Active Detection and Localization of Textureless Objects in Cluttered Environments

Marco Imperoli[a], Alberto Pretto[a,*]

[a]*Department of Computer, Control and Management Engineering*
Sapienza University of Rome
Via Ariosto 25, 00185 Rome, Italy

Abstract

This paper introduces an active object detection and localization framework that combines a robust untextured object detection and 3D pose estimation algorithm with a novel next-best-view selection strategy. We address the detection and localization problems by proposing an edge-based registration algorithm that refines the object position by minimizing a cost directly extracted from a 3D image tensor that encodes the minimum distance to an edge point in a joint direction/location space. We face the next-best-view problem by exploiting a sequential decision process that, for each step, selects the next camera position which maximizes the mutual information between the state and the next observations. We solve the intrinsic intractability of this solution by generating observations that represent scene realizations, i.e. combination samples of object hypothesis provided by the object detector, while modeling the state by means of a set of constantly resampled particles.
Experiments performed on different real world, challenging datasets confirm the effectiveness of the proposed methods.

Keywords: Object Detection and Localization, Next-Best-View, Bin-Picking, Chamfer Distance

1. Introduction

The capability to detect and accurately localize objects randomly placed inside an unstructured environment is an essential requirement for highly autonomous robots that need to identify, grasp and manipulate objects in an accurate and reliable way. Many object detection and localization systems have been studied and developed to meet this requirement in the last decades and, also thanks to the increasing interest on mobile manipulators[1], such systems are recently becoming more and more important. Actually, there is a wide range of applications that can highly benefit from a robust and reliable object detection and localization system, from service robotics to robot-aided manufacturing applications. Noteworthy examples include the pick&place and the random bin-picking problems, in which an industrial robot has to grasp and manipulate parts randomly placed inside a bin or in a conveyor belt.

Image-based object detection systems [e.g. 1, 2] usually assume that the searched objects are characterized by salient visual patterns or textures. Unfortunately, these methods can't naturally handle untextured, non-Lambertian objects: this often prevents the use of these methods for industrial and domestic applications, where objects are often untextured and made with non-Lambertian materials as metal or glass. Moreover they usually provide only

Figure 1: The KUKA youBot mobile manipulator used in the experiments.

a rough 2D localization of the objects inside the image. Thanks to the availability of many commercial depth sensors as RGB-D cameras, laser triangulation systems and 3D laser range finders, many vision systems are currently mainly based on 3D measurements [e.g. 3, 4, 5]. Although these systems benefit of a full 3D representation of the workspace, they still have some important limitations: current depth sensors such as time-of-light cameras and structured-light 3D cameras can't easily handle reflective nei-

---

*Corresponding author

[1]We use the definition "mobile manipulator" to refer to a mobile robot with high-level manipulation capabilities, i.e. a mobile base equipped with one ore multiple robotic arms (e.g., Fig. 1).
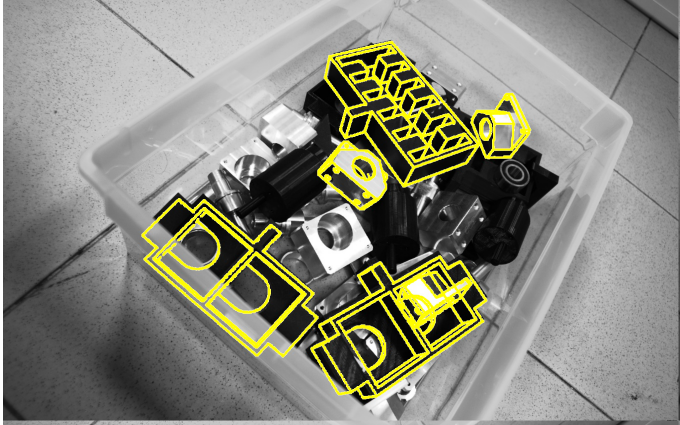
Figure 2: Some object detection and localization results obtained with the proposed active perception strategy (in this examples we are detecting three types of objects).

ther very small objects[2] due to their technology constraints, while high precision laser triangulation systems require expensive devices to move the scanning head over the area of interest, and the acquisition process can take several seconds.

We believe that vision, possibly coupled with depth information used to provide scale and location priors, still remains the primary source of information for detection and localization of objects in challenging environments. In many cases, edge-based algorithms still provide superior performances. In this context, state-of-the-art methods [e.g. 6, 7] usually perform an efficient and exhaustive template matching over the whole image. Templates usually represents shapes extracted from the 3D CAD of the object, seen from a number of viewpoints. Unfortunately, in our experience we found that: (i) The huge 6D searching space imposes a coarse-grained viewpoint discretization, so it is usually required to perform many time-consuming object registration steps over a large set of object candidates in order to accurately detect the true best matches; (ii) Given as input a single view of the scene, often none of the tested state-of-the-art matching algorithms provide as first output the best, true-positive, matches.

To address these problems, in this work we present an effective active perception framework based on the Direct Directional Chamfer Optimization registration method (D^2CO) we recently proposed in [8]. Our method provides: (i) Fast and robust 3D object registration: (ii) An effective active perception strategy able to solve the detection ambiguities and to improve the object localization accuracy.

D^2CO is model-based and works on grey level images: the backbone of our method is represented by the 3D distance transform proposed in [7] we call here Directional Chamfer Distance (DCD) (Sec. 3.3). The DCD is com-

puted using a 3D image tensor that, for each image pixel coordinates and for each (discretized) direction, provides the minimum distance from a template (e.g., a projection of a CAD model) in a joint direction/location space. In our experience, DCD-based object detection (Sec. 3.4) often provides better detection results compared with other state-of-the-art matching algorithms in case of clutter and undetected image edges. The key idea of the D^2CO registration algorithm is to refine the parameters (i.e., the object pose) using a cost function that exploits the DCD tensor in a *direct way*, i.e. by retrieving the costs and the derivatives directly from the tensor (Sec. 4). Being a piecewise smooth function of both the image translation and the (edge) orientation, the DCD ensures a wide basin of convergence. Differently from other registration algorithms based on the iterative closest point method [ICP, 9], D^2CO does not require to re-compute the point-to-point correspondences, since the data association is implicitly encoded in the DCD tensor. We will show that D^2CO outperforms other methods in almost all tests (Sec. 6), while getting a gain in speed of a factor 10 compared to the second most competitive approach.

In many cases, a single view of a scene does not provide sufficient information to detect and accurately locate the objects of interest: objects in the working area can be mutually occluded, moreover different objects may look very similar from different viewpoints. We address these problems firstly by introducing a very simple but effective multi-view extension of the D^2CO algorithm, then by proposing a novel solution to the next-best-view (NBV) problem that aims to solve the detection ambiguities while maximizing the confidence and the localization accuracy: the latter represents the main contribution of this work (Sec. 5).

The NBV problem refers to the sensor placement problem that, given the previous sensor measurements, asks for the next sensor position that results in a better understand-

---

[2]We found that the reflective and the smaller objects used in our experiments are sometimes not even perceived from RGB-D cameras as the Microsoft Kinect, even at small distances.

ing of the scene. Many solution of the NBV problem have been proposed to perform 3D object modeling and scene reconstruction tasks, while surprisingly only a few works deal with both the detection and the localization tasks [e.g. 10, 11] and only a few of these works proposed *practical* strategies to be employed in real robots for everyday applications. Active 3D object detection and localization still remains a very challenging task: it is generally NP-Hard [12] and, without any efficient approximation, it is likely intractable.

