






ContourPose: Monocular 6-D Pose Estimation Method for Reflective Textureless Metal Parts

Zaixing He , Senior Member, IEEE, Quanzhi Li, Xinyue Zhao , Jin Wang , Huarong Shen, Shuyou Zhang , and Jianrong Tan 

Abstract—Pose estimation is an essential technology for industrial robots to perform precise gripping and assembly. The state-of-the-art deep learning-based approach uses an indirect strategy, i.e., first finding local correspondence between the 2-D image and 3-D model, and then using the perspective-n-point and RANSAC methods to calculate the poses of ordinary objects. However, the metal parts in industry are reflective and textureless, making it difficult to identify distinguishable point features to establish 2-D–3-D correspondences. To address this problem, in this article, we propose a novel deep learning based two-stage method for pose estimation of reflective textureless metal parts, which accurately estimates the target pose using monocular red green blue (RGB) images. Since contours play an important role in both keypoints prediction and pose estimation stages, our method is named ContourPose. First, an additional contour decoder is adopted to implicitly constrain the keypoints prediction in the former stage, which improves the accuracy of the keypoints prediction. Then, the predicted contour of the previous stage is taken as geometric prior that is used to iteratively solve for the optimal pose. Experiments indicate that the proposed approach for reflective textureless metal parts has a significant improvement over the state-of-the-art approaches.

Index Terms—Deep learning, reflective textureless metal parts, six degrees of freedom (6-D) pose estimation.

I. INTRODUCTION

SIX degrees of freedom (6-D) pose estimation, i.e., recovering the rotation and translation of an object in the 3-D Euclidean space, is a key task for robotic vision. With the fast development of intelligent manufacturing, pose estimation of industrial objects has become a crucial technology for tasks such as part gripping, unit assembly, and human–machine collaboration [1].

Manuscript received 20 March 2023; revised 19 May 2023; accepted 9 June 2023. Date of publication 20 July 2023; date of current version 4 October 2023. This article was recommended for publication by Associate Editor Huichan Zhao and Editor Sven Behnke upon evaluation of the reviewers' comments. This work was supported in part by the National Natural Science Foundation of China under Grant 52275514, Grant 52275547, and Grant 52175032 and in part by the Zhejiang Provincial Natural Science Foundation of China under Grant LY21E050021. (Corresponding author: Xinyue Zhao.)

Zaixing He, Quanzhi Li, Xinyue Zhao, Jin Wang, Shuyou Zhang, and Jianrong Tan are with the School of Mechanical Engineering, State Key Lab of Fluid Power and Mechatronic Systems, Zhejiang University, Hangzhou 310058, China (e-mail: zaixinghe@zju.edu.cn; liquanzhi@zju.edu.cn; zhaoxinyue@zju.edu.cn; dwjcom@zju.edu.cn; zsy@zju.edu.cn; egi@zju.edu.cn).

Huarong Shen is with Zhejiang Feihang Intelligent Technology, Co., Ltd., Huzhou 313200, China (e-mail: noodlesh@163.com).

This article has supplementary material provided by the authors and color versions of one or more figures available at <https://doi.org/10.1109/TRO.2023.3290300>.

Digital Object Identifier 10.1109/TRO.2023.3290300

Metal parts are the main components of machines and industrial products. Although many pose estimation methods have achieved promising results in recent years, the reflective and textureless nature of metal parts poses a great challenge to them due to their heavy reliance on surface features.

The existing methods to handle this task can be classified as RGB image-based or red green blue–depth (RGB-D) image-based methods. In terms of ordinary textured and nonreflective objects, the RGB-D based method achieves higher accuracy and robustness because they use additional depth information to estimate the pose. However, RGB-D cameras have limited ability to accurately capture depth information for non-Lambertian materials, such as metal parts, which often produce fragmented depth images [2]. In addition, it is expensive to acquire high precision point clouds of objects with industrial-grade depth camera cameras or 3-D scanners, while the cost of consumer-grade depth cameras is considerably low. Therefore, it is more appropriate to use RGB monocular cameras for 6-D pose estimation of such objects.

The RGB-based methods also have limitations. Traditional pose estimation methods use various types of effective image feature descriptors, such as SIFT [3], SURF [4], and ORB [5], to establish 2-D–3-D correspondences based on the similarity of these descriptors. Then, the target pose can be obtained by solving the Perspective-n-Point (PnP) problem. However, these methods are only applicable to objects with rich textures. Reflective textureless metal parts barely show apparent gradient variations in images and these descriptors will not be able to find an exact correspondence. To address this issue, some methods use geometric features such as lines [6], moments [7], or edges [8] as templates to retrieve the best matching result. However, these methods cannot achieve high accuracy because the matched result is chosen from a limited number of samples. Increasing the number of templates to enhance accuracy requires more computational resources and may result in matching errors when too many templates are used, consequently leading to pose estimation errors.

In recent years, deep learning has been widely used in pose estimation for its significant improvements compared to traditional methods. For example, some methods use convolutional neural networks (CNNs) to regress 2-D keypoints, establish 2-D–3-D correspondences, and then use the PnP algorithm to calculate the pose. If only sparse keypoints, e.g., several or tens of keypoints, are selected to predict, such as BB8 [9], PVNet [10], and YOLO-6D [11], these methods will be referred to

as pose estimation methods using sparse keypoint predictions. On the contrary, if it needs to establish pixelwise 2-D–3-D correspondences of the objects in the image, such as CDPN [12] and DPOD [13], these methods will be instead indicated as pose estimation methods using dense keypoint predictions. Although these methods work well on public benchmarks (e.g., LINEMOD [14] and YCB-Video [15]) of ordinary objects, they have limitations in measuring metal parts. For example, PVNet predicts the location of 2-D keypoints by pixelwise voting, but the textureless surface of metal parts does not provide effective semantic information. In addition, due to the reflective nature of metal parts, slight changes in shooting angles can lead to large color differences in the images, which can cause incorrect correspondences with dense methods, such as DPOD and CDPN.

Although the current method of dense keypoints prediction achieves superior performance for the pose estimation of ordinary objects, the surface of reflective and textureless metal parts can provide little semantic information. To address this issue, we propose a novel method, called ContourPose. The proposed method can be divided into two stages.

In the first stage, the proposed method uses the object contours to implicitly constrain the prediction of keypoints. We call this stage ContourNet. In the second stage, the proposed method utilizes the contours predicted in the previous stage as geometric priors to achieve higher accuracy in solving the pose. Experiments show that it improves the accuracy of the final pose estimation compared to the widely used RANSAC + PnP [16] method.

In summary, this work has the following contributions.

- 1) A novel pose estimation approach is proposed for reflective textureless metal parts, which achieves superior performance on industrial metal parts.
- 2) A new deep learning framework is proposed for the pose estimation of reflective textureless metal parts. This framework adopts an additional contour decoder to implicitly constrain the prediction of keypoints. The proposed method does not need a predetector such as YOLO [17] or RetinaNet [18] to locate the target object.
- 3) A new method is proposed to iteratively solve the optimal pose using the contour as a geometric prior. The proposed method improves the accuracy of final pose estimation compared to widely used RANSAC-based methods.

II. RELATED WORK

This section introduces recent related deep learning-based 6-D pose estimation methods dealing with RGB images, which can be divided into the following four categories.

A. Direct Pose Regression Methods

Given an input image, these methods output the 6-D pose of the target object directly using CNNs; these methods are also referred to as single-stage methods in some studies. Kendall et al. [19] introduces a CNN architecture that regresses the camera pose directly from an RGB image, which is similar to object pose estimation. SSD-6D [20] extends the popular single-shot multibox detector (SSD) [21] object detection framework to cover the full 6-D pose space by adding a

translation and orientation regression module. However, SSD-6D predicts rotation by scoring the discrete views, so it needs to add the pose refinement process to get better results. PoseCNN [15] decouples translation and rotation prediction to predict 6-D object pose. The 3-D rotation of the object is represented using unitary quaternions. However, predicting 3-D rotations is difficult because the nonlinearity of the rotation space makes CNNs less generalizable. The methods mentioned above all rely on the feature extraction ability of CNNs, which cannot achieve good results on textureless industrial objects and are difficult to apply to more challenging reflective metal parts.

DenseFusion [22] designs a dense pixel-level fusion method that integrates features of RGB data and point clouds in an appropriate way to obtain good results. But these point cloud based methods are not applicable to metal parts because it is difficult to obtain reliable point clouds for this kind of objects due to their reflective nature. EfficientPose [23] achieves end-to-end multiobject pose estimation by extending an efficient and accurate 2-D object detection method EfficientDet [24]. However, pose estimation methods based on object detection usually depend on the results of prior detection, but the features of reflective metal parts are different under different lighting and viewing angles, making it difficult to train a detector for these reflective parts.

B. Template Matching-Based Methods

These methods discretize the rotation space into sampled templates. Given an input image, they match the closest template from the database, thus bypassing the explicit labeling of pose ambiguity. Some traditional methods [14] create features by hand, which is very time consuming. Spares templated-based (STB) method [6] uses high-level geometric features and the correlation of straight contours to estimate the 6-D pose of metal parts. But this method is sensitive to background noise. AAE [25] trains an autoencoder with synthetic data, which encodes the rotation of the object into a latent space. There, pose ambiguity is implicitly handled through similar latent codes for symmetric poses. Multipath-AAE [26] extends the AAE [25] to multiple objects. However, these methods have some limitations, such as the domain gap between real and synthetic data and the regularity of the latent space. Although Wen et al. [27] introduce edges to supervise the regularity of the latent space, addressing the domain gap problem remains challenging, particularly for metal parts, due to the difficulty in rendering realistic reflections of these parts. The method [28] proposes an approach to estimate pose by using YOLO [17] and LINEMOD [14]. But the accuracy of this method is not high and can only be used for the bin picking task. GFI [29] proposes a generative feature-to-image network to generate edge templates for pose estimation by matching the most consistent edges. This method achieves good results for pose estimation of metal parts but is computationally slow.

