

Symmetry Aware Evaluation of 3D Object Detection and Pose Estimation in Scenes of Many Parts in Bulk

Romain Brégier^{1,2}, Frédéric Devernay², Laetitia Leyrit¹, James L. Crowley²

¹ Siléane, Saint-Étienne

² Univ. Grenoble Alpes, Inria, CNRS, Grenoble INP, LIG, F-38000 Grenoble

{r}.{bregier} at sileane.com – {frederic}.{devernay} at inria.com

Abstract

While 3D object detection and pose estimation has been studied for a long time, its evaluation is not yet completely satisfactory. Indeed, existing datasets typically consist in numerous acquisitions of only a few scenes because of the tediousness of pose annotation, and existing evaluation protocols cannot handle properly objects with symmetries. This work aims at addressing those two points. We first present automatic techniques to produce fully annotated RGBD data of many object instances in arbitrary poses, with which we produce a dataset of thousands of independent scenes of bulk parts composed of both real and synthetic images. We then propose a consistent evaluation methodology suitable for any rigid object, regardless of its symmetries. We illustrate it with two reference object detection and pose estimation methods on different objects, and show that incorporating symmetry considerations into pose estimation methods themselves can lead to significant performance gains. The proposed dataset is available at <http://rbregier.github.io/dataset2017>.

1. Introduction

Detecting instances of 3D rigid objects and estimating their poses given visual data is of major interest for applications such as augmented reality, scene understanding and robotics, and has been an open field of research since the early days of computer vision. Nonetheless and despite important progresses in this field, existing 3D pose estimation methods still fundamentally cannot deal with any kind of rigid object, since they rely on the assumption that the pose of a rigid object corresponds to a single 6-degrees-of-freedom rigid transformation [1, 2, 3, 4, 5, 6]. This hypothesis leads to ambiguities when dealing with objects showing some proper symmetries – *i.e.* invariances under some rigid transformations – notably for evaluation purposes, as the pose of such an object can be represented by multiple rigid

transformations. Symmetries are actually common among manufactured objects, and while ad hoc solutions to this issue have been proposed for performance evaluation [4, 6], they consist in relaxing the pose estimation validation criteria and are therefore not suited for applications requiring precise positioning.

On an other level, a scenario of particular practical interest for object detection and pose estimation consists in the so-called *bin-picking* problem, where instances of a rigid object have to be detected and localized within a bin containing many instances in bulk. Despite the active research on 3D pose estimation, the current state of the art for bin-picking is relatively unknown. Secrecy is indeed the norm regarding the performances of industrial solutions, and available object recognition datasets are not representative of this scenario, as they typically consider only a limited number of object instances lying on a flat surface on a given face, because of the cost of pose annotation.

This article tries to address those issues. In section 3, we propose automatic techniques for generating fully annotated real range image and synthetic RGBD datasets of many instances of objects in arbitrary poses, with no redundancy between the acquisitions, contrary to existing approaches. We suggest in section 4 to take the object symmetries into account for both evaluating accuracy and improving performances of existing pose estimation methods. Section 5 is focused on the evaluation problem, for which we formalize performance metrics in order to better deal with scenes containing many object instances. These developments are supported by experiments on two well-established object detection and pose estimation methods, in section 6.

2. Related work on 3D object pose estimation evaluation

Several approaches have been proposed to evaluate 3D object detection and pose estimation for manipulation tasks based on robotic experiments [7, 8, 9, 10], but pose accuracy is difficult to evaluate in such scenario due to the lack




Figure 1: Samples exhibiting the strong correlation of data within existing datasets. From left to right: datasets of Tejani *et al.* [5], Hinterstoisser *et al.* [4, 13], T-LESS [16] and Desk3D [14].

of ground truth, and its online nature makes reproducibility difficult. The use of datasets annotated with ground truth poses of every visible object instance instead enables a more quantitative and reproducible offline evaluation. Over the years, several publicly available datasets emerged and we summarize the characteristics of major ones in table 1. Other 3D object recognition datasets of interest obviously exist (see notably Firman [11]) but we focus here on the case of rigid objects without intra-class variability. Those datasets (except for T-LESS [16]) consist in views of scenes of a few objects lying on a given face on a table, acquired from pan-tilt viewpoints with only limited roll along the sensor axis. As such, they provide useful material for working on tasks such as indoor scene understanding, but their limited variability of poses relative to the camera is not representative of the problem of localizing objects in arbitrary poses.

Moreover, and because manually annotating objects poses in each image is a tedious process, datasets of more than a few tens of images [4, 5, 14, 16, 15] rely on automatic annotation techniques. Objects instances are typically placed at given poses relatively to some fiducial markers, whose automatic detection in the scene enable to compute the poses of instances relatively to the camera. This approach enables to easily generate large annotated datasets, however consisting in numerous acquisitions of a few scenes from different viewpoints (see figure 1). This strong correlation between data samples makes those datasets less representative of the data distribution in a genuine application, and training or performance evaluation on such datasets is therefore likely to suffer from some overfitting effects.

