

# Focal Length and Object Pose Estimation via Render and Compare

Georgy Ponimatkin<sup>1,2</sup>    Yann Labbé<sup>3</sup>    Bryan Russell<sup>4</sup>  
 Mathieu Aubry<sup>1</sup>    Josef Sivic<sup>2</sup>

<sup>1</sup>LIGM, École des Ponts, Univ Gustave Eiffel, CNRS    <sup>2</sup>CIIRC CTU    <sup>3</sup>ENS/Inria    <sup>4</sup>Adobe Research

georgy.ponimatkin@enpc.fr

<https://ponimatkin.github.io/focalpose>

## Abstract

We introduce *FocalPose*, a neural render-and-compare method for jointly estimating the camera-object 6D pose and camera focal length given a single RGB input image depicting a known object. The contributions of this work are twofold. First, we derive a focal length update rule that extends an existing state-of-the-art render-and-compare 6D pose estimator to address the joint estimation task. Second, we investigate several different loss functions for jointly estimating the object pose and focal length. We find that a combination of direct focal length regression with a re-projection loss disentangling the contribution of translation, rotation, and focal length leads to improved results. We show results on three challenging benchmark datasets that depict known 3D models in uncontrolled settings. We demonstrate that our focal length and 6D pose estimates have lower error than the existing state-of-the-art methods.

## 1. Introduction

The projection of a 3D object into an image depends not only on the object’s relative pose to the camera, but also on the camera’s intrinsic parameters. While it is possible to capture objects in a controlled environment where the camera’s intrinsic parameters are known (*e.g.*, a calibrated camera on a robot), for many “in-the-wild” images we do not have control over the capture process and these parameters are unknown, *e.g.*, Internet pictures or archival photographs.

Given an input image, we seek to retrieve a 3D model of a depicted object from a model library and estimate the relative camera-object 6D pose jointly with the camera’s focal length (depicted in Figure 1). This problem has its origins in the early days of computer vision [26, 27, 36] and has important modern-day applications in augmented reality and computer graphics, such as applying in situ object overlays or editing the position of an object via 3D compositing in uncontrolled consumer-captured images.

The problem of 6D object pose estimation in an uncal-



Figure 1. Given a single input photograph (left) and a known 3D model, our approach accurately estimates the 6D camera-object pose together with the focal length of the camera (right), here shown by overlaying the aligned 3D model over the input image. Our approach handles a large range of focal lengths and the resulting perspective effects.

ibrated setting is, by its nature, challenging. First, it is difficult to distinguish subtle changes of the camera’s focal length from changes in an object’s depth. Second, including the camera’s focal length increases the number of parameters that must be estimated and hence increases the optimization complexity. Finally, “in-the-wild” consumer-captured images may depict large appearance variation for a particular object instance in the model library. Variation may be due to differences in illumination and the depicted object having slightly different, non-identical shapes or surface appearance in different real-world instance captures. For example, consider different instances of the same car model that have a similar overall shape but may have different color, wear and tear, or customizable features (*e.g.*, additional headlights, alloy wheels, or a spoiler).

Previous approaches for this task primarily rely on establishing local 2D-3D correspondences between an image

and a 3D model using either hand-crafted [2, 3, 7, 8, 17, 27] or CNN features [12, 19, 20, 31, 32, 34, 35, 38, 41, 42, 47, 48], followed by robust camera pose estimation using PnP [23]. These approaches often fail on scenes with large textureless areas where local correspondences cannot be reliably established. In contrast, the recent best-performing 6D object pose estimation methods are based on the render-and-compare strategy [22, 24, 28, 30, 48], which performs a dense alignment over all pixels of rendered views of the 3D model to its depiction in the input image. However, all prior render-and-compare methods fall short of handling the aforementioned desired uncontrolled, uncalibrated setting as they assume a controlled environment where the camera intrinsic parameters are fixed and known *a priori*. Also, these prior methods typically operate over only a handful of known objects.

To address these challenges, we build on the strengths of render and compare and extend it to handle our desired uncontrolled, uncalibrated setting. We introduce FocalPose, a novel render-and-compare approach for jointly estimating an object’s 6D pose and camera focal length based on a monocular image input. Our contributions are twofold. First, we extend a recent state-of-the-art [18] method for 6D pose estimation (CosyPose [22]) by deriving and integrating focal length update rules in a differentiable manner, which allows our method to overcome the added complexity of including the focal length. Second, we investigate several different loss functions for jointly estimating object pose and focal length. We find that a combination of direct focal length regression with a reprojection loss disentangling the contribution of translation, rotation, and focal length leads to the best performance and allows our method to distinguish subtle differences due to the focal length and the object’s depth. We apply our method to three real-world consumer-captured image datasets with varying camera focal lengths and show that our focal length and 6D pose estimates have lower error compared to the state of the art. As an added benefit, our work is the first render-and-compare method applied to a large collection of 3D meshes (20-200 meshes for Pix3D [39],  $\sim 150$  for the car datasets [45]).

## 2. Related Work

### 6D pose estimation of rigid objects from RGB images.

This task is one of the oldest problems in computer vision [26, 27, 36] and has been successfully approached by estimating the pose from 2D-3D correspondences obtained via local invariant features [3, 7, 8, 27], or by template-matching [17]. Both of these strategies rely on shallow hand-designed image features and have been revisited with learnable deep convolutional neural networks (CNNs) [19, 20, 31, 32, 34, 35, 38, 41, 42, 47, 48]. The best-performing methods for 6D pose estimation from RGB images are now based on variants of the deep *render-and-compare* strat-

egy [22, 24, 28, 30, 48]. However, these methods assume the full perspective camera model is known so that the object can be rendered and compared with the input image. We build on the state-of-the-art render-and-compare approach of Labbé *et al.* [22] and extend it to the “in-the-wild” uncontrolled set-up where the focal length of the camera is not known and has to be estimated together with the object’s 6D pose directly from the input image.

