



HAL
open science

3D-Aware Ellipse Prediction for Object-Based Camera Pose Estimation

Matthieu Zins, Gilles Simon, Marie-Odile Berger

► **To cite this version:**

Matthieu Zins, Gilles Simon, Marie-Odile Berger. 3D-Aware Ellipse Prediction for Object-Based Camera Pose Estimation. 3DV 2020 - International Virtual Conference on 3D Vision, Nov 2020, Fukuoka / Virtual, Japan. hal-02975379v1

HAL Id: hal-02975379

<https://inria.hal.science/hal-02975379v1>

Submitted on 23 Oct 2020 (v1), last revised 24 May 2021 (v2)

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.

3D-Aware Ellipse Prediction for Object-Based Camera Pose Estimation

Matthieu Zins Gilles Simon Marie-Odile Berger
Inria, Université de Lorraine, LORIA, CNRS

matthieu.zins@inria.fr gilles.simon@loria.fr marie-odile.berger@inria.fr

Abstract

In this paper, we propose a method for coarse camera pose computation which is robust to viewing conditions and does not require a detailed model of the scene. This method meets the growing need of easy deployment of robotics or augmented reality applications in any environments, especially those for which no accurate 3D model nor huge amount of ground truth data are available. It exploits the ability of deep learning techniques to reliably detect objects regardless of viewing conditions. Previous works have also shown that abstracting the geometry of a scene of objects by an ellipsoid cloud allows to compute the camera pose accurately enough for various application needs. Though promising, these approaches use the ellipses fitted to the detection bounding boxes as an approximation of the imaged objects. In this paper, we go one step further and propose a learning-based method which detects improved elliptic approximations of objects which are coherent with the 3D ellipsoid in terms of perspective projection. Experiments prove that the accuracy of the computed pose significantly increases thanks to our method and is more robust to the variability of the boundaries of the detection boxes. This is achieved with very little effort in terms of training data acquisition – a few hundred calibrated images of which only three need manual object annotation.

1. Introduction

In the last few years, deep learning has invested all the fields of computer vision. That of visual localization has not escaped this wave, but the specific nature of this problem has led to it being approached from a wide variety of angles. On one side, convolutional neural networks (CNNs) have been trained to detect visual beacons independently of the camera viewpoint and environmental conditions [4, 23, 31]. These methods require an accurate geometric model of the scene for both training and actual pose computation, e.g. by PnP [16]. At the opposite side of the spectrum, the six pose parameters are directly obtained from end-to-end CNN regression, thanks to prior training on data associating images




Figure 1. Camera pose estimation from objects.

with ground truth poses [15, 13, 20]. Intermediate or alternative approaches have also emerged, making more or less intensive use of geometric models [14, 30, 34].

While model-free methods perform poorly when computing the pose for images distant from the training data [29], obtaining an accurate geometric model of a multi-object scene can be challenging especially when specular or textured-less objects are considered.

In this paper, we propose a method for coarse pose computation which is robust to large changes in motion and environmental conditions and does not require a detailed model of the scene. This method meets the growing need of easy deployment of robotics or augmented reality applications in any environments, especially those for which no accurate model nor huge amount of ground truth data are available. Our method builds on [7, 8] where a model-based method which operates at the level of objects is proposed. Considering objects as features for pose computation allows us to take advantage of the performance of recognition systems, especially their good invariance to pose and environmental condition changes, while not necessitating systematic retraining of the system when a new scene is considered. In [7, 8], pose is computed from 2D/3D object correspondences with 2D objects detected by any object detection algorithm such as YOLO [26]. Approximating 2D objects with ellipses and 3D objects with ellipsoids allowed the authors to obtain direct solutions for pose computation without initial guess (Fig. 1).

Though the method is very promising, the main source of inaccuracy originates from a poor approximation of objects with an ellipse aligned with the image axes and inscribed in the detection bounding box (BB). In this paper, we go one step further and propose a learning-based method which detects improved elliptic approximations of objects which are coherent with the 3D ellipsoid i.e. that are likely to be the projection of the ellipsoid. This way of detecting elliptic abstractions of objects significantly improves the accuracy of the recovered pose. Our main contributions are as follows:

- A network for predicting ellipses around objects that are coherent with their 3D ellipsoidal abstractions. Its goal is to overcome the weaknesses of directly fitting the ellipses to the axis-aligned BBs. We propose a data augmentation procedure that allows for robustness to box boundaries variability. We also show that, even if the box boundaries contain a part of background that might be learnt by the network, our method is almost unaffected in typical scenarios, where the scene does not change between training and testing.
- We show how the concept of ellipsoidal abstractions of objects and 3D-coherent ellipse predictions can be used for robust pose computation when only a small amount of data is available on the scene. We show that the pose accuracy little depends on the choice of this ellipsoidal abstraction, which makes the method flexible and easy to use in practice. Only three calibrated images need to be annotated by hand to build the ellipsoid cloud. Annotations of the object are then obtained by projection in the training images.

Experiments prove that the accuracy of the computed pose dramatically increases thanks to our 3D-aware ellipse detection procedure and is close to the accuracy obtained with learning-based methods that require an accurate 3D model for training.

2. Background and related works

Pose computation from monocular RGB images is an important problem in computer vision which witnessed a complete renaissance with the emergence of deep learning. Thanks to the ability of such methods to detect features across a wide range of viewpoints, largely independently from environmental conditions, this opens the way towards more robust localization and matching methods, especially able to handle few-textured scenes.