Pose accuracy An essential part of the evaluation methodology consists in selecting a criterion to decide whether or not a pose hypothesis matches a ground truth pose. Such matching criterion is typically defined based on a measure of similarity between the pose hypothesis and the ground truth. Several similarity measures have been considered for this task, such as the *Intersection over Union* ratio of 2D silhouettes, the translational and rotational errors [3], and the

average displacement of vertices of a 3D model of the object [4]. The first criterion, while widely used for 2D object localization, is however not suited for 3D pose estimation, as multiple poses may have similar projected silhouette on the image plane. The latter two are not well defined for objects with symmetry properties, since there is no unique displacement between two poses of a symmetric object. The most widespread workaround to this issue consists in the use of an *ad hoc* dissimilarity measure [4], based on the distance between vertices of a 3D model \mathcal{M} of the object at those given poses

$$\text{avg}_{\mathbf{x}_1 \in \mathcal{M}} \min_{\mathbf{x}_2 \in \mathcal{M}} \|\mathbf{T}_1(\mathbf{x}_1) - \mathbf{T}_2(\mathbf{x}_2)\|, \quad (1)$$

where $\mathbf{T}_1, \mathbf{T}_2 \in SE(3)$ are active rigid transforms representing the poses. However, such criterion remains problematic, as it cannot distinguish between poses of similar 3D shapes, such as the flipped poses of coffee cups depicted figure 3a. Hodan *et al.* [6] recently suggested the use of an *ambiguity-invariant pose error function*, considering a pose hypothesis as valid if and only if it is plausible given the available data. While not denying the interest of dealing with ambiguities, especially in the context of active vision, we consider this approach problematic for evaluation, since numerous applications rely on a precise pose estimation and cannot be satisfied by plausible hypotheses. This is all the more true when dealing with highly occluded objects, for which the number of plausible hypotheses can be infinitely large.

Precision and recall Once able to classify pose results, one can define metrics to quantify the performances of a given algorithm. In the case of pose estimation of a single object instance per scene, performances are typically described by the recognition rate [1, 3, 4], that is the fraction of scenes for which the pose returned is the correct one on the whole dataset. This metric is however insufficient for scenes containing an unknown number of instances, as it provides no information on the false positive results produced. In such case, the performance is typically described in terms of precision and recall [5], averaged over all the data samples. Such approach however focuses on the exhaustive retrieval of every instance in each scene, which might not be the actual objective of use cases dealing with scenes of many instances.

In this article, we propose a methodology to generate both real and synthetic data of many object instances in arbitrary poses, with no dependency between data samples, in order to overcome the limitations of existing datasets, and metrics adapted for performance evaluation in scenes of many object instances, with potential symmetries.

Table 1: Characteristics of typical rigid object detection and pose estimation datasets.

We distinguish between *localization* and *detection* [6], depending on whether the number of instances to retrieve is constant or not. A dataset is considered *redundant* if it contains numerous acquisitions of the same scene.

Dataset	Modality	Problem class	Multiple instances	Multiple objects	Absence of data redundancy	Pose variability	Clutter
Mian <i>et al.</i> [2]	Point cloud	Localization	no	yes	yes	limited	no
Aldoma <i>et al.</i> [12]	Colored point cloud	Detection	no	yes	partial	limited	no
Hinterstoisser <i>et al.</i> [4, 13]	RGBD	Localization	no	yes	no	limited	yes
Desk3D [14]	Point cloud	Detection	no	yes	no	limited	yes
Tejani <i>et al.</i> [5]	RGBD	Localization	yes	no	no	limited	yes
Doumanoglou <i>et al.</i> [15]	RGBD	Detection	yes	yes	no	high	limited
T-LESS [16]	RGBD	Detection	yes	yes	no	high	limited

3. Generating annotated datasets of many instances

We propose in this section two methods to generate datasets of independent scenes of many object instances in arbitrary poses, automatically annotated with the poses of those objects. The first method is based on the annotation of real range data thanks to the use of tagged object instances. The second one relies on computer simulations. We consider in our experiments multiple instances of a single object because it is representative of the bin-picking problem, however multiple objects could be considered in a similar fashion.

3.1. Real data

We cover the surface of each instance of object with a set of unique fiducial markers [17], densely enough such that at least one marker is always visible from any point of view. The pose of a marker relative to its corresponding object instance is assumed to be known, and in our experiments, we ensured the precise positioning of markers by carving their locations on the 3D model of the object, and by producing its instances by 3D printing. This shape modification is not required but is merely performed here for convenience, and precise positioning of markers could be achieved *e.g.* through the use of a specific assembly jig. Using a binocular stereoscopic system, we produce a range image for each scene through the use of a pseudo-random pattern projector and an off-the-shelf stereo matching algorithm. We also acquire intensity images from both camera with a diffuse lighting, depicted in figure 2. This second modality is exploited in order to automatically recover the ground truth poses.

Visible markers can indeed be quite reliably detected, identified and localized within the intensity images. Because each marker is uniquely assigned to a given instance of object, object detection is straightforward: a given instance is considered present in the scene if and only if at least one of its marker is detected. The potential occlusion of every marker of an instance is not a major issue for this approach, as it only occurs when the instance is nearly en-

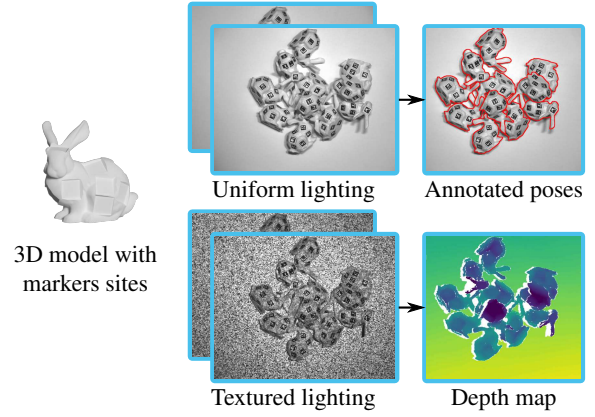


Figure 2: Automatic production of annotated range images of many object instances in bulk. Range images are produced by active stereo matching (left), while ground truth poses are annotated by detecting, in intensity images, markers placed on object instances (right).

tirely occluded – provided that the object is densely enough covered by small enough markers – and pose annotation is of limited interest such case.