**Camera calibration.** Camera calibration techniques [1, 9, 10, 29, 33, 40, 43, 49] recover the camera model (intrinsic parameters) and its pose (extrinsic parameters) jointly. A limitation is that they require estimating 2D-3D correspondences in multiple images using structured object patterns [11, 14, 40, 43], identifying specific image elements such as lines or vanishing points [6, 9, 40] or structured features (*e.g.*, human face landmarks [4]). These requirements limit their applicability to unconstrained images where these structures are not present. Other works [46] have considered in-the-wild images, but only focus on recovering the focal length of the camera. In contrast, our approach recovers both components of the camera calibration (focal length and 6D camera pose) given a single image of a known object.

### Joint 6D pose and focal length estimation from a single in-the-wild image.

The prior work closest to our approach establishes point correspondences, followed by robust fitting of the camera model [12, 13, 45]. Wang *et al.* [45] uses Faster R-CNN with a scalar regression head and L1 loss for estimating the focal length, and the 6D pose is estimated by predicting 2D-3D correspondences followed by PnP. GP2C [12] extends this approach via a two-step procedure that predicts initial 2D-3D correspondences and focal length with a similar direct regression, followed by applying a PnP solver to refine jointly the 6D pose and the focal length. The model cannot be trained end-to-end as it relies on a separate non-differentiable optimizer. GCVNet [13] uses an approximation of the PnP solver for differentiability, but its results are limited by this approximation. In contrast, our work builds on the success of the recent render-and-compare methods [22, 24] for 6D rigid pose estimation. Our 6D pose and focal length updates are learned end-to-end using our novel focal length update parameterization coupled with a disentangled training loss. Our approach produces lower-error focal length and pose estimates compared to the two-step approach of GP2C [12] and the prior one-shot end-to-end approaches [13, 45].

## 3. Approach

Our goal is to estimate the 6D pose of objects in a photograph taken with unknown focal length. To achieve this goal, we use a render-and-compare strategy where we estimate jointly the camera focal length with the 6D pose. We

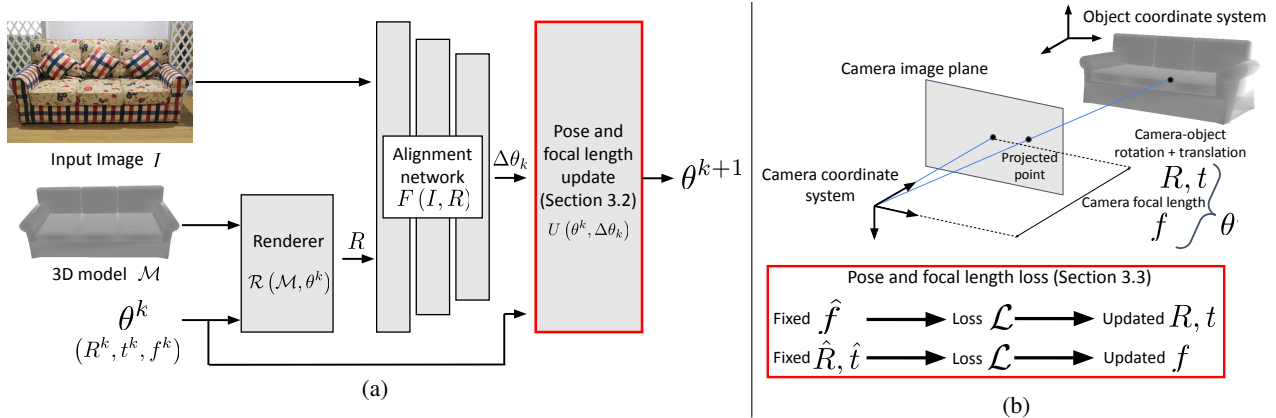


Figure 2. **FocalPose overview.** (a) Given a single in-the-wild RGB input image  $I$  of a known object 3D model  $\mathcal{M}$ , parameters  $\theta^k$  composed of focal length  $f^k$  and the object 6D pose (3D translation  $t^k$  and 3D rotation  $R^k$ ) are iteratively updated using our render-and-compare approach. Rendering  $R$ , together with the input image  $I$ , are given to a deep neural network  $F$  that predicts update  $\Delta\theta_k$ , which is then converted into parameter update  $\theta^{k+1}$  using a non-linear update rule  $U$ . (b) Illustration of the camera-object setup with parameters  $\theta$  composed of 3D translation  $t$ , 3D rotation  $R$  and focal length  $f$ . The alignment network is trained using a novel pose and focal length loss that disentangles the focal length and pose updates. The two main contributions of this work are highlighted by red boxes in the figure.

assume knowledge of a database of 3D models that may appear in the image, but our results show that the approach is effective even if the 3D models are only approximate.

### 3.1. Approach Overview

The first step of our approach, illustrated in Fig. 2, identifies the object location in the input image and retrieves a 3D model from the database that matches the depicted object instance. We use an object detector [15] trained on real images of these known objects. At test time, we run this detector on the test image to obtain a 2D bounding box of the object and its corresponding 3D model  $\mathcal{M}$ . We describe a *render and compare* approach, which iteratively estimates the focal length and 6D pose of the identified object. We denote the current estimate of focal length and 6D pose at iteration  $k$  jointly as  $\theta^k$ . The object model is first *rendered* using the current estimates  $\theta^k$  into an image  $\mathcal{R}(\mathcal{M}, \theta^k)$  using a renderer  $\mathcal{R}$ . The rendering  $\mathcal{R}(\mathcal{M}, \theta^k)$  and observed input image  $I$  are given to a deep neural network  $F$  which predicts the pose and focal length update  $\Delta\theta_k$ :

$$\Delta\theta_k = F(I, \mathcal{R}(\mathcal{M}, \theta^k)). \quad (1)$$

The intuition is that the neural network *compares* the input image  $I$  with the rendering  $\mathcal{R}(\mathcal{M}, \theta^k)$  and based on their (potentially subtle) differences predicts the update in the rendering parameters  $\Delta\theta_k$ . The pose and focal length updates  $\Delta\theta_k$  are designed to be, as much as possible, free of non-linearities and thus easy to predict by the neural network  $F$ . The pose and focal length at the next iteration  $k+1$  is then computed by a non-linear update rule  $U$ :

$$\theta^{k+1} = U(\theta^k, \Delta\theta_k), \quad (2)$$

where  $\theta^k$  is the current estimate of the pose and focal length,  $\Delta\theta_k$  is the prediction by the network  $F$  given by eq. (1), and

