

HccePose(BF): Predicting Front & Back Surfaces to Construct Ultra-Dense 2D-3D Correspondences for Pose Estimation

Yulin Wang
Southeast University
yulinwang@seu.edu.cn

Mengting Hu
Southeast University
220240361@seu.edu.cn

Hongli Li
Purdue University
li5125@purdue.edu

Chen Luo*
Southeast University
chenluo@seu.edu.cn

Abstract

In pose estimation for seen objects, a prevalent pipeline involves using neural networks to predict dense 3D coordinates of the object surface on 2D images, which are then used to establish dense 2D-3D correspondences. However, current methods primarily focus on more efficient encoding techniques to improve the precision of predicted 3D coordinates on the object’s front surface, overlooking the potential benefits of incorporating the back surface and interior of the object. To better utilize the full surface and interior of the object, this study predicts 3D coordinates of both the object’s front and back surfaces and densely samples 3D coordinates between them. This process creates ultra-dense 2D-3D correspondences, effectively enhancing pose estimation accuracy based on the Perspective-n-Point (PnP) algorithm. Additionally, we propose Hierarchical Continuous Coordinate Encoding (HCCE) to provide a more accurate and efficient representation of front and back surface coordinates. Experimental results show that, compared to existing state-of-the-art (SOTA) methods on the BOP website, the proposed approach outperforms across seven classic BOP core datasets. Code is available at <https://github.com/WangYuLin-SEU/HCCePose>.

1. Introduction

As a fundamental task in computer vision, pose estimation [1, 2] holds significant value across various applications, including automated product picking in warehouses [3], sheet metal assembly in industrial settings [4], and object tracking in virtual reality [5]. In pose estimation for seen objects, some studies [6–13] predefine 3D keypoints on the surface of a given object and use neural networks to predict the 2D projections of them on 2D images, establishing sparse 2D-3D correspondences. These correspondences are then used with the Perspective-n-Point (PnP) algorithm [14] to compute the object pose. Subsequently, some methods [15–

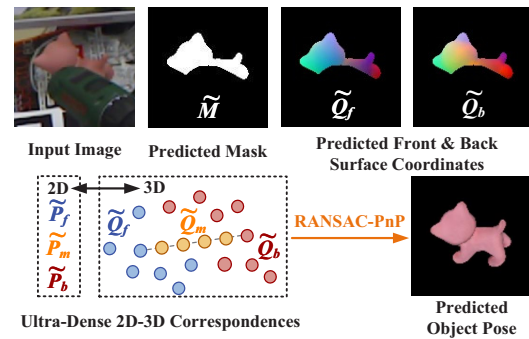


Figure 1. Pose estimation based on ultra-dense 2D-3D correspondences. For the input image, the network predicts the object mask and the coordinates of the object’s front and back surfaces. Based on the predicted coordinates, we perform dense sampling of 3D coordinates between the front and back surfaces to construct ultra-dense 2D-3D correspondences, and use the RANSAC-PnP solver [14] to compute the object pose. Here, \tilde{Q}_f , \tilde{Q}_b , and \tilde{Q}_m represent the sets of 3D coordinates for the object’s front surface, back surface, and interior of the object, respectively. Their corresponding sets of 2D coordinates are denoted as \tilde{P}_f , \tilde{P}_b , and \tilde{P}_m .

[24] leverage neural networks to predict the 3D coordinates of the object’s front surface from 2D images, constructing denser 2D-3D correspondences that significantly enhance the pose estimation accuracy achieved by PnP. However, existing studies typically focus only on the front surface of the object, overlooking the potential benefits of incorporating the back surface and interior of the object. Therefore, this paper proposes utilizing neural networks to simultaneously predict the coordinates of both the object’s front and back surfaces, along with densely sampling 3D coordinates between these surfaces, further enhancing the density of the 2D-3D correspondences. By leveraging this ultra-dense 2D-3D correspondence, the RANSAC-PnP solver can achieve more precise object poses (see Fig. 1).

The accuracy of object pose estimation using the PnP algorithm heavily depends on the precision of the predicted surface coordinates. To enhance the accuracy of both front and back surface coordinate predictions, this paper proposes

*Corresponding author.

an efficient Hierarchical Continuous Coordinate Encoding (HCCE). HCCE encodes x , y , and z components of each surface coordinate separately and employs mirroring operations to transform these components into multi-level continuous codes. Neural networks are then used to predict these hierarchical continuous codes based on the input 2D image, enabling the calculation of the object’s front and back surface coordinates. Furthermore, during network training, our method calculates a separate histogram for each coordinate component, which records the proportion of mispredicted codes at different hierarchical levels. This reflects the learning difficulty of the network at various levels. Based on these error histograms, we adjust the weights of different hierarchical codes in the loss function, facilitating effective hierarchical learning. By utilizing weight adjustments based on multiple histograms rather than a single histogram [21], the stability of training and the accuracy of the predicted surface coordinates can be effectively enhanced.

Experimental results show that on the seven classic BOP core datasets [1, 4, 15, 25–28], our method outperforms SOTA RGB-based approach [21] by 2.4% in BOP score. When trained on RGB but tested on RGB-D data, it achieves a 4.7% improvement over the SOTA method [23]. Additionally, in the 2D segmentation task, our method surpasses the best existing approach [21] by 3.7%, further demonstrating its effectiveness.

Our contributions are summarized as follows:

- To the best of our knowledge, this study is the first to demonstrate a neural network’s ability to simultaneously predict both the front and back surfaces of an object for pose estimation.
- Our method constructs ultra-dense 2D-3D correspondences by densely and uniformly sampling 3D coordinates across the front and back surfaces, which effectively enhances pose accuracy.
- Our approach introduces a novel Hierarchical Continuous Coordinate Encoding (HCCE), which effectively enhances the accuracy of the predicted coordinates.
- We propose a hierarchical learning approach based on multiple histograms, which enables more efficient and stable training of networks based on HCCE.

2. Related Work

To estimate the pose of seen objects, early methods [26, 28–36] directly regressed 3D translation and rotation using neural networks. However, such approaches are less robust under occlusion or background clutter. To overcome these limitations, later works [6–13, 15–23] adopted a two-stage pipeline that first constructs 2D-3D correspondences and then computes the pose using a PnP algorithm [14], leading to more accurate pose estimation.