For each detected instance, marker detection in intensity image $j \in \{1, 2\}$ provides us with the 2D corners coordinates $\mathbf{p}_{i,j} \in \mathbb{R}^2, i \in M_j$ of the detected markers associated with the instance. Because the pose of those markers relative to the object instance is known, each $\mathbf{p}_{i,j}$ can be put in relation with its corresponding 3D point $\mathbf{X}_{i,j} \in \mathbb{R}^3$ within the object frame. The pose $[\hat{\mathbf{R}}, \hat{\mathbf{t}}]$ of the considered instance – described by a 3×3 rotation matrix $\hat{\mathbf{R}}$ and a 3D translation vector $\hat{\mathbf{t}}$ – can therefore be estimated by solving the multiview perspective-n-point problem

$$[\hat{\mathbf{R}}, \hat{\mathbf{t}}] = \underset{\mathbf{R}, \mathbf{t}}{\operatorname{argmin}} \sum_{j \in \{1, 2\}} \sum_{i \in M_j} \|\pi_j(\mathbf{R}\mathbf{X}_{i,j} + \mathbf{t}) - \mathbf{p}_{i,j}\|^2, \quad (2)$$

where π_j represents the projection from a 3D point in the reference coordinate system onto the image plane of camera j . The annotation procedure is therefore automatic, but we proceed in our experiments to a visual validation stage to

ensure the quality of the ground truth. Validation of about 700 scenes required less than 40min for a single person, and 6% of images were discarded due to failures of the marker detector used, consisting either in false positives, false negatives, mislabeled markers or imprecision in the localization of markers' corners. This failure rate could probably be diminished significantly through more robust markers detection, but we considered it acceptable given that scene acquisition and validation of automatic annotations are quite fast to perform. Additional annotations such as the occlusion rate of each instance are produced by comparison with CGI-renderings of the instances at the annotated poses.

3.2. Synthetic data

While effective, this marker-based approach is intrusive, in that markers remain visible in intensity images, which limit the suitability of this approach for object recognition techniques based on such modality. Those issues can be avoided through the use of synthetic data, for which ideal ground truth can be produced. It is indeed possible to generate synthetic datasets reasonably representative of controlled environments such as an industrial setup for bin-picking, and for our experiments we produce synthetic scenes of bulk objects through the physical simulation of drop of instances into a bin, from which we synthesize top-view RGBD images (see examples figure 4).

Sensor simulation Producing plausible synthetic RGB images is relatively simple thanks to existing ray-tracing renderers, but while range data simulators [18] or depth noise models (e.g. [19]) have been proposed, they remain quite sensor-dependent and their tuning is not trivial. To overcome those issues, we render stereo image pairs of the synthetic scene lit by a virtual pseudo-random pattern projector using the Blender Cycles renderer [20], and perform on those images the same 3D stereo reconstruction as in our real experiments. This approach produces range images with a 3D reconstruction noise visually similar to the one observed in our real data, without the need for an explicit noise model, and despite a coarse modeling of the scene (we discuss this question further in section 6.2). The reconstruction technology we simulate here is similar to the one used in industrial cameras such as the Ensenso N10, and other triangulation-based sensors such as LASER scanners or Kinect v1 could be simulated in a similar fashion, given access to the processing algorithms used by these devices to transform raw input data into depth images.

Based on those approaches, we generated both real and synthetic datasets, considering objects with various symmetry properties, in scenes of various number of instances piled up in bulk (see fig. 4 and supplementary material). The *clutter* dataset is an exception depicting very cluttered sce-

nes, to illustrate the suitability of our annotation procedure in such scenario.

4. Dealing with symmetric objects

Being able to quantify the accuracy of a pose estimate is a prerequisite for evaluation, but usual measures are not fit to deal with symmetric objects. In this section, we suggest the use of a distance suited to any bounded rigid object, and discuss how it can be used within 3D object detection and pose estimation techniques themselves, in order to increase performances.

4.1. Pose distance

Brégier *et al.* [21] recently proposed a pose definition valid for any rigid object – including those with symmetries – as a distinguishable static state of the object. Within their framework, a pose can be identified to an equivalency class $\{\mathbf{T} \circ \mathbf{G} | \mathbf{G} \in G\}$ of $SE(3)$, defined by an active rigid transformation \mathbf{T} and up to any rigid transformation the object is invariant to (the *proper symmetry group* $G \subset SE(3)$).

Given a bounded rigid object, they propose a distance between two poses consisting in the length of the smallest displacement from one pose to an other, the length of a displacement being defined as the RMS displacement of surface points of the object

$$d(\mathcal{P}_1, \mathcal{P}_2) \triangleq \min_{\mathbf{G}_1, \mathbf{G}_2 \in G} \sqrt{\frac{1}{S} \int_S \|\mathbf{T}_2 \circ \mathbf{G}_2(\mathbf{x}) - \mathbf{T}_1 \circ \mathbf{G}_1(\mathbf{x})\|^2 ds}, \quad (3)$$

where S is the surface area of the object. This distance is physically meaningful, and accounts properly for object's symmetries contrary to the widespread measure (1) as illustrated figure 3. Moreover, it can be estimated efficiently in closed form, whereas the measure (1) requires to perform systematically a computation on every vertex of the model. To this aim, the authors propose a representation of a pose \mathcal{P} as a finite set of points $\mathcal{R}(\mathcal{P})$ of at most 12 dimensions, depending on the proper symmetry class of the object. Those representations are listed in table 2 and in particular,

- A pose of an object without proper symmetries (such as the bunny covered with markers used in our experiments) is represented by a 12D point, consisting in the concatenation of the 3D position of its centroid \mathbf{t} and of its rotation matrix \mathbf{R} , anisotropically scaled by a matrix $\mathbf{\Lambda}$ to account for the object's geometry.
- A pose of an object with a finite non-trivial proper symmetry group is represented by several of those 12D points, two in the case of the *brick* object figure 4 which is invariant under the group of proper symmetries $G = \{\mathbf{I}, \mathbf{R}_z^\pi\}$ consisting in the identity transformation and a rotation of $1/2$ turn around a given axis.