C. Render and Compare Methods

Another category of methods is based on rendering synthetic images of the object model under different poses and comparing them with the observed image. These methods rely on iterative procedures that consider previous estimate of a pose p_k and

predict an update Δp_{k+1} that will make current internal mesh state to better fit the observed object. For example, BB8 [9] designs an iterative refinement process step to compare the input image and the rendering of the initial pose object using another CNN, thereby improving the prediction of 2-D projections. Given an initial pose estimation, DeepIM [30] designs a network to iteratively refine the pose by matching the rendered image against the observed image, and gradually improves the matching result by iteratively updating the rendered image. CosyPose [31] optimizes the pose hypotheses using a multiview consistency module that resolves conflicts and inconsistencies and estimates the camera viewpoints. However, these methods are not suitable for reflective metal parts because they rely on renderers, which have difficulty in simulating the surface color and texture changes of reflective parts under different lighting conditions [32]. This variation results in significant differences between rendered and real images, making it difficult to make effective comparisons.

D. Keypoints Localization-Based Methods

These methods follow a two-stage pipeline for pose estimation. They establish 2-D–3-D correspondence using a neural network in the first stage, and then calculate the target pose using the PnP and RANSAC methods in the second stage. The methods for keypoints localization can be divided into two categories, one is to select only several keypoints, e.g., several or tens of keypoint for prediction, and this type of method is classified as the sparse methods. The other category is to establish pixelwise 2-D–3-D correspondences of object in an image, and this type of method is classified as the dense methods.

BB8 [9] uses CNN as a keypoints detector to output 2-D coordinates of eight corner points of the object’s 3-D bounding box. YOLO-6D [11] uses a YOLO [17] architecture to regress keypoints. These methods output the keypoints directly through a fully connected layer, which depend on the distribution of the training images and lead to problems such as poor spatial generalization and proneness to overfitting. Inspired by the success of 2-D human pose estimation [33], another category of methods output pixelwise heatmaps of keypoints to improve accuracy. The method [34] predicts heatmaps of multiple small blocks independently and then overlays them to obtain accurate and robust keypoints. PVNet [10] regresses pixelwise vectors point at the keypoints and uses these vectors to vote for the location of the keypoints. HybridPose [35] extends the approach of PVNet [10] by utilizing a hybrid intermediate representation to express different geometric information in the input image, including keypoints, edge vectors, and symmetry correspondences. All the above-mentioned methods preselect only a few keypoints, classified as sparse methods. These methods are designed for ordinary objects, and the selected keypoints usually do not have semantic information. It is not appropriate for metal parts with many semantic points, such as corner points and circle centers.

The methods where each pixel is used as a keypoint to generate the corresponding prediction, we will refer to this as a dense method. On some irregularly curved objects, the accuracy and robustness of dense prediction are usually higher than that of sparse prediction. Pix2Pose [40] outputs a 3-D coordinate

for each pixel of the object in the picture, thus establishing dense 2-D–3-D correspondences. DPOD [13] considers that it is difficult to directly predict 3-D coordinates, so it predicts UV maps instead, and then establishes 2-D–3-D correspondence through UV maps. CDPN [12] decouples the prediction of rotation and translation; rotation is obtained by PnP calculation and translation is determined by 2-D bounding box and depth estimation, but the method requires a predetector. PSGMN [37] utilizes the information from the 3-D model. First extracting the features of the 3-D model using a graph convolutional neural network, and then establishing the 2-D–3-D correspondences through a pseudo-Siamese neural network.

However, due to the reflective nature of metal parts, the colors of the object’s pixels in the image vary significantly with the lighting conditions and the shooting angle. Moreover, the surfaces of metal parts are less distinguishable than that of textured objects. For this reason, methods regressing dense keypoints might be less effective with this category of objects. To handle this problem, we first predict sparse semantic points, and then we use dense contour points for validation in the pose estimation stage to iteratively solve for the optimal pose. Our method combines the advantages of sparse and dense prediction.

III. PROPOSED APPROACH

In this article, we propose a novel 6-D pose estimation method for reflective textureless metal parts in industry. The task of 6-D pose estimation is to detect the object and estimate its 3-D rotation and translation in the given image. The 6-D pose represents a rigid transformation of the object coordinate system to the camera coordinate system, where R denotes the 3-D rotation and t denotes the 3-D translation.

At once at least three pairs of 2-D–3-D correspondences are determined, $[R|t]$ can be computed by solving the PnP problem. Therefore, the key step of pose estimation is to obtain the precise 2-D–3-D correspondences. We estimate the object pose using a two-stage pipeline: The first stage is a neural network that predicts a heatmap representing the 2-D keypoints that we call “ContourNet.” The second stage is an iterative pose optimization algorithm that uses the object contours as prior information. We call this stage Pose Evaluation using Contour as a Prior (PECP). We observe that the contour of objects with sharp edges is invariant to different viewing angles and lighting conditions. In the first stage, with the constraint of contour, the keypoints are constrained to be within the region of the target object in the image, which allows our method to achieve high accuracy without a predetector. In the second stage, the contour of the previous stage is used as a geometric prior to iteratively solve the optimal pose.

The architecture of ContourPose, which consists of two stages, is shown in Fig. 1.

A. Keypoint Prediction With Implicit Constraints on Contour

1) *Network Framework*: The basic framework of the ContourNet is shown in Fig. 1. In this stage, the input is an RGB image of an object and a modified FCN model is proposed to encode the image. To predict keypoints, the first step is to locate the position of the target object in the image. Some methods

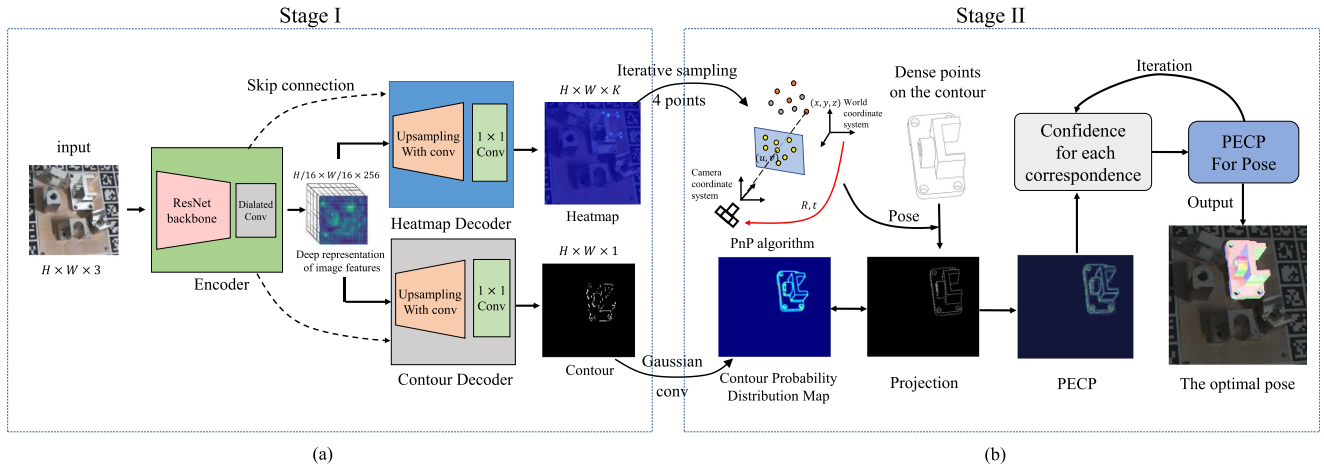


Fig. 1. Overview of ContourPose. (a) Keypoints prediction with implicit constraints on contour, named ContourNet. In this stage, we predict the keypoints and contours of the target object. (b) Pose estimation using contour as a priori. In this stage, the contour predicted in the previous stage is used as geometric priors to eliminate outlier poses, and finally output the optimal pose in the result set.

[12], [13] use the result of the predictor as input so that the input image contains only the target object. Other methods [10], [35], [37] use masks and then locate keypoints within the mask. We observe that for metal parts, the contour has richer semantic information compared to mask. The contour can act as a mask by restricting keypoints within the contour and also establish implicit constraints with the predicted semantic points. Therefore, we add a decoder for predicting contour to implement implicit constraints on keypoints prediction. The predictions of contour and keypoints heatmaps cotrain the parameters of the encoder of this network, so the prediction of contour actually affects the prediction of keypoints as well.

Many methods such as PVNet [10] and PSGMN [37] will output both the keypoints information and the mask using only one decoder. In other words, one part of the output tensor represents the keypoints and the other part represents the mask. However, our experiments in Section IV demonstrate that this approach leads to the prediction of keypoints being completely dependent on the contour. This means that the contours have a very strong constraint on the keypoints. For example, in some extreme cases, imprecise predictions of contours can result in completely erroneous predictions of keypoints. Therefore, to decouple the contour and the keypoints, we use one decoder to output the contour and another decoder to output the heatmap of the keypoints. This structure ensures that the contour has an implicit constraint on the keypoints and prevents the keypoints from being overly dependent on the contour.