One of the pioneering method in the use of deep learning for pose computation is poseNet [15] where absolute camera pose is regressed. Besides the difficulty of balancing rotation and translation errors in the cost function, a major default of such approaches, pointed out by Sattler [29], is that direct pose regression is more closely related to pose

approximation via image retrieval than to accurate pose estimation via 3D structure. As a result, such methods generalize poorly to trajectories far from the training sequences. Other methods propose to regress dense 3D scene coordinates [2, 3] and estimate the camera pose by solving a PnP problem. These methods were coupled with advanced versions of RANSAC [1] and obtain very good performance, but only on small-scale scenes.

In order to scale to large environments, [33] propose a method that relies on the detection on a set of objects of interest and exploits dense 2D/3D correspondences between them and reference images to estimate the pose. However, their method is limited to planar objects.

Finding the pose of the camera from general shape objects can also be viewed as estimating the objects pose in the camera frame. Many works exist on this subjects [12, 18, 24, 25, 30, 31, 36]. They usually assume to have access to a detailed model of the object in order to generate synthetic training data and, thereby enlarge the 6D pose space where pose can be correctly recovered. SSD-6D [12] extends the idea of 2D object detection and infers 6D pose based on a discrete viewpoint classification while an auto-encoder is used in [30] to recover the object orientation. Another way to infer object pose is by predicting the 2D projections of the corners of the bounding box of the 3D object with a CNN. This avoids the need for a meta-parameter to balance the position and orientation error since the 6D pose can be estimated with PnP from 2D/3D correspondences. In BB8 [25], segmentation is first performed to detect the objects and a CNN then infers the projection of the BBs. Data augmentation with a random background is performed during training to reduce the influence of the scene context. Building on YOLO and BB8, Tekin *et al.* [31] proposed a single-shot deep CNN architecture which directly predicts the 2D projections of the 3D BB vertices in the image.

The BB of detected objects can be also used for localization issues with the idea of decoupling the rotation and the translation. In PoseRBPF [5], and PoseCNN [34] the translation is constrained to be along the vector from the camera center to the center of the BB. T_z is then approximately determined from the ratio between 3D ROI and 2D BB sizes. An auto-encoder trained on the synthetic rendering of the 3D object is used for rotation estimation in PoseRBF, while PoseCNN regresses a quaternion for rotation and a mask, both with the need of an accurate model of the object. Without accurate knowledge of the objects, but under some assumptions on the scene, e.g. objects lying on a flat scene, Li *et al.* [17] used object detections to estimate relative camera poses for viewpoints changes. Wang *et al.* [32] proposed a category-level approach to predict normalized objects coordinates without the need of detailed 3D models. However, the method is not straightforward for new objects from unseen categories and depth information is needed to com-

pute the pose. In the context of autonomous driving, [21] estimates the pose and the dimensions of an object’s 3D BB from a 2D BB and the surrounding image pixels under additional constraints on the box orientation. However, modelling objects as 3D cuboids and the 2D detections by rectangles does not allow to derive closed-form solutions to projection equations and leads to solutions with a high combinatorics.

Modeling 2D/3D objects correspondences based on ellipses/ellipsoids was already used by [28] in the context of multiview reconstruction and by Nicholson *et al.* in [22] in the context of SLAM. Resolution was based on the minimization of a geometric cost function over the six camera parameters, using odometry sensors for initial position and orientation. Recent works have proposed direct solutions without the need of initial prior for pose computation from ellipse-ellipsoid matches: [7] show that the problem of estimating the camera pose from ellipse-ellipsoid correspondences has at most 3 degrees of freedom, since the position can be obtained from its orientation. Direct closed form solution can thus be estimated once the orientation is known. In [8], a direct method for full recovery of the pose from more than 2 ellipse-ellipsoid correspondences was proposed under assumptions satisfied by many robotics applications. In practical experiments, axis-aligned ellipses are inferred from the BBs detected by YOLO. The authors however note that such an elliptic 2D approximation is not always sufficiently accurate and may lead to a significant error on the estimated pose.

Instead of just detecting a rectangular BB of 2D objects, it is also possible to detect elliptic objects. An immediate idea is to perform instance segmentation, as in Mask R-CNN [9] and directly fit an ellipse to the output mask. A recent work [6] shows that such a strategy fails to capture the ellipse orientation and proposes a CNN for ellipse regression. Such a work is typically dedicated to the detection of objects with elliptic shapes. Besides the fact that the detection is not aware of a particular 3D model, another difference with our work is that we are interested in detecting ellipsoids which approximate an object, whose shape may be quite different from a perfect ellipsoid.

3. System overview

Our method uses an ellipsoid cloud as a light scene model, where each ellipsoid approximates an object. The pose computation system consists of three main components (Fig. 2, bottom part). The first component is a CNN designed to detect objects as BBs. The second component is another CNN aimed at predicting, for each BB, the ellipse corresponding to the projection of the object’s ellipsoid according to the camera pose. Unlike global approaches, such as PoseNet, which apply to the whole image, our prediction network uses local patches around each detected object.

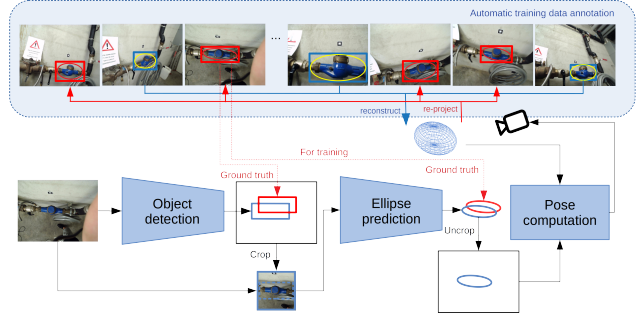


Figure 2. Overview of our full system. The new ellipse prediction block is the key part of our work.

