606–619. Springer, 2016. 5, 6
- [26] Tomas Hodan, Frank Michel, Eric Brachmann, Wadim Kehl, Anders GlentBuch, Dirk Kraft, Bertram Drost, Joel Vidal, Stephan Ihrke, Xenophon Zabulis, et al. Bop: Benchmark for 6d object pose estimation. In *Proceedings of the European Conference on Computer Vision (ECCV)*, pages 19–34, 2018. 1, 6
- [27] Tomáš Hodaň, Xenophon Zabulis, Manolis Lourakis, Štěpán Obdržálek, and Jiří Matas. Detection and fine 3d pose estimation of texture-less objects in rgb-d images. In *2015 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, pages 4421–4428. IEEE, 2015. 6
- [28] Qingyong Hu, Bo Yang, Linhai Xie, Stefano Rosa, Yulan Guo, Zhihua Wang, Niki Trigoni, and Andrew Markham. Randla-net: Efficient semantic segmentation of large-scale point clouds. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 11108–11117, 2020. 6
- [29] Yinlin Hu, Pascal Fua, Wei Wang, and Mathieu Salzmann. Single-stage 6d object pose estimation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 2930–2939, 2020. 7
- [30] 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
- [31] Xiaoke Jiang, Donghai Li, Hao Chen, Ye Zheng, Rui Zhao, and Liwei Wu. Uni6d: A unified cnn framework without projection breakdown for 6d pose estimation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 11174–11184, 2022. 7, 8
- [32] Xingyu Jiang, Jiayi Ma, Guobao Xiao, Zhenfeng Shao, and Xiaojie Guo. A review of multimodal image matching: Methods and applications. *Information Fusion*, 73:22–71, 2021. 2
- [33] Wadim Kehl, Fabian Manhardt, Federico Tombari, Slobodan Ilic, and Nassir Navab. Ssd-6d: Making rgb-based 3d detection and 6d pose estimation great again. In *Proceedings of the IEEE international conference on computer vision*, pages 1521–1529, 2017. 2
- [34] Jason Ku, Ali Harakeh, and Steven L Waslander. In defense of classical image processing: Fast depth completion on the cpu. In *2018 15th Conference on Computer and Robot Vision (CRV)*, pages 16–22. IEEE, 2018. 5
- [35] Jason Ku, Melissa Mozifian, Jungwook Lee, Ali Harakeh, and Steven L Waslander. Joint 3d proposal generation and object detection from view aggregation. In *2018 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, pages 1–8. IEEE, 2018. 2
- [36] Chi Li, Jin Bai, and Gregory D Hager. A unified framework for multi-view multi-class object pose estimation. In *Proceedings of the european conference on computer vision (eccv)*, pages 254–269, 2018. 1, 2
- [37] Yiming Li, Tao Kong, Ruihang Chu, Yifeng Li, Peng Wang, and Lei Li. Simultaneous semantic and collision learning for 6-dof grasp pose estimation. In *2021 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, pages 3571–3578. IEEE, 2021. 1
- [38] 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. 7
- [39] 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. 7
- [40] Ming Liang, Bin Yang, Shenlong Wang, and Raquel Urtasun. Deep continuous fusion for multi-sensor 3d object detection. In *Proceedings of the European Conference on Computer Vision (ECCV)*, pages 641–656, 2018. 2, 6
- [41] Tsung-Yi Lin, Priya Goyal, Ross Girshick, Kaiming He, and Piotr Dollár. Focal loss for dense object detection. In *Proceedings of the IEEE international conference on computer vision*, pages 2980–2988, 2017. 5
- [42] Yuanpeng Liu, Laishui Zhou, Hua Zong, Xiaoxi Gong, Qiaoyun Wu, Qingxiao Liang, and Jun Wang. Regression-

- based three-dimensional pose estimation for texture-less objects. *IEEE Transactions on Multimedia*, 21(11):2776–2789, 2019. [2](#)
- [43] Jiayi Ma, Xingyu Jiang, Aoxiang Fan, Junjun Jiang, and Junchi Yan. Image matching from handcrafted to deep features: A survey. *International Journal of Computer Vision*, 129:23–79, 2021. [2](#)
- [44] Van Nguyen Nguyen, Yinlin Hu, Yang Xiao, Mathieu Salzmann, and Vincent Lepetit. Templates for 3d object pose estimation revisited: Generalization to new objects and robustness to occlusions. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 6771–6780, 2022. [7](#)
- [45] Keunhong Park, Arsalan Mousavian, Yu Xiang, and Dieter Fox. Latentfusion: End-to-end differentiable reconstruction and rendering for unseen object pose estimation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 10710–10719, 2020. [2](#)
- [46] 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. [7](#)