In our approach, we assume to use a robot provided with high-level manipulation capabilities (e.g., a mobile manipulator) equipped with an RGB camera mounted on the robot arm end effector (e.g., Fig. 1). The robot is looking for one or more objects of interest using the detection and localization algorithms presented in Sec. 3 and 4, respectively. Our method is inspired by the active perception framework introduced in [13]: we employs a sequential decision process that, for each step, selects the next camera position which maximizes the mutual information between the state (i.e., the object position) and the observations (Sec. 5.1). The active perception formulation provided in [13] is very attractive since: (i) It uses a probabilistic framework, i.e. sensor measurements and placements are not considered noiseless or ideal; (ii) It leads naturally to an iterative algorithm for state estimation; (iii) The convergence of the sequential decision process can be proven. Unfortunately, a direct computation of the mutual information is often intractable since it requires to iterate over the whole observations space and over the whole state space. We address this problem by introducing a novel model-based observations sampling algorithm (Sec. 5.2). The idea is to generate observations by means of "scene realizations", i.e. by sampling *combinations of object hypothesis* provided by the object detector. The next observations can be synthesized by projecting the scene realizations in an efficient way. If a scene is composed by many know objects, as in the case of the random bin-picking problem in an industrial scenario, this method enables to well model the mutual occlusions between objects, since some of the scene realizations could well approximate the real scene. Following the sequential decision framework proposed in [13], we select the next view as the one that maximizes the mutual information (MI) between the system state and the synthesized observations. We model the probability density function over the state by means of a set of particles that represent objects positions: this allows to represent multi-modal probability functions that implicitly enables our system to detect multiple instances of an object type (e.g. Fig. 2). At each new scene observation, we improve the localization accuracy employing the multi-view D^2CO algorithm (Sec. 4.3) over all the collected images. Particles are then resampled in order to remove particles with low weights and to condense them around areas where they get high weights (Sec. 5.3). Experiments and quantitative evaluations (Sec. 6) performed on challenging scenes show the effectiveness of our active detection and localization framework.

2. Related Work

2.1. Object Detection and Localization

Vision systems for learning, detecting and localizing objects have been widely studied in the computer vision and robotics communities for many years.

Early object detection systems [14, 15] relied on 3D object models matched against the input images: due to the reduced computational resources available at that time, these approaches were forced to adopt strong assumptions that limited the effectiveness when applied to real images. More recently, the research in this field moved toward image-based systems: objects models in this case are learnt directly by images of the target objects seen by a number of viewpoints. Viola and Jones [16] used an Adaptive Boosting algorithm applied to a set of Haar-like features efficiently extracted from images; Yeh *et al.* [2] presented a localization and recognition algorithm based on a feature based branch-and-bound approach. Distinctive local features [17, 18] has been used for object detection in [1] and [19], organized in a connected object parts graph and in a deformable configurations of object parts, respectively. An overview of general image-based object recognition and localization techniques can be found in [20], along with a performance evaluation of many types of visual local descriptors used for 6 DoF pose estimation. Very recent advances in deep learning achieve state-of-the-art object detection results, e.g. employing features extracted from a deep, pre-trained convolutional neural network [21].

Even though these methods have shown improving performances in standard computer vision datasets [e.g. 22], they are mainly designed to provide only a rough 2D localization of the object inside the image, without any deeper interpretation about the 3D scene.

Recently, 3D model-based systems are gaining popularity again. Compared to image-based systems, besides ensuring greater accuracy in the localization task, they can generally deal with untextured objects.

A common approach is to apply standard image-based techniques to rendered CAD images. In [23] HOG features [18] learned from a set of synthetic images that represent projections of the models seen from different viewpoints are matched against 2D natural images. Aubry *et al.* [24] used part-based correspondences between 3D CAD models and the test image to recover the object class and its pose. These approaches require the off-line computation of many exemplar models starting from the 3D CAD: Choy *et al.* [25] tried to solve this problem introducing a method for generating 3D CAD model exemplar templates on-the-fly. Another class of model-based systems is based on template matching: models can be represented by means of lines or points to be matched with lines or points extracted from the images. Recent template-based object matching algorithms use spread image gradient orientations saved in

a cache memory-friendly way [6] and multi-path edgelet constellations [26]. We presented a model-based vision system for 3D localization of planar textureless objects [27] that exploits a modified Generalized Hough Transform to select object candidates, and a constrained optimize-and-score registration procedure. In template-matching, the iterative closest point [ICP, 9] is probably the best known point registration method: at each iteration, given the current parameters (i.e., a rigid-body transformation), ICP re-computes the correspondences between points and then updates the parameters as a solution of a least square problem. Fitzgibbon [28] proposed to use the Levenberg-Marquardt algorithm to solve the ICP inner loop, while employing a fast distance lookup based on the Chamfer Distance Transform. Jian and Vemuri [29] proposed a generalization of the ICP algorithm that represents the input point sets using Gaussian mixture models.

The Chamfer Distance Transform [30, 31] has played an important role in many template-based detection and matching algorithms. Even if the original formulation suffers from not being robust to outliers, Chamfer matching and especially its variations still remain powerful tools used for edge-based object detection and matching. Choi and Christensen [32] employed Chamfer matching inside a particle filtering framework for textureless object detection and tracking. Shotton *et al.* [33] presented a matching scheme called Oriented Chamfer Matching (OCM): they proposed to augment the Chamfer distance with an additional channel that encodes the edge points orientations. Cai *at al.* [34] used sparse edge-based image descriptor to efficiently prune object-pose hypotheses, and OCM for hypotheses verification. Recently, Liu *et al.* [7] extended this idea proposing the Fast Directional Chamfer Matching (FDCM) scheme, that exploits a 3D distance transform that provides the minimum distance to an edge point in a joint direction/location space: reported results show that FDCM outperforms OCM. The approach presented in [8] and extended in this paper takes inspiration from [7] and [28].

2.2. Active Object Detection

The concept of *active perception* was introduced by Bajcsy [35], where the term "active" refers to the task of actively change the state of the sensor (e.g., its position, focus, etc...) in order to improve the data interpretation process. Several active perception methods have been proposed to address many robotics problems, e.g. search [36, 37, 38], exploration [39], object modeling [40, 41, 42] and object detection and recognition. A comprehensive survey on many recent works in the filed of active vision in robotics systems can be found in [43].

Active perception applied to the object detection and localization task aims to provide the robot with a sequence of sensor positions and orientations that allows to acquire the most informative observations able to solve the detection ambiguities while maximizing the confidence and the localization accuracy. This application is generally defined next-best-view planning: traditionally developed for the object modeling task [e.g. 44, 40], NBV has been extended to deal with more complex tasks such as the object detection and recognition problems.

The concept of active object recognition has been introduced in [45], where a mobile manipulator has been exploited to move the camera using a motion strategy based on low-level image features (e.g., lines) in order to infer the framed object type. Schiele and Crowley [46] presented an NBV algorithm for the vision based object recognition task where the next view is selected using the mutual information (MI) between the state and the observations. In [47] the next view planning is based on the maximization of the average reduction of the entropy over object hypotheses. Denzler and Brown [13] extended the works presented in [47, 46] by modeling the next view selection task as a sequential decision problem. They proposed to select the action that maximizes the mutual information between the state and the observation, and to update the posterior probability of the state using the Bayes rule. Farshidi *et al.* [48] proposed a multi-camera probabilistic approach for increasing the confidence level in the object recognition task. The next camera positions at each recognition step are selected based on statistical metrics quantifying the quality of the observations, the mutual information and the Cramér-Rao lower bound.