Fig. 2 provides more detailed information about ContourNet. Residual blocks [38] are utilized to convolute the input features. To downsample the inputs and extract features at different scales, we employ max pooling and stride convolutional layers. Dilated convolution [39] can expand the receptive field and capture multiscale contextual information. The bilinear upsampling layer is employed to increase the resolution of the low-resolution feature map to the same size as the original input image. The skip connection concatenates the feature maps of the residual blocks before downsampling and after upsampling to preserve the raw information of the image.

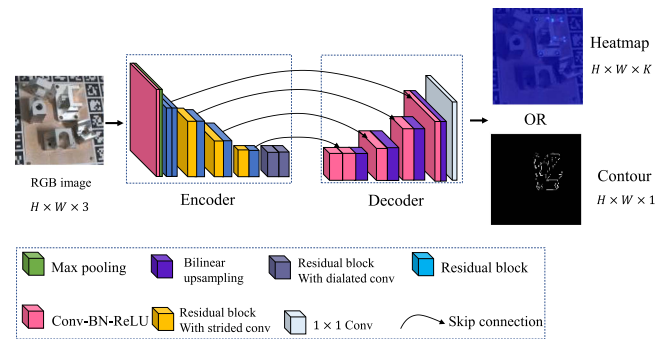


Fig. 2. ContourNet: keypoints prediction network with implicit constraints on contours. Input a picture containing the target object, after the encoding-decoding process, output the heatmap and contour of the keypoints of the object. The design of decoder for heatmap and contour is identical, only the output channels are different.

Given C classes of objects and K keypoints for each class, our network takes the $H \times W \times 3$ RGB image as input, then processes it with a fully convolutional architecture with ResNet-18 [38] as the backbone. When the resolution of the feature map is downsampled to $H/16 \times W/16$, we discard the subsequent downsampling and convolution steps of ResNet-18. Instead, we add dilated convolutions to improve the receptive field of the network. After that, we repeatedly apply upsampling and convolution the feature map until its size is the same as that of the input image, i.e., $H \times W$. Finally, we apply 1×1 convolution on the feature map to obtain the contour or keypoints heatmap. We implement a skip connection between the decoder and the encoder at the same resolution of the feature map. The predictions of keypoints and contours are obtained by two independent decoders.

ContourNet can be seen as a regression model for two tasks, one is the regression of the keypoints heatmap, which we can consider as a submodel of ContourNet, denoted by Φ , and the other is the regression of the contour of the target object, which is denoted by Λ . The heatmap follows a Gaussian distribution, and each pixel of the heatmap represents the probability of the keypoint being present at that location. The pixel with the highest

probability is considered as the keypoint. To learn the heatmap, we use ℓ_2 loss. The corresponding loss function is defined as

$$l(w_E, w_{HD}) = \sum_{k=1}^K \ell_2(\tilde{H}_k(w_E, w_{HD}) - H_k) \quad (1)$$

$$\tilde{H}_k = \Phi_{w_E, w_{HD}}(I) \quad (2)$$

where w_E represents the parameters of the encoder, w_{HD} represents the parameters of the heatmap decoder, K is the number of keypoints for each object, H represents the groundtruth of the heatmap, \tilde{H} represents the predicted heatmap, and I represents the input image.

In the contour regression task, there is a problem of positive and negative category imbalance in learning because the contour of the object accounts for a small part. For this reason, we use the weighted cross entropy as the loss function. The corresponding loss function is defined as

$$l(w_E, w_{CD}) = -\beta \sum_{p \in Y^+} \log(\tilde{Y}_p = 1; w_E, w_{CD}) - (1 - \beta) \sum_{p \in Y^-} \log(\tilde{Y}_p = 0; w_E, w_{CD}) \quad (3)$$

$$\tilde{Y} = \Lambda_{w_E, w_{CD}}(I) \quad (4)$$

where w_{CD} represents the parameters of the contour decoder, \tilde{Y} is the predicted contour, and p denotes each pixel in the contour map. $\beta = |Y^-|/|Y^+ + Y^-|$ and $1 - \beta = |Y^+|/|Y^+ + Y^-|$, where $|Y^+|$ and $|Y^-|$ denote the edge and nonedge in the contour ground truth. The final loss function of the ContourNet is defined as

$$l(w) = l(w_E, w_{HD}) + \rho l(w_E, w_{CD}) \quad (5)$$

where w is the parameter of the whole network, which contains the encoder w_E , the heatmap decoder w_{HD} , and the contour decoder w_{CD} . This shows that the contour and keypoints heatmap regression tasks jointly train the network encoder, while their respective decoders do not interfere with each other. ρ is a hyperparameter to balance the two parts of the loss.

2) *Keypoints Generation*: Keypoints need to be defined based on the model of the 3-D object. One simple approach is to directly select the eight corner points of the 3-D bounding box as keypoints. However, most of these points are not on the objects and as such are less suited as prediction targets [13], [15]. Common methods use the farthest point sampling (FPS) [40] algorithm to select several points at the farthest Euclidean distance on the model, such as PVNet [10]. However, these methods are designed for ordinary object datasets, such as LINEMOD [14] and YCB-Video [15], because most of these objects are curved and it is difficult to manually select keypoints. In contrast, Pavlakos [41] proposes to select semantic points as keypoints, which is very suitable for metal parts, but the article does not mention the specific details of how to select these points.

In fact, selecting appropriate keypoints is crucial for accurate pose estimation. Unlike ordinary curved objects, metal parts have more semantic points, such as corner points of contours, the center of circles, midpoints of lines and arcs, among others. Fig. 3 illustrates an intuitive and effective algorithm

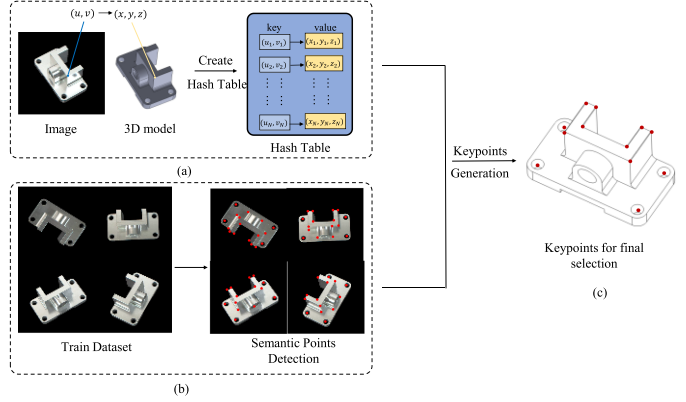


Fig. 3. Keypoints generation process. (a) Establish 2-D-3-D correspondence of candidate semantic points and store them in a hash table, N represents the number of candidate semantic points. (b) Iterate through the training images and record the frequency of candidate points with the semantic point detection algorithm. (c) K points with the highest confidence are chosen as keypoints.

proposed for generating keypoints for metal parts. For each class of parts, we traverse all its training images $TS = \{ts_i | i = 1, 2, \dots, n\}$ and then build a 2-D-3-D hash table for each image $H = \{h_i | i = 1, 2, \dots, n\}$, where n denotes the number of all training images for each part. This hash table contains the 2-D-3-D correspondence of candidate semantic points S_{cand} . After that, semantic point detection algorithms are applied to each image to detect semantic points, such as circle centers, corner points, and midpoints. Specifically, Hough transform is used to detect circle centers and midpoints, while Shi-Tomasi corner detection [42] is used to detect corner points. The semantic points that appear most frequently in the perspective of the training images are selected. Using the established hash table, we locate the 3-D points corresponding to the 2-D semantic points, and finally select K points into S_{pred} as keypoints. The keypoints selected by our method are evenly dispersed on the object surface, which makes the PnP algorithm more stable. It is worth noting that some previous methods have selected a fixed number of keypoints, but since each metal part has a different shape and its semantic points are not the same, K is determined according to the part shape in our method. If the object has less than eight semantic points, our method uses the FPS algorithm to define the keypoints. Algorithm 1 shows the flow of the keypoints generation algorithm.

B. Pose Estimation Using Contours as Geometric Priors

After obtaining the keypoints in the first stage, a common approach is to iteratively solve for the optimal pose by eliminating the incorrect correspondence based on the RANSAC and PnP methods. However, since RANSAC only uses the distribution of points to distinguish between inner and outer points, the RANSAC-based method works well for dense points but may not be applicable when there are few points. To solve this issue, we propose a method that utilizes the contour as geometric priors to iteratively solve for the optimal pose.

Our approach is shown in Fig. 4. First, the proposed method randomly samples four points from K points and uses the EPnP [43] algorithm to calculate a temporary pose T in each iteration.