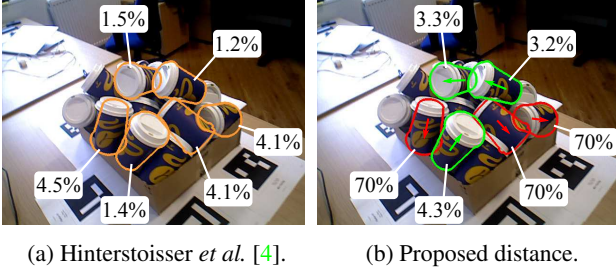


Figure 3: Dissimilarity (in % of the object’s diameter D) between pose hypotheses (overlaid contours) and the corresponding ground truth poses. In this example from Doumanoglou’s dataset [15], each pose hypothesis would be considered as valid according to the widespread criterion of Hinterstoisser *et al.* [4] for symmetric objects (dissimilarity below $10\% \cdot D$). The proposed distance on the other hand properly accounts for the object’s symmetries, which enables to discriminate well true positives (green) and false positives (upside down hypotheses, in red).

- Poses of revolution objects (such as *pepper* and *gear* figure 4) are represented by 6D vectors consisting in the concatenation of the 3D position of their centroids and the direction of their revolution axis, scaled to account for the object’s geometry.

Using this representation, they show that the distance between two poses $\mathcal{P}_1, \mathcal{P}_2$ can be evaluated as the Euclidean distance between any given point $\mathbf{p}_1 \in \mathcal{R}(\mathcal{P}_1)$ and the pointset $\mathcal{R}(\mathcal{P}_2)$:

$$\forall \mathbf{p}_1 \in \mathcal{R}(\mathcal{P}_1), d(\mathcal{P}_1, \mathcal{P}_2) = \min_{\mathbf{p}_2 \in \mathcal{R}(\mathcal{P}_2)} \|\mathbf{p}_2 - \mathbf{p}_1\|. \quad (4)$$

This “nearly Euclidean” structure enables to perform neighborhood queries such as nearest-neighbor or radius searches in an exact or approximate fashion using existing algorithms developed for Euclidean spaces (regular grids, kD-trees, *etc.*). It also enables to average poses of an object with potential symmetries efficiently. Thanks to those possibilities, high-level techniques can be developed in the space of poses, such as the Mean Shift algorithm, which finds local maxima of density distributions.

Given its advantages, we suggest to use this distance for quantifying the accuracy of a pose estimation. It expresses the RMS error in positioning of surface points of the object, and we define a criterion $m(p, t) \triangleq (d(p, t) < \delta)$ to decide whether or not a pose hypothesis p matches a ground truth pose t based on a hard threshold δ , in length unit. The choice of such threshold is very application-dependent, but without further information we chose arbitrarily a value of 10% of the diameter of the smallest sphere enclosing the object and centered on the centroid of object’s surface.

4.2. Considering symmetries for 3D pose estimation


Finding modes An important class of 3D detection and pose estimation methods are based on the aggregation of multiple votes for poses [3, 22, 10, 5, 23, 24] in order to generate pose hypotheses. They typically rely on density estimation, modes finding techniques such as Mean Shift, or clustering operations. Those operations depend on the choice of metric over the pose space, and can actually be performed efficiently considering the distance (3) (see [21] for details). Contrary to a metric suited for $SE(3)$, the suggested one enables to take object symmetries into account directly in the 3D pose estimation method, and to better exploit the aggregate of votes. We test this hypothesis in our experiments by adapting to this approach the method of Drost *et al.* [3], to which we will refer as PPF.

Filtering duplicates More generally, considering symmetries can improve the performance of any 3D pose estimation method thanks to early duplicate removal, as one can ignore pose hypotheses too similar to an already retained one, according to the distance (3). Removing duplicates is a simple way to increase the precision in object detection and pose estimation, and performing it at an early stage benefits to computation time, by limiting the number of pose hypotheses to refine and validate. Equivalently given a constant number of pose hypotheses retained, duplicates removal enables to consider a larger number of truly different pose hypotheses, which should benefit to the recall. We experiment this approach on two methods: PPF evoked previously; and the sliding window technique of Hinterstoisser *et al.* [4] (referred to as LINEMOD+).

5. Performance metrics

In scenes containing many parts, retrieving the pose of every single object instance is not always required. The pose of very occluded objects is indeed often ambiguous, and the retrieval of a limited number of instances is sufficient for many applications – *e.g.* in robotic manipulation. We propose here some adaptations of the usual performance metrics – precision and recall – to take those aspects into account. We discuss the case of a scene containing potentially multiple instances of a rigid object, but multiple scenes should obviously be considered for statistical robustness through the use of aggregated metrics, such as the mean precision and recall. The multi-object case can be handled similarly, as long as no specific relations between the objects are implied (*e.g.* no object categories). Let T be the set of ground truth poses of object instances within the scene, and P the set of result poses retrieved via an object detection and pose estimation method.

Table 2: Classification of every potential group of proper symmetry for a 3D bounded physical object, and expression of pose representatives enabling fast distance computations.