This has the advantage of restricting a lot the change of appearance and should enable our prediction network to better interpolate new views. Finally, the last component aims at calculating the pose from the ellipses thus obtained, the ellipsoid cloud and the labels associated with these features.

The object detection network takes as input an image and returns the BBs of the detected objects, with their class. Many networks have already been proposed to solve this task very efficiently [26, 27] and many types of objects are already recognized by these networks. It is, moreover, easy to fine-tune a pre-trained network in order to make it recognize new objects, as only a few images are required for that purpose. In order to preserve the aspect ratio of the detected object, the crop box is completed with the content of the image to meet the ellipse prediction network’s input requirement. From these cropped images, the network predicts the ellipse parameters expressed in the BB frame. All the ellipses are then transferred to the initial image frame according to the coordinates of the BBs. The transferred ellipses, associated with the detection labels, are then used along with the ellipsoid cloud to compute the pose.

Pose computation from one or more ellipse-ellipsoid pairs is based on the work described in [7, 8]. The method used depends on whether or not external orientation data are available. Measuring camera orientation from image content or sensor acquisitions is generally easier to do than estimating the position. Indeed, position sensors are often inaccurate (GPS) or have a limited range and are sensitive to environmental disturbances (magnetic, acoustic sensors, etc.). By contrast, automatic vanishing point detection may lead to accurate orientation measures as long as the scene satisfies the Manhattan world assumption [7]. On the other hand, IMU sensors can be used, limiting drift and expressing orientations in an absolute reference frame by taking into account the Earth’s magnetic field and gravity. A very interesting result from a practical point of view is that knowing the orientation of the camera, a closed form solution makes it possible to deduce its position from a single ellipse-ellipsoid pair [7]. If the orientation is available, we therefore use a RANSAC procedure with minimal sets




Figure 3. Bottom: Images and detections used for reconstruction. Middle: Reconstructed ellipsoid. Top: Projection in new views.

of one ellipse-ellipsoid pair [7]. Multiple association hypotheses may occur in case objects of the same class appear in the image, but this only adds as many correspondences to be included in the RANSAC draws, preserving a linear algorithmic complexity. If orientation data is not available, two ellipse-ellipsoid pairs are at least required to compute the pose. We then use RANSAC with minimal sets of two correspondences [8].

4. Ellipse prediction

4.1. Motivation for 3D supervision

We give here an intuition of why a 2D-only supervision is not sufficient for detecting good ellipses. Pose estimation from ellipse-ellipsoid pairs comes down to aligning cones, which means that the accuracy is heavily influenced by how much the ellipse detection satisfies the real projection of the ellipsoid. In Fig. 3, we simulated a "perfect" 2D detector by finding the minimum-area ellipse containing the object and reconstructed an ellipsoid from three views. We can observe that the 2D detections in other views are not fully coherent with the projection of this ellipsoid. In fact, the deformations caused by perspective, the viewpoints used for reconstruction and the shape of the object have a significant impact. This is why our method relies on 3D supervision and manages to be almost unaffected by these factors.

4.2. Ellipse parameterization

Different parameterizations of an ellipse can be used, among which the following were considered: (i) three points (center and two points corresponding to the axes ends), (ii) five parameters of the quadratic form matrix, (iii) center, axes lengths and orientation angle.

Using the five parameters of the quadratic form gives a minimal representation, because an ellipse has five degrees of freedom. However, these parameters have no physical interpretation which makes the associated L2 distance loss hard to optimize. Parameterizing an ellipse by three points

has the advantage of having a real physical sense. The first point is chosen at the center of the ellipse and the other two points at the end of each axis. An inconvenient of this representation is that there is no implicit constraint to force orthogonality. Instead, we could choose to represent the ellipse with its center, axes lengths and orientation angle. Both parameterizations have ambiguities due to the ellipse symmetries. A canonical form can be obtained by enforcing the first axis to be the largest one, its endpoint to be in the right half-circle and the second axis endpoint to be in the bottom half-circle. Also, the angle is measured between the horizontal axis and the largest axis. This parameterization, however, implies discontinuities in the orientation angle when the ellipse is almost vertical. Moreover, L2 loss functions with non-normalized heterogeneous quantities are hard to optimize (e.g. the camera position and orientation parameters in PoseNet [15]) are hard to optimize.

To tackle these issues, we use the center, axes and orientation parameterization but with a classified, rather than regressed, orientation angle [21]. More specifically, this angle is discretized into n overlapping bins distributed between $-\frac{\pi}{2}$ and $\frac{\pi}{2}$. Our network predicts the corresponding bin but also regresses a correction angle as a value added to the mid-angle of the predicted bin.

4.3. MultiBin Loss

The center and dimensions of the ellipse are both optimized using a L2 loss, respectively

$$\mathcal{L}_{center} = \frac{1}{N} \|c_{gt} - c_{pred}\|^2 \quad (1)$$

$$\mathcal{L}_{dim} = \frac{1}{N} \|d_{gt} - d_{pred}\|^2 \quad (2)$$

where c_{gt} and c_{pred} are the ground truth and predicted centers, d_{gt} and d_{pred} are the ground truth and predicted dimensions, and N is the number of images in the batch.