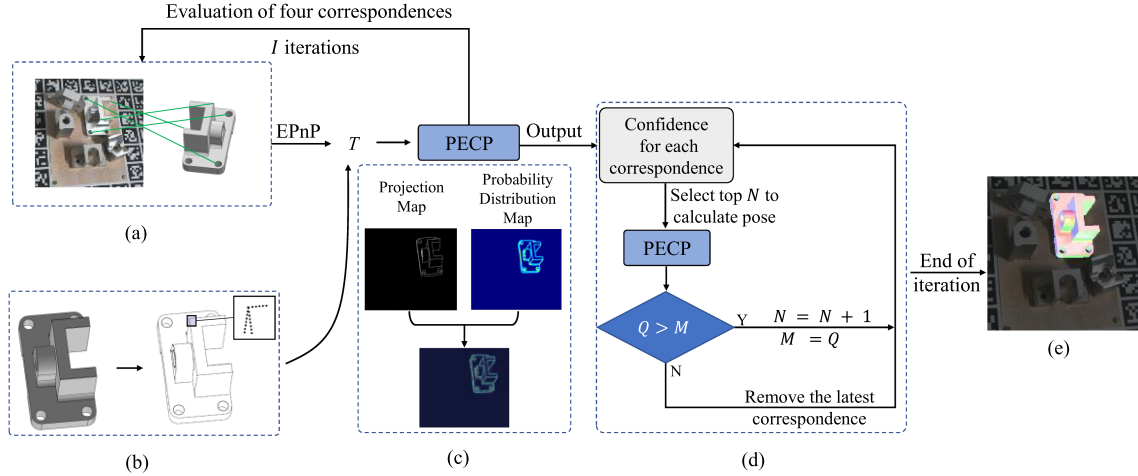


Fig. 4. Pose estimation using contours as geometric priors. (a) Randomly sample four correspondences. (b) Sampling dense contour points from the 3-D model. (c) PECP. (d) Progressive sampling to solve the pose. (e) Output the optimal pose.

Algorithm 1: Keypoints Generation Process.

Input: S_{cand}, TS
Output: S_{pred}

- 1: **if** $Len(S_{cand}) \leq 8$ **then**
- 2: $S_{pred} \leftarrow Use_FPS_Algorithm$
- 3: Return S_{pred}
- 4: **end if**
- 5: **for** i **in** n **do**
- 6: $h_i \leftarrow Create_Hashtable$
- 7: $S_{2d} \leftarrow Projection(S_{cand})$
- 8: **for** j **in** $Len(S_{cand})$ **do**
- 9: $h_i[S_{2d_j}] \leftarrow S_{cand_j}$
- 10: **end for**
- 11: **end for**
- 12: $arr \leftarrow Create_Array(Len(S_{cand}))$
- 13: **for** ts_i **in** TS **do**
- 14: $S_{2d}^s \leftarrow Semantic_Point_Detection(ts_i)$
- 15: **for** j **in** $Len(S_{2d}^s)$ **do**
- 16: **if** $h_i[S_{2d_j}^s] \neq null$ **then**
- 17: $C_{3d} \leftarrow h_i[S_{2d_j}^s]$
- 18: $idx \leftarrow Get_Index(C_{3d})$
- 19: $arr[idx]++$
- 20: **end if**
- 21: **end for**
- 22: **end for**
- 23: $S_{pred} \leftarrow Get_Top_K_Points(arr)$
- 24: Return S_{pred} ;

To validate our pose estimation, we sample dense points on the contour of the 3-D model, which are referred to as contour validation points S , as shown in Fig. 4(b). Specifically, S_i represents the set of all 3-D contour points. $S = \{s_i | i = 1, 2, \dots, m\}$, $s_i = [x_i, y_i, z_i, 1]^T$, where m represents the number of points. We calculate the image coordinate system coordinates P of these points using the image projection x model

$$P = \kappa[R|t]S \quad (6)$$

$$\kappa = \begin{bmatrix} f_x & 0 & c_x \\ 0 & f_y & c_y \\ 0 & 0 & 1 \end{bmatrix} \quad (7)$$

where P denotes the projection of the contour validation points S on the 2-D image, which is referred to as the projection map. $P = \{p_i | i = 1, 2, \dots, m\}$, $p_i = [u_i, v_i, 1]^T$, u, v denotes its position in the pixel coordinate system. κ denotes the intrinsic parameters of the camera, which represents the transformation between the camera coordinate system to the image coordinate system. R and t denote the 3-D rotation and translation of the object coordinates, and $[R|t]$ represents the object's pose.

The accuracy of the pose is positively associated with the overlap of the projection map and the contour image from ContourNet. We refer to this algorithmic process as PECP. Specifically, we convolve the contours predicted in the previous stage. The convolution function is defined as

$$G(x, y) = \frac{1}{2\pi\sigma^2} e^{-\frac{x^2+y^2}{2\sigma^2}} \quad (8)$$

$$H(u, v) = (G * \tilde{Y})(u, v) = \sum_x \sum_y G(x, y) \tilde{Y}(u-x, v-y). \quad (9)$$

We use a Gaussian convolution kernel, where the size of the convolution kernel size is 3, $\delta = 9$. The probability distribution map is modeled using a Gaussian function, where the closer a pixel is to the contour, the higher its probability. With a given pose, the confidence score can be calculated using the PECP method

$$\mu_T = \sum_{p \in P(T)} H(u_i(p), v_i(p)), \quad (10)$$

where μ represents the confidence score of the pose, H is the probability distribution map, and P is the projection of the contour validation points. If the confidence score is higher than the threshold that we indicate with θ , we assign this score to the

TABLE I
SYMBOLS AND DESCRIPTION IN THE POSE ESTIMATION STAGE

Name	Description
T	Temporary pose
S	Set of 3D contour points
P	Projection of 3D contour points in 2D plane
m	Number of contour points
μ	Confidence of temporary pose ρ
θ	Threshold of correct pose
I	Iterations
M	Maximum confidence of the pose obtained by progressive sampling
Q	Confidence of current pose
λ	A set of correspondence which is ordered by PECP

four correspondences which compute the pose

$$\lambda_i = \begin{cases} \lambda_i + \mu_T - \theta, & \mu_T > \theta \\ \lambda_i, & \mu_T \leq \theta \end{cases} \quad (11)$$

where $\lambda = \{\lambda_i | i = 1, 2, \dots, K\}$ represents each corresponding confidence. The value of θ is determined based on the number of contour validation points, and in our experiments, it was set to $1/3 m$. We select four points from the correspondence set obtained in the previous stage to calculate a temporary pose T . Next, we calculate the confidence score μ_T of pose T using PECP and (11), and assigned this score to the four correspondences used to compute the pose. After I iterations, the correspondence set sorted by confidence is obtained. The number of iterations I can be calculated as follows:

$$I = \frac{\log(1 - \text{Pr})}{\log(1 - r^4)} \quad (12)$$

where r denotes the probability that the sampled correspondences are the correct ones. It is assumed that only four correspondences are correct in the extreme cases, $r = 4/K$. Pr is the probability of being correct at least once after I iterations. After the iterations are completed, an initial pose is calculated with the four highest confidence correspondences. PECP is used to calculate the confidence of this pose, denoted as M . Then the next highest confidence point is selected to solve for the pose, and (11) is used to calculate its confidence, denoted as Q . If $Q > M$, the point is kept and the value of M is updated; otherwise, this point is discarded. When all the points in the set λ are sampled, the optimal pose is obtained. All symbols in this stage can be found in Table I. Our method improves the robustness and accuracy of pose estimation compared to the RANSAC-based method.

IV. EXPERIMENTAL RESULTS

In this section, we present the performance of ContourPose compared to other models for 6-D pose estimation methods of metal parts and describe the ablation experiments.




Fig. 5. Training dataset. Ten parts with different shapes are included.




Fig. 6. Test dataset with four background scenes. (a) Black background. (b) Black background with texture. (c) Simulated rust background. (d) Reflective metal background.

A. Experimental Setup

1) *Dataset*: We demonstrate our method on two datasets: a dataset of reflective, textureless metal parts created by our team, and the T-LESS [44] dataset. The reflective metal dataset includes ten metal parts with varying shapes, as shown in Fig. 5. Each part comprises 660 training images, and each image contains only one object part. The test dataset comprises images with multiple parts stacked up randomly, including some unrelated parts, which increases the difficulty of pose estimation for the target object. The scenes in the test dataset are based on real industrial scenes and can be divided into four cases: black background, black textured background, rust background, and reflective metal background, as shown in Fig. 6. Each type of background is shot under four lighting cases, including bright artificial lighting, bright natural lighting, dim lighting, and artificial lighting. Each category contains four scenes with different lighting and different parts distribution, so we tested our method on a total of 16 different scenes. The T-LESS dataset [44] contains 30 objects that lack color and texture information. Moreover, most test images exhibit obvious occlusion or stacking.

The dataset used for evaluation resembles the actual environment of an industrial site and presents a challenging task for pose estimation. Obtaining good results on these datasets would demonstrate the feasibility of our approach in real industrial settings.

2) *Data Augmentation*: Since there were only 660 training images in reflective metal dataset, we rendered 6600 realistic reflective and textureless synthetic images of the metal parts with Blender [10]. Our network was trained at a ratio of 1:10 between real and synthetic images. To create these synthetic images, we first rendered the 3-D model surface to the glossy metal part in the Blender GUI, and then distributed the sampled cameras with viewpoints in a Fibonacci arrangement, with the camera




Fig. 7. Synthetic picture of reflective metal parts. (a) Real images in the training dataset. (b) Synthetic image rendered with blender.

pitch and azimuth angles covering the range of possible poses. After that, we arrange multiple light sources randomly and vary the light intensity randomly to simulate various scenarios of possible reflections of metal parts. The rendered image is shown in Fig. 7(b), and we render it as similar as possible to the real image in Fig. 7(a). To prevent overfitting, we employed a cut-and-paste strategy during training [45], i.e., the target object was extracted and pasted onto a random background, which was randomly sampled from SUN397 [46]. We also added more data augmentation to the original image including random cropping during training, random scaling, color dithering, and 3-D rotation.