Unfortunately, none of the systems presented above have been tested in realistic scenarios or using challenging datasets.

More recently, Eidenberger and Scharinger [10] proposed an active object detection and localization system that uses both SIFT features and stereo information; the active perception strategy is realized in form of an approximated partially observable Markov decision process (POMDP). Being based on visual features, this method is well suited only for textured objects such as the household items used in the experiments. The POMDP framework has been exploited also in [49]: differently from [10], authors used a non-myopic strategy, and the object detection system is based on 3D features extracted from point clouds acquired with an RGB-D camera. Depth sensors are used also in [11], where the bin-picking problem is faced in an active way: the next views are estimated as a trade off between high information gain and low traveling cost. The information gain is computed by considering the number of discovered unknown 3D cells along the traveled path. Wu *et al.* [50] proposed a recognition system for volumetric shapes that recovers the object category using a Convolutional Deep Belief Network applied to view-based 2.5D depth maps: in order to increase the recognition confidence, this system exploits a next-best-view strategy inspired by [13]. The mutual information to be maximized is computed in an approximated way by sampling shapes to generate hypotheses of the current shape, and then rendering each hypothesis to obtain the depth maps for different viewpoints. Although this system is very promising, it

seems not well suited to deal with cluttered environments.

3. Object Detection

In this section, we describe our object detection approach, that exploits the DCD tensor (Sec. 3.3) in order to extract a set of object candidates (i.e., their 'rough' 3D locations) from an input image (Sec. 3.4), based on the given model template (Sec. 3.1).

3.1. Object Model

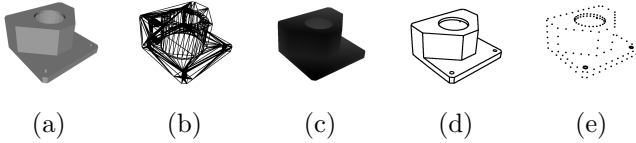(a)   (b)   (c)   (d)   (e)

Figure 3: An example of template extraction from a 3D CAD model: (a) Original Stl 3D CAD model; (b) Wireframe; (c) OpenGL z-buffer; (d) 3D edges template (only visible edges); (d) Raster template.

Edges represent the most informative image features that characterize untextured objects: edges are usually generated by occlusions (depth edges) and high curvature. Given a 3D CAD model of an object (e.g., a Stl file, Fig. 3(a)), we need to extract a 3D template that includes only the *visible* edges (Fig. 3(d)). We start from the 3D model wireframe (Fig. 3(b)), preserving only edges that belong to high curvature parts or to the external object shape, while using the OpenGL z-buffer (i.e., the depth buffer, Fig. 3(c)) to deal with occlusions. Some results of this procedure are shown in the second row of Fig. 9. It is important to note that, in the general case, this procedure should be repeated for each viewpoint, i.e., for each object position.

We finally produce a *rasterization* of this template (Fig. 3(e)), i.e. we extract from the template a set of m sample points $\mathcal{M} = \{\mathbf{o}_1, \ldots, \mathbf{o}_m\} \in \mathbb{R}^3$, in the object reference frame. Typically we employ a rasterization step of $1 - 2$ mm. We also collect another set of m points, $\mathcal{M}' = \{\mathbf{o}'_1, \ldots, \mathbf{o}'_m\} \in \mathbb{R}^3$, where:

$$\mathbf{o}'_i = \mathbf{o}_i + \hat{\tau}(\mathbf{o}_i) \cdot dr \tag{1}$$

$\hat{\tau}(\mathbf{o}_i)$ is a function that provides the unit tangent vector, i.e. the unit direction vector of the 3D edge the raster point belongs to, while dr is a small scalar increment, $dr \ll 1$. Given a transformation $\mathbf{g}_{cam,obj} \in \mathbb{SE}(3)$ from the object frame to the camera frame and $\pi : \mathbb{R}^3 \to \mathbb{R}^2$ a general projection function, we can project the points $\mathbf{o}_i$ and $\mathbf{o}'_i$ on the image plane as:

$$\mathbf{x}_i = \pi(\mathbf{o}_i, \mathbf{g}_{cam,obj}) \qquad \mathbf{x}'_i = \pi(\mathbf{o}'_i, \mathbf{g}_{cam,obj}) \tag{2}$$
$$\mathbf{x}_i, \mathbf{x}'_i \in \mathbb{R}^2, \ i = \{1, \ldots, m\}$$

The idea behind these two set of 3D points is simple: by projecting on the image plane the raster points $\mathcal{M}$ along with the points $\mathcal{M}'$, it is possible to easily recover also the 2D local directions (orientations) of the projected edge points in the image plane.

3.2. Edge Points Extraction

In order to match the edge template extracted from the CAD model, we need to detect edges in the input image. We adopt here the concept of *edgelet*, a straight segment that can be part of a longer, possibly curved, line, extracted using a state-of-the-art detection algorithm for line segment detector called LSD [51]. This algorithm searches the input image for edgelets starting from pixels with higher gradient magnitude, looking in the neighbourhood for pixels with similar gradient directions. A line-support region is therefore selected and validated in order to discard the outliers.

We employ the LSD algorithm on a Gaussian pyramid of the input image, enabling to include in the final set also edgelets that appear in higher scales. This technique increases the sensitivity of the edge detector, at the cost of a reduced accuracy in the localization of the edgelets. For each detected edgelet, we also compute its orientation in the image reference frame. We define as $\mathcal{E} = \{\mathbf{e}_1, \ldots, \mathbf{e}_n\} \in \mathbb{R}^2$ the set of pixels (*edgels*) that belong to edgelets.

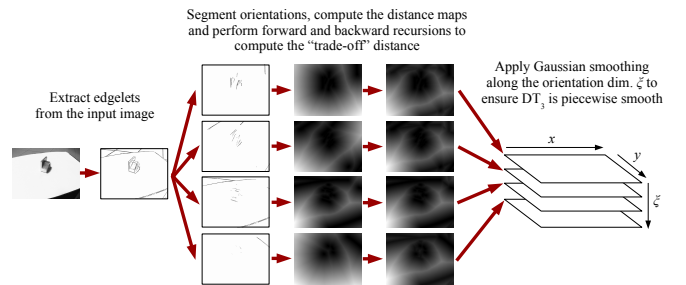### 3.3. Directional Chamfer Distance

Figure 4: Computation of the Directional Chamfer Distance tensor. The ξ coordinate represents the edge discretized directions (orientations).

As introduced above, our approach leverages the Directional Chamfer Distance tensor in both the detection and registration steps. The DCD tensor ($\mathcal{DT}_3$) is represented by an ordered sequence of distance transforms[3] $\mathcal{DT}$, each one representing a *discretized edge direction* ϕ_i, $i = 1, \ldots, q$. The basic idea behind the Directional Chamfer Distance is simple:

1. Divide the set of edgelets computed in Sec. 3.2 into q subsets by quantizing their directions;

---

[3]A distance transform, also called distance map, is an image where each pixel reports the distance to the closest edge pixel.