Proper symmetry class	Revolution without rotoreflection invariance	Revolution with rotoreflection invariance	Spherical	Finite
Proper symmetry group G	$\{\mathbf{R}_z^\alpha \alpha \in \mathbb{R}\}$	$\{\mathbf{R}_x^\delta \mathbf{R}_z^\alpha \delta \in \{0, \pi\}, \alpha \in \mathbb{R}\}$	$SO(3)$	$G \subset SO(3)$
Pose representatives $\mathcal{R}(\mathcal{P})$	$(\lambda(\mathbf{R}e_z)^\top, \mathbf{t}^\top)^\top \in \mathbb{R}^6$	$\{(\pm\lambda(\mathbf{R}e_z)^\top, \mathbf{t}^\top)^\top\} \subset \mathbb{R}^6$	$\mathbf{t} \in \mathbb{R}^3$	$\{(\text{vec}(\mathbf{R}\mathbf{G}\mathbf{\Lambda})^\top, \mathbf{t}^\top)^\top \mathbf{G} \in G\} \subset \mathbb{R}^{12}$

Assumptions: Object frame $(\mathbf{O}, e_x, e_y, e_z)$ chosen such as \mathbf{O} to be the center of mass of the object and the e_z axis to be aligned with the symmetry axis for revolution objects.

Notations: \mathbf{I} is the identity rotation. $\mathbf{R}_x^\alpha, \mathbf{R}_y^\alpha, \mathbf{R}_z^\alpha$ represent rotations of angle $\alpha \in \mathbb{R}$ around respectively e_x, e_y, e_z axes.

$$\mathbf{\Lambda} \triangleq (\frac{1}{S} \int_S \mathbf{x} \mathbf{x}^\top ds)^{1/2}, \text{ and } \lambda \triangleq \sqrt{\lambda_r^2 + \lambda_z^2} \text{ for revolution objects where } \mathbf{\Lambda} = \text{diag}(\lambda_r, \lambda_r, \lambda_z).$$

Instances of interest Only a subset $T_o \subset T$ of object instances present in the scene might be of interest to retrieve. In our experiments, we choose $T_o \triangleq \{t \in T | o(t) < \delta_o\}$ as the subset of instances t with an occlusion rate $o(t)$ smaller than $\delta_o = 50\%$. Given T_o , we define the notions of true positives (TP), false positives (FP), and false negatives (FN)

$$\begin{aligned} TP &= \{(p, t) \in P \times T_o | m(p, t) \wedge p = n_P(t) \wedge t = n_T(p)\} \\ FP &= \{p \in P | \neg m(p, n_T(p)) \vee n_P(n_T(p)) \neq p\} \\ FN &= \{t \in T_o | \neg m(t, n_P(t)) \vee n_T(n_P(t)) \neq t\}, \end{aligned} \quad (5)$$

where $n_S(q) \triangleq \text{argmin}_{r \in S} d(q, r)$ is the nearest pose within a set S from a pose q . Duplicates are considered as false positives with this definition, and a result corresponding to an instance whose retrieval is of no interest (within $T \setminus T_o$) is neither considered as a true nor a false positive. The complementary notions of precision and recall can then be derived as

$$\begin{aligned} \text{precision} &= |TP| / (|TP| + |FP|), \\ \text{recall} &= |TP| / (|FN| + |TP|). \end{aligned} \quad (6)$$

Limited number of retrievals To properly handle the case where the number of results $|P|$ is restricted to $n \in \mathbb{N}^*$, we propose to alter the definition of recall as follows

$$\text{recall}_{\leq n \text{ results}} = |TP| / \min(n, |FN| + |TP|). \quad (7)$$

6. Experiments

This section presents experiments performed on object detection and pose estimation using the PPF and LINEMOD+ methods, adapted to deal with symmetric objects (see section 4.2). Through these, we aim to illustrate our different evaluation proposals suited for scenes depicting multiple instances of any rigid object; and also to bring out the benefits of considering symmetries for object detection and pose estimation.

6.1. Protocol

Our experiments are based on the LINEMOD implementation of Stefan Holtzer available in PCL [25], and on our own implementation of PPF. Prior to evaluation, we refine pose hypotheses thanks to a projective ICP algorithm – an additional step often performed to improve pose estimation accuracy [4, 24, 26]. The evaluation therefore focuses on the ability to generate pose hypotheses within the convergence basin of poses of actual instances. We perform experiments on both our real and synthetic datasets (some of those generated using object models from T-LESS [16]), as well as the two bin-picking datasets from Doumanoglou *et al.* [15]. Computation time is not evaluated here, as we used unoptimized code unrepresentative of the original methods.

Post-processing step The addition of a post-processing step (PP) for objects pose hypotheses has been shown to substantially improve performances [4, 27]. We evaluate this effect by considering variants of PPF and LINEMOD+ consisting in keeping the 20 best hypotheses returned by the method, scoring them according to their consistency with the input data and filtering duplicates.

Considering symmetries As discussed in section 4.2, we also evaluate the impact of considering the proper symmetries of the object if any (suffixed by *sym* in table 3). We use the same set of templates for LINEMOD+ in both cases, so as not to bias the comparison.