$\theta^{k+1}$  are the updated pose and focal length. Note that  $U$  is not learnt but derived from the 3D to 2D projection model and takes into account the nonlinearities of the imaging process. The neural network  $F$  is trained in such a way that the updated pose and focal length  $\theta^{k+1}$  are progressively closer to their ground truth. Our approach is summarized in Fig. 2

**Discussion.** Existing render-and-compare estimators [22, 24] require knowledge of the camera intrinsic parameters. In our scenario, the problem is more challenging because the rendering also depends on the unknown focal length. We address this challenge by proposing an update rule for the focal length as well as a modification of the update rules for 6D pose parameters accounting for the unknown focal length (Sec 3.2). We then introduce a novel loss function adapted for joint focal length and 6D pose estimation, which disentangles the effects of the pose and focal length updates for better end-to-end training of the network (Sec. 3.3). Please see the **supp. materials** for details of our implementation,  $\theta^0$  parameter initialization, and our training data.

### 3.2. Update rules with focal length estimation

The standard render-and-compare approach to 6D pose estimation [22,24] considers only translation  $t^k$  and rotation  $R^k$  as parameters  $\theta^k$ . We additionally estimate the focal length  $f^k$  as an unknown, and thus need to build an appropriate rule  $U$  (as defined in eq. (2)) for updating jointly all parameters. In detail, we assume a pinhole camera model with focal length  $f_x^k = f_y^k = f^k$  in which the optical center is set at the center of the image. We define the 6D pose of the object with respect to the camera by a 3D rotation  $R^k$  and a 3D translation  $t^k = [x^k, y^k, z^k]$ . Next, we describe our updates for focal length and 6D pose.

**Focal length update.** To build an appropriate focal length

update rule, we take into account the fact that it should remain strictly positive throughout the update iterations. We consider update rules that are multiplicative, *i.e.*, they scale an initial guess  $f^0$  by a sequence of multiplications. Let  $f^k$  be the current estimate of the focal length at iteration  $k$  and  $v_f^k$  be the focal length update predicted by the network  $F$  (see eq. (1)). We define the updated focal length  $f^{k+1}$  as the multiplication,

$$f^{k+1} = e^{v_f^k} f^k. \quad (3)$$

The sequence of multiplicative updates can be written as  $f^{k+1} = e^{\sum_{i=1}^k v_f^i} f^0$ , where  $f^0$  is the initial focal length and  $v_f^i, i = 0, \dots, k-1$  are the individual updates. An alternative to the above strategy would be enforcing positivity of the focal length update via a sigmoid function instead of an exponential function. We found the exponential and sigmoid functions to behave similarly, but the sigmoid update requires setting an additional scale parameter. Hence, we opted for the simpler exponential updates as described in eq. (3).

**6D pose update.** For the update of the 6D pose, we build on the update rule introduced in DeepIM [24] that disentangles 3D rotation and 3D translation updates. In more detail, the network  $F$  is trained to predict a translation of the projected object center into the image  $[v_x^k, v_y^k]$  (measured in pixels), and a ratio  $v_z^k$  of the camera-to-object depth between the observed and the rendered image. The 3D translation of the object is then updated from the quantities  $[v_x^k, v_y^k, v_z^k]$  predicted by network  $F$ , taking into account the nonlinear projection equations derived from the camera model. In [24] the focal length is known and fixed. In our scenario the focal length is not fixed and we replace the known fixed focal length with the *predicted* focal length  $f^{k+1}$ . In detail, the updated 3D translation  $[x^{k+1}, y^{k+1}, z^{k+1}]$  of the object with respect to the camera is obtained as :

$$x^{k+1} = \left( \frac{v_x^k}{f^{k+1}} + \frac{x^k}{z^k} \right) z^{k+1} \quad (4)$$

$$y^{k+1} = \left( \frac{v_y^k}{f^{k+1}} + \frac{y^k}{z^k} \right) z^{k+1} \quad (5)$$

$$z^{k+1} = v_z^k z^k, \quad (6)$$

where  $[v_x^k, v_y^k, v_z^k]$  are the object translation updates predicted by network  $F$  as part of  $\Delta\theta$  (eq. 1),  $[x^k, y^k, z^k]$  is the 3D translation vector of the relative camera-object pose at iteration  $k$ ,  $[x^{k+1}, y^{k+1}, z^{k+1}]$  is the new updated 3D translation vector, and  $f^{k+1}$  is the updated focal length of the camera given by eq. (3).

To obtain the update of the rotation component of the object pose we use directly the prediction of the alignment network  $F$  in a multiplicative update, which does not depend on the focal length. In particular, we parametrize the

rotation update using two 3-vectors  $v_{R,1}^k, v_{R,2}^k$  that define the rotation matrix  $R(v_{R,1}^k, v_{R,2}^k)$  by Gram-Schmidt orthogonalization as described in [50]. This parametrization was found to work well for different prediction tasks [50] including 6D object pose estimation [22]. The resulting update rule is then written as

$$R^{k+1} = R(v_{R,1}^k, v_{R,2}^k) R^k, \quad (7)$$

where  $R^{k+1}$  is the new updated object rotation,  $R^k$  is the current object rotation, and  $R(v_{R,1}^k, v_{R,2}^k)$  is the rotation matrix obtained by Gram-Schmidt orthogonalization from the two 3-vectors  $v_{R,1}^k, v_{R,2}^k$  predicted by the alignment network  $F$  as part of  $\Delta\theta_k$ . Note that this rotation update is similar to the one used in DeepIM [24].