2. Draw each edgelets set in a different binary image (i.e., an *edge map*);

3. Compute one distance map for each subset using the edge map computed above.

In this way, each map reports the minimum distance from a set of edges that share the same discretized direction. Liu *et al.* [7] extended this idea enabling the DCD tensor $\mathcal{DT}_3$ to encode the minimum distance to an edge point in a joint location and orientation space. Be $\mathbf{x}_i$ and $\mathbf{x}'_i$ the 2D projections on the image plane of the 3D points $\mathbf{o}_i$ and $\mathbf{o}'_i$, respectively (see Eq. 2). For each raster point projection $\mathbf{x}_i$, we can compute its scalar direction (orientation) in the image reference frame as:

$$\xi_i = \Xi(\mathbf{x}_i, \mathbf{x}'_i) = atan\left(\frac{\mathbf{d}_i(1)}{\mathbf{d}_i(0)}\right) \qquad (3)$$

where $\mathbf{d}_i \triangleq \mathbf{x}'_i - \mathbf{x}_i,\ \mathbf{d}_i \in \mathbb{R}^2$. The distance to the closest edge point (edgel) $\mathbf{e}_j$ in a joint location-orientation space can be recovered as:

$$\mathcal{DT}_3(\mathbf{x}_i, \xi_i) = \\ \min_{\mathbf{e}_j \in \mathcal{E}} \left(\|\mathbf{x}_i - \mathbf{e}_j\| + \lambda\|\Phi(\xi_i) - \Phi(o(\mathbf{e}_j))\|_\pi\right) \quad (4)$$

where $o(\mathbf{e}_j)$ is the edgel orientation, $\Phi(.)$ is an operator that provides the nearest quantized orientation ϕ and λ is a weighting factor between the location and orientation distances. The tensor $\mathcal{DT}_3$ that provides the distance reported in Eq. 4 can be easily computed pixel-wise by applying a forward recursion followed by a backward recursion to the sequence of distance maps described above. In the forward recursion, for each pixel $\mathbf{x}_i$ and for each discretized direction ϕ_j, the tensor is updated as:

$$\mathcal{DT}_3(\mathbf{x}_i, \phi_j) = \\ \min\left(\mathcal{DT}_3(\mathbf{x}_i, \phi_j), \mathcal{DT}_3(\mathbf{x}_i, \phi_{j-1}) + \lambda\|\phi_{j-1} - \phi_j\|_\pi\right) \quad (5)$$

similarly, in the backward recursion the tensor is updated as:

$$\mathcal{DT}_3(\mathbf{x}_i, \phi_j) = \\ \min\left(\mathcal{DT}_3(\mathbf{x}_i, \phi_j), \mathcal{DT}_3(\mathbf{x}_i, \phi_{j+1}) + \lambda\|\phi_{j+1} - \phi_j\|_\pi\right) \quad (6)$$

Note that both the recursions should continue also over a full cycle, i.e. until the value for a pixel is not changed. Since our optimization framework employs the tensor $\mathcal{DT}_3$ in a direct way, we need to ensure that $\mathcal{DT}_3 : \mathbb{R}^3 \to \mathbb{R}$ is piecewise smooth. To this end, we smooth the tensor along the direction (orientation) dimension using a simple Gaussian filter. We illustrate each step used to compute the tensor $\mathcal{DT}_3$ in Fig. 4. In all our experiments, we use 60 discretized orientations and we set λ to 100 and σ^2 (the variance the Gaussian filter) to 1.

Since we perform the object detection task without knowing any accurate scale prior, the huge 6D searching space imposes a coarse-grained viewpoint discretization. In order to speedup the process, for each object in the dataset we pre-compute the (projected) raster templates along with their image orientations for a large number of possible 3D locations. Each template includes in such a way a set of image points along with their orientation: by performing a set of lookups on the tensor $\mathcal{DT}_3$, we can compute the average distance as:

$$\mu(\mathcal{DT}_3) = \frac{1}{m}\sum_{i=1}^{m}\mathcal{DT}_3(\mathbf{x}_i, \xi_i) \qquad (7)$$

Finally, we sort the templates for increasing distances: the top rated templates (e.g., see Fig. 5) represent our objects hypothesis (or "object candidates"), to be registered and validated.

4. Object Registration

Once we have obtained a set of object candidates, we need to precisely locate each true positive object, while discarding the outliers. D²CO refines the object position employing a non-linear optimization procedure that minimizes a tensor-based cost function. In Sec. 4.3 we introduce a natural extension of the D²CO algorithm for multi-view object localization.

4.1. Single-view D²CO

Let us express the transformation $\mathbf{g}_{cam,obj}$ in terms of a translation vector $\mathbf{T} = [t_x\ t_y\ t_z]^T$ and, using the axis-angle representation, an orientation vector $\mathbf{\Omega} = [r_x\ r_y\ r_z]^T$, both in $\mathbb{R}^3$. We make explicit this fact using the notation $\mathbf{g}_{cam,obj} = \mathbf{g}(\mathbf{T}, \mathbf{\Omega})$. $\mathbf{R}(\mathbf{\Omega}) \doteq \exp(\widehat{\mathbf{\Omega}})$ is the rotation matrix corresponding to the rotation vector $\mathbf{\Omega}$, where $\widehat{\mathbf{\Omega}}$ is the skew-symmetric matrix corresponding to $\mathbf{\Omega}$ [52]. Given a set of m raster points $\mathbf{o}_i$ extracted from the 3D CAD model, from Eq. 2 we can obtain the corresponding image projections $\mathbf{x}_i$ and from Eq. 3 we can compute their scalar orientations ξ_i. Our optimization procedure aims to find the parameters $(\tilde{\mathbf{T}}, \tilde{\mathbf{\Omega}}) \in \mathbb{R}^6$ that minimize:

$$E(\mathbf{T}, \mathbf{\Omega}) = \frac{1}{2}\sum_{i=1}^{m}\mathcal{DT}_3(\mathbf{x}_i, \xi_i)^2 \qquad (8)$$

While we can assume that, for small viewpoint transformations, the 3D raster points $\mathbf{o}_i$ do not change (i.e., we can neglect changes in the occlusions), this fact does not apply for their image projections $\mathbf{x}_i$. Moreover, Eq. 8 also requires to constantly update the projected (edge) point orientations (Eq. 3).

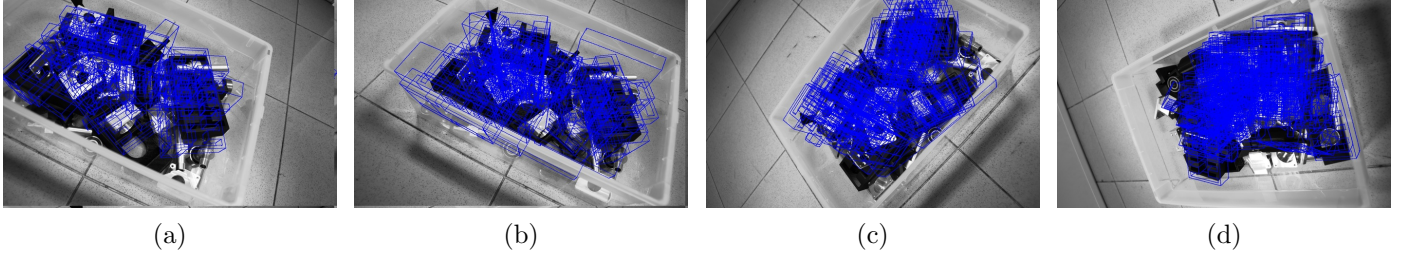In order to apply a non-linear minimization on $E(\mathbf{T}, \mathbf{\Omega})$,

Figure 5: Some examples of object candidates extraction, looking for a single object type, with: (a) top 40 candidates; (b) top 50 candidates; (c) top 60 candidates; (d) top 100 candidates.