3) *Generating Contour Groundtruth*: Since our network is required to regress the contours, we need the groundtruth of the contours to calculate the loss. However, the groundtruth of the contours is not provided within the dataset. Therefore, we use OpenGL to render the synthetic image corresponding to each training image, and then apply the Canny edge detection algorithm to extract the contours and obtain the ground truth.

4) *Training Strategy*: We set the initial learning rate as 0.1 and halve it every 20 epochs. All models are trained for 150 epochs using the AdamW optimizer with 2 Nvidia RTX 3090 GPUs. The weight decay of the optimizer is set to 0.1. The hyperparameter ρ in the loss function is set to 100 to balance the magnitudes of the two parts in the loss function. We train a specific model for each metal part the same as in [10] and [37].

B. Metric

We evaluate our method using three standard metrics: the pointwise mean 2-D projection metric [47], the average 3-D distance of model points (ADD) metric [15], and the average R/t error of the valid poses. The T-LESS dataset specifies that the dataset is evaluated with the e_{VSD} metric [48]. Therefore, our experiments on the T-LESS dataset are evaluated with e_{VSD} .

2-D Projection metric: This metric computes the mean distance in the 2-D image between the projections of the 3-D mesh model from the estimated pose and the ground truth pose. It is generally accepted that if this distance is less than 5 pixels, the pose is considered correct. However, this assessment criterion is not rigorous enough for accurate attitude estimation in industrial environments. Thus, we gradually reduced the constraint from 5 pixels to 0 pixels and compared the performance variations of each method.

ADD(-S) metric: This metric computes the mean distance between two transformed model points using the estimated pose




Fig. 8. Ablation studies with iterations and threshold. (a) Accuracy and computational speed under different iterations, where the value of θ is set to $1/3 m$. (b) Accuracy under different threshold θ , where the number of iterations T is calculated by (12).

and the groundtruth pose. It is claimed that the estimated pose is correct if the distance is less than 10% of the model diameter. For symmetric objects, we use the ADD-S metric, where the mean distance is computed based on the closest point distance.

R/t error metric: We specify that if the pose is correct under the ADD metric, it is a valid pose, and the R/t error metric is intended to quantify the actual error in the valid pose, where R error denotes the error α, β, γ in 3-D rotation, while the t error denotes the error x, y, z in 3-D translation.

e_{VSD} metric: This metric is proposed by Hodan et al. [48], and it evaluates the visible surface discrepancy (VSD) between the target pose and the predicted pose. This metric was used in most experiments evaluating the T-LESS dataset. In [48], an estimated pose with $e_{\text{VSD}} < 0.3$ is defined as a valid pose.

C. Parameter Analysis

We conduct the parameter analysis experiment to investigate the effect of two parameters: the number of iterations I and the correct pose threshold θ . We gradually increase the number of iterations I from 0 to 1000 while keeping θ set to $1/3 m$, and calculate the accuracy and computational speed at different iterations. The accuracy is defined by the ADD(-S) metric, and the computational speed is defined by FPS, as shown in Fig. 8(a). When the number of iterations is 0, it corresponds to no confidence ranking of the correspondence obtained in the previous stage, which may lead to errors in the initial four correspondences used to solve the pose, resulting in decreased accuracy. As the number of iterations increases, the accuracy rate gradually increases, but the computational speed of the method decreases. To balance accuracy and computational speed, we set the optimal number of iterations between 300 and 500. In fact, the number of iterations calculated by (12) falls within this range.

θ is the threshold value of the correct pose, as shown in Fig. 8(b). We evaluate the impact of different θ on accuracy, where accuracy is defined by the ADD(-S) metric. Generally, the value of θ is not sensitive to the accuracy of the final pose estimation. As θ increases, accuracy and precision first slightly increase and then decrease. This is because when θ is small, it follows from (11) that the pose with higher accuracy contributes to higher confidence. As the threshold value increases, the influence of some incorrect poses on the confidence can be avoided, thus improving accuracy and precision. However, if θ is too large, many correct poses may be judged as incorrect, ultimately leading to a decrease in accuracy and precision of the

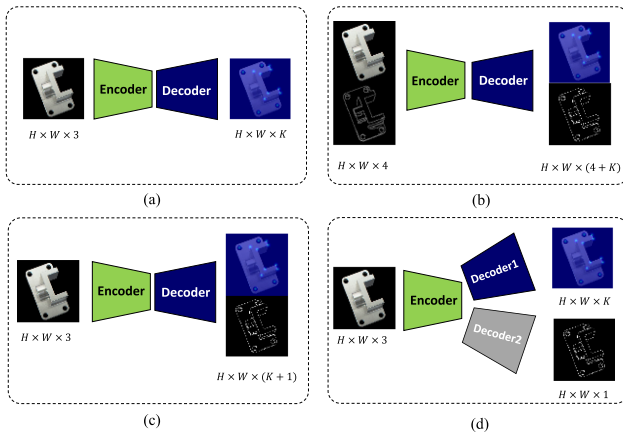


Fig. 9. Ablation study of network design. (a) Network1. Input an image and output the heatmap of semantic points. (b) Network2. Input a four-channel image, the first three channels is RGB channels, and the last channel is the contour detected with canny. Output heatmap and noise reduced contour. (c) Network3. Input an RGB image. Output a feature map with $K + 1$ channels, K represents the number of keypoints, and the last channel is the predicted contour. (d) Proposed method. Input an RGB image, output the heatmap by one decoder, and the contour by another decoder. (a) Network 1. (b) Network 2. (c) Network 3. (d) ContourNet.




Fig. 10. Pose estimation results in terms of ADD metric. Gradually increasing the error correspondence, PECP is more robust than RANSAC-based methods.




Fig. 11. Evaluation on reflective metal parts dataset in terms of 2-D projection metric. (a) Mean 2-D projection metric compared with other methods. (b) Performance of each object in terms of 2-D projection metric using our method.

final pose estimation. As shown in Fig. 8(b), when the value of θ is less than 0.4 m, the proposed method performs well.

D. Ablation Studies

1) *Network Analysis*: We conduct ablation studies in network design schemes to demonstrate the effectiveness of our introduction of contour decoders, as shown in Fig. 9. Table II lists the performance of different networks for different inputs

TABLE II
COMPARING THE PERFORMANCE OF DIFFERENT NETWORK SCHEMES FOR INPUT IMAGES IN TERMS OF ADD (-S) METRIC

Method	GT Mask	GT BBOX	Yolov3	w/o pre-detector
Network1	91.35	80.53	77.16	64.50
Network2	86.79	76.49	54.96	24.04
Network3	98.55	88.55	72.11	48.31
ContourNet	99.59	94.62	92.22	90.79

The bold values indicate the best performance.

in terms of ADD(-S) metrics, using obj5 as an example. In our initial design, as shown in Fig. 9(a), only the semantic points of the object are regressed by the heatmap. However, since the test image contains multiple metal parts, each of which also has many semantic points, it needs a predetector to achieve good results.

A common practice is to localize the target object by introducing the mask. However, for metal parts, the contour provides more advanced mask with more semantic information, so we chose to use the contour to constrain the prediction region of keypoints. As shown in Fig. 9(b), the input is a four-channel image, where the first three channels are RGB channels, and the fourth channel is the contour obtained by Canny [49] detection. Since Canny does not detect contours well, we want the network to implement a denoising autoencoder function, i.e., input a mutilated noisy contour and output a complete contour. Through experiments, we found that this approach leads to a “lazy” network because the contour and semantic points are strongly associated, causing the network to only focus on the input contour and ignore the RGB images. When the accuracy of the input contour is high, the output keypoints can achieve high accuracy, and once the input contour is less effective, the performance of the method is degraded significantly.

The third scheme is to implement the prediction of contours with a network, as shown in Fig. 9(c). In this scheme, the contour and keypoints cotrain the model, and we expect to establish the relationship between the keypoint prediction and the contour in this way, thus achieving a constraint of the contour on the keypoints prediction. However, experiments show that this scheme creates a strong constraint between contours and keypoints. If the contour prediction is accurate, the keypoints are also accurate. But in extreme cases where the contour prediction accuracy decreases, it results in decreased accuracy of keypoints prediction.

To decouple the contour and keypoints, the proposed network architecture uses a separate decoder to predict the contour. The predictions of contours and heatmaps train the encoder together, but each trains its corresponding decoder. This scheme makes the contour implicitly constrain the prediction of keypoints instead of guiding it directly. When the contours are entirely incorrect, the prediction of keypoints degrades to the network1. As given in Table II, the ContourNet architecture significantly improves the accuracy and robustness of pose estimation.

2) *Comparison With the RANSAC-Based Pose Calculation*: To demonstrate the effectiveness of PECP in solving the optimal pose, we compare it with RANSAC-based methods. Table III presents the results of the comparison in terms of different




Fig. 12. Qualitative results for reflective metal parts dataset, where the green edge line is a visualization of our predicted pose. We place each part in two–three scenes, and each scene also contains one–three target parts. (a) obj1. (b) obj2. (c) obj3. (d) obj4. (e) obj5. (f) obj6. (g) obj7. (h) obj8. (i) obj9. (j) obj10.




Fig. 13. Qualitative results for different scenes. The scenes are divided into A, B, C, and D according to the different backgrounds. Each scene has a different lighting situation. (a) Scene A. (b) Scene B. (c) Scene C. (d) Scene D.