2.1. PnP-based Object Pose Estimation

In PnP-based object pose estimation, related work can be categorized into keypoint-based methods and surface coordinate-based methods. Keypoint-based methods [6–13] typically predefine a set of 3D keypoints on the object model, such as the corner points of the object’s 3D bounding box (BB8 [6]) or keypoints selected using farthest point sampling [37, 38] (PvNet [11]). These methods utilize neural networks to predict the 2D projections of these 3D keypoints on the input 2D image, thereby constructing sparse 2D-3D correspondences. To establish dense 2D-3D correspondences, surface coordinate-based methods [17–24, 39] predict the 3D coordinates of the object’s front surface corresponding to each 2D pixel coordinate on the input 2D image. The poses computed via the PnP algorithm [14] using these dense 2D-3D correspondences are generally more accurate than those derived from keypoint-based methods.

2.2. Surface Coordinate Encoding

To accurately predict the 3D coordinate of an object’s surface, surface coordinate-based methods have proposed various encoding techniques, generally categorized as continuous encoding [17, 19, 22] and discrete encoding [20, 21, 24]. In continuous encoding, methods such as CDPN [17], GDR-Net [22], and Pix2Pose [19] normalize the surface coordinates and utilize neural networks to predict the normalized coordinates corresponding to 2D pixel coordinates on the input 2D image. In discrete encoding, methods like DPOD [20] and ZebraPose [21] encode different regions of the object surface into discrete codes and then use neural networks to predict the discrete codes corresponding to 2D pixel coordinates on the input image. Subsequently, these methods decode the surface regions using the predicted discrete codes, transforming them into surface coordinates. For instance, when encoding object surface regions, DPOD [20] employs a two-channel UV map to segment the surface into different regions, while ZebraPose [21] uses clustering algorithms to encode the surface into hierarchical binary codes. This hierarchical encoding structure enables ZebraPose [21] to achieve SOTA performance among RGB-based methods.

The methods discussed above [17, 19–24] develop various encoding techniques to predict surface coordinates and have achieved accurate pose estimation. However, they share a common limitation: only the front surface of the object is considered for prediction. Although StereoPose [40] also predicts both front and back surface coordinates, it is specifically designed for transparent objects and relies on stereo images. In contrast, our method simultaneously predicts both front and back surface coordinates, and performs dense sampling between them to construct ultra-dense 2D-3D correspondences. Additionally, we introduce an effective encoding method, HCCE, which significantly outperforms

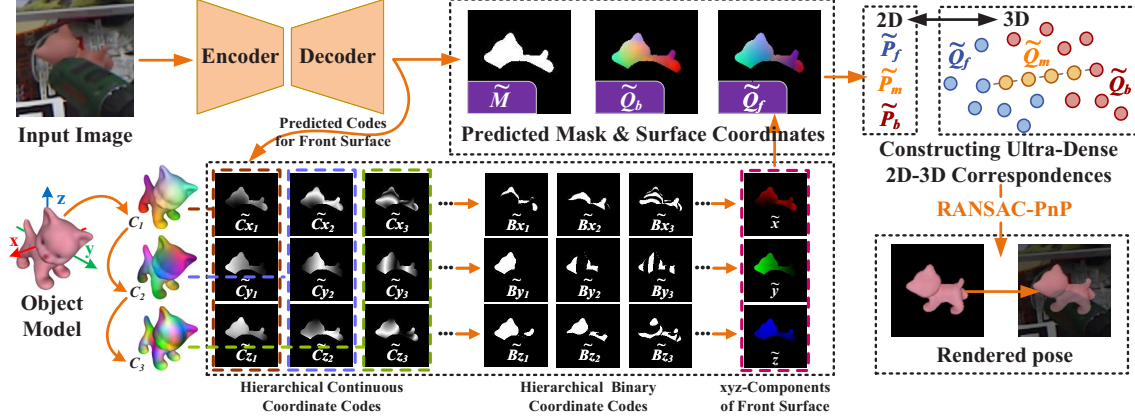


Figure 2. The inference pipeline. Our method begins by cropping a raw image based on the results of 2D object detection. The cropped image is then used as input to neural networks, which predict both the object mask and the coordinates of the front and back surfaces. To efficiently and accurately represent surface coordinates, we propose HCCE that encodes surface coordinates as multi-level codes, which are then learned by neural networks. During inference, the predicted multi-level continuous codes are first converted into binary codes, which are subsequently used to decode the surface coordinates (illustrated for the front surface coordinates in the figure). Using these predicted front and back surface coordinates, we densely sample 3D coordinates between the two surfaces to construct ultra-dense 2D-3D correspondences. Based on these ultra-dense correspondences, the method applies the RANSAC-PnP [14] to compute the object’s pose.

traditional hierarchical binary encoding [21] in surface coordinate prediction accuracy. Leveraging high-precision coordinates and ultra-dense 2D-3D correspondences, our approach achieves improved accuracy in pose estimation.

3. Methodology

Our method simultaneously predict the front and back surface coordinates through a neural network. Based on these predictions, dense sampling is performed to construct ultra-dense 2D-3D correspondences, leveraging potential benefits of the object’s back and intermediate regions. These correspondences are then used by the RANSAC-PnP algorithm [14] for pose estimation. To further enhance performance, we introduce HCCE to encode both surface coordinates, while a hierarchical learning approach improves training stability. The overall method is illustrated in Fig. 2.

Next, this section introduces the construction of ultra-dense 2D-3D correspondences and pose estimation (Section 3.1), the Hierarchical Continuous Coordinate Encoding method (Section 3.2), and the loss function for the neural network (Section 3.3).

3.1. Ultra-Dense 2D-3D Correspondences

To harness the potential benefits of incorporating the object’s back surface and interior, we construct Ultra-Dense 2D-3D correspondences. The method first utilizes a neural network to simultaneously predict the 3D coordinates of both the object’s front and back surfaces, denoted as \tilde{Q}_f and \tilde{Q}_b , respectively. Their corresponding 2D coordinates are represented as \tilde{P}_f and \tilde{P}_b , where $\tilde{P}_f = \tilde{P}_b$.

To further enrich the correspondences, we perform dense and uniform sampling between the front and back surfaces. Before sampling, a k-d tree is used to compute the average distance between the nearest point pairs in the predicted 3D coordinates, denoted as \bar{d} . Given the front and back 3D coordinates \tilde{q}_1 and \tilde{q}_2 corresponding to the same 2D coordinate, the number of interpolation points n is determined by:

$$n = \lfloor \frac{\|\tilde{q}_1 - \tilde{q}_2\|_2}{\bar{d}} \rfloor \quad (1)$$

Based on the sampling quantity n , uniform dense sampling between the front and back surfaces can be achieved using the following function:

$$s(\tilde{q}_1, \tilde{q}_2, a) = a \times \tilde{q}_1 + (1 - a) \tilde{q}_2 \quad (2)$$

where $a \in \left\{ \frac{t}{n+1} \mid t = 1, 2, \dots, n \right\}$.