### 3.3. Pose and focal length training loss

We now present our network training loss, where we assume the training data consist of image and aligned model pairs. Note that a training pair may be a real image with a manually aligned model or a rendered image of a model under a specified 6D pose and focal length. Given input parameters  $\theta^k$ , the output parameters  $\theta^{k+1}$  are fully defined by the network outputs  $\Delta\theta$  given by eq. (1) and the differentiable update rules described by eqs. (3)-(7) in the previous section. In the following, we consider a single network iteration and denote  $\theta = \{R, t, f\}$  as the estimated parameters. For jointly learning to estimate the 6D pose and the focal length, we use the following loss that penalizes errors in the output 6D pose predictions  $(R, t)$  and the estimated focal length  $f$ :

$$\begin{aligned} \mathcal{L}(\theta, \hat{\theta}) &= \mathcal{L}_{pose}((R, t), (\hat{R}, \hat{t})) \\ &+ \alpha \mathcal{L}_{focal}((R, t, f), (\hat{R}, \hat{t}, \hat{f})), \end{aligned} \quad (8)$$

where  $\theta = \{R, t, f\}$  are the estimated pose and focal length parameters,  $\hat{\theta} = \{\hat{R}, \hat{t}, \hat{f}\}$  are the ground truth pose and focal length parameters,  $\mathcal{L}_{pose}$  is a loss that penalizes errors in the 6D pose estimate,  $\mathcal{L}_{focal}$  is our novel loss function that jointly takes into account the errors in the focal length and the 6D predicted pose, and  $\alpha$  is a scalar hyper-parameter. This loss is written for a single instance, but our model is trained to minimize the average loss over all training images. We now describe the individual losses  $\mathcal{L}_{focal}$  and  $\mathcal{L}_{pose}$ .

**Focal length loss.** We use the following focal length loss:

$$\mathcal{L}_{focal} = \beta \mathcal{L}_H(f, \hat{f}) + \mathcal{L}_{DR}((R, t, f), (\hat{R}, \hat{t}, \hat{f})), \quad (9)$$

where  $\mathcal{L}_H$  is Huber regression loss,  $\mathcal{L}_{DR}$  is disentangled reprojection loss and  $\beta$  is a scalar hyper-parameter. The individual terms are explained next. The Huber regression loss  $\mathcal{L}_H$  measures the errors between the estimated and the



ground truth focal length using a logarithmic parametrization of the focal length following the recommendations from Grabner *et al.* [12] for better training:

$$\mathcal{L}_H(f, \hat{f}) = \|\log(f) - \log(\hat{f})\|_H, \quad (10)$$

where again  $\hat{f}$  is the ground truth focal length and  $f$  is the focal length estimated by our model.

While using only the loss  $\mathcal{L}_H$  is possible for training our model, we found better results are obtained by also considering the 2D errors of the projected 3D model in the image using the current estimates of the focal length and object 6D pose. We first define the reprojection error:

$$\mathcal{L}_{proj}((R, t, f), (\hat{R}, \hat{t}, \hat{f})) = \sum_{p \in \mathcal{M}} \|\pi(K(f), R, t, p) - \pi(K(\hat{f}), \hat{R}, \hat{t}, p)\|_1, \quad (11)$$

where  $K(f)$  is the intrinsic camera matrix of our camera model with focal length  $f$ ,  $p \in \mathcal{M}$  are 3D points sampled on the object model,  $\pi(K(f), R, t, p)$  is the projection of a 3D point  $p$  using the current estimates of all the parameters, and  $\pi(K(\hat{f}), \hat{R}, \hat{t}, p)$  is the projection of the same 3D point  $p$  using ground truth parameters. This loss can be seen as the counterpart of the pose loss  $\mathcal{L}_{pose}$  (defined below): instead of penalizing errors in 3D space, it penalizes reprojection errors in the image while also taking into account the estimated focal length  $f$ . However, this loss does not disentangle the effects of the pose and focal length predictions. We thus introduce our disentangled reprojection loss:

$$\mathcal{L}_{DR} = \frac{1}{2} \mathcal{L}_{proj}((R, t, \hat{f}), (\hat{R}, \hat{t}, \hat{f})) \quad (12)$$

$$+ \frac{1}{2} \mathcal{L}_{proj}((\hat{R}, \hat{t}, f), (\hat{R}, \hat{t}, \hat{f})), \quad (13)$$

where each term separately measures the 2D reprojection errors resulting from errors in the 6D pose (the first term) and in the focal length (the second term). This disentanglement leads to faster convergence and better model accuracy, as we show in our ablation results.

**6D pose loss.** For  $\mathcal{L}_{pose}$  (in equation (8)), we build on the loss used in CosyPose [22]. This loss is based on the point-matching loss [24, 47] that measures the error between the alignment of the points on the 3D model  $\mathcal{M}$  transformed with the predicted pose  $(R, t)$  and the ground truth pose  $(\hat{R}, \hat{t})$ . CosyPose [22] extends this loss to take into account object symmetries, and uses the disentanglement ideas of [37] to separate the influence of translation errors along the camera axis, image plane, and rotations. In our approach, we do not consider object symmetries as they are nontrivial to obtain for 3D models in the wild considered in this work. In detail, for the pose loss we utilize the following distance metric between two poses specified by  $\{R_1, t_1\}$

and  $\{R_2, t_2\}$ :

$$D(\{R_1, t_1\}, \{R_2, t_2\}) = \frac{1}{|\mathcal{M}|} \sum_{p \in \mathcal{M}} \| (R_1 p + t_1) - (R_2 p + t_2) \|_1, \quad (14)$$

where  $\|\cdot\|_1$  denotes  $L_1$  norm,  $R_i$  is a rotation matrix,  $t_i$  is a translation vector and  $p \in \mathcal{M}$  is a point sampled from the mesh  $\mathcal{M}$ . Following [22], we disentangle the pose loss as