we have to compute its derivatives ∇E. Application of one step of the chain rule yields:

$$\nabla E = \sum_{i=1}^{m} \nabla \mathcal{DT}_3 \, \nabla \left[\mathbf{x}_i, \Xi(\mathbf{x}_i, \mathbf{x}_i') \right] \qquad (9)$$

Since $\mathcal{DT}_3$ is only defined at discrete points, its derivatives $\nabla \mathcal{DT}_3$ should be computed in a numerical, approximate way. To this end, we compute the x and y derivatives as the image derivatives of the currently selected distance map and, in a similar way, the derivative along the orientation direction ξ as:

$$\frac{\delta \mathcal{DT}_3}{\delta \xi}(\mathbf{x}, \xi) = \frac{\mathcal{DT}_3(\mathbf{x}, \xi+1) - \mathcal{DT}_3(\mathbf{x}, \xi-1)}{2} \qquad (10)$$

We lookup the $\mathcal{DT}_3$ tensor employing a bilinear interpolation operator, adapted to the 3D nature of $\mathcal{DT}_3$: this enables to improve the level of smoothness of the cost function.
We perform the optimization using the Levenberg - Marquardt algorithm and, as suggested in [28], a Huber loss function in order to reduce the influence of outliers.
Some registration results are reported in Fig. 11.

4.2. The Scoring Function

Some of the selected hypothesis used as initial guess for the registration may represent false positive objects: after the position refinement presented above, we need to employ a metrics that allows us to discard the outliers and to select the best matches. We use a simple but effective scoring function based on local image gradient directions. For each $\mathbf{x}_i$, we can compute its direction ξ_i with Eq. 3: in the ideal case of a perfect match, this direction should correspond to the local gradient direction $\mathcal{I}_\theta(\mathbf{x}_i)$ (up to a rotation of π *radians*), where $\mathcal{I}_\theta$ is the gradient direction image computed directly from the input image. We define the scoring function as:

$$\Psi(\mathbf{g}_{cam,obj}) = \frac{1}{m} \sum_{1}^{m} \left| \cos\left(\mathcal{I}_\theta(\mathbf{x}_i) - \xi_i \right) \right| \qquad (11)$$

Clearly $\Psi(\mathbf{g}_{cam,obj})$ can get values from 0 to 1, where 1 represents the score for a perfect, ideal match.
Differently from the average Directional Chamfer Distance

(Eq. 7): (i) the scoring function in Eq. 11 is not affected by undetected edges and (ii) small errors in the object localization lead to large variations in the score. The latter observation can be interpreted with the fact that the average DCD is more "smooth" respect to Eq. 11.
In our experiments, all good matches (inliers) with no or small occlusions usually obtain a score greater than 0.8.

4.3. Multi-view D^2CO

Assume that we have collected N_v images $\mathcal{I}_j$ of the scene, $j = 1, \ldots, N_v$, taken from different viewpoints. We define $\mathbf{g}_{j,1} = \mathbf{g}_j \in \mathbb{SE}(3)$ as the transformation from the camera frame in position 1 to the camera frame in position j (i.e., the camera position where the first image $\mathcal{I}_1$ has been acquired), with $\mathbf{g}_1 = \mathbf{I}$ ($\mathbf{I}$ is the identity matrix). We assume here that the transformations $\mathbf{g}_j$ are ideal, i.e. error-free. This assumption can be applied also to real systems: industrial robot manipulators, for instance, provide a superior accuracy, thus the error in positioning is negligible. The cost function of the optimization procedure (Eq. 8) can be easily extended to the multi-view case as:

$$E(\mathbf{T}, \mathbf{\Omega}) = \frac{1}{2} \sum_{j=1}^{N_v} \sum_{i=1}^{m^{(j)}} \mathcal{DT}_{3_j} \left(\mathbf{x}_i^{(j)}, \xi_i^{(j)} \right)^2 \qquad (12)$$

$\mathcal{DT}_{3_j}$ in this case is the DCD tensor computed from the image $\mathcal{I}_j$, while $\mathbf{x}_i^{(j)}$ and $\xi_i^{(j)}$, $i^{(j)} = 1, \ldots, m^{(j)}$, are the raster points projections and their directions, respectively, computed at the j-th position. The scoring function in Eq. 11 scales for the multi-view settings in a similar way. Some registration results obtained with the multi-view D^2CO algorithm are reported in Fig. 13.

5. Active Detection and Localization

The proposed object recognition and localization system behaves as a passive receiver of information, since the processed image is not actively selected. Unfortunately, there are many facts that may prevent a single, random view of the scene to provide enough confidence to the object identification and localization algorithms, among others:

- Objects with specific geometry may be highly self occluded for specific viewpoints (e.g., Fig. 9 (c));

- On a cluttered environment, objects are often mutually occluded;

- Different objects look very similar from many viewpoints (e.g., Fig. 9 (a),(b)).

In the next sections, we describe our active perception strategy that, after a first random view of the working space, it sequentially plans the sensing process in order to increase the detection confidence and the accuracy of the localization.

5.1. Next-Best-View Probabilistic Framework

The next-best-view problem can be formulated as a repeated sequence of three steps: action, observation, and state estimation. This sequence is repeated until some termination criteria are met, e.g. when the system reaches a desired detection confidence. Neither actions nor observations (i.e., sensor data) are ideal, so it is natural to use a probabilistic framework: we denote the state of the system at time t with the random variable $\mathbf{l}_t$, the action with the random variable $\mathbf{a}_t$ and the observation with the random variable $\mathbf{z}_t$. In the object detection and localization problem the state could be the objects position, the action a sensor placement (e.g., a manipulator movement) and the observation an image of the working area.

Suppose that at time step $t-1$ we have actively selected the sequence of actions (e.g, the camera placements) up to $\mathbf{a}_{t-1}$. Once the observation $\mathbf{z}_{t-1}$ is made, we can perform the state estimation, updating a *posterior* probability density function (pdf) $p(\mathbf{l}_{t-1}|\mathbf{z}_{t-1}, \ldots, \mathbf{z}_0)$ conditioned on the sensor readings up to time $t-1$. Given the state transition model $p(\mathbf{l}_t|\mathbf{l}_{t-1})$, we can obtain the *prior* pdf at time t as:

$$p(\mathbf{l}_t|\mathbf{z}_{t-1}, \ldots, \mathbf{z}_0) = \int_{l_{t-1}} p(\mathbf{l}_t|\mathbf{l}_{t-1})(\mathbf{l}_{t-1}|\mathbf{z}_{t-1}, \ldots, \mathbf{z}_0)dl_{t-1} \quad (13)$$