The resulting densely sampled 3D points are denoted as \tilde{Q}_m with corresponding 2D coordinates \tilde{P}_m , where $\tilde{P}_m = \tilde{P}_f = \tilde{P}_b$. Ultimately, the Ultra-Dense 2D-3D correspondences are constructed as:

$$\tilde{P}_u = \left\{ \tilde{P}_b, \tilde{P}_f, \tilde{P}_m \right\}, \tilde{Q}_u = \left\{ \tilde{Q}_b, \tilde{Q}_f, \tilde{Q}_m \right\} \quad (3)$$

After constructing ultra-dense 2D-3D correspondences, we observe that multiple 3D points share the same 2D projection ($\tilde{P}_m = \tilde{P}_f = \tilde{P}_b$). Directly applying RANSAC-PnP [14] may therefore select multiple 3D points corresponding to the same 2D pixel in a single iteration, which can lead to unreliable pose estimates. To address this, we constrain each RANSAC-PnP iteration to sample only one 3D point

per 2D pixel. The RANSAC process is repeated 150 times, and in each iteration, the reprojection error of the predicted front and back 3D points is computed using a 2-pixel threshold. The pose with the lowest error is finally selected as the result.

3.2. Hierarchical Coordinate Encoding

For encoding object surface coordinates, hierarchical surface encoding [21] has demonstrated greater efficiency compared to non-hierarchical approaches [17, 19, 20, 22]. Existing approaches [21] typically use hierarchical binary codes to represent different regions of an object’s surface. However, experiments in the related work have not explored whether it is more effective to encode surface regions or coordinate components. To investigate this, we introduce Hierarchical Binary Coordinate Encoding (HBCE), which we find improves pose estimation accuracy over hierarchical binary surface encoding. Nevertheless, during the training of HBCE-based methods, we observed that the neural network faced challenges in learning binary codes near the edges of dark and light stripes, as illustrated in Fig. 3. To address this, we propose Hierarchical Continuous Coordinate Encoding (HCCE), which utilizes multi-level continuous codes in place of binary codes to eliminate these stripes. Next, we will introduce HBCE and HCCE in detail.

3.2.1. Hierarchical Binary Coordinate Encoding (HBCE)

In HBCE, each component of a surface coordinate is encoded as multi-level binary codes, where the coordinate components are normalized to the range of 0 to 1. The maximum of k corresponds to the number of pixels occupied by the object in the input image. The traditional binary encoding function for floating-point numbers is employed to encode coordinate components. Based on the binary encoding function at level i , the components x , y , and z of the surface coordinates can be encoded as binary code $Bx_{i,k}$, $By_{i,k}$, and $Bz_{i,k}$, respectively, as illustrated in Fig. 2.

Using neural networks, multi-level binary codes are predicted from the input 2D image. These predicted binary codes enable the calculation of the surface coordinate components. For example, the decoding of the component x can be expressed as:

$$x_k \approx \sum_{i=1}^8 2^{-k} \times Bx_{i,k} \quad (4)$$

Since excessively high encoding levels do not significantly enhance the precision of the represented coordinate components, this paper sets the upper limit of the encoding level to 8.

3.2.2. Hierarchical Continuous Coordinate Encoding (HCCE)

In HCCE, hierarchical continuous codes are generated through mirroring operations. For surface coordinate pre-

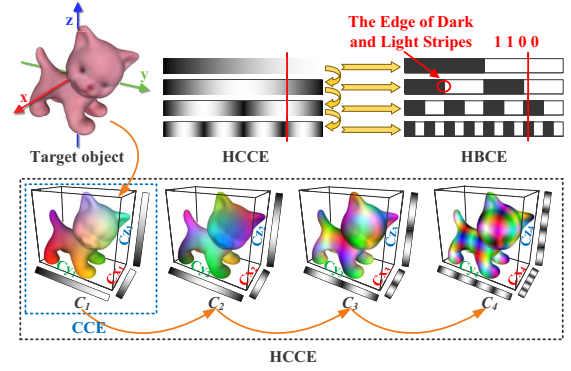


Figure 3. The hierarchical coordinate encoding method. In Hierarchical Binary Coordinate Encoding (HBCE), we observed that neural networks struggle to learn codes near the edges between dark and light stripes. To eliminate these stripes, we propose Hierarchical Continuous Coordinate Encoding (HCCE), which encodes coordinate components as multi-level continuous codes. Here, C_1 to C_4 represent the hierarchical continuous encoding of the object from the first to the fourth level, respectively. In addition, Continuous Coordinate Encoding (CCE) is an encoding method widely used by approaches such as CDPN [17], GDR-Net [22], and Pix2Pose [19].

diction, the method first predicts the hierarchical continuous codes from the input image. The continuous codes are subsequently converted into binary codes, which are then used to calculate the surface coordinates, following a process consistent with HBCE. Next, we will describe the specific steps involved in HCCE and how the continuous codes are converted into binary codes.

In encoding the surface coordinates, HCCE encodes each component of the coordinates into multi-level continuous codes. Given a coordinate component x_k , its continuous code at the first level is simply $Cx_{1,k} = f_1(x_k) = x_k$. To ensure the continuity of the higher-level codes, the $(i-1)$ -th level code is mirrored to compute the i -th level code, where $i > 1$, as follows:

$$Cx_{i,k} = f_i(x_k) = \begin{cases} f_{i-1}(2x_k), & x_k < 0.5 \\ f_{i-1}(2 - 2x_k), & x_k \geq 0.5 \end{cases} \quad (5)$$

In this process, $f_{i-1}(\cdot)$ and $f_i(\cdot)$ are the continuous encoding functions for the $(i-1)$ -th and i -th levels, respectively. When $x_i < 0.5$, the continuous code at the i -th level is obtained by copying the continuous code from the $(i-1)$ -th level; when $x_i \geq 0.5$, the continuous code at the i -th level is derived by mirroring the continuous code from the $(i-1)$ -th level.

In decoding the surface coordinates, the method first converts the continuous codes into binary codes and then calculates the surface coordinates according to Eq. (4). In the conversion from continuous codes to binary codes, a binarization function $g(t)$ is defined such that $g(t) = 0$ when

$t < 0.5$ and $g(t) = 1$ when $t \geq 0.5$. Thus, the continuous code at the first level $Cx_{1,k}$ can be converted to the binary code $Bx_{1,k} = g(Cx_{1,k})$.