$$\begin{aligned} \mathcal{L}_{pose} = & D(U(\theta^k, \{v_x^k, v_y^k, v_z^k, \hat{R}^k, \hat{v}_f^k\}), \hat{R}, \hat{t}) \\ & + D(U(\theta^k, \{\hat{v}_x^k, \hat{v}_y^k, v_z^k, \hat{R}^k, \hat{v}_f^k\}), \hat{R}, \hat{t}) \\ & + D(U(\theta^k, \{\hat{v}_x^k, \hat{v}_y^k, \hat{v}_z^k, R^k, \hat{v}_f^k\}), \hat{R}, \hat{t}), \end{aligned} \quad (15)$$

where  $\theta^k$  are the pose and focal length parameters at iteration  $k$ ,  $\hat{R}$  is a ground truth rotation,  $\hat{t}$  is a ground truth translation,  $D$  is a distance defined by Eq. (14) and  $U$  is an update function defined by (2). The main idea of this loss is to separate the influence of translation errors in the  $x - y$  plane, depth alignment errors along the  $z$  axis, and rotation errors. In Eq. (15) the terms  $\{v_x^k, v_y^k, v_z^k, \hat{R}^k, \hat{v}_f^k\}$ ,  $\{\hat{v}_x^k, \hat{v}_y^k, v_z^k, \hat{R}^k, \hat{v}_f^k\}$  and  $\{\hat{v}_x^k, \hat{v}_y^k, \hat{v}_z^k, R^k, \hat{v}_f^k\}$  represent the necessary updates that lead to such loss disentanglement. Here  $[v_x^k, v_y^k, v_z^k]$  are translation updates at iteration  $k$  as predicted by the network  $F$ ,  $R^k$  is a rotation update at iteration  $k$  predicted by network  $F$  and  $\hat{v}_f^k$  is a focal length update at iteration  $k$ .  $\hat{v}_i^k$  and  $\hat{R}^k$  then represent the updates needed to transform the current parameters to the ground truth values, which leads to the disentanglement along each of the dimensions. The first term in Eq. (15) leads to the disentanglement along the  $x - y$  axis, since this term provides the gradients resulting from the  $x - y$  alignment errors. Analogously, the second and third terms provide gradients that arise from depth and rotation alignment errors.

## 4. Experiments

We evaluate our method for focal length and 6D pose estimation on three challenging benchmarks: the Pix3D [39], CompCars [45] and StanfordCars [45] datasets. In the remainder of this section we first introduce the benchmark datasets and give details of the full pose estimation pipeline. Then, in Sec. 4.1 we present the ablation of the main components of the proposed loss function. In Sec. 4.2 we compare our method with the state of the art [12, 13, 44] addressing the same task. Finally, in Sec. 4.3 we discuss the main limitations of our approach.

**Datasets and evaluation criteria.** We consider three real-world in-the-wild datasets depicting objects with known 3D models annotated with ground truth focal length and 6D pose of the object. Following Grabner *et al.* [12], we consider the *bed*, *chair*, *sofa*, and *table* classes in the Pix3D

dataset [39]. The images for each object class are considered as separate datasets. The Stanford cars and CompCars datasets [45] contain images of different car instances. Note that for the Pix3D chair images and both cars datasets, there are hundreds of different object instances in the dataset, which makes the task of recognizing the object instance challenging. We use the standard set of evaluation criteria used by prior work [12, 13, 44] that include the detection accuracy and several 6D pose metrics. The results are reported as median errors (smaller is better) between the prediction and ground truth (e.g., the  $MedErr_R$  is the median rotation error) and accuracies (higher is better), which report the percentage of images with an error smaller than a certain threshold (e.g.,  $Acc_{R\frac{\pi}{6}}$  reports the percentage of test images with the rotation error smaller than  $\frac{\pi}{6}$ ). See the supplemental for a detailed description of all the evaluation criteria.

**The complete pose estimation pipeline.** The first step of our pipeline returns bounding box coordinates for depicted model instances in the input image via a Mask R-CNN detector. One detector is trained for each object class. For each detected instance, we crop the input image given the bounding box and apply an instance classifier to obtain which 3D model instance to align. In our case we finetune the DINO model [5] as the instance classifier. We align the 3D model instance corresponding to the top classifier score. The classifier achieves top-1 retrieval accuracy of 62.1% for Pix3D, 71.2% for Stanford Cars and 79.0% for CompCars datasets. Next, we estimate the coarse 6D pose and focal length using the full image, bounding box, and retrieved 3D model instance. Finally, the refiner FocalPose model iteratively refines the estimates for  $N$  iterations given the coarse estimates.

#### 4.1. Loss ablation study

In this section we ablate the different components of our proposed loss function. We train the coarse and refinement networks with the three different losses introduced in Section 3.3. We report the results in Table 1. First, our solution (c.) combining the Huber regression loss with the 2D reprojection error taking into account the object 3D model and its 6D pose results in significantly lower errors than simply using the regression loss (a.) used in Grabner *et al.* [12]. Second, our new loss (c.), which disentangles the effects of focal length and pose, results in lower median errors compared to the standard reprojection loss that does not disentangle pose and focal length (b.).

#### 4.2. Comparison to the state-of-the-art

Below we report the results of our approach on the three different datasets and compare with other methods for 6D object pose and focal length estimation [12, 13, 44].

**Pix3D dataset.** We report the average for the four classes

Loss	$MedErr_R$	$MedErr_t \cdot 10$	$MedErr_f \cdot 10$
a. $\mathcal{L}_H$	6.61	1.51	4.17
b. $\mathcal{L}_H + \mathcal{L}_{proj}$	3.28	1.42	1.45
c. $\mathcal{L}_H + \mathcal{L}_{DR}$	<b>2.98</b>	<b>1.29</b>	<b>1.36</b>

Table 1. **Training loss ablation on Pix3D sofa.** The median alignment errors for refinement models trained using different loss functions. Our proposed combination of Huber regression loss with a disentangled reprojection loss (c.) performs best.

(bed, chair, sofa, table) in Table 2 (top). The per-class results are in the **supplementary material**. On average over all classes, our method significantly outperforms the other methods in 5 out of the 8 metrics. In particular, we see a clear improvement in the estimated focal length (almost 11% relative reduction in the median focal length error, from 0.172 to 0.155). We see also a clear improvement in the estimated 3D translation (20% relative reduction in the median 3D translation error, from 0.185 to 0.148). Please note that the 3D translation is related to the focal length because of the focal length/depth ambiguity. These improvements are significant and validate the contribution of our method.