6.2. Suitability of synthetic data for evaluation

The use of synthetic data raises the question of its suitability for evaluation purposes. To assess the usefulness of our depth sensor simulation procedure for this task, we generate a synthetic dataset of 308 images depicting virtual copies of real scenes of object instances lying on a flat

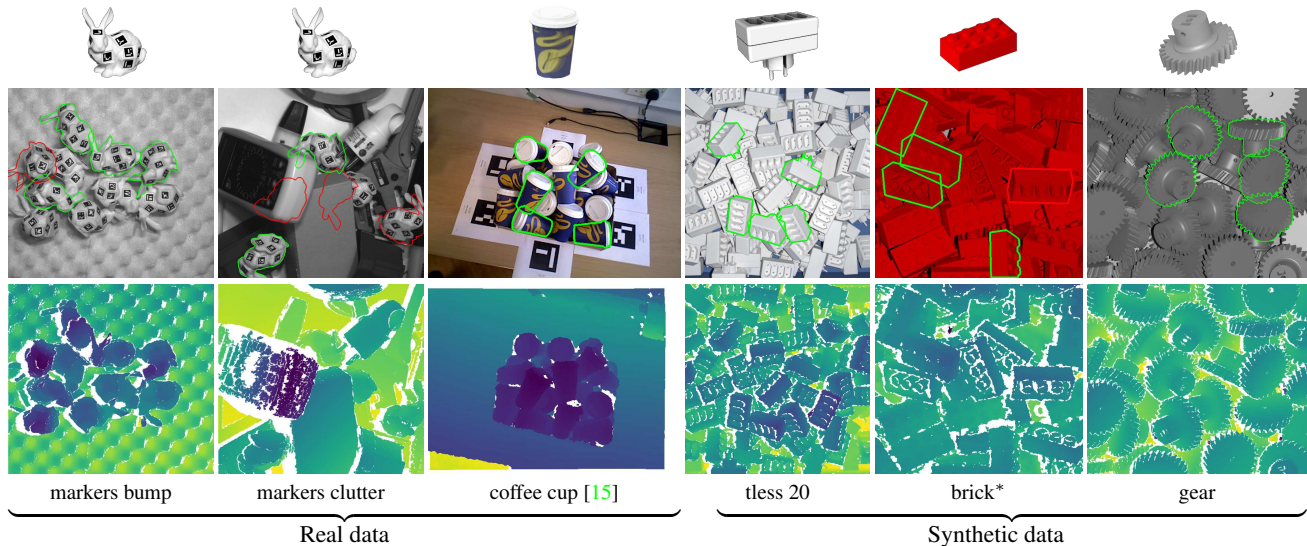


Figure 4: Examples of results from our experiments in object detection and pose estimation. Top: considered objects. Middle: RGB images, with silhouettes contours of the first 5 results returned by the PPF *sym* method with post-processing, classified as true positives (green) or false positives (red). Bottom: corresponding range data. *: the false positive brick highlighted in red is flipped upside down compared to the corresponding ground truth.




Figure 5: Synthetic dataset depicting virtual copies of real scenes for comparison purposes (depth and intensity).

background (*markers flat*), and we perform evaluations on both datasets. These virtual scenes are synthesized automatically based on the pose annotations of real data produced using markers, and a plane detection algorithm for the background. Parameters of the virtual cameras used for range data generation match those of real cameras, and figure 5 depicts an example of synthesized data.

Because we used simple and ideal material and lighting in our CGI renderings, the synthetic data produced is different from the real one, and of slightly better quality (96.6% of the pixels with depth information, against 96.2% for the real data). This quantitatively affects performances and every pose estimation method we tested showed slightly better performances on the synthetic dataset than on the real one (see *markers flat* results in table 3). Nonetheless, synthetic data remains plausible thanks to our depth sensor simulation procedure, and the performances of the evaluated methods compare to one another very similarly on both datasets, as illustrated figure 6. We therefore consider the synthetic data produced to be *realistic enough* for comparative evaluation purposes.

6.3. Effect of accounting for object symmetries

While precision-recall curves such as the ones figure 6 provide a finer-grained understanding of performances, we synthesize our results for the sake of readability in table 3 through different metrics. We consider the *Average Precision* (AP) [28], a usual metric consisting in the area under the precision-recall curve, given the goal of retrieval of instances less than 50% occluded. One might only be interested in the retrieval of a few object instances – *e.g.* for robotics manipulation – therefore we also present the Average Precision given at most $n \in \{1, 3\}$ results returned (AP_n). The formalism enabling to compute those metrics is defined in section 5.

Table 3: Performances of two object detection and pose estimation methods, and their variants exploiting object symmetries (*sym*) if any, and post-processing of the best 20 pose hypotheses. Accounting for symmetries improves performances of both methods.

Dataset		Raw method												With post-processing of the best 20 pose hypotheses (PP)											
		PPF			PPF <i>sym</i>			LINEMOD+			LINEMOD+ <i>sym</i>			PPF			PPF <i>sym</i>			LINEMOD+			LINEMOD+ <i>sym</i>		
		AP	AP1	AP3	AP	AP1	AP3	AP	AP1	AP3	AP	AP1	AP3	AP	AP1	AP3	AP	AP1	AP3	AP	AP1	AP3	AP	AP1	AP3
Real data	markers bump	.35	.66	.43	–	–	–	.85	1.00	.96	–	–	–	.56	.97	.84	–	–	–	.91	1.00	.99	–	–	–
	markers clutter	.34	.36	.31	–	–	–	.57	.67	.53	–	–	–	.52	.70	.52	–	–	–	.68	.83	.69	–	–	–
	markers flat	.26	.54	.31	–	–	–	.83	.99	.97	–	–	–	.46	.94	.76	–	–	–	.90	1.00	.99	–	–	–
	juice [15]	.04	.15	.07	–	–	–	.01	.01	.01	–	–	–	.07	.29	.11	–	–	–	.06	.24	.10	–	–	–
	coffee cup [15]	.16	.76	.53	.28	.96	.85	.03	.37	.10	.08	.37	.17	.23	.98	.90	.30	1.00	.92	.10	.95	.61	.20	1.00	.93
Synthetic data	markers flat	.29	.55	.36	–	–	–	.87	.99	.97	–	–	–	.50	.94	.79	–	–	–	.91	.99	.99	–	–	–
	tless 22	.08	.52	.34	–	–	–	.19	.63	.54	–	–	–	.12	.89	.76	–	–	–	.21	.81	.81	–	–	–
	bunny	.29	.83	.66	–	–	–	.39	.97	.94	–	–	–	.37	.99	.97	–	–	–	.45	.99	.98	–	–	–
	tless 20	.10	.49	.35	.20	.82	.64	.17	.81	.44	.25	.81	.75	.14	.92	.84	.23	.98	.94	.24	1.00	.97	.31	1.00	.99
	tless 29	.15	.69	.40	.19	.76	.56	.14	.71	.34	.20	.71	.50	.21	.90	.76	.23	.91	.79	.20	.88	.84	.26	.92	.86
	brick	.05	.24	.13	.08	.35	.22	.20	.97	.47	.31	.97	.76	.10	.68	.47	.13	.77	.59	.32	.98	.95	.39	.99	.96
	gear	.24	.42	.30	.62	.94	.89	.15	.93	.31	.44	.95	.84	.30	.81	.76	.63	.99	.97	.25	.99	.92	.50	.99	.98
	candlestick	.09	.32	.22	.16	.60	.47	.17	.86	.29	.38	.92	.78	.15	.85	.75	.22	.85	.78	.26	1.00	.96	.49	1.00	1.00
	pepper	.04	.08	.06	.06	.25	.13	.03	.11	.05	.04	.11	.08	.08	.68	.38	.12	.85	.57	.03	.13	.07	.03	.14	.08