Subsequently, based on the binary code $Bx_{i-1,k}$ from the $(i-1)$ -th level, the continuous code at the i -th level $Cx_{i,k}$ can be converted to $Bx_{i,k}$:

$$Bx_{i,k} = \begin{cases} g(Cx_{i,k}), Bx_{i-1,k} = 0 \\ 1 - g(Cx_{i,k}), Bx_{i-1,k} = 1 \end{cases} \quad (6)$$

Here, $Bx_{i-1,k} = 1$ indicates that $Cx_{i,k}$ was generated through mirroring, which means that the binarized $g(Cx_{i,k})$ does not correspond to the hierarchical binary code $Bx_{i,k}$. Therefore, mirroring is reversed in the decoding process, resulting in $Bx_{i,k} = 1 - g(Cx_{i,k})$.

3.3. The Loss Function of Neural Network

In neural network training, the proposed method utilizes a neural network to learn the object mask and the hierarchical continuous codes of front and back surface coordinates. With the number of encoding levels for each coordinate component set to 8, the neural network outputs 49 channels, comprising 3×8 channels for the front surface coordinates, 3×8 channels for the back surface coordinates, and 1 channel for the object mask. Consequently, the loss function of networks can be divided into mask loss L_M and hierarchical loss (front loss L_{xyz}^{Front} and back loss L_{xyz}^{Back}). The total loss of the neural network can be expressed as: $L = L_M + \gamma \times (L_{xyz}^{Front} + L_{xyz}^{Back})$, where γ is the weighting coefficient for the hierarchical loss in the total loss function.

3.3.1. Mask Loss

During training, the proposed method utilizes the L1 metric to compute the loss L_M between the ground-truth object mask M and the predicted object mask \widetilde{M} , $L_M = \sum_{i=1}^m \|M_i - \widetilde{M}_i\|_1$. \widetilde{M}_i denotes the output on a pixel and m

describes the number of outputs. Pixels with $\widetilde{M} > 0$ correspond to the predicted object mask, while those with $\widetilde{M} \leq 0$ are considered background. The total number of pixels classified as the object mask is denoted as n .

3.3.2. Hierarchical Loss

During training, the proposed method uses histograms to quantify the proportion of erroneous predictions at different hierarchical levels. This approach effectively measures the difficulty the neural network encounters in learning across various levels. Based on these error histograms, the method dynamically adjusts the weights of the continuous codes in the loss function at each level. This adjustment facilitates hierarchical learning and enhances the stability of neural network training.

For different coordinate components, the method computes the corresponding histograms separately. For in-

stance, consider the calculation of the histogram $h_{f,x}$ for the x component of the front surface coordinates. The process begins by obtaining the predicted hierarchical continuous code $\widetilde{Cx}_{f,i}$ at the i -th level, which is then converted into the corresponding hierarchical binary code $\widetilde{Bx}_{f,i}$. Next, the proportion of erroneous binary codes at the i -th level, denoted as $r_{f,x,i}$, is calculated. Consequently, the intensity $h_{f,x,i}$ of the x component in the histogram $h_{f,x}$ at the i -th level can be computed as follows:

$$h_{f,x,i} = \exp(\sigma \cdot \min\{r_{f,x,i}, 0.5 - r_{f,x,i}\}) \quad (7)$$

To reduce the weight of codes at levels with a high proportion of errors in the loss function, $r_{f,x,s}$ is truncated at 0.5. An exponential term $\exp(\cdot)$ is used to smooth the abrupt changes in intensity caused by the truncation. Additionally, σ is a constant employed to adjust the intensity range in the histogram. By normalizing $h_{f,x,i}$ in the histogram $h_{f,x}$, the weight $w_{f,x,i}$ for the i -th level in the loss function can be obtained.

Based on the weights $w_{f,x,i}$ for each level, the loss function L_x^{Front} can be expressed as follows:

$$L_x^{Front} = \sum_{i=1}^8 \left(w_{f,x,i} \cdot \sum_{j=1}^n \|Cx_{f,i,j} - \widetilde{Cx}_{f,i,j}\|_1 \right) \quad (8)$$

Here, $\widetilde{Cx}_{f,i,j}$ denotes the j -th code at the i -th level, where the upper limit for j is the area n of the predicted mask \widetilde{M} .

Similarly, the histograms for the y and z components of the object's front surface, along with the histogram-based loss functions L_y^{Front} and L_z^{Front} , can also be calculated. Thus, the hierarchical loss for the object's front surface can be expressed as $L_{xyz}^{Front} = L_x^{Front} + L_y^{Front} + L_z^{Front}$. Likewise, the hierarchical loss L_{xyz}^{Back} for the object's back surface can be computed in the same manner.

4. Experiments

In this section, we first introduce the experimental setup, including implementation details, datasets, and evaluation metrics. Next, we conduct ablation studies on the encoding method, loss function weight adjustment, and 2D-3D correspondences. Finally, we compare our experimental results with SOTA methods [21, 23] on the BOP website across seven classic BOP core datasets [1, 4, 15, 25–28].

4.1. Experimental Setup

4.1.1. Implementation Details

Network Setup. The network architecture in this paper is adapted from ZebraPose [21], with a ResNet34 [41] version for ablation experiments and an EfficientNet-B4 [42] version for comparison experiments. Both network versions take RGB images as input and output the object mask along

with the 3D surface coordinates. The input and output resolutions are 256×256 and 128×128, respectively.

Preprocessing of Training Data. The DZI strategy, proposed by CDPN [17], is employed to simulate randomly detected 2D bounding boxes. Additionally, the random enhancement strategy, introduced by GDR-Net [22], is incorporated to simulate real-world RGB image noise.

Training Setup. Separate neural networks are trained for each object, with each training process taking approximately 24 hours on a single NVIDIA RTX 4090 GPU. For the ablation experiments, the ResNet34-based version is used, with a maximum of 100k training epochs, a batch size of 24, and a constant learning rate of 0.0002. For the comparative experiments, the EfficientNet-B4 version is used, with a maximum of 600k training epochs and a batch size of 12. Pose estimation for a single object within a bounding box takes approximately 30 ms during inference.

Label Preparation. Following the BOP toolkit’s practice of rendering front surface depth maps using VisPy to generate coordinate labels, we adopt a similar approach with PyOpenGL. By setting the depth test to `GL_LESS` and `GL_GREATER`, we respectively obtain front and back surface depth maps, which are then used to generate the corresponding labels.