**CompCars and Stanford cars.** A similar pattern of results is shown in Table 2 (middle, bottom) also for the CompCars and Stanford cars datasets that contain hundreds of different car models. Our approach obtains the best results in 4 (CompCars) and 5 (Stanford cars) of the 8 reported metrics. In particular, our method significantly improves the focal length estimates (11% relative reduction on CompCars and 54% relative reduction on Stanford cars) and the 3D translation estimates (10% relative reduction on CompCars and 52% relative reduction on StanfordCars). Again, these improvements are significant and validate the contribution of our method.

**Qualitative results.** We report examples of qualitative results for our method on the four classes of the Pix3D dataset in Fig. 4 and qualitative results on Stanford cars and CompCars datasets in Fig. 5. Please note that the renderings of the predictions (taking into account focal length and object 6D pose) show precise alignment with the observed image for in-the-wild photographs. Notably, these qualitative results demonstrate the robustness of our approach to large object truncation and strong perspective effects. Please see the **supplementary material** for additional qualitative results and comparisons.

#### 4.3. Limitations

There are three main failure modes of our approach, illustrated in Fig. 3. First, we observe high rotation errors for *symmetric objects* such as tables or stools, where the correct orientation is ambiguous. Please note that none of the used evaluation criteria take into account the symmetries of

Method	Dataset	Detection	Rotation		Translation	Pose	Focal	Projection	
		$Acc_{D_{0.5}}$	$MedErr_R$ .1	$Acc_{R_{\frac{\pi}{6}}}$	$MedErr_t$ $\cdot 10^1$	$MedErr_{R,t}$ $\cdot 10^1$	$MedErr_f$ $\cdot 10^1$	$MedErr_P$ $\cdot 10^2$	$Acc_{P_{0.1}}$
[45]	Pix3D	96.0%	7.25	87.8%	2.52	1.76	2.41	6.33	71.5%
[12]-LF		96.2%	6.92	88.4%	1.85	1.30	1.72	3.85	85.5%
[12]-BB		<b>97.7%</b>	6.89	<b>90.8%</b>	1.94	1.30	1.75	3.66	<b>88.0%</b>
Ours		95.5%	<b>4.92</b>	84.1%	<b>1.49</b>	<b>1.09</b>	<b>1.53</b>	<b>2.97</b>	79.2%
[45]	CompCars	<b>98.9%</b>	5.24	97.6%	3.30	2.35	3.23	7.85	73.7%
[12]-LF		98.8%	5.23	97.9%	2.61	1.86	2.97	4.21	95.1%
[12]-BB		<b>98.9%</b>	4.87	98.1%	2.55	1.84	2.95	3.87	<b>95.7%</b>
[13]-TwoStep		-	4.37	98.1%	3.22	1.90	3.79	4.54	90.2%
[13]-GCVNet		-	<b>3.99</b>	<b>98.4%</b>	3.18	1.89	3.76	4.31	90.5%
Ours		98.2%	<b>3.99</b>	98.4%	<b>2.35</b>	<b>1.67</b>	<b>2.65</b>	<b>2.95</b>	93.0%
[45]	Stanford	99.6%	5.43	98.0%	2.33	1.80	2.34	7.46	76.4%
[12]-LF		<b>99.6%</b>	5.38	<b>98.3%</b>	1.93	1.51	2.01	3.72	96.2%
[12]-BB		<b>99.6%</b>	5.24	<b>98.3%</b>	1.92	1.47	2.07	3.25	<b>96.5%</b>
[13]-TwoStep		-	5.09	97.5%	2.29	1.52	2.52	3.78	93.6%
[13]-GCVNet		-	4.92	97.5%	2.20	1.46	2.43	3.65	94.6%
Ours		99.5%	<b>4.44</b>	95.1%	<b>1.00</b>	<b>0.84</b>	<b>1.09</b>	<b>2.55</b>	93.8%

Table 2. **Comparison with the state of the art for 6D pose and focal length prediction** on the Pix3D, CompCars and Stanford cars datasets. **Bold** denotes the best result among directly comparable methods. Our approach outperforms other competing methods in 4/5 out the 8 reported metrics on all three datasets. Clear improvements (ranging from 10% to 50% relative reduction in the median error) are obtained in the focal length (“Focal”) and 3D translation (“Translation”) estimates (shaded columns) on all three datasets validating our approach and demonstrating our method deals well with the focal length/depth ambiguity.




Figure 3. **Main failure modes** are: (a) symmetric objects, (b) local minima, and (c) incorrect 3D models identified by the object detector.

objects. Second, our iterative alignment procedure can get stuck into a *local minima* where the predicted object model in the predicted configuration is reasonably aligned but the errors are still high, *e.g.*, because the object is flipped upside down. This failure could be mitigated by running our approach from multiple initializations or running our refinement network on better coarse estimates. Finally, we ob-

serve that in some situations the 3D model retrieved by our pipeline is incorrect. These failure modes lead to large errors, which explains the lower accuracies measured by the  $Acc_{R_{\frac{\pi}{6}}}$  and  $Acc_{P_{0.1}}$  metrics. Nevertheless, our approach achieves significantly lower median errors (5 out of the 8 reported metrics) compared to the current state-of-the-art methods, which demonstrates the high precision of our approach outside of these failure modes.

**Broader impact.** Our work has the potential to positively impact practical applications in augmented reality and robotics, among them overlaying artistic effects on viewed objects or for a robotic assistant that can manipulate real-world objects. However, our work could also potentially be used as a component for 3D-assisted manipulation of an image or video via object compositing to create misinformation.