Fig. 14. Qualitative results for T-LESS dataset, where the green edge line is a visualization of our predicted pose.

metrics. The proposed method outperforms the RANSAC-based method in both accuracy and precision.

While RANSAC relies only on the point distribution to eliminate outlier correspondences, our method uses the contour as a geometric prior to compute the confidence score of each correspondence, and then iteratively solve for the optimal pose with progressive sampling. To investigate the robustness of our method against outlier correspondences, we conducted an

TABLE III
COMPARE THE PERFORMANCE OF RANSAC AND PECP FOR SOLVING POSE IN TERMS OF DIFFERENT METRICS

Method	ContourNet + RANSAC	ContourNet+ PECP
2D Projection	95.81%	97.06%
ADD	90.15%	94.1%
R	2.16 °	0.84 °
t	5.92 mm	4.44 mm

The bold values indicate the best performance.

experiment where we systematically introduced outlier correspondences and computed the pose. We altered the location of the predicted keypoint so that it became an outlier, and then solved the pose using the RANSAC-based and PECP-based methods, respectively. We evaluated the results using the ADD metric.

As shown in Fig. 10, we used obj5 as an example and sequentially increased the number of error points. As the proportion of errors gradually increased, PECP demonstrated higher accuracy than RANSAC-based methods. Particularly, when the number of error points exceeded half of all points, the performance of RANSAC dropped sharply, whereas our method still achieved good results.

E. Comparison With the State-of-the-Art Methods

We compared our method with BB8 [9], AAE [25], STB [6], PSGMN [37], and GFI [29] on the reflective metal parts dataset, while STB, PSGMN, and GFI were designed specifically for pose estimation on metal parts. BB8 improves the end branch of VGG [50] by directly outputting the 2-D coordinates of the eight corner points of the 3-D bounding box. We adopted the optimization method of BB8, which regresses the keypoints by heatmap. AAE encodes object rotation into the latent space




Fig. 15. Pose estimation of target parts in complex scenes. Each scene contains five target parts and several unrelated parts. (a) White background. (b) Wooden background. (c) Metal background. (d) Black background. Our method can accurately estimate the pose of each target part. For the convenience of display, we combine the pose estimation results of multiple parts in one figure.




Fig. 16. Illustration of grasping task of reflective metal parts in complex scenes. (a) Metal part to grasp. (b) Initial position of the industrial robot. (c) Image taken by the camera at the initial position. The proposed method successfully estimates the 6-D pose of the target part. (d) Successful grasping of the metal part.

using an augmented autoencoder, which only requires synthetic data for training. This method can effectively learn the 6-D pose of objects under different environmental backgrounds and occlusion conditions. Given an image, AAE only predicts the object's rotation, while the object's translation is estimated using a 2-D bounding box. PSGMN is a dense matching method that establishes the pixelwise correspondence between 3-D models and 2-D images using graph neural networks, and then calculates the pose by RANSAC and PnP algorithms. STB and GFI were specifically designed for pose estimation of metal parts. STB uses high-level geometric features and linear contours to represent metal part templates requiring only sparse templates to obtain highly accurate poses. GFI proposed a feature-image generation model. Given a feature representing a pose, it can generate an image of the object in the exact same pose. This method avoids extracting features from reflective metal parts.

We compared the 2-D projection metric with the other methods, as shown in Fig. 11. Our method is significantly better than other methods. In addition to that, our method shows more performance when the constraints are more rigorous. With a pixel threshold of only 2 pixels, the proposed method achieves 87.46% in the 2-D projection metric. GFI [29] achieves better results when the pixel threshold is less than 2 pixels. The method achieves high accuracy by matching the most similar templates.

Table IV summarizes the comparison with other methods using the ADD(-S) metric. BB8 [9] estimates the pose by regressing the eight corner points of the object BBOX, which can easily result in incorrect keypoints due to complex backgrounds or interference from other parts. AAE [25] only encodes

TABLE IV
COMPARISON WITH DIFFERENT METHODS ON REFLECTIVE METAL PARTS DATASET USING THE ADD(-S) METRIC

Methods	BB8	AAE	STB	PSGMN	GFI	Ours
Obj 1	71.93	76.96	64.21	94.23	95.32	100.00
Obj 2	43.49	76.43	66.49	70.86	96.77	97.54
Obj 3	43.44	84.32	54.65	82.45	92.16	95.35
Obj 4	26.72	32.42	48.90	74.95	91.49	88.14
Obj 5	48.80	64.77	36.46	79.57	87.85	90.70
Obj 6	68.42	54.32	62.36	84.34	85.03	96.71
Obj 7	18.73	49.33	29.45	74.11	76.31	91.82
Obj 8	19.21	72.12	45.49	75.89	84.22	95.31
Obj 9	44.95	67.09	61.26	79.94	89.92	93.50
Obj 10	39.18	71.32	59.23	87.30	85.11	92.30
Average	42.60	64.91	52.85	80.36	88.42	94.14

The bold values indicate the best performance.

rotation information of the object and will suffer from decreased accuracy when dealing with complex parts. STB [6] performs well for individual parts but can encounter errors due to template matching issues when there are interferences from other parts in the dataset, as is the case with similar geometric features across all parts. PSGMN [37] matches the object pixels in the image with the nodes of the 3-D model. However, in the test dataset, there are multiple metal parts with similar surface features, which can cause semantic segmentation errors and incorrect matching results. Therefore, PSGMN requires a predetector to achieve accurate pose estimation. GFI [29] has to do edge detection on the target object in the image, and then generate the image that is most similar to the edge detection result. After that, this method can get the pose by the feature corresponding to the image. The method relies on the result of edge detection, and the performance of the method will be degraded when the edges are inaccurate. STB [6] and GFI [29] rely on contours to estimate the poses, and the accuracy of the pose estimation decreases when the contours are not clear or accurate. In contrast, our method only takes the contour as an implicit constraint, establishes 2-D-3-D correspondence by predicting semantic points, and then solves the poses with PnP. The contour only

TABLE V
COMPARING THE PERFORMANCE OF DIFFERENT METHODS IN DIFFERENT SCENARIOS IN TERMS OF ADD(-S) METRICS

Methods	BB8	AAE	STB	PSGMN	GFI	Ours
Scene A	44.82	69.77	54.43	84.75	90.34	96.75
Scene B	50.95	57.82	50.09	83.27	89.49	91.24
Scene C	20.54	49.19	31.39	75.09	86.77	89.01
Scene D	46.42	64.23	49.78	72.41	88.25	95.60
Average	40.68	60.26	46.42	78.88	88.71	93.15

The bold values indicate the best performance.

TABLE VI
COMPARISON OF DIFFERENT METHODS USING THE R/t ERROR METRIC IN VALID POSES

Method	x/mm	y/mm	z/mm	$\alpha/^\circ$	$\beta/^\circ$	$\gamma/^\circ$
BB8	1.87	1.51	7.61	3.10	2.89	0.85
AAE	1.48	1.1	7.92	5.25	4.99	2.23
STB	1.47	0.94	7.49	1.11	1.44	0.85
PSGMN	2.50	1.97	8.45	4.00	3.74	1.52
GFI	2.47	1.94	6.59	1.85	1.91	0.97
Ours	0.71	0.79	4.31	1.00	1.10	0.42

The bold values indicate the best performance.

serves to improve the accuracy of the pose estimation not to determine it. The comparison shows that our approach is superior to the state-of-the-art method. Some qualitative results of pose estimation are shown in Fig. 12.

Table V lists the performance of each method in different scenes in terms of ADD(-S) metric. Scene A is a black background, Scene B is a black textured background, Scene C is a rusted background, and Scene D is a reflective metallic background. Our method demonstrates robustness in various scenes, including those with different lighting conditions. Fig. 13 provides some qualitative results of pose estimation in different scenes.

Table VI summarizes the comparison of other methods in terms of the R/t error metric. We only considered valid poses, i.e., those that were correct under the ADD(-S) metric. The R error represents the 3-D rotational error of the object, while the t error represents the 3-D translation error. Our method achieved an average error of less than 1 mm in 2-D translation and only about 1° for Euler angles α , β , and γ . This shows that our method can achieve high accuracy for reflective textureless metal parts, which can be implemented for practical applications in industrial scenes.

To evaluate the effectiveness of our proposed method on general industrial objects, we conducted experiments on the T-LESS [44] dataset. We used the predetector provided by Pix2Pose [36] to identify the target object and trained our model using real and synthetic images, as was done in CosyPose [31], SurfEmb [51], and CDPN [12]. We then estimated the 6-D pose of the target object using our method. Fig. 14 shows the qualitative results of

TABLE VII
COMPARISON WITH DIFFERENT METHODS ON THE T-LESS DATASET

Method	$e_{vSD} < 0.3$
AAE	18.2
Pix2Pose	34.4
CDPN	47.8
CosyPose	72.8
SurfEmb	77.0
Ours	72.4

The bold values indicate the best performance.

our method on the T-LESS dataset, demonstrating its applicability even for objects without sharp edges. Table VII provides a summary comparison of our method with other methods on the T-LESS dataset, where all methods used RGB images as input. Our method is designed for reflective metal parts with sharp contours that can significantly improve the accuracy of the proposed method. While our method did not achieve the best performance on the T-LESS dataset, the experimental results confirm that it can handle general industrial objects.

F. Pose Estimation for Metal Part Grasping

To further demonstrate the effectiveness of our proposed method, we design some more challenging scenes. These scenes are prevalent with parts stacked and placed in different backgrounds, as well as interference from other similar parts. As shown in Fig. 15, our method can accurately estimate the pose of the target part in such complex scenes.