As proposed in [21], the angular part has been split into both classification and regression. We use a standard *cross-entropy* loss for the classification part (\mathcal{L}_{bin}) and a correction loss is added to reduce the difference between the predicted angle and the ground truth angle in each of the bins that overlaps with this angle (one or two bins).

$$\mathcal{L}_{correction} = -\frac{1}{n_{\theta^*}} \sum_i^n o_i \cdot \cos(\theta^* - \sigma_i - \Delta\theta_i) \quad (3)$$

where θ^* is the ground truth angle, n_{θ^*} is the number of bins that overlap the ground truth angle (1 or 2), o_i is either 0 or 1 and indicates if bin i overlaps the ground truth angle, σ_i is the mid-angle of the i -th bin and $\Delta\theta_i$ is the predicted correction angle that is applied to the mid-angle of the bin. This loss has effects only on the correction angles of the bins that actually overlap with the ground truth angle.




Figure 4. Ellipse prediction network.

The final loss of our network is a weighted combination of each individual loss.

$$\mathcal{L} = \alpha(\mathcal{L}_{center} + \mathcal{L}_{dim}) + (\mathcal{L}_{bin} + \beta\mathcal{L}_{correction}) \quad (4)$$

The parameters α and β are used to balance the loss intensities and were determined empirically. We used the same values for all experiments (respectively 0.01 and 1.0).

4.4. Network architecture

Our architecture (Fig. 4) includes a base convolutional network used for features extraction, followed by four parallel branches. We used a pre-trained VGG-19 network as base and just cut off the final *softmax* part, which is traditionally used for classification. Then, the parallel branches take these features as input and consist of three-layers fully connected networks with *ReLU* as non-linear activations.

The *center* branch predicts two values corresponding to the center coordinates. The *dimensions* branch regresses the width and height of the ellipse. The *bin* branch predicts n values, where n is the number of bins. These values correspond to scores or confidences of belonging to a specific bin and are then passed to a softmax layer. Finally, the *correction* branch regresses $2n$ values, $\cos(\Delta\theta_i)$ and $\sin(\Delta\theta_i)$ for each bin. To get valid cosine and sine values, a normalization step was added at the end of this branch.

4.5. Network architecture validation

We validated our network architecture and loss function on synthetic data, consisting of multiple views of an object with the camera placed on a semi-sphere around it and random background from COCO dataset [19]. We sampled 100 values of azimuth and 20 of elevation for training and test data were generated at intermediate positions. For evaluation, we computed the *IoU* between the predicted ellipse and the ground truth one. We obtained very high *IoUs* (more than 0.95) on the test images, keeping in mind that the viewpoints coverage of the training was quite large. Figure 5 also shows the advantage of the multibin strategy over the direct regression.




Figure 5. Left: Mean *IoUs* obtained on test images for an increasing number of training epochs. Right: Example of predicted ellipse (green). The ground truth ellipse is in red.

5. Data acquisition and augmentation

The only prerequisite for our system is that images of the scene have been acquired and calibrated from various viewpoints, i.e. we know for each training image the intrinsic parameters of the camera as well as its position and orientation in a global frame. This is the minimum prerequisite of any visual pose computation methods using deep learning. In practice, relatively few images are needed (around 200) both to fine-tune the detection network (if required) and to train the prediction network. Obtaining these data can be done in several ways, e.g. by using a differential GPS coupled with orientation sensors, a Structure-from-Motion (SFM) technique, or a marker [35].

5.1. 3D model and training data generation

Once the calibrated images have been obtained, we propose a practical method to get both the ellipsoid cloud and the ground truth boxes and ellipses in all training images by just labeling a few boxes in a small number of them. It consists of the following steps (Fig. 2, top): (i) choose a minimum of three images of the scene showing the object(s) from various viewing angles, (ii) define boxes around the objects visible in these images (blue boxes Fig. 2, top) and associate a label to each box, (iii) fit ellipses to these boxes, with axes parallel to those of the boxes, (iv) build the ellipsoid cloud from the 2D ellipses and projection matrices; an analytic solution was obtained by Rubino *et al.* to solve this problem from only three calibrated views [28], (v) reproject the ellipsoid cloud in all training images by using the ground truth projection matrices. The obtained ellipses (in red in Fig. 2) will be used as ground truth data for training the ellipse prediction CNN. (vi) If needed, the BBs of the projected ellipses (red boxes in Fig. 2) will be used as ground truth BBs to fine-tune the object detection CNN.

The ellipsoids obviously depend on the images chosen for reconstruction. Fortunately, we show in Section 6 that their size and orientation may vary significantly without degrading the method performance. This is due to the fact that the ellipsoids are only intermediate data used to calculate the pose from the detected ellipses, themselves predicted to

match the perspective projection of the ellipsoids. The only critical point is that the ellipsoids used to generate the training data must be those used at runtime to compute the pose. Also, the internal calibration of the camera does not change between training and testing.

5.2. Data augmentation

Data augmentation plays an important role in the training of our prediction network and its generalization with a relatively limited number of annotated images. Several strategies were performed online (during training): (i) Color jittering randomly changes the brightness, contrast and saturation of an image in order to simulate illumination changes. (ii) Blurring filters the images with a randomly-sized Gaussian kernel in order to accommodate different resolutions caused by the object distance. (iii) Shifting randomly translates the images so that the object is not always perfectly centered, which should accommodate noisy object crops. (iv) In-plane rotations as well as (v) perspective deformations (homographies) were added to generate new views of the object. The last two can, for example, simulate a camera which is not held upright, or not aiming at the object center.

6. Experiments

6.1. Metrics

In our experiments, we used three common metrics: reprojection error (in pixels), ADD (as in [10, 12, 25]) and camera pose error (either only the position or the full 6D pose). Unless otherwise specified, the percentages in the tables represent the proportion of test images that satisfy a certain threshold.