4.1.2. Datasets

The performance of the method was tested on seven classic BOP core datasets: LM-O [15], YCB-V [28], IC-BIN [43], TUD-L [2], HB [27], T-LESS [25], and ITODD [4]. During training, we used the PBR [44, 45] training data provided by BOP [46, 47] for all seven datasets. For the YCB-V, T-LESS, and TUD-L datasets, real data was also used in training. During testing, we utilized the detection sequences and 2D detection results provided by the BOP 2023 challenge [46, 47].

4.1.3. Evaluation Metrics

ADD(-S) Accuracy [48]. In the ablation experiments, we use the widely adopted ADD(-S) metric to evaluate the accuracy of object pose predictions. It includes ADD for non-symmetric objects and ADD-S for rotationally symmetric objects. For objects without rotational symmetry, ADD computes the average distance between the transformed vertices of the object based on the predicted pose and the ground truth pose. For objects with rotational symmetry, ADD-S calculates the average distance between the nearest vertices, rather than the same vertex. A predicted pose is considered correct if the ADD(-S) error is less than 10% of the object’s diameter. ADD(-S) accuracy, also referred to as average recall (AR), represents the proportion of correctly predicted poses across all test sequences. Additionally, we also report the Area Under the Curve (AUC) for ADD(-S) and ADD-S [28], with a maximum threshold of 10 cm.

Coordinate Accuracy. In addition to reporting pose accuracy, the ablation experiments also assess the accuracy of predicted surface coordinates. A prediction is considered correct if the Euclidean distance between the predicted and ground truth coordinates is below a given threshold. Coordinate accuracy is defined as the ratio of correctly predicted coordinates to the total number of predicted coordinates.

BOP Score. In the comparison experiments, the BOP score of the proposed method is presented on seven classic BOP core datasets [1, 4, 15, 25–28]. The BOP score is the average of three evaluations: VSD (Vertex Symmetric Distance), MSSD (Mean Symmetric Surface Distance), and MSPD (Mean Symmetric Projection Distance) [1, 2]. These evaluations have complex definitions, and the specific details can be found on the BOP website.

4.2. Ablation Experiments

Ablation experiments on encoding methods and loss weight adjustments were conducted on the IC-BIN dataset [43]. The results, presented in Fig. 4 and Tab. 1, include the prediction accuracy of object surface coordinates and object pose accuracy, with a primary focus on ADD(-S) accuracy. Additionally, ablation studies on 2D-3D correspondences were performed on three datasets: LM-O [15], TUD-L [2], and IC-BIN [43], with the average precision (AP) of BOP scores reported in Tab. 2. To ensure fair comparisons, OpenCV’s RANSAC-PnP [14] is consistently used across all methods, regardless of the original PnP algorithms employed by other approaches.

4.2.1. Encoding Methods

The current SOTA RGB-based pose estimation method, ZebraPose [21], employs hierarchical binary surface encoding, which represents different regions of the object surface as multi-level binary codes. However, we find that directly encoding coordinate components is more efficient and accurate than encoding surface regions. Experimental results show that our proposed hierarchical binary surface encoding (HBCE + f) improves surface coordinate accuracy (see Fig. 4) and enhances pose estimation accuracy by 0.97% (see Tab. 1). However, it was observed that the network struggled to learn the binary codes near the edges of the dark and light stripes, as shown in Fig. 3. To eliminate these stripes, this paper introduces hierarchical continuous coordinate encoding (HCCE), which encodes coordinate components as multi-level continuous codes. Compared to HBCE, HCCE (HCCE + f(h3)) further improves the prediction accuracy of object surface coordinates and increases pose estimation accuracy by 5.13%. This indicates that the neural network finds it easier to learn continuous codes than binary codes, which effectively enhances the accuracy of pose estimation.

Table 1. The ablation experiments on the IC-BIN dataset [43]. h0 signifies the absence of histogram-based weight adjustment; h1 denotes weight adjustment using a single histogram; h3 represents weight adjustment using multiple histograms; f indicates the prediction of only the object’s front surface. Cell colors indicate column-wise values from low (dark blue) to high (dark red).

Method		AR of ADD(-S)	AR of ADD-S	AUC of ADD(-S)	AUC of ADD-S
Zebra-Pose	f	55.85	61.82	72.91	76.40
CCE	f	55.42	62.15	73.58	78.14
HBCE	f	56.82	63.50	73.15	76.56
HCCE	f(h0)	61.35	67.00	76.35	79.82
	f(h1)	60.44	65.23	74.92	78.68
	f(h3)	61.95	68.30	77.67	81.11

4.2.2. Weight Adjustment

In hierarchical learning, higher-level codes are more complex; for example, C_4 in Fig. 3 is much more complex than C_1 . The more complex the code, the more difficult it is for neural networks to learn. To measure the learning difficulty at different levels, ZebraPose [21] uses a single histogram to track the proportion of incorrect codes at each level. Subsequently, ZebraPose adjusts the weights of different levels in the loss function to facilitate hierarchical learning. However, hierarchical learning based on a single histogram is not suitable for HCCE. Compared to not adjusting the weights (HCCE + f(h0)), using dynamic weights based on a single histogram (HCCE + f(h1)) reduces the prediction accuracy of object surface coordinates, resulting in a 0.91% decrease in ADD(-S) accuracy. Therefore, our method generates a separate histogram for each coordinate component and proposes dynamic weights based on multiple histograms. In Fig. 5, we illustrate the histogram intensities and weight coefficients at different encoding levels for the x -component of the 3D coordinates on the front surface. As training progresses, the peak of the weight coefficients




Figure 4. Accuracy of surface coordinates. Percentage of correctly predicted coordinates under thresholds of 2%, 5%, and 10% of the object diameter.

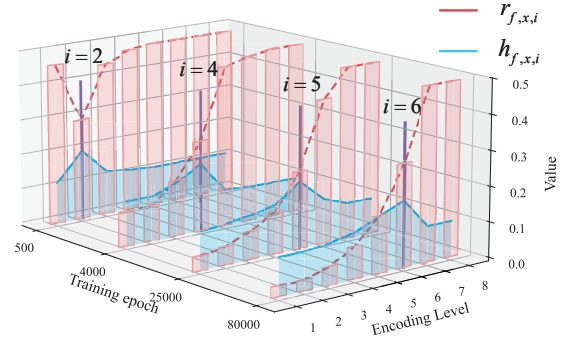


Figure 5. Weight adjustments across different training epochs.

shifts from lower to higher encoding levels, enabling hierarchical learning from low- to high-level encodings. Weight adjustment based on multiple histograms (HCCE + f(h3)) achieves more accurate coordinate predictions, improving the ADD(-S) accuracy by 0.6% compared to not adjusting the weights. This demonstrates that weight adjustment based on multiple histograms can enhance the stability of network training and improve the accuracy of pose estimation.