## 5. Conclusion

We have demonstrated successful joint estimation of camera-object 6D pose and camera focal length given a single still image. Key to our success was our extension of render and compare that incorporated the estimated focal length in the iterative update rules and a disentangled loss for training. We have shown that our approach produces lower-error focal length and pose estimates compared to prior art. Our approach can be extended to other camera






Figure 4. **Pix3D qualitative results.** For each example (each row), we show the input image (left), ground truth focal length and pose annotation (center) and our prediction (right). We overlay a rendering of the detected 3D model with the jointly estimated 6D pose and focal length. Notice how our method produces precise alignments for truncated objects (rows 1, 2, 8, 9) and handles large perspective effects (rows 3, 5, 6). Notice also that in row 8 our prediction is better than the manually annotated ground truth.

intrinsic parameters besides focal length, including differ-




Figure 5. Example qualitative results on the **CompCars** (rows 1-4) and **Stanford cars** (rows 5-8) datasets.

ent forms of camera distortions, provided they can be reliably rendered. This work opens up the possibility of downstream applications in augmented reality/computer graphics and reasoning over “in-the-wild” articulated and interacted objects in video.

**Acknowledgements.** This work was partly supported by the European Regional Development Fund under the project IMPACT (reg. no. CZ.02.1.01/0.0/0.0/15\_003/0000468), the Ministry of Education, Youth and Sports of the Czech Republic through the e-INFRA CZ (ID:90140), the French government under management of Agence Nationale de la Recherche as part of the “Investissements d’avenir” program, reference ANR19-P3IA-0001 (PRAIRIE 3IA Institute).



## References

- [1] Alex M Andrew. Multiple view geometry in computer vision. *Kybernetes*, 2001. 2
- [2] Mathieu Aubry, Daniel Maturana, Alexei Efros, Bryan Russell, and Josef Sivic. Seeing 3D chairs: exemplar part-based 2D-3D alignment using a large dataset of CAD models. In *CVPR*, 2014. 2
- [3] Herbert Bay, Tinne Tuytelaars, and Luc Van Gool. Surf: Speeded up robust features. In *ECCV*, pages 404–417. Springer, 2006. 2
- [4] Xavier P Burgos-Artizzu, Matteo Ruggero Ronchi, and Pietro Perona. Distance estimation of an unknown person from a portrait. In *ECCV*, pages 313–327. Springer, 2014. 2
- [5] Mathilde Caron, Hugo Touvron, Ishan Misra, Hervé Jégou, Julien Mairal, Piotr Bojanowski, and Armand Joulin. Emerging properties in self-supervised vision transformers. In *ICCV*, 2021. 6
- [6] Qian Chen, Haiyuan Wu, and Toshikazu Wada. Camera calibration with two arbitrary coplanar circles. In *ECCV*, pages 521–532. Springer, 2004. 2
- [7] Alvaro Collet, Manuel Martinez, and Siddhartha S Srinivasa. The moped framework: Object recognition and pose estimation for manipulation. *The international journal of robotics research*, 30(10):1284–1306, 2011. 2
- [8] A Collet and S S Srinivasa. Efficient multi-view object recognition and full pose estimation. In *ICRA*, pages 2050–2055, May 2010. 2
- [9] Markéta Dubská, Adam Herout, Roman Juránek, and Jakub Sochor. Fully automatic roadside camera calibration for traffic surveillance. *IEEE Transactions on Intelligent Transportation Systems*, 16(3):1162–1171, 2014. 2
- [10] Olivier Faugeras. *Three-dimensional computer vision: a geometric viewpoint*. MIT press, 1993. 2
- [11] David A. Forsyth and Jean Ponce. *Computer Vision - A Modern Approach, Second Edition*. Pitman, 2012. 2
- [12] Alexander Grabner, Peter M Roth, and Vincent Lepetit. Gp2c: Geometric projection parameter consensus for joint 3d pose and focal length estimation in the wild. In *CVPR*, pages 2222–2231, 2019. 2, 5, 6, 7, 12, 14
- [13] Yaohang Han, Huijun Di, Hanfeng Zheng, Jianyong Qi, and Jianwei Gong. Gcvnet: Geometry constrained voting network to estimate 3D pose for fine-grained object categories. In *PRCV*, pages 180–192. Springer, 2020. 2, 5, 6, 7, 12
- [14] R. I. Hartley and A. Zisserman. *Multiple View Geometry in Computer Vision*. Cambridge University Press, ISBN: 0521540518, second edition, 2004. 2
- [15] Kaiming He, Georgia Gkioxari, Piotr Dollar, and Ross Girshick. Mask R-CNN. In *ICCV*, Oct 2017. 3, 12
- [16] Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. Deep residual learning for image recognition. In *CVPR*, pages 770–778, 2016. 11, 12
- [17] S Hinterstoisser, S Holzer, C Cagniart, S Ilic, K Konolige, N Navab, and V Lepetit. Multimodal templates for real-time detection of texture-less objects in heavily cluttered scenes. In *ICCV*, pages 858–865, Nov. 2011. 2
- [18] Tomáš Hodaň, Martin Sundermeyer, Bertram Drost, Yann Labbé, Eric Brachmann, Frank Michel, Carsten Rother, and Jiří Matas. BOP challenge 2020 on 6D object localization. *ECCVW*. 2
- [19] Yinlin Hu, Joachim Hugonot, Pascal Fua, and Mathieu Salzmann. Segmentation-driven 6D object pose estimation. In *CVPR*, pages 3385–3394, 2019. 2
- [20] Wadim Kehl, Fabian Manhardt, Federico Tombari, Slobodan Ilic, and Nassir Navab. SSD-6D: Making RGB-based 3D detection and 6d pose estimation great again. In *ICCV*, pages 1521–1529, 2017. 2
- [21] Diederik P. Kingma and Jimmy Ba. Adam: A method for stochastic optimization. In Yoshua Bengio and Yann LeCun, editors, *ICLR*, 2015. 11
- [22] Y. Labbe, J. Carpentier, M. Aubry, and J. Sivic. CosyPose: Consistent multi-view multi-object 6D pose estimation. In *ECCV*, 2020. 2, 3, 4, 5, 11
- [23] Vincent Lepetit, Francesc Moreno-Noguer, and Pascal Fua. Eppn: An accurate O(n) solution to the PnP problem. *IJCV (International Journal of Computer Vision)*, 81:155–166, 2009. 2
- [24] Yi Li, Gu Wang, Xiangyang Ji, Yu Xiang, and Dieter Fox. DeepIM: Deep iterative matching for 6D pose estimation. In *ECCV*, pages 683–698, 2018. 2, 3, 4, 5, 12
- [25] Tsung-Yi Lin, Piotr Dollár, Ross Girshick, Kaiming He, Bharath Hariharan, and Serge Belongie. Feature pyramid networks for object detection. In *CVPR*, 2017. 12
- [26] David G Lowe. Three-dimensional object recognition from single two-dimensional images. *Artif. Intell.*, 31(3):355–395, 1987. 1, 2
- [27] D G Lowe. Object recognition from local scale-invariant features. In *CVPR*, volume 2, pages 1150–1157 vol.2, Sept. 1999. 1, 2
- [28] Fabian Manhardt, Wadim Kehl, Nassir Navab, and Federico Tombari. Deep model-based 6D pose refinement in RGB. In *ECCV*, pages 800–815, 2018. 2
- [29] Gaku Nakano. A versatile approach for solving pnp, pnpf, and pnpfr problems. In *ECCV*, pages 338–352. Springer, 2016. 2
- [30] Markus Oberweger, Paul Wohlhart, and Vincent Lepetit. Generalized feedback loop for joint hand-object pose estimation. *TPAMI (IEEE Transactions on Pattern Analysis and Machine Intelligence)*, 2019. 2
- [31] 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
- [32] Georgios Pavlakos, Xiaowei Zhou, Aaron Chan, Konstantinos G Derpanis, and Kostas Daniilidis. 6-DoF object pose from semantic keypoints. In *ICRA*, pages 2011–2018. IEEE, 2017. 2
- [33] Adrian Penate-Sanchez, Juan Andrade-Cetto, and Francesc Moreno-Noguer. Exhaustive linearization for robust camera pose and focal length estimation. *TPAMI (IEEE Transactions on Pattern Analysis and Machine Intelligence)*, 35(10):2387–2400, 2013. 2
- [34] Sida Peng, Yuan Liu, Qixing Huang, Xiaowei Zhou, and Huijun Bao. PVNet: Pixel-wise voting network for 6DoF pose estimation. In *CVPR*, pages 4561–4570, 2019. 2