6.2. Camera position estimation

We evaluated the accuracy of the camera position estimated using our predicted ellipses on the LINEMOD dataset [10], which provides RGBD images of 15 objects in cluttered environments with ground truth pose. Note that we did not perform specific retraining for object detection in LINEMOD and T-LESS and used the 2D boxes given by the projected 3D object.

Using an initial ellipsoid, obtained by reconstruction from a few images, we generated ellipse annotations and trained our prediction network. We limit this experiment to the camera position because only one object can be used per sequence (the rest of the scene changes). Thus, the camera orientation is assumed to be known. To be realistic, we added a random noise uniformly sampled in $[-2^\circ, 2^\circ]$ on each of its Euler angle at test time. Figure 6 and Tables 1 and 2 show the results obtained in terms of reprojection error, position error and ADD. The experiments show a significant improvement of the position estimation, compared to Gaudilliere *et al.* [7] (inscribed ellipse) and the obtained

Method	Tekin <i>et al.</i> [31]	Gaudilliere <i>et al.</i> [7]				Ours		
		5 px	10 px	15 px	20 px	5 px	10 px	15 px
ape	92.10	95.39	100.0	100.0	100.0	100.0	100.0	100.0
cam	93.24	49.77	94.47	100.0	100.0	94.47	100.0	100.0
can	97.44	57.60	79.26	98.62	100.0	99.54	100.0	100.0
cat	97.41	68.20	98.62	100.0	100.0	90.32	100.0	100.0
driller	79.41	16.13	61.75	90.32	98.62	96.77	100.0	100.0
duck	94.65	89.40	100.0	100.0	100.0	100.0	100.0	100.0
eggbox	90.33	97.70	100.0	100.0	100.0	100.0	100.0	100.0
glue	96.53	54.38	88.02	95.85	99.54	94.93	99.08	100.0
holepunc	92.86	83.41	100.0	100.0	100.0	99.54	100.0	100.0
iron	82.94	17.05	51.15	78.34	93.09	90.78	99.54	100.0
lamp	76.87	18.43	60.37	84.79	97.24	92.63	99.54	100.0
phone	86.07	34.56	70.97	88.48	97.24	90.78	100.0	100.0

Table 1. Reprojection error.

Method	Tekin <i>et al.</i> [31]	Gaudilliere <i>et al.</i> [7]			Ours		
		10%	15%	25%	10%	15%	25%
ape	21.62	18.43	35.94	56.68	64.98	80.65	93.09
cam	36.57	34.10	56.68	84.33	47.47	69.12	87.56
can	68.80	12.90	18.43	31.34	72.81	87.56	98.16
cat	41.82	26.27	37.33	55.30	24.42	41.01	61.29
driller	63.51	42.86	57.14	76.04	88.48	95.85	99.54
duck	27.23	31.34	47.00	67.28	62.67	78.80	93.09
eggbox	69.85	16.59	22.58	40.09	78.34	91.71	98.16
glue	80.02	11.98	23.04	32.72	42.40	57.14	79.72
holepunc	42.63	12.90	20.74	30.88	76.50	88.94	97.70
iron	74.97	16.59	25.81	40.55	76.50	91.24	97.70
lamp	71.11	23.04	35.48	58.99	84.79	94.47	97.24
phone	74.74	22.12	29.03	42.86	64.98	82.03	96.31

Table 2. ADD.




Figure 6. Proportion of positions correctly estimated wrt. the error threshold (in cm) on the horizontal axis. Left: using the inscribed ellipse. Right: using the predicted ellipse.

accuracies are comparable with Tekin *et al.* [31], a state-of-the-art method trained using a detailed model of the object.

6.3. Full 6D pose estimation

We evaluated the performance of our method for the full 6D pose estimation on T-LESS [11], another RGBD dataset providing 20 scenes of texture-less objects. In each scene, a few objects are placed on a board and images are acquired with the camera placed on a semi-sphere of radius 75 cm. Compared to LINEMOD, several objects can be used in a scene, allowing us to recover the full 6D pose of the camera. We tested our method on a scene with 6 objects and 500 images, split in two halves for training and tests. We obtained the ellipsoid of each object directly from the CAD model and generated our training images. Table 3 shows the 6D-pose errors obtained with our method and compares

Method	Gaudilliere <i>et al.</i> [8]	Ours
Error in orientation (in degrees)	3.15 (± 1.96)	2.65 (± 1.40)
Error in position (in cm)	4.09 (± 2.42)	2.23 (± 2.05)

Table 3. Median error and standard deviation of the orientation and position estimates.




Figure 7. Inscribed ellipses (left) vs Predicted ellipses (right). Noisy BBs used for cropping and extracting the ellipses used for pose computation are in green. Ground truth projection of the ellipsoids are in red and ground truth objects BBs are in blue. Note that, despite noisy crops, the predicted ellipses fit much better to the ground truth projections.

it with Gaudilliere *et al.* [8] (inscribed ellipses). This confirms the improvement made by our predicted ellipses, that was previously demonstrated on the position only.

6.4. Important points analysis

6.4.1 Robustness to detection noise

We show here that a major strength of our learning-based approach is its robustness to noisy object detections. As explained before, we use a two-step approach for detecting an object in the form of an ellipse. First, the object is detected with a 2D box, and then, this box is transformed into an oriented ellipse. The object detections step is not the focus of this work, but is nevertheless crucial. In the previous experiments, this step was completely simulated using a 2D box given by the projection of the 3D model. This is not totally realistic and, thus, we measured the influence of noisy objects detections on the pose estimation. On the one hand, it is easy to see that spatial noise on the BB has a direct impact on its inscribed ellipse. On the other hand, our decoupling between the box and the ellipse provides a much better robustness to variance in the detection (Fig. 7 and 8). This is especially true for the objects marked with the arrows in Fig. 7. Even though the crop passed to the prediction network does not contain the whole object, the inferred ellipses are still correct. This robustness is mostly achieved thanks to our data augmentation which randomly shifts the image (equivalent to shifting the BB before cropping).