4.2.3. 2D-3D correspondences

In this section, we examine the average precision (AP) of BOP scores for all objects in three datasets (LM-O [15], TUD-L [2], and IC-BIN [43]) to assess how different surface information impacts pose estimation accuracy.

Tab. 2 presents the impact on the AP of BOP scores across three datasets when 2D-3D correspondences are con-

Table 2. Effect of front and back surface information on the Average Precision (AP) of BOP scores. bf denotes the simultaneous prediction of both the object’s front and back surfaces, while bfu refers to joint prediction of the object’s front and back surfaces, along with uniform sampling between them. Cell colors indicate row-wise values from low (dark blue) to high (dark red).

Dataset	Object	f	b	bf	bfu
LM-O [15]	ape	84.9	85.7	85.7	85.8
	can	94.8	94.8	95.2	95.5
	cat	87.2	85.9	87.6	87.7
	driller	93.1	93.1	94.5	94.5
	duck	78.7	82.1	80.5	82.4
	eggbox	56.6	57.1	56.0	57.1
	glue	84.4	83.9	85.8	86.1
	holepuncher	82.2	75.8	80.1	82.3
	TUD-L [2]	obj01	95.0	95.4	95.7
obj02	85.4	86.3	86.1	86.8	
obj03	93.2	92.2	92.6	93.3	
IC-BIN [43]	obj01	62.9	58.5	63.6	63.7
	obj02	70.1	65.9	70.7	71.5
Mean		82.2	81.3	82.6	83.3

Table 3. Comparison with SOTA methods on seven classic BOP core datasets in 6D localization and 2D segmentation tasks. Scores are expressed as percentages.

Task	6D Localization			2D Segmentation			6D Localization		
Test Data	RGB			RGB			RGB-D		
Method	GPose [49]	ZebraPose [21]	Ours	DLZDet [50]	ZebraPose [21]	Ours	PFA [51]	GDRNPP [23]	Ours
$AR(AP)_{LM-O}$	69.9	72.9	75.5	46.0	51.6	52.7	79.2	79.2	80.5
$AR(AP)_{T-LESS}$	79.9	82.1	85.6	58.4	72.1	76.8	84.9	87.2	87.9
$AR(AP)_{TUD-L}$	83.1	85.0	86.9	73.2	71.8	76.7	96.3	93.6	94.4
$AR(AP)_{IC-BIN}$	62.6	59.2	63.5	31.6	49.3	51.2	70.6	70.2	72.4
$AR(AP)_{ITODD}$	46.0	50.4	54.2	23.9	46.2	51.3	52.6	58.8	73.4
$AR(AP)_{HB}$	87.6	92.2	91.9	60.0	68.9	76.1	86.7	90.9	93.1
$AR(AP)_{YCB-V}$	80.9	82.8	83.9	66.9	73.1	74.5	89.9	83.4	91.1
$AR(AP)_{Core}$	72.9	74.9	77.3	51.4	61.9	65.6	80.0	80.5	84.7

structured based on the front surface, the back surface, and the intermediate 3D coordinates sampled between them. In the LM-O dataset, the contributions of front and back surface information to the improvement in BOP scores are roughly equivalent; in the TUD-L dataset, front surface information is less effective than back surface information; whereas in the IC-BIN dataset, front surface information proves more advantageous. These experimental results indicate that pose estimation based on front surface information and that based on back surface information each have their own strengths and weaknesses. Therefore, when using only front surface information yields unsatisfactory performance, incorporating back surface information can effectively enhance the accuracy of pose prediction.

Furthermore, the experimental results (see the ‘‘Mean’’ row in Tab. 2) show that simultaneously leveraging both front and back surface information increases the BOP scores by 0.4% and 1.3%, respectively. Among the 13 objects evaluated, 8 objects achieved BOP scores that were equal to or higher when using both surfaces compared to using only front or only back information. This suggests that for most objects, a hybrid approach that combines 2D-3D correspondences from both the front and back surfaces is more effective in extracting predictive information, thereby improving pose estimation performance.

To further capitalize on the predicted front and back surface coordinates, we uniformly sample new 3D coordinates between the two surfaces to construct ultra-dense 2D-3D correspondences. Compared to using only front or back surface information, the adoption of ultra-dense correspondences improves the BOP scores by 1.1% and 2.0%, respectively, further demonstrating that this approach more effectively integrates information from both surfaces, thereby enhancing the accuracy of pose estimation.

4.3. Comparison with the SOTA Methods

The comparison experiments were conducted on the seven classic BOP core datasets [1, 4, 15, 25–28], and the BOP

scores of our method can be viewed on the BOP website. To ensure a fair comparison, only methods based on the BOP 2023 official detection results were considered. The scores for the methods listed in Tab. 3, including ZebraPose [21], GPose [49], MegaPose [52], and GDRNPP [22, 23], are also obtained from the BOP website. Notably, if a method achieves higher scores with alternative detection results, those scores are not reported here.

When tested on RGB data, our method outperforms the best existing approach [21] by 2.4% in the BOP score, as shown in Tab. 3. Among the methods trained on RGB data but tested with RGB-D data, the current top-performing approach refines poses using Iterative Closest Point (ICP [53]). In contrast, our method leverages the more efficient FoundationPose [54], leading to a 4.7% improvement in the BOP score (see Tab. 3). Furthermore, in the 2D segmentation task, our method surpasses existing approaches [21] by 3.7%, further demonstrating its effectiveness.

5. Conclusion

In this study, we propose an ultra-dense 2D-3D correspondence to enable more accurate pose estimation via RANSAC-PnP. It is constructed by simultaneously predicting the coordinates of both the front and back object surfaces, followed by dense sampling of intermediate 3D points. We further introduce an efficient Hierarchical Continuous Coordinate Encoding (HCCE) and a weight adjustment approach based on multiple histograms, which together facilitate accurate predictions of both surfaces. Compared to SOTA methods [21, 23], our method achieves competitive scores and rankings on the BOP Challenge. However, due to varying learning difficulties across objects, the trained network is inherently object-specific, making it infeasible to train a single unified model for multiple objects.

Acknowledgments

This work was supported by the National Natural Science Foundation of China under grant No. 52375487 and the special fund of Jiangsu Province for Key Research and Development No. BE2023041. Additionally, the third author carried out this work during her undergraduate and master's studies at Southeast University under the supervision of Dr. Luo, prior to commencing her doctoral studies at Purdue University.