Moreover, we implemented our method in a real grasping task for reflective textureless metal parts, as illustrated in Fig. 16. The grasping task was performed using a six-axis industrial robot with a CCD camera mounted on its end-effector, and the intrinsic parameters of the camera and hand-eye matrix are well calibrated. We place several parts in a plane with some stacking as the target. The robot's end-effector is first moved to the initial position, where the picture is taken and the pose of the target part is calculated, and then it is moved to the corresponding position for grasping. Our proposed pose estimation method achieves a real-time performance of 20 FPS on a single Nvidia RTX 3090 GPU.

V. CONCLUSION AND FUTURE WORK

In this article, we proposed a 6-D pose estimation method of reflective textureless metal parts with contour constraints and geometric prior. The ContourPose consists of two stages, namely, the keypoints prediction stage and the pose estimation stage. We achieve an implicit constraint on the keypoints prediction by adding a decoder to predict the contour in the first stage. In the pose estimation stage, we utilized the predicted contours from the previous stage as a priori to iteratively solve the optimal pose, which greatly enhances the accuracy and precision of the pose. Our proposed method for reflective textureless

metal parts outperforms current state-of-the-art pose estimation methods. Furthermore, we demonstrated the effectiveness of our method in an industrial application for grasping tasks. The dataset and more information in this article can be found on https://github.com/ZJU-IVIR/RT-Less_10parts.

Our proposed method trains a specific network to estimate a single object similar to PVNet [10], CDPN [12], and PSGMN [37]. However, unlike these methods, our method does not require a predetector if there is only one target object in the scene. Currently, for multiple identical objects in one scene, our method still relies on a predetector to locate each instance. In future work, we plan to leverage topological constraints to group the keypoints of each instance and enable detection of multiple objects of the same class.

REFERENCES

- [1] Z. He, W. Feng, X. Zhao, and Y. Lv, "6D pose estimation of objects: Recent technologies and challenges," *Appl. Sci.*, vol. 11, no. 1, 2020, Art. no. 228.
- [2] H. Zhang and Q. Cao, "Detect in RGB, optimize in edge: Accurate 6D pose estimation for texture-less industrial parts," in *Proc. Int. Conf. Robot. Automat.*, 2019, pp. 3486–3492, doi: [10.1109/ICRA.2019.8794330](https://doi.org/10.1109/ICRA.2019.8794330).
- [3] Y. Ke and R. Sukthankar, "PCA-SIFT: A more distinctive representation for local image descriptors," in *Proc. IEEE Comput. Soc. Conf. Comput. Vis. Pattern Recognit.*, 2004, pp. 506–513.
- [4] H. Bay, A. Ess, T. Tuytelaars, and L. Van Gool, "Speeded-up robust features (SURF)," *Comput. Vis. Image Understanding*, vol. 110, no. 3, pp. 346–359, 2008.
- [5] E. Rublee, V. Rabaud, K. Konolige, and G. Bradski, "ORB: An efficient alternative to SIFT or SURF," in *Proc. Int. Conf. Comput. Vis.*, 2011, pp. 2564–2571.
- [6] Z. He, Z. Jiang, X. Zhao, S. Zhang, and C. Wu, "Sparse template-based 6-D pose estimation of metal parts using a monocular camera," *IEEE Trans. Ind. Electron.*, vol. 67, no. 1, pp. 390–401, Jan. 2020.
- [7] O. Tahri and F. Chaumette, "Complex objects pose estimation based on image moment invariants," in *Proc. IEEE Int. Conf. Robot. Automat.*, 2005, pp. 436–441.
- [8] S. Hinterstoisser et al., "Gradient response maps for real-time detection of textureless objects," *IEEE Trans. Pattern Anal. Mach. Intell.*, vol. 34, no. 5, pp. 876–888, May 2012.
- [9] M. Rad and V. Lepetit, "BB8: A scalable, accurate, robust to partial occlusion method for predicting the 3D poses of challenging objects without using depth," in *Proc. IEEE Int. Conf. Comput. Vis.*, 2017, pp. 3828–3836.
- [10] S. Peng, Y. Liu, Q. Huang, X. Zhou, and H. Bao, "PVNet: Pixel-wise voting network for 6DoF pose estimation," in *Proc. IEEE/CVF Conf. Comput. Vis. Pattern Recognit.*, 2019, pp. 4556–4565.
- [11] B. Tekin, S. N. Sinha, and P. Fua, "Real-time seamless single shot 6D object pose prediction," in *Proc. IEEE/CVF Conf. Comput. Vis. Pattern Recognit.*, 2018, pp. 292–301, doi: [10.1109/CVPR.2018.00038](https://doi.org/10.1109/CVPR.2018.00038).
- [12] Z. Li, G. Wang, and X. Ji, "CDPN: Coordinates-based disentangled pose network for real-time RGB-based 6-DoF object pose estimation," in *Proc. IEEE/CVF Int. Conf. Comput. Vis.*, 2019, pp. 7677–7686, doi: [10.1109/ICCV.2019.00777](https://doi.org/10.1109/ICCV.2019.00777).
- [13] S. Zakhharov, I. Shugurov, and S. Ilic, "DPOD: 6D pose object detector and refiner," in *Proc. IEEE/CVF Int. Conf. Comput. Vis.*, 2019, pp. 1941–1950, doi: [10.1109/ICCV.2019.00203](https://doi.org/10.1109/ICCV.2019.00203).
- [14] S. Hinterstoisser et al., "Model based training, detection and pose estimation of texture-less 3D objects in heavily cluttered scenes," in *Proc. Asian Conf. Comput. Vis.*, 2012, pp. 548–562.
- [15] Y. Xiang, T. Schmidt, V. Narayanan, and D. Fox, "PoseCNN: A convolutional neural network for 6D object pose estimation in cluttered scenes," *Robot.: Sci. Syst. (RSS)*, 2018.
- [16] L. Ferraz, X. Binefa, and F. Moreno-Noguer, "Very fast solution to the PnP problem with algebraic outlier rejection," in *Proc. IEEE Conf. Comput. Vis. Pattern Recognit.*, 2014, pp. 501–508.
- [17] J. Redmon, S. Divvala, R. Girshick, and A. Farhadi, "You only look once: Unified, real-time object detection," in *Proc. IEEE Conf. Comput. Vis. Pattern Recognit.*, 2016, pp. 779–788.
- [18] T.-Y. Lin, P. Goyal, R. Girshick, K. He, and P. Dollár, "Focal loss for dense object detection," in *Proc. IEEE Int. Conf. Comput. Vis.*, 2017, pp. 2999–3007.
- [19] A. Kendall, M. Grimes, and R. Cipolla, "PoseNet: A convolutional network for real-time 6-DOF camera relocalization," in *Proc. IEEE Int. Conf. Comput. Vis.*, 2015, pp. 2938–2946.
- [20] W. Kehl, F. Manhardt, F. Tombari, S. Ilic, and N. Navab, "SSD-6D: Making RGB-based 3D detection and 6D pose estimation great again," in *Proc. IEEE Int. Conf. Comput. Vis.*, 2017, pp. 1530–1538.
- [21] W. Liu et al., "SSD: Single shot multibox detector," in *Proc. Eur. Conf. Comput. Vis.*, 2016, pp. 21–37.
- [22] C. Wang et al., "DenseFusion: 6D object pose estimation by iterative dense fusion," in *Proc. IEEE/CVF Conf. Comput. Vis. Pattern Recognit.*, 2019, pp. 3338–3347, doi: [10.1109/CVPR.2019.00346](https://doi.org/10.1109/CVPR.2019.00346).
- [23] Y. Bukschat and M. Vetter, "Efficientpose: An efficient, accurate and scalable end-to-end 6D multi object pose estimation approach," 2020, *arXiv:2011.04307*.
- [24] M. Tan, R. Pang, and Q. V. Le, "EfficientDet: Scalable and efficient object detection," in *Proc. IEEE/CVF Conf. Comput. Vis. Pattern Recognit.*, 2020, pp. 10778–10787, doi: [10.1109/CVPR42600.2020.01079](https://doi.org/10.1109/CVPR42600.2020.01079).
- [25] M. Sundermeyer, Z.-C. Marton, M. Durner, M. Brucker, and R. Triebel, "Implicit 3D orientation learning for 6D object detection from RGB images," in *Proc. Eur. Conf. Comput. Vis.*, 2018, pp. 699–715.
- [26] M. Sundermeyer et al., "Multi-path learning for object pose estimation across domains," in *Proc. IEEE/CVF Conf. Comput. Vis. Pattern Recognit.*, 2020, pp. 13913–13922.
- [27] Y. Wen, H. Pan, L. Yang, and W. Wang, "Edge enhanced implicit orientation learning with geometric prior for 6D pose estimation," *IEEE Robot. Automat. Lett.*, vol. 5, no. 3, pp. 4931–4938, Jul. 2020.
- [28] A. Blank et al., "6DoF pose-estimation pipeline for texture-less industrial components in bin picking applications," in *Proc. Eur. Conf. Mobile Robots*, 2019, pp. 1–7, doi: [10.1109/ECMR.2019.8870920](https://doi.org/10.1109/ECMR.2019.8870920).
- [29] Z. He, M. Wu, X. Zhao, S. Zhang, and J. Tan, "A generative feature-to-image robotic vision framework for 6D pose measurement of metal parts," *IEEE/ASME Trans. Mechatron.*, vol. 27, no. 5, pp. 3198–3209, Oct. 2022, doi: [10.1109/TMECH.2021.3109344](https://doi.org/10.1109/TMECH.2021.3109344).
- [30] Y. Li, G. Wang, X. Ji, Y. Xiang, and D. Fox, "DeepIM: Deep iterative matching for 6D pose estimation," *Int. J. Comput. Vis.*, vol. 128, pp. 657–678, 2020.
- [31] Y. Labbé, J. Carpentier, M. Aubry, and J. Sivic, "CosyPose: Consistent multi-view multi-object 6D pose estimation," in *Proc. Eur. Conf. Comput. Vis.*, 2020, pp. 574–591.
- [32] Z. Fan, Y. Zhu, Y. He, Q. Sun, H. Liu, and J. He, "Deep learning on monocular object pose detection and tracking: A comprehensive overview," *ACM Comput. Surv.*, vol. 55, no. 4, pp. 1–40, 2022.
- [33] A. Newell, K. Yang, and J. Deng, "Stacked hourglass networks for human pose estimation," in *Proc. Eur. Conf. Comput. Vis.*, 2016, pp. 483–499.
- [34] M. Oberwiger, M. Rad, and V. Lepetit, "Making deep heatmaps robust to partial occlusions for 3D object pose estimation," in *Proc. Eur. Conf. Comput. Vis.*, Jul. 2018, pp. 125–141.
- [35] C. Song, J. Song, and Q. Huang, "Hybridpose: 6D object pose estimation under hybrid representations," in *Proc. IEEE/CVF Conf. Comput. Vis. Pattern Recognit.*, 2020, pp. 428–437.
- [36] K. Park, T. Patten, and M. Vincze, "Pix2Pose: Pixel-wise coordinate regression of objects for 6D pose estimation," in *Proc. IEEE/CVF Int. Conf. Comput. Vis.*, 2019, pp. 7667–7676, doi: [10.1109/ICCV.2019.00776](https://doi.org/10.1109/ICCV.2019.00776).
- [37] C. Wu, L. Chen, Z. He, and J. Jiang, "Pseudo-Siamese graph matching network for textureless objects' 6D pose estimation," *IEEE Trans. Ind. Electron.*, vol. 69, no. 3, pp. 2718–2727, Mar. 2022, doi: [10.1109/TIE.2021.3070501](https://doi.org/10.1109/TIE.2021.3070501).
- [38] K. He, X. Zhang, S. Ren, and J. Sun, "Deep residual learning for image recognition," in *Proc. IEEE Conf. Comput. Vis. Pattern Recognit.*, 2016, pp. 770–778.
- [39] F. Yu and V. Koltun, "Multi-scale context aggregation by dilated convolutions," in *Proc. Int. Conf. Learn. Representations*, 2016.
- [40] C. R. Qi, L. Yi, H. Su, and L. J. Guibas, "PointNet++: Deep hierarchical feature learning on point sets in a metric space," in *Proc. Adv. Neural Inf. Process. Syst.*, 2017, vol. 30, pp. 5099–5108.
- [41] G. Pavlakos, X. Zhou, A. Chan, K. G. Derpanis, and K. Daniilidis, "6-DoF object pose from semantic keypoints," in *Proc. IEEE Int. Conf. Robot. Automat.*, 2017, pp. 2011–2018.
- [42] J. Shi and Tomasi, "Good features to track," in *Proc. IEEE Conf. Comput. Vis. Pattern Recognit.*, 1994, pp. 593–600, doi: [10.1109/CVPR.1994.323794](https://doi.org/10.1109/CVPR.1994.323794).