6.4.2 Influence of the reconstructed ellipsoid

In this experiment, we illustrate the influence of the ellipsoid chosen as approximation of the target object. For that, we used three different ellipsoids (Fig. 9). The first




Figure 8. Influence of noisy BB detections. Left: Orientation error (in degrees). Right: Position error (in cm). Both horizontal axes represent the half-range of the noisy horiz. and vert. shifts applied on the 2 corners defining the BB.

Metric Threshold	Reprojection error 5 pixels error	Position error 5 cm	ADD 10% of diameter
Ellipsoid 1	96.77	94.47	84.79
Ellipsoid 2	96.77	94.93	84.33
Ellipsoid 3	96.31	92.17	82.49

Table 4. Results obtained with different ellipsoids.




Figure 9. The three different ellipsoids used in our experiment.

one (green) was obtained by reconstruction from five images, the second one (blue) by directly fitting to the CAD model and the third one (red) was the same ellipsoid manually stretched and rotated. The results, available in Table 4, are very similar, which indicates that the choice of ellipsoid has no real influence and thus, that the network is able to learn the projection of more-or-less any ellipsoid. This is an important point of our method, because this decoupling between the ellipse and the object box detection allows us to use, for example, only a sub-part of the object.

6.4.3 Comparison with instance segmentation

The previous experiments clearly show the superiority of our method compared to the axis-aligned ellipses. To get around this axes alignment problem, an alternative could be to fit an ellipse to a mask predicted by an instance-segmentation network. We chose Mask R-CNN [9], which obtains state-of-the-art results for this task and fine-tuned it on the *driller*. Table 5 shows that the ellipses fitted to the predicted masks (*Ell. Mask R-CNN*) or even to the ground truth masks (*Ell. GT masks*) do not reach the same level of accuracy as our predicted ellipses. This confirms our intuition presented in section 4.1.

Method	Position error			ADD			
	5 cm	10 cm	15 cm	10%	15%	20%	25%
Ell. Mask R-CNN	59.91	94.93	99.08	38.71	55.76	68.20	81.57
Ell. GT masks	74.65	100.0	100.0	47.00	66.36	82.95	95.39
Ours	94.47	100.0	100.0	84.79	94.01	98.62	99.54

Table 5. Comparison with an ellipse directly fitted to an object segmentation.

Metric	Reprojection error 5 pixels		Position error 5cm		ADD 10% of diameter	
	inscr.	pred.	inscr.	pred.	inscr.	pred.
ape	94.99	96.74	89.72	93.99	29.07	46.87
driller	21.17	17.15	47.81	52.19	25.91	24.45

Table 6. Comparison between the use of inscribed or predicted ellipses in case of a new scene configuration, without retraining.

6.4.4 New scene configuration

We explain here how our method may adapt to new scene configurations, where some objects are moved between training and runtime, without the need to retrain the system. The underlying idea is to exploit the capabilities of the initially trained predictor to infer non axis-aligned ellipses which are then used to build the ellipsoid from a few images. By reconstructing the scene from these ellipses we should get ellipsoids close to the initial ones and thus avoid having to train the network again. This strategy has been tested on LINEMOD test images of the same sequence, in which the scene has been changed and the target object rotated. An "easy" (the ape) and a "hard" (the driller) objects were chosen in this experiment, depending on how much background is visible in the object crop and how much the object shape changes with a change of viewpoint. Table 6 shows that our method still outperforms the direct BB fitting for the easy object (ape) and reaches similar results for the hard ones (driller). Even if our approach of cropping the object allows us to be insensitive to most of the background, some of it is still visible in the crop. This shows a limit to our light 3D supervision. Having access to a 3D model (even an approximated one, but not just an ellipsoid) would enable us to be more independent to the background by generating random background during training.

6.5. Application on a real scenario

We demonstrate our full system on a real case using the WatchPose dataset [35], which provides pose annotated images of 10 objects in industrial environments. For each object, around 200 images are provided. Half of them were acquired "near" the object (around 60 cm) and the rest farther away (1.4 m). This dataset differs from the previous ones and is more representative of a real application. Indeed, the images were acquired with a hand-held smartphone without specific care about motion noise, reflections and shadows.

We follow the protocol explained in Section 5.1 to build the ellipsoids and generate the ellipse annotations. The BBs of the ellipses were used to fine-tune our object detection network whereas the ellipses were used to train our ellipse

Method	Median error (in mm)	Threshold			
		5 cm	10 cm	15 cm	20 cm
Gaudilliere <i>et. al.</i> [7]	84.11	27.87	57.38	77.05	83.61
Ours	30.75	76.27	96.61	98.31	100

Table 7. Position error in the *easy* case.

Method	Median error (in mm)	Threshold			
		5 cm	10 cm	15 cm	20 cm
Gaudilliere <i>et. al.</i> [7]	120.31	14.61	38.20	66.29	85.39
Ours	63.90	34.01	75.00	90.91	96.59

Table 8. Position error in the *hard* case.



Figure 10. Predicted ellipses and ground truth projections of the ellipsoid in WatchPose.

prediction network. Finally, the two trained networks are used to relocalize the camera on test images. We tested two scenarios: an easy one, where a subset of *near* and *far* images were used for training, and a hard one, where training was done only on *near* images and testing on *far* images.