For the sake of simplicity, we rewrite the prior pdf at time t in a compact form, $p(\mathbf{l}_t) = p(\mathbf{l}_t|\mathbf{z}_{t-1}, \ldots, \mathbf{z}_0)$. When the state is static[4], e.g. the position of the object does not change over time, the prior pdf at time t is equal to the posterior pdf computed at time $t-1$, $p(\mathbf{l}_t) = p(\mathbf{l}_{t-1}|\mathbf{z}_{t-1}, \ldots, \mathbf{z}_0)$. The next-best-view problem asks at this point: what is the best action $\mathbf{a}_t$ to perform in order to acquire the most informative observation $\mathbf{z}_t$? The solution proposed in [13] is to select the action $\mathbf{a}_t^*$ such that the reduction in the uncertainty of $\mathbf{l}_t$ due to the observation $\mathbf{z}_t$[5] is maximized. In information theory, this reduction of uncertainty is measured by the mutual information (MI) between the prior

pdf $p(\mathbf{l}_t)$ and the pdf $p(\mathbf{z}_t|\bar{\mathbf{a}}_t)$[53], i.e.:

$$I(\mathbf{l}_t; \mathbf{z}_t|\bar{\mathbf{a}}_t) = \\ \int_{\mathbf{l}_t} \int_{\mathbf{z}_t} p(\mathbf{l}_t)p(\mathbf{z}_t|\mathbf{l}_t, \bar{\mathbf{a}}_t) log \left(\frac{p(\mathbf{z}_t|\mathbf{l}_t, \bar{\mathbf{a}}_t)}{p(\mathbf{z}_t|\bar{\mathbf{a}}_t)} \right) d\mathbf{z}_t d\mathbf{l}_t \quad (14)$$

where here and in the rest of the paper we use the notation $\bar{\mathbf{x}}$ to define a realization of a random variable $\mathbf{x}$. Note that Eq. 14 *is not* the conditional mutual information $I(\mathbf{l}_t; \mathbf{z}_t|\mathbf{a}_t)$: $\bar{\mathbf{a}}_t$ is a realization of $\mathbf{a}_t$ and thus it is a parameter of the MI, i.e. $I(\mathbf{l}_t; \mathbf{z}_t|\bar{\mathbf{a}}_t) = I(\mathbf{l}_t; \mathbf{z}_t|\mathbf{a}_t = \bar{\mathbf{a}}_t)$. The optimal action $\mathbf{a}_t^*$ is finally calculated as:

$$\mathbf{a}_t^* = \underset{\bar{a}_t}{\operatorname{argmax}} \, I(\mathbf{l}_t; \mathbf{z}_t|\bar{\mathbf{a}}_t) \quad (15)$$

Once the action has been taken, the new observation $\mathbf{z}_t$ is used to update the posterior pdf at time t using the Bayes rule:

$$p(\mathbf{l}_t|\mathbf{z}_t, \ldots, \mathbf{z}_0) = \frac{p(\mathbf{z}_t|\mathbf{l}_t, \mathbf{a}_t)p(\mathbf{l}_t)}{p(\mathbf{z}_t|\mathbf{a}_t)} \quad (16)$$

where at every iteration the previous posterior is interpreted as the prior used in the current state update and the density $p(\mathbf{z}_t|\mathbf{a}_t)$ (also called *evidence*) can be calculated as:

$$p(\mathbf{z}_t|\mathbf{a}_t) = \int_{\mathbf{l}_t} (\mathbf{z}_t|\mathbf{l}_t, \mathbf{a}_t)p(\mathbf{l}_t)d\mathbf{l}_t \quad (17)$$

The convergence of this sequential decision process can be formally proven. Unfortunately, a direct computation of Eq. 15 is often practically intractable: it is required to compute the MI (Eq. 14) for each possible action. Even if we can choose between a small number of actions, the probability density function $p(\mathbf{z}_t|\mathbf{l}_t, \bar{\mathbf{a}}_t)$ (often called *likelihood function*) is defined in the space of *all* observations and states, for each action $\bar{\mathbf{a}}_t$. A solution would be to use a parametric density function (e.g., a Gaussian distribution) to represent the likelihood, or to discretize the observations space: unfortunately in the general case this function is multi-modal while an effective discretization of the observations space leads often to intractable solutions.

5.2. Observations as Objects Combinations

Our solution to address the intractability of the next-best-view problem is to represent the observations space as a discrete set of *combinations of hypothetical scene elements*. For the sake of simplicity, suppose that we are looking for a single instance of an object with a mobile manipulator equipped with an RGB camera mounted on the robot arm end effector (Fig. 1). Assuming that the scene contains at least one instance of the searched object, let us specialize the framework presented in the previous section:

- The state $\mathbf{l}_t$ defines the 3D position $\mathbf{g}_{cam,obj} \in \mathbb{SE}(3)$ of the object with respect to the reference frame defined by the initial camera position;

---

[4]In this work, we always assume to deal with a static system.

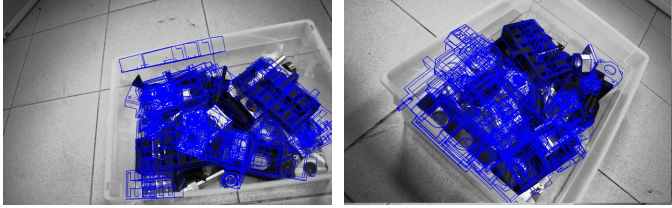[5]$\mathbf{z}_t$ clearly depends on the chosen action.

Figure 6: Object candidates refined using D^{2}CO: in these examples we are looking for three different types of objects. We use combinations of these "objects" as observations.

- The actions $\mathbf{a}_t$ define a discretized set of N_{pos} camera positions $A = \{a_1, \ldots, a_{N_{pos}}\}$ in a neighborhood of the working area (e.g, on a hemisphere that encloses the working area).

- The observations $\mathbf{z}_t$ are edge maps extracted from the images, e.g. using LSD (Sec. 3.2) or other edge detectors.

The likelihood function $p(\mathbf{z}_t|\mathbf{l}_t, \bar{\mathbf{a}}_t)$ should tell us what is the probability of getting a specific edge map $\mathbf{z}_t$, computed from the current image, given the object position represented by the state $\mathbf{l}_t$. Even for a given state, this function should be defined for *all* natural images: this facts clearly prevents the direct application of the next-best-view formulation presented above. To introduce our solution, we assume for now that:

1 The scene contains only known objects, and the object CAD models are given;
2 The object detection algorithm presented in Sec. 3 provides as output a set of N_{obj} candidates that, among a number of false positives, includes at least one true positive for each object in the scene.

The assumption (2) means that for each object in the scene, at least one object candidate is detected "close" to its real position. If now we apply D^{2}CO (Sec. 4) to each candidate, most of the true positives will converge toward the correct positions (e.g., Fig. 6). The idea is to build a set of synthetic scenes, or "scene realizations", by means of combinations without repetitions of the N_{obj} object candidates (Fig. 7 illustrates the concept of our procedure). For each combination, we can generate an edge map with the projections of the 3D edges templates that belong to the subset of object candidates. We select a set of *plausible k-combination without repetitions*, $k = 1, \ldots, N_{obj}$: *plausible* means here that we discard impossible combinations, e.g. combinations with intersections between objects. It is clear that, given the assumption (1) a (2), it exists at least a combination of candidates such that its edge map well approximates the edge map extracted from the *real* image (e.g., the edge map highlighted in red in Fig. 7), up to the inaccuracies in the objects localization and the nuisance factors in the image.

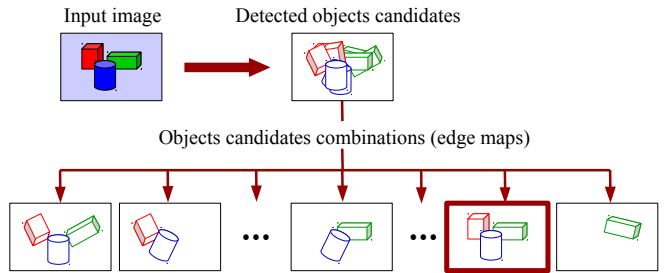Of course we don't know what are the combinations that

Figure 7: An overview of the concept behind our model-based approximation of the likelihood function. The edge maps in the second row represent scene realizations (e.g., candidates combinations) seen from a given viewpoint. The edge map that better approximates the ground truth is highlighted in red. (Objects in the edge maps are depicted with different colors for clarity).

best match the current edge map. Moreover, the assumptions (1) and (2) are probably never met in practice. However, we can observe that:

- In many object detection tasks (e.g., the random bin-picking problem in an industrial scenario, Fig. 6) often the scene contains many instances of the same object, so many object candidates combinations may represent good scene approximations;

- By increasing the number of the searched objects and the number of object candidates, we can obtain a better approximation of the real scene;

- From a given scene realization, it is possible to generate the related edge map for *any* camera position.