- [35] 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 *CVPR*, pages 3828–3836, 2017. [2](#)
- [36] Lawrence G Roberts. *Machine perception of three-dimensional solids*. PhD thesis, Massachusetts Institute of Technology, 1963. [1](#), [2](#)
- [37] Andrea Simonelli, Samuel Rota Bulo, Lorenzo Porzi, Manuel López-Antequera, and Peter Kotschieder. Disentangling monocular 3D object detection. In *ICCV*, 2019. [5](#)
- [38] Chen Song, Jiaru Song, and Qixing Huang. HybridPose: 6D object pose estimation under hybrid representations. In *CVPR*, pages 431–440, 2020. [2](#)
- [39] Xingyuan Sun, Jiajun Wu, Xiuming Zhang, Zhoutong Zhang, Chengkai Zhang, Tianfan Xue, Joshua B Tenenbaum, and William T Freeman. Pix3D: Dataset and methods for single-image 3D shape modeling. In *CVPR*, 2018. [2](#), [5](#), [6](#)
- [40] Richard Szeliski. *Computer vision: algorithms and applications*. Springer Science & Business Media, 2010. [2](#)
- [41] Bugra Tekin, Sudipta N Sinha, and Pascal Fua. Real-time seamless single shot 6D object pose prediction. In *CVPR*, pages 292–301, 2018. [2](#)
- [42] Jonathan Tremblay, Thang To, Balakumar Sundaralingam, Yu Xiang, Dieter Fox, and Stan Birchfield. Deep object pose estimation for semantic robotic grasping of household objects. In *CoRL*, 2018. [2](#)
- [43] Roger Tsai. A versatile camera calibration technique for high-accuracy 3D machine vision metrology using off-the-shelf tv cameras and lenses. *IEEE Journal on Robotics and Automation*, 3(4):323–344, 1987. [2](#)
- [44] Yaming Wang, Xiao Tan, Yi Yang, Ziyu Li, Xiao Liu, Feng Zhou, and Larry S Davis. Improving annotation for 3D pose dataset of fine-grained object categories. *arXiv preprint arXiv:1810.09263*, 2018. [5](#), [6](#), [13](#)
- [45] Yaming Wang, Xiao Tan, Yi Yang, Xiao Liu, Errui Ding, Feng Zhou, and Larry S Davis. 3D pose estimation for fine-grained object categories. In *ECCVW*, pages 0–0, 2018. [2](#), [5](#), [6](#), [7](#), [12](#), [13](#), [14](#)
- [46] Scott Workman, Connor Greenwell, Menghua Zhai, Ryan Baltenberger, and Nathan Jacobs. DeepFocal: A method for direct focal length estimation. In *ICIP*, pages 1369–1373. IEEE, 2015. [2](#)
- [47] Yu Xiang, Tanner Schmidt, Venkatraman Narayanan, and Dieter Fox. PoseCNN: A convolutional neural network for 6D object pose estimation in cluttered scenes. In *Robotics: Science and Systems (RSS)*, 2018. [2](#), [5](#)
- [48] Sergey Zakharov, Ivan Shugurov, and Slobodan Ilic. DPOD: 6D pose object detector and refiner. In *CVPR*, pages 1941–1950, 2019. [2](#)
- [49] Yinqiang Zheng, Shigeki Sugimoto, Imari Sato, and Masatoshi Okutomi. A general and simple method for camera pose and focal length determination. In *CVPR*, pages 430–437, 2014. [2](#)
- [50] Yi Zhou, Connelly Barnes, Jingwan Lu, Jimei Yang, and Hao Li. On the continuity of rotation representations in neural networks. In *CVPR*, 2019. [4](#)