Our experiments show that our method is able to estimate the camera position with a median error around 3 cm in the *easy* case and 6.4 cm for the *hard* one. We compared our approach to Gaudilliere's method [7]. Our 3D-supervised method shows a significant improvement (see Tables 7, 8 and Fig. 10). For comparison with a global method (where the network uses the whole image), we trained PoseNet [15] during 5000 epochs on the *near* images. While the mean position error obtained by testing also on *near* images was quite good (around 3 cm), the generalization to more distant images was not satisfactory at all (a mean position error of 80 cm on *far* images). This confirms that our learning on local object patches has a better generalization ability to handle new viewpoints.

7. Conclusion

In this paper, we proposed a method for object-based pose estimation which does not require an accurate model of the scene. Its main component is a new 3D-aware ellipse prediction network. By learning from different viewpoints, the network is able to map the object appearance to ellipse parameters which are coherent with the projection of the object ellipsoidal abstraction, and thus, improves the computed pose a lot. Three key aspects of the method are its robustness to variance in the box detection boundaries, its good invariance to the chosen ellipsoidal abstraction and its minimal amount of manual annotations required, making the method of large practical interest.

References

- [1] E. Brachmann, A. Krull, S. Nowozin, J. Shotton, F. Michel, S. Gumhold, and C. Rother. DSAC - differentiable RANSAC for camera localization. In *2017 IEEE Conference on Computer Vision and Pattern Recognition, CVPR 2017, Honolulu, HI, USA, July 21-26, 2017*, pages 2492–2500. IEEE Computer Society, 2017.
- [2] E. Brachmann and C. Rother. Learning less is more - 6d camera localization via 3d surface regression. In *2018 IEEE Conference on Computer Vision and Pattern Recognition, CVPR 2018, Salt Lake City, UT, USA, June 18-22, 2018*, pages 4654–4662. IEEE Computer Society, 2018.
- [3] M. Bui, S. Albarqouni, S. Ilic, and N. Navab. Scene coordinate and correspondence learning for image-based localization. In *British Machine Vision Conference 2018, BMVC 2018, Newcastle, UK, September 3-6, 2018*, page 3. BMVA Press, 2018.
- [4] A. Crivellaro, M. Rad, Y. Verdie, K. M. Yi, P. Fua, and V. Lepetit. A Novel Representation of Parts for Accurate 3D Object Detection and Tracking in Monocular Images. In *iccv*, volume 00, pages 4391–4399, Dec. 2015.
- [5] X. Deng, A. Mousavian, Y. Xiang, F. Xia, T. Bretl, and D. Fox. Poserbpf: A rao-blackwellized particle filter for 6d object pose estimation. In A. Bicchi, H. Kress-Gazit, and S. Hutchinson, editors, *Robotics: Science and Systems XV, University of Freiburg, Freiburg im Breisgau, Germany, June 22-26, 2019*, 2019.
- [6] W. Dong, P. Roy, C. Peng, and V. Isler. Ellipse r-cnn: Learning to infer elliptical object from clustering and occlusion, 2020.
- [7] V. Gaudillière, G. Simon, and M.-O. Berger. Camera Relocalization with Ellipsoidal Abstraction of Objects. In *ISMAR 2019 - 18th IEEE International Symposium on Mixed and Augmented Reality*, pages 19–29, Beijing, China, Oct. 2019.
- [8] V. Gaudillière, G. Simon, and M.-O. Berger. Perspective-2-Ellipsoid: Bridging the Gap Between Object Detections and 6-DoF Camera Pose. *IEEE Robotics and Automation Letters*, 5(4):5189–5196, 2020.
- [9] K. He, G. Gkioxari, P. Dollár, and R. B. Girshick. Mask R-CNN. *IEEE Trans. Pattern Anal. Mach. Intell.*, 42(2):386–397, 2020.
- [10] S. Hinterstoisser, V. Lepetit, S. Ilic, S. Holzer, G. R. Bradski, K. Konolige, and N. Navab. Model based training, detection and pose estimation of texture-less 3d objects in heavily cluttered scenes. In K. M. Lee, Y. Matsushita, J. M. Rehg, and Z. Hu, editors, *Computer Vision - ACCV 2012 - 11th Asian Conference on Computer Vision, Daejeon, Korea, November 5-9, 2012, Revised Selected Papers, Part I*, volume 7724 of *Lecture Notes in Computer Science*, pages 548–562. Springer, 2012.
- [11] T. Hodaň, 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. *IEEE Winter Conference on Applications of Computer Vision (WACV)*, 2017.
- [12] 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 *IEEE International Conference on Computer Vision, ICCV 2017, Venice, Italy, October 22-29, 2017*, pages 1530–1538. IEEE Computer Society, 2017.
- [13] A. Kendall and R. Cipolla. Modelling uncertainty in deep learning for camera relocalization. In *IEEE international conference on Robotics and Automation*, pages 4762–4769, 2016.
- [14] A. Kendall and R. Cipolla. Geometric loss functions for camera pose regression with deep learning. In *IEEE Conference on Computer Vision and Pattern Recognition*, pages 5974–5983, 2017.
- [15] A. Kendall, M. Grimes, and R. Cipolla. Posenet: A convolutional network for real-time 6-dof camera relocalization. In *2015 IEEE International Conference on Computer Vision, ICCV 2015, Santiago, Chile, December 7-13, 2015*, pages 2938–2946. IEEE Computer Society, 2015.
- [16] V. Lepetit, F. Moreno-Noguer, and P. Fua. EPnP: An Accurate O(n) Solution to the PnP Problem. *Int. J. Comput. Vis.*, 81(2):155–166, 2009.
- [17] J. Li, D. Meger, and G. Dudek. Context-coherent scenes of objects for camera pose estimation. In *IROS*, 2017.
- [18] Z. Li, G. Wang, and X. Ji. CDPN: coordinates-based disentangled pose network for real-time rgb-based 6-dof object pose estimation. In *2019 IEEE/CVF International Conference on Computer Vision, ICCV 2019, Seoul, Korea (South), October 27 - November 2, 2019*, pages 7677–7686. IEEE, 2019.
- [19] T. Lin, M. Maire, S. J. Belongie, J. Hays, P. Perona, D. Ramanan, P. Dollár, and C. L. Zitnick. Microsoft COCO: common objects in context. In D. J. Fleet, T. Pajdla, B. Schiele, and T. Tuytelaars, editors, *Computer Vision - ECCV 2014 - 13th European Conference, Zurich, Switzerland, September 6-12, 2014, Proceedings, Part V*, volume 8693 of *Lecture Notes in Computer Science*, pages 740–755. Springer, 2014.
- [20] I. Melekhov, J. Ylioinas, J. Kannala, and E. Rahtu. Image-based localization using hourglass networks. In *IEEE International Conference on Computer Vision*, pages 879–886, 2017.
- [21] A. Mousavian, D. Anguelov, J. Flynn, and J. Kosecka. 3d bounding box estimation using deep learning and geometry. In *2017 IEEE Conference on Computer Vision and Pattern Recognition, CVPR 2017, Honolulu, HI, USA, July 21-26, 2017*, pages 5632–5640. IEEE Computer Society, 2017.
- [22] L. Nicholson, M. Milford, and N. Sünderhauf. Quadricslam: Dual quadrics from object detections as landmarks in object-oriented slam. *IEEE RA-L*, 2019.
- [23] M. Oberweger, M. Rad, and V. Lepetit. Making Deep Heatmaps Robust to Partial Occlusions for 3D Object Pose Estimation. In *The European Conference on Computer Vision (ECCV)*, September 2018.
- [24] S. Peng, Y. Liu, Q. Huang, X. Zhou, and H. Bao. Pvnnet: Pixel-wise voting network for 6dof pose estimation. In *IEEE Conference on Computer Vision and Pattern Recognition, CVPR 2019, Long Beach, CA, USA, June 16-20, 2019*, pages 4561–4570. Computer Vision Foundation / IEEE, 2019.
- [25] 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 *IEEE International Conference on Computer Vision, ICCV 2017, Venice,*