AP: Average Precision for the retrieval of instances less than 50% occluded.

AP_n ($n \in \mathbb{N}^*$): Average Precision given at most n results returned for the retrieval of instances less than 50% occluded.

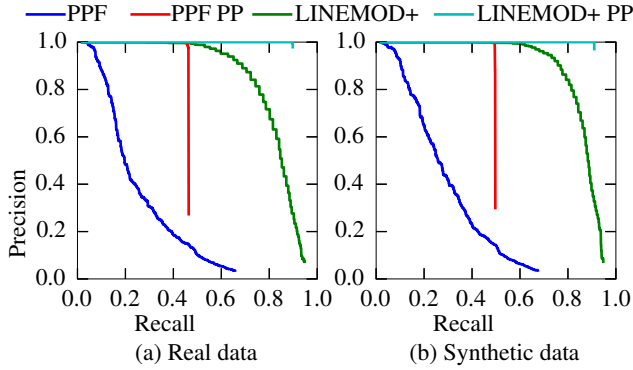


Figure 6: Comparison of performances obtained with real and synthetic data. Precision-recall curves for the retrieval of less than 50% occluded instances.

Results discussion As expected, adding a basic post-processing step (PP) to both PPF and LINEMOD+ significantly improves performances on every dataset (with an average AP improvement of respectively +59% and +60%). Taking object symmetries into consideration also leads to great improvements for the PPF method, with or without postprocessing and regardless of the considered metrics (e.g. respectively +18% and +99% increases of AP3). The PPF method indeed relies on the aggregation of multiple weak pose hypotheses in order to generate stronger ones, and considering symmetries significantly helps in this aggregation, as discussed in section 4.2.

LINEMOD+ also benefits from considering symmetries (+54% and +93% increases of AP respectively with and

without post-processing). Since symmetry considerations are only used to filter out duplicates prior to pose refinement, we should not observe performance improvements when retrieving at most one pose hypothesis without post-processing for this method. The small difference observed here (+1% increase of AP1) is merely an artifact of our implementation, which includes a basic filtering step of pose hypotheses outside the frustum of the camera after pose refinement. Accounting for symmetries however benefits in the other cases to both precision, by removing duplicates, and recall, as it allows to consider more truly different poses for a given number of pose hypotheses, a point we observe in the global performance improvements of LINEMOD+ PP (+54% AP and +11% AP3) which already includes a duplicates filtering step, even without symmetry considerations.

7. Conclusion

We focused in this article on the evaluation of 3D rigid object detection and pose estimation techniques, in scenes containing an arbitrary number of instances, in arbitrary poses. We proposed two methods to generate automatically annotated datasets of such scenes, and metrics suited for the performance evaluation of this generic scenario, even in the case of symmetric objects. We showed how those symmetry considerations could be adapted within existing pose estimation methods themselves, and our experimental results suggest that it leads to significant performance improvements.