These observations suggest that such object candidates combinations can be used, for each camera position (i.e., for each action $\mathbf{a}_t$), as generators of observation samples. Formally speaking, for a given state $\bar{\mathbf{l}}_t$ and a given camera position $\bar{\mathbf{a}}_t$, we approximate the likelihood $p(\mathbf{z}_t|\bar{\mathbf{l}}_t, \bar{\mathbf{a}}_t)$ ($\bar{\mathbf{l}}_t$ and $\bar{\mathbf{a}}_t$ are both parameters in this case) with a probability mass function defined in a discrete sample space $Z_t = \{z_{1,t}, \ldots, z_{N_{comb},t}\}$, with N_{comb} the number of combinations. The $i - th$ sample $z_{i,t}$ represents an edge map generated from the $i - th$ object candidates combination, given the camera position $\bar{\mathbf{a}}_t$. In order to assign a probability to each sample, we could define the probability mass function to be proportional to a function of the average Directional Chamfer Distance (Eq. 7), e.g.:

$$p(\mathbf{z}_{i,t}|\bar{\mathbf{l}}_t, \bar{\mathbf{a}}_t) \propto e^{-\mu\left(\mathcal{DT}_3(\mathbf{z}_{i,t}, \bar{\mathbf{l}}_t, \bar{\mathbf{a}}_t)\right)^2} \qquad (18)$$

computed on the points $\mathbf{x}_i \in \mathbb{R}^2$ that are the projections of the 3D raster points of the searched object: the positions of these image points depend on both the current state $\bar{\mathbf{l}}_t$ (i.e., the object position with respect to the reference frame defined by the initial camera position) and the current camera position $\bar{\mathbf{a}}_t$. In Eq. 18 we make explicit the fact that the DCD depends on the state, on the scene

9

realization and on the camera position. For the sake of efficiency, during the next-best-view selection we prefer to use the conventional Chamfer Distance in place of the Directional Chamfer Distance, so we need only to compute one distance map $\mathcal{DT}$ for each combination rather than the full DCD tensor, i.e.:

$$p(\mathbf{z}_{i,t}|\bar{\mathbf{l}}_t, \bar{\mathbf{a}}_t) \propto e^{-\mu\left(\mathcal{DT}(\mathbf{z}_{i,t},\bar{\mathbf{l}}_t,\bar{\mathbf{a}}_t)\right)^2} \qquad (19)$$

Unfortunately, the solution proposed above is still not applicable in practice: the number of combinations is exponential with respect to the number of object candidates N_{obj}. To overcome this problem, we propose to sample a fixed number of $N_{comb} \ll 2^{N_{obj}}$ combinations using the sampling technique reported in Algorithm 1. We firstly define a probability (line 1) for each object candidate (e.g. Fig. 6), derived from its score (Eq. 11). For each combination, we draw from a binomial distribution[6] the maximum number of objects K that will be included in the combination (line 4). Up to K objects are then selected by drawing the object candidates with probability $p(obj_i)$, thus candidates with a high score will be easily included in the combinations (lines 5-13).

Algorithm 1: Objects Combinations Sampling

Data: The N_{obj} object candidates with their scores $\Psi(\mathbf{g}_{cam,obj_i})$ (Eq. 11), the required number of combinations N_{comb}

Result: The set of sampled combinations $\mathcal{S}_{comb}$

1 For each object candidate obj_i, compute a probability $p(obj_i) \propto e^{-(1-\Psi(\mathbf{g}_{cam,obj_i}))^2}$;

2 $\mathcal{S}_{comb} \leftarrow \{\}$;

3 **for** $j \leftarrow 1$ **to** N_{comb} **do**

4   Sample K from the binomial distribution $Bin(K|n = N_{obj}, p = 0.5)$;

5     $\mathcal{C}_j \leftarrow \{\}$;

6   **for** $k \leftarrow 1$ **to** K **do**

7   Sample an object $obj_i, obj_i \notin \mathcal{C}_j$, where the probability of drawing obj_i is given by $p(obj_i)$;
        `/* Add the current object if it does not intersect any previously added object */`

8   **if** $\mathcal{C}_j \cup \{obj_i\}$ *is a plausible combination* **then**

9             $\mathcal{C}_j \leftarrow \mathcal{C}_j \cup \{obj_i\}$;

10   **end if**

11   **end for**

12     $\mathcal{S}_{comb} \leftarrow \mathcal{S}_{comb} \cup \{\mathcal{C}_j\}$;

13 **end for**

The average Chamfer Distance required in Eq. 19 can be computed in an efficient way, avoiding the explicit computation of the N_{comb} edge maps and the related distance

---

[6]A binomial distribution $B(n, p)$ can be approximated with a Gaussian distribution $N(\mu, \sigma)$ with $\mu = np$ and $\sigma^2 = np(1 - p)$.

---

maps: we just compute the distance maps for the N_{obj} object candidates, where $N_{obj} \ll N_{comb}$. Given a camera position $\bar{\mathbf{a}}_t$, for each object candidate we generate the related distance map for that specific camera position. We also store a depth buffer that for each pixel provides the depth of the nearest edgel. The average Chamfer Distance is computed using for each image point $\mathbf{x}_i$ the minimum value between the distances extracted from the distance maps of the objects that belong to the current combination, while using the corresponding depth buffers to handle the occlusions.

5.3. The Proposed Algorithm

Algorithm 2: Model-Based Next-Best-View Detection and Localization

1 Move the camera toward the area of interest and acquire an image;

2 Detect up to N_{obj} object candidates using the procedure described in Sec. 3;

3 Refine the position of the N_{obj} objects using D²CO (Sec. 4);

4 Sample N_{comb} plausible combinations $\mathcal{S}_{comb}$ as described in Algorithm 1;

5 **for** $i \leftarrow 1$ **to** N_{pos} **do**

6   **for** $j \leftarrow 1$ **to** N_{comb} **do**

7   Project the combination $\mathcal{C}_j$ into the image plane of a camera placed in the position a_i;

8   Compute the distance map $\mathcal{DT}$ for that projection;

9   **end for**

10 **end for**

11 Sample N_{part} particles $\mathbf{l}_t = \{\mathbf{l}_t^{[1]}, \ldots, \mathbf{l}_t^{[N_{part}]}\}$ in the neighborhoods of the object candidates;

12 **while** *termination criteria are not met* **do**

13   Find $\mathbf{a}_t^* = \underset{\bar{a}_t}{\operatorname{argmax}}\, I(\mathbf{l}_t; \mathbf{z}_t|\bar{\mathbf{a}}_t)$ where the MI is computed using Eq. 21 and Eq. 24 ;

14   Move the camera in the position defined by $\mathbf{a}_t^*$ and acquire a new image;

15   **foreach** $\mathbf{l}_t^{[i]}$ **do**

16   Refine the particle positions using multi-view D²CO (Sec. 4.3);

17   Compute the particle weight w_i using the Bayes rule (Eq. 16) where the likelihood is computed as in Eq. 24, using only the *real* observations up to time t;