References

- [1] Tomáš Hodaň, Martin Sundermeyer, Bertram Drost, Yann Labbé, Eric Brachmann, Frank Michel, Carsten Rother, and Jiří Matas. Bop challenge 2020 on 6d object localization. In *ECCVW*, pages 577–594, 2020. 1, 2, 5, 6, 8
- [2] Tomáš Hodaň, Frank Michel, Eric Brachmann, Wadim Kehl, Anders Glent Buch, Dirk Kraft, Bertram Drost, Joel Vidal, Stephan Ihrke, Xenophon Zabulis, Caner Sahin, Fabian Manhardt, Federico Tombari, Tae-Kyun Kim, Jiří Matas, and Carsten Rother. Bop: Benchmark for 6d object pose estimation. In *ECCV*, pages 19–34, 2018. 1, 6, 7
- [3] Stephen Tyree, Jonathan Tremblay, Thang To, Jia Cheng, Terry Mosier, Jeffrey Smith, and Stan Birchfield. 6-dof pose estimation of household objects for robotic manipulation: An accessible dataset and benchmark. In *IROS*, pages 13081–13088, 2022. 1
- [4] Bertram Drost, Markus Ulrich, Paul Bergmann, Philipp Hartinger, and Carsten Steger. Introducing mvtec itodd-a dataset for 3d object recognition in industry. In *ICCVW*, pages 2200–2208, 2017. 1, 2, 5, 6, 8
- [5] Prithviraj Banerjee, Sindi Shkodrani, Pierre Moulon, Shreyas Hampali, Fan Zhang, Jade Fountain, Edward Miller, Selen Basol, Richard Newcombe, Robert Wang, Jakob Julian Engel, and Tomas Hodan. Introducing hot3d: An egocentric dataset for 3d hand and object tracking, 2024. Preprint at <https://arxiv.org/abs/2406.09598>. 1
- [6] 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 *ICCV*, pages 3828–3836, 2017. 1, 2
- [7] Kartik Gupta, Lars Petersson, and Richard Hartley. Cullnet: Calibrated and pose aware confidence scores for object pose estimation. In *ICCVW*, pages 2758–2766, 2019.
- [8] Zelin Zhao, Gao Peng, Chengkun Li, Haoyu Wang, Cewu Lu, and Hao-Shu Fang. Estimating 6d pose from localizing designated surface keypoints, 2018. Preprint at <https://arxiv.org/abs/1812.01387>.
- [9] Chen Song, Jiaru Song, and Qixing Huang. Hybridpose: 6d object pose estimation under hybrid representations. In *CVPR*, pages 431–440, 2020.
- [10] Markus Oberweger, Mahdi Rad, and Vincent Lepetit. Making deep heatmaps robust to partial occlusions for 3d object pose estimation. In *ECCV*, pages 119–134, 2018.
- [11] Sida Peng, Yuan Liu, Qixing Huang, Xiaowei Zhou, and Huijun Bao. Pvnnet: Pixel-wise voting network for 6dof pose estimation. In *CVPR*, pages 4561–4570, 2019. 2
- [12] Bo Chen, Tat-Jun Chin, and Marius Klimavicius. Occlusion-robust object pose estimation with holistic representation. In *WACV*, pages 2929–2939, 2022.
- [13] Bugra Tekin, Sudipta N. Sinha, and Pascal Fua. Real-time seamless single shot 6d object pose prediction. In *CVPR*, pages 292–301, 2018. 1, 2
- [14] Vincent Lepetit, Francisc Moreno-Noguer, and Pascal Fua. Ep n p: An accurate o (n) solution to the p n p problem. *IJCV*, 81:155–166, 2009. 1, 2, 3, 6
- [15] Eric Brachmann, Alexander Krull, Frank Michel, Stefan Gumhold, Jamie Shotton, and Carsten Rother. Learning 6d object pose estimation using 3d object coordinates. In *ECCV*, pages 536–551, 2014. 1, 2, 5, 6, 7, 8
- [16] 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 *CVPR*, pages 3364–3372, 2016.
- [17] Zhigang Li, Gu Wang, and Xiangyang Ji. Cdpn: Coordinates-based disentangled pose network for real-time rgb-based 6-dof object pose estimation. In *ICCV*, pages 7678–7687, 2019. 2, 4, 6
- [18] He Wang, Srinath Sridhar, Jingwei Huang, Julien Valentin, Shuran Song, and Leonidas J Guibas. Normalized object coordinate space for category-level 6d object pose and size estimation. In *CVPR*, pages 2642–2651, 2019.
- [19] Kiru Park, Timothy Patten, and Markus Vincze. Pix2pose: Pixel-wise coordinate regression of objects for 6d pose estimation. In *ICCV*, pages 7668–7677, 2019. 2, 4
- [20] Sergey Zakharov, Ivan Shugurov, and Slobodan Ilic. Dpod: 6d pose object detector and refiner. In *ICCV*, pages 1941–1950, 2019. 2, 4
- [21] 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 *CVPR*, pages 6738–6748, 2022. 2, 3, 4, 5, 6, 7, 8
- [22] Gu Wang, Fabian Manhardt, Federico Tombari, and Xiangyang Ji. Gdr-net: Geometry-guided direct regression network for monocular 6d object pose estimation. In *CVPR*, pages 16611–16621, 2021. 2, 4, 6, 8
- [23] Xingyu Liu, Ruida Zhang, Chenyangguang Zhang, Bowen Fu, Jiwen Tang, Xiquan Liang, Jingyi Tang, Xiaotian Cheng, Yukang Zhang, Gu Wang, and Xiangyang Ji. Gdrnpp, 2022. https://github.com/shanice-1/gdrnpp_bop2022. 2, 5, 8
- [24] Shun Iwase, Xingyu Liu, Rawal Khirodkar, Rio Yokota, and Kris M Kitani. Repose: Fast 6d object pose refinement via deep texture rendering. In *ICCV*, pages 3303–3312, 2021. 1, 2
- [25] 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 *WACV*, pages 880–888, 2017. 2, 5, 6, 8
- [26] Thanh-Toan Do, Ming Cai, Trung Pham, and Ian Reid. Deep-6dpose: Recovering 6d object pose from a single rgb image, 2018. Preprint at <https://arxiv.org/abs/1802.10367>. 2