References

- [1] A. E. Johnson and M. Hebert, "Using spin images for efficient object recognition in cluttered 3D scenes," *Pattern Analysis and Machine Intelligence, IEEE Transactions on*, vol. 21, no. 5, p. 433–449, 1999. 1, 2
- [2] A. S. Mian, M. Bennamoun, and R. Owens, "Three-dimensional model-based object recognition and segmentation in cluttered scenes," *Pattern Analysis and Machine Intelligence, IEEE Transactions on*, vol. 28, no. 10, p. 1584–1601, 2006. 1, 3
- [3] B. Drost, M. Ulrich, N. Navab, and S. Ilic, "Model globally, match locally: Efficient and robust 3D object recognition," in *Computer Vision and Pattern Recognition (CVPR), 2010 IEEE Conference on*. IEEE, 2010, p. 998–1005. 1, 2, 5
- [4] S. Hinterstoisser, V. Lepetit, S. Ilic, S. Holzer, G. Bradski, K. Konolige, and N. Navab, "Model based training, detection and pose estimation of texture-less 3D objects in heavily cluttered scenes," in *Computer Vision—ACCV 2012*. Springer, 2013, p. 548–562. 1, 2, 3, 5, 6
- [5] A. Tejani, D. Tang, R. Kouskouridas, and T. Kim, "Latent-Class hough forests for 3D object detection and pose estimation," in *Computer Vision—ECCV 2014*. Springer, 2014, p. 462–477. 1, 2, 3, 5
- [6] T. Hodan, J. Matas, and v. Obdržálek, "On evaluation of 6D object pose estimation," in *European Conference on Computer Vision Workshops (ECCVW) 2016, 2016*. 1, 2, 3
- [7] B. K. Horn and K. Ikeuchi, "Picking parts out of a bin," Massachusetts Institute of Technology Artificial Intelligence Laboratory, Tech. Rep., 1983. 1
- [8] M. Liu, O. Tuzel, A. Veeraraghavan, Y. Taguchi, T. K. Marks, and R. Chellappa, "Fast object localization and pose estimation in heavy clutter for robotic bin picking," *The International Journal of Robotics Research*, vol. 31, no. 8, p. 951–973, 2012. 1
- [9] D. Buchholz, S. Winkelbach, and F. M. Wahl, "RANSAM for industrial Bin-Picking," in *Robotics (ISR), 2010 41st International Symposium on and 2010 6th German Conference on Robotics (ROBOTIK)*, Jun. 2010, pp. 1–6. 1
- [10] J. J. Rodrigues, J. Kim, M. Furukawa, J. Xavier, P. Aguiar, and T. Kanade, "6D pose estimation of textureless shiny objects using random ferns for bin-picking," in *Intelligent Robots and Systems (IROS), 2012 IEEE/RSJ International Conference on*. IEEE, 2012, p. 3334–3341. 1, 5
- [11] M. Firman, "RGBD datasets: Past, present and future," in *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition Workshops*, 2016, p. 19–31. 2
- [12] A. Aldoma, F. Tombari, L. Di Stefano, and M. Vincze, "A global hypotheses verification method for 3d object recognition," in *European Conference on Computer Vision*. Springer, 2012, p. 511–524. 3
- [13] E. Brachmann, A. Krull, F. Michel, S. Gumhold, J. Shotton, and C. Rother, "Learning 6D object pose estimation using 3D object coordinates," in *Computer Vision – ECCV 2014*. Springer, 2014, pp. 536–551. 2, 3
- [14] U. Bonde, V. Badrinarayanan, and R. Cipolla, "Robust instance recognition in presence of occlusion and clutter," in *European Conference on Computer Vision*. Springer, 2014, p. 520–535. 2, 3
- [15] A. Doumanoglou, R. Kouskouridas, S. Malassiotis, and T.-K. Kim, "Recovering 6d Object Pose and Predicting Next-Best-View in the Crowd," in *The IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2016. 2, 3, 5, 6, 7, 8
- [16] 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," in *Applications of Computer Vision (WACV), 2017 IEEE Winter Conference on*. IEEE, 2017, pp. 880–888. 2, 3, 6
- [17] S. Garrido-Jurado, R. Muñoz-Salinas, F. J. Madrid-Cuevas, and M. J. Marín-Jiménez, "Automatic generation and detection of highly reliable fiducial markers under occlusion," *Pattern Recognition*, vol. 47, no. 6, p. 2280–2292, 2014. 3
- [18] M. Gschwandtner, R. Kwitt, A. Uhl, and W. Pree, "BlenSor: Blender sensor simulation toolbox," in *Advances in Visual Computing*. Springer, 2011, p. 199–208. 4
- [19] A. Handa, T. Whelan, J. McDonald, and A. J. Davison, "A benchmark for RGB-D visual odometry, 3D reconstruction and SLAM," in *Robotics and automation (ICRA), 2014 IEEE international conference on*. IEEE, 2014, p. 1524–1531. 4
- [20] Blender Online Community, *Blender - a 3D modelling and rendering package*, Blender Foundation, Blender Institute, Amsterdam, 2016. [Online]. Available: <http://www.blender.org> 4
- [21] R. Brégier, F. Devernay, L. Leyrit, and J. Crowley, "Defining the pose of any 3D rigid object and an associated distance," 2016, manuscript submitted for publication. [Online]. Available: <https://hal.inria.fr/hal-01415027> 4, 5
- [22] E. Kim and G. Medioni, "3D object recognition in range images using visibility context," in *Intelligent Robots and Systems (IROS), 2011 IEEE/RSJ International Conference on*. IEEE, 2011, p. 3800–3807. 5
- [23] T. Birdal and S. Ilic, "Point pair features based object detection and pose estimation revisited," in *3D Vision (3DV), 2015 International Conference on*. IEEE, 2015, p. 527–535. 5
- [24] W. Kehl, F. Tombari, N. Navab, S. Ilic, and V. Lepetit, "Hashmod: A hashing method for scalable 3D object detection," in *British Machine Vision Conference*, 2015. 5, 6
- [25] R. B. Rusu and S. Cousins, "3D is here: Point cloud library (PCL)," in *Robotics and Automation (ICRA), 2011 IEEE International Conference on*. IEEE, 2011, p. 1–4. 6

- [26] W. Kehl, F. Milletari, F. Tombari, S. Ilic, and N. Navab, “Deep learning of local RGB-D patches for 3D object detection and 6D pose estimation,” in *European Conference on Computer Vision*. Springer, 2016, p. 205–220. 6
- [27] A. Aldoma, F. Tombari, L. Di Stefano, and M. Vincze, “A global hypothesis verification framework for 3D object recognition in clutter,” *IEEE transactions on pattern analysis and machine intelligence*, vol. 38, no. 7, p. 1383–1396, 2016. 6
- [28] E. Zhang and Y. Zhang, “Average precision,” in *Encyclopedia of Database Systems*, L. Liu and M. T. Özsu, Eds. Boston, MA: Springer US, 2009, pp. 192–193. 7