- [47] Sida Peng, Yuan Liu, Qixing Huang, Xiaowei Zhou, and Hujun Bao. Pvnnet: Pixel-wise voting network for 6dof pose estimation. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 4561–4570, 2019. [2](#), [5](#), [6](#), [7](#)
- [48] Nuno Pereira and Luís A. Alexandre. MaskedFusion: Mask-based 6d object pose estimation. In *19th IEEE International Conference on Machine Learning and Applications (ICMLA 2020)*, December 2020. [6](#)
- [49] Charles R Qi, Or Litany, Kaiming He, and Leonidas J Guibas. Deep hough voting for 3d object detection in point clouds. In *proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 9277–9286, 2019. [2](#)
- [50] 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](#)
- [51] Chen Song, Jiaru Song, and Qixing Huang. Hybridpose: 6d object pose estimation under hybrid representations. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 431–440, 2020. [7](#)
- [52] Yongzhi Su, Jason Rambach, Nareg Minaskan, Paul Lesur, Alain Pagani, and Didier Stricker. Deep multi-state object pose estimation for augmented reality assembly. In *2019 IEEE International Symposium on Mixed and Augmented Reality Adjunct (ISMAR-Adjunct)*, pages 222–227. IEEE, 2019. [1](#)
- [53] Yongzhi Su, Mahdi Saleh, Torben Fetzner, Jason Rambach, Nassir Navab, Benjamin Busam, Didier Stricker, and Federico Tombari. ZebraPose: Coarse to fine surface encoding for 6dof object pose estimation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 6738–6748, 2022. [7](#)
- [54] Bugra Tekin, Sudipta N Sinha, and Pascal Fua. Real-time seamless single shot 6d object pose prediction. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 292–301, 2018. [2](#)
- [55] Shinji Umeyama. Least-squares estimation of transformation parameters between two point patterns. *IEEE Transactions on Pattern Analysis & Machine Intelligence*, 13(04):376–380, 1991. [5](#)
- [56] Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N Gomez, Łukasz Kaiser, and Illia Polosukhin. Attention is all you need. *Advances in neural information processing systems*, 30, 2017. [3](#), [4](#)
- [57] Kentaro Wada, Edgar Sucar, Stephen James, Daniel Lenton, and Andrew J Davison. Morefusion: Multi-object reasoning for 6d pose estimation from volumetric fusion. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 14540–14549, 2020. [2](#), [6](#)
- [58] 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. [1](#), [2](#), [3](#), [4](#), [5](#), [6](#), [7](#)
- [59] Jiase Wang, Kai Chen, and Qi Dou. Category-level 6d object pose estimation via cascaded relation and recurrent reconstruction networks. In *2021 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, pages 4807–4814. IEEE, 2021. [1](#)
- [60] Yangzheng Wu, Mohsen Zand, Ali Etemad, and Michael Greenspan. Vote from the center: 6 dof pose estimation in rgb-d images by radial keypoint voting. In *European Conference on Computer Vision*, pages 335–352. Springer, 2022. [6](#), [7](#)
- [61] Yu Xiang, Tanner Schmidt, Venkatraman Narayanan, and Dieter Fox. Posecnn: A convolutional neural network for 6d object pose estimation in cluttered scenes. *arXiv preprint arXiv:1711.00199*, 2017. [1](#), [2](#), [5](#), [6](#), [7](#)
- [62] Qian Xie, Yu-Kun Lai, Jing Wu, Zhoutao Wang, Yiming Zhang, Kai Xu, and Jun Wang. Mlcvnet: Multi-level context votenet for 3d object detection. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 10447–10456, 2020. [2](#)
- [63] Danfei Xu, Dragomir Anguelov, and Ashesh Jain. Pointfusion: Deep sensor fusion for 3d bounding box estimation. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 244–253, 2018. [1](#), [2](#), [6](#), [7](#)
- [64] Sergey Zakharov, Ivan Shugurov, and Slobodan Ilic. Dpod: 6d pose object detector and refiner. In *Proceedings of the IEEE International Conference on Computer Vision*, pages 1941–1950, 2019. [7](#)
- [65] Hengshuang Zhao, Jianping Shi, Xiaojuan Qi, Xiaogang Wang, and Jiaya Jia. Pyramid scene parsing network. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 2881–2890, 2017. [6](#)
- [66] Guangyuan Zhou, Huiqun Wang, Jiabin Chen, and Di Huang. Pr-gcn: A deep graph convolutional network with point refinement for 6d pose estimation. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 2793–2802, 2021. [6](#), [7](#)
- [67] Guangliang Zhou, Yi Yan, Deming Wang, and Qijun Chen. A novel depth and color feature fusion framework for 6d object pose estimation. *IEEE Transactions on Multimedia*, 2020. [2](#)