- [43] V. Lepetit, F. Moreno-Noguer, and P. Fua, "EPNP: An accurate O(n) solution to the PNP problem," *Int. J. Comput. Vis.*, vol. 81, no. 2, pp. 155–166, 2009.
- [44] T. Hodan, P. Haluza, Š. Obdržálek, J. Matas, M. Lourakis, and X. Zabulis, "T-LESS: An RGB-D dataset for 6D pose estimation of texture-less objects," in *Proc. IEEE Winter Conf. Appl. Comput. Vis.*, 2017, pp. 880–888.
- [45] D. Dwibedi, I. Misra, and M. Hebert, "Cut, paste and learn: Surprisingly easy synthesis for instance detection," in *Proc. IEEE Int. Conf. Comput. Vis.*, 2017, pp. 1310–1319.
- [46] J. Xiao, J. Hays, K. A. Ehinger, A. Oliva, and A. Torralba, "Sun database: Large-scale scene recognition from abbey to zoo," in *Proc. IEEE Comput. Soc. Conf. Comput. Vis. Pattern Recognit.*, 2010, pp. 3485–3492.
- [47] E. Brachmann, F. Michel, A. Krull, M. Y. Yang, and S. Gumhold, "Uncertainty-driven 6D pose estimation of objects and scenes from a single RGB image," in *Proc. IEEE Conf. Comput. Vis. Pattern Recognit.*, 2016, pp. 3364–3372.
- [48] T. Hodan et al., "Bop: Benchmark for 6D object pose estimation," in *Proc. Eur. Conf. Comput. Vis.*, 2018, pp. 19–34.
- [49] W. Rong, Z. Li, W. Zhang, and L. Sun, "An improved CANNY edge detection algorithm," in *Proc. IEEE Int. Conf. Mechatron. Automat.*, 2014, pp. 577–582.
- [50] K. Simonyan and A. Zisserman, "Very deep convolutional networks for large-scale image recognition," 2014, *arXiv:1409.1556*.
- [51] R. L. Haugaard and A. G. Buch, "SurfEmb: Dense and continuous correspondence distributions for object pose estimation with learnt surface embeddings," in *Proc. IEEE/CVF Conf. Comput. Vis. Pattern Recognit.*, 2022, pp. 6739–6748.



Zaixing He (Senior Member, IEEE) received the B.Sc. and M.Sc. degrees in mechanical engineering from Zhejiang University, Hangzhou, China, in 2006 and 2008, respectively, and the Ph.D. degree in information science from the Graduate School of Information Science and Technology, Hokkaido University, Sapporo, Japan, in 2012.

He is currently an Associate Professor with the School of Mechanical Engineering, Zhejiang University. He has authored or coauthored more than 40 peer-reviewed papers in prestigious journals, such as IEEE

TRANSACTIONS ON ROBOTICS, IEEE TRANSACTIONS ON INDUSTRIAL ELECTRONICS, IEEE TRANSACTIONS ON INDUSTRIAL INFORMATICS, IEEE TRANSACTIONS ON INSTRUMENTATION AND MEASUREMENT, IEEE/ASME TRANSACTIONS ON MECHATRONICS, *Pattern Recognition*, and *Neurocomputing*. He served as a Lead Guest Editor or Guest Editor for several journals, including IEEE TRANSACTIONS ON CONSUMER ELECTRONICS and *Mathematics*, and a Program Chair or TPC for more than ten international conferences. His research interests include robotic vision, visual intelligence of manufacturing equipment, and optical-based measurement.



Quanzhi Li received the B.S. degree in mechanical and electronic engineering from Chongqing University, Chongqing, China, in 2021. He is currently working toward the M.S. degree in mechanical engineering with Zhejiang University, Hangzhou, China.

His research interests include robotic vision, pose estimation, and artificial intelligence.



Xinyue Zhao received the M.S. degree in mechanical engineering from Zhejiang University, Hangzhou, China, in 2008, and the Ph.D. degree in information science from the Graduate School of Information Science and Technology, Hokkaido University, Sapporo, Japan, in 2012.

She is currently an Associate Professor with the School of Mechanical Engineering, Zhejiang University, China. Her research interests include machine vision and image processing. She has authored or coauthored nearly 50 peer-reviewed journal papers.



Jin Wang received the B.Sc. and Ph.D. degrees in mechanical engineering from Zhejiang University, Hangzhou, China, in 2003 and 2008, respectively.

He is currently an Associate Professor with the College of Mechanical Engineering, Zhejiang University, China. His research interests include design automation and optimization, industrial robotic off-line planning, industrial robot intelligent control, and multirobot coordinated control and application.



Huarong Shen received the B.Sc. degree in electronic information engineering and the M.Sc. degree in agriculture from Zhejiang University, Hangzhou, China, in 2009 and 2014, respectively.

He is currently a Chief Executive Officer of Zhejiang Feihang Intelligent Technology, Co., Ltd., Huzhou, China.



Shuyou Zhang received the M.S. degree in mechanical engineering and the Ph.D. degree in applied mathematics from the State Key Laboratory of CAD&CG, Zhejiang University, Hangzhou, China, in 1991 and 1999, respectively.

He is currently a Professor with the Department of Mechanical Engineering, Zhejiang University, Hangzhou, China. His research interests include product digital design, design and stimulation for complex equipment, and engineering and computer graphics.



Jianrong Tan received the B.S. degree in mechanical engineer and electronic engineering from The Open University of China, Beijing, China, in 1982, the M.S. degree in engineering from the Huazhong University of Science and Technology, Wuhan, China, in 1985, and the Ph.D. degree in mathematics from Zhejiang University, Hangzhou, China, in 1987.

He is an academian of China Engineering Academy, and is currently a Professor with the State Key Laboratory of CAD&CG, Zhejiang University. His main research interests include virtual-reality-based simulation, machine learning, CAX, and robotics.