- Italy, October 22-29, 2017*, pages 3848–3856. IEEE Computer Society, 2017.
- [26] J. Redmon and A. Farhadi. YOLO9000: better, faster, stronger. In *CVPR*, 2017.
- [27] S. Ren, K. He, R. Girshick, and J. Sun. Faster R-CNN: Towards Real-Time Object Detection with Region Proposal Networks. In *Advances in Neural Information Processing Systems 28*, pages 91–99. 2015.
- [28] C. Rubino, M. Crocco, and A. D. Bue. 3d object localisation from multi-view image detections. *IEEE TPAMI*, 2018.
- [29] T. Sattler, Q. Zhou, M. Pollefeys, and L. Leal-Taixé. Understanding the limitations of cnn-based absolute camera pose regression. In *IEEE Conference on Computer Vision and Pattern Recognition, CVPR 2019, Long Beach, CA, USA, June 16-20, 2019*, pages 3302–3312. Computer Vision Foundation / IEEE, 2019.
- [30] M. Sundermeyer, Z.-C. Marton, M. Durner, M. Brucker, and R. Triebel. Implicit 3D Orientation Learning for 6D Object Detection from RGB Images. In V. Ferrari, M. Hebert, C. Sminchisescu, and Y. Weiss, editors, *Computer Vision – ECCV 2018*, pages 712–729, Cham, 2018. Springer International Publishing.
- [31] B. Tekin, S. N. Sinha, and P. Fua. Real-Time Seamless Single Shot 6D Object Pose Prediction. In *2018 IEEE Conference on Computer Vision and Pattern Recognition, CVPR 2018, Salt Lake City, UT, USA, June 18-22, 2018*, pages 292–301, 2018.
- [32] H. Wang, S. Sridhar, J. Huang, J. Valentin, S. Song, and L. J. Guibas. Normalized object coordinate space for category-level 6d object pose and size estimation. In *IEEE Conference on Computer Vision and Pattern Recognition, CVPR 2019, Long Beach, CA, USA, June 16-20, 2019*, pages 2642–2651. Computer Vision Foundation / IEEE, 2019.
- [33] P. Weinzaepfel, G. Csurka, Y. Cabon, and M. Humenberger. Visual localization by learning objects-of-interest dense match regression. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, June 2019.
- [34] Y. Xiang, T. Schmidt, V. Narayanan, and D. Fox. PoseCNN: A Convolutional Neural Network for 6D Object Pose Estimation in Cluttered Scenes. In *Proceedings of Robotics: Science and Systems*, Pittsburgh, Pennsylvania, June 2018.
- [35] C. Yang, G. Simon, J. See, M.-O. Berger, and W. Wang. WatchPose: A View-Aware Approach for Camera Pose Data Collection in Industrial Environments. *Sensors*, 20(11), May 2020.
- [36] S. Zakharov, I. Shugurov, and S. Ilic. DPOD: 6d pose object detector and refiner. In *2019 IEEE/CVF International Conference on Computer Vision, ICCV 2019, Seoul, Korea (South), October 27 - November 2, 2019*, pages 1941–1950. IEEE, 2019.