- [27] Roman Kaskman, Sergey Zakharov, Ivan Shugurov, and Slobodan Ilic. Homebreweddb: Rgb-d dataset for 6d pose estimation of 3d objects. In *ICCVW*, pages 2767–2776, 2019. **6**
- [28] Yu Xiang, Tanner Schmidt, Venkatraman Narayanan, and Dieter Fox. Posecnn: A convolutional neural network for 6d object pose estimation in cluttered scenes. In *RSS*, 2018. **2, 5, 6, 8**
- [29] Catherine Capellen, Max Schwarz, and Sven Behnke. Conposecnn: Dense convolutional 6d object pose estimation. In *VISIGRAPP*, pages 162–172, 2020.
- [30] Alex Kendall, Matthew Grimes, and Roberto Cipolla. Posenet: A convolutional network for real-time 6-dof camera relocalization. In *ICCV*, pages 2938–2946, 2015.
- [31] Torsten Sattler, Qunjie Zhou, Marc Pollefeys, and Laura Leal-Taixe. Understanding the limitations of cnn-based absolute camera pose regression. In *CVPR*, pages 3302–3312, 2019.
- [32] Alex Kendall and Roberto Cipolla. Geometric loss functions for camera pose regression with deep learning. In *CVPR*, pages 5974–5983, 2017.
- [33] Siddharth Mahendran, Haider Ali, and René Vidal. 3d pose regression using convolutional neural networks. In *CVPRW*, pages 2174–2182, 2017.
- [34] Hao Su, Charles R Qi, Yangyan Li, and Leonidas J Guibas. Render for cnn: Viewpoint estimation in images using cnns trained with rendered 3d model views. In *ICCV*, pages 2686–2694, 2015.
- [35] Zhigang Li, Yuwang Wang, and Xiangyang Ji. Monocular viewpoints estimation for generic objects in the wild. *IEEE Access*, 7:94321–94331, 2019.
- [36] Thanh-Toan Do, Trung Pham, Ming Cai, and Ian Reid. Real-time monocular object instance 6d pose estimation. 2018. **2**
- [37] Jingtao Li, Jian Zhou, Yan Xiong, Xing Chen, and Chaitali Chakrabarti. An adjustable farthest point sampling method for approximately-sorted point cloud data. In *SiPS*, pages 1–6, 2022. **2**
- [38] Meng Han, Liang Wang, Limin Xiao, Hao Zhang, Chenhao Zhang, Xiangrong Xu, and Jianfeng Zhu. Quickfps: Architecture and algorithm co-design for farthest point sampling in large-scale point clouds. *TCAD*, 42(11):4011–4024, 2023. **2**
- [39] Zong-Wei Hong, Yen-Yang Hung, and Chu-Song Chen. Rdpn6d: Residual-based dense point-wise network for 6dof object pose estimation based on rgb-d images. In *CVPRW*, pages 5251–5260, 2024. **2**
- [40] 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, 2023. **2**
- [41] Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. Deep residual learning for image recognition. In *CVPR*, pages 770–778, 2016. **5**
- [42] Mingxing Tan and Quoc Le. Efficientnet: Rethinking model scaling for convolutional neural networks. In *ICML*, pages 6105–6114, 2019. **5**
- [43] Andreas Drounoglou, Rigas Kouskouridas, Sotiris Malassiotis, and Tae-Kyun Kim. Recovering 6d object pose and predicting next-best-view in the crowd. In *CVPR*, pages 3583–3592, 2016. **6, 7**
- [44] Maximilian Denninger, Dominik Winkelbauer, Martin Sundermeyer, Wout Boerdijk, Markus Wendelin Knauer, Klaus H Strobl, Matthias Humt, and Rudolph Triebel. Blenderproc2: A procedural pipeline for photorealistic rendering. *JOSS*, 8(82):4901, 2023. **6**
- [45] Tomáš Hodaň, Vibhav Vineet, Ran Gal, Emanuel Shalev, Jon Hanzelka, Treb Connell, Pedro Urbina, Sudipta N Sinha, and Brian Guenter. Photorealistic image synthesis for object instance detection. In *ICIP*, pages 66–70, 2019. **6**
- [46] Tomas Hodan, Martin Sundermeyer, Yann Labbe, Van Nguyen Nguyen, Gu Wang, Eric Brachmann, Bertram Drost, Vincent Lepetit, Carsten Rother, and Jiri Matas. Bop challenge 2023 on detection segmentation and pose estimation of seen and unseen rigid objects. In *CVPRW*, pages 5610–5619, 2024. **6**
- [47] Martin Sundermeyer, Tomáš Hodaň, Yann Labbe, Gu Wang, Eric Brachmann, Bertram Drost, Carsten Rother, and Jiří Matas. Bop challenge 2022 on detection, segmentation and pose estimation of specific rigid objects. In *CVPRW*, pages 2785–2794, 2023. **6**
- [48] 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 *ACCV*, pages 548–562, 2013. **6**
- [49] Ruida Zhang, Ziqin Huang, Gu Wang, Xingyu Liu, Chenyangguang Zhang, and Xiangyang Ji. Gpose, 2023. https://bop.felk.cvut.cz/method_info/373/. **8**
- [50] Anonymity. Dlzdet, 2022. https://bop.felk.cvut.cz/method_info/304/. **8**
- [51] Yinlin Hu, P. Fua, and Mathieu Salzmann. Perspective flow aggregation for data-limited 6d object pose estimation. In *ECCV*, 2022. **8**
- [52] Yann Labbé, Lucas Manuelli, Arsalan Mousavian, Stephen Tyree, Stan Birchfield, Jonathan Tremblay, Justin Carpentier, Mathieu Aubry, Dieter Fox, and Josef Sivic. MegaPose: 6D Pose Estimation of Novel Objects via Render & Compare. In *CoRL*, 2022. **8**
- [53] Paul J Besl and Neil D McKay. Method for registration of 3-d shapes. In *SPIE*, volume 1611, pages 586–606, 1992. **8**
- [54] Bowen Wen, Wei Yang, Jan Kautz, and Stan Birchfield. Foundationpose: Unified 6d pose estimation and tracking of novel objects. In *CVPR*, pages 17868–17879, 2024. **8**
- [55] Tsung-Yi Lin, Michael Maire, Serge Belongie, James Hays, Pietro Perona, Deva Ramanan, Piotr Dollár, and C. Lawrence Zitnick. Microsoft coco: Common objects in context. In David Fleet, Tomas Pajdla, Bernt Schiele, and Tinne Tuytelaars, editors, *ECCV*, pages 740–755, 2014.
- [56] Martin Sundermeyer, Tomáš Hodaň, Yann Labbé, Gu Wang, Eric Brachmann, Bertram Drost, Carsten Rother, and Jiří Matas. Bop challenge 2022 on detection, segmentation and pose estimation of specific rigid objects. In *2023 IEEE/CVF Conference on Computer Vision and Pattern Recognition Workshops (CVPRW)*, pages 2785–2794, 2023.