18   **end foreach**

19   Resample a new generation of N_{part} particles according to their weights $w_i, i = 1 \ldots N_{part}$;

20 **end while**

We represent the pdf $p(\mathbf{l}_t)$ over the state $\mathbf{l}_t$ by means of a set of N_{part} particles $\mathbf{l}_t = \{\mathbf{l}_t^{[1]}, \ldots, \mathbf{l}_t^{[N_{part}]}\}$, each one denoting an object position. In this way we can model
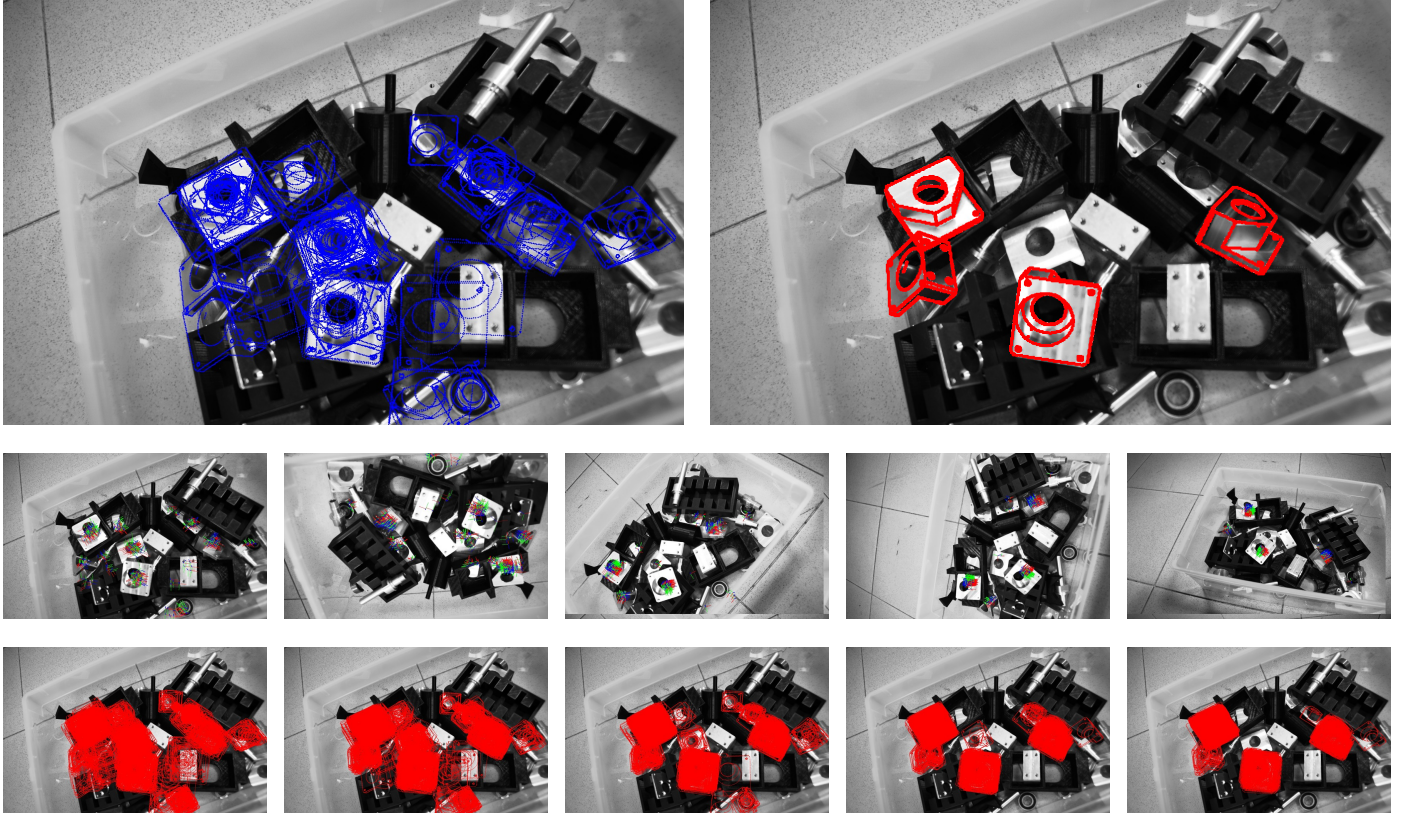
Figure 8: An example of the proposed next-best-view strategy: in this example we are looking for instances of the object with model reported in Fig. 9(b). [First row]: The detected object candidates (left) and the final object detection and localization result obtained after 5 views. [Second row]: The sequence of five images acquired during the active perception process, particles are represented by means of their positions and orientations. [Third row]: The particles (represented in this case by objects projections) seen from the first view during the five steps: thanks to the importance resampling step, particles with low weights are discarded.

multi-modal probability functions that implicitly enables our system to detect multiple instance of an object, each one represented by a mode in the (approximated) probability function.

We can rewrite the mutual information of Eq. 14 as:

$$I(\mathbf{l}_t; \mathbf{z}_t | \bar{\mathbf{a}}_t) =$$
$$E_{p(\mathbf{l}_t)} \left[E_{p(\mathbf{z}_t | \mathbf{l}_t, \bar{\mathbf{a}}_t)} \left[log \left(\frac{p(\mathbf{z}_t | \mathbf{l}_t, \bar{\mathbf{a}}_t)}{p(\mathbf{z}_t | \bar{\mathbf{a}}_t)} \right) \right] \right] \quad (20)$$

where $E[\cdot]$ is the expectation of a random variable. If the pdf over the state is represented by a set of N_{part} particles and the observations space has been discretized by a set of object candidates combinations (see Sec. 5.2), the mutual information can be computed as:

$$I(\mathbf{l}_t; \mathbf{z}_t | \bar{\mathbf{a}}_t) =$$
$$\frac{1}{N_{part}} \sum_{i=1}^{N_{part}} \sum_{j=1}^{N_{comb}} \left(p(\mathbf{z}_{j,t} | \mathbf{l}_t^{[i]}, \bar{\mathbf{a}}_t) \, log \left(\frac{p(\mathbf{z}_{j,t} | \mathbf{l}_t^{[i]}, \bar{\mathbf{a}}_t)}{p(\mathbf{z}_{j,t} | \bar{\mathbf{a}}_t)} \right) \right)$$
$$(21)$$

$p(\mathbf{z}_{j,t} | \mathbf{l}_t^{[i]}, \bar{\mathbf{a}}_t)$ could be computed using Eq. 19, while the evidence $p(\mathbf{z}_{j,t} | \bar{\mathbf{a}}_t)$ is calculated applying the law of total probability (Eq. 17).

Here we introduce a more effective way to compute the likelihood $p(\mathbf{z}_{j,t} | \mathbf{l}_t^{[i]}, \bar{\mathbf{a}}_t)$. At time step t we are looking for a "good" action a_t to take, but we have already collected $t-1$ real views of the scene along with the related DCD tensors. Our goal is to compute a likelihood that considers both the past *real* observations up to time step $t-1$ and the next observations that are synthesized by Algorithm 1. Given a particle $\mathbf{l}_t^{[i]}$, let us define the summation of the average Directional Chamfer Distances up to time $t-1$ as:

$$\Upsilon_{\mathcal{DT}_3}(t-1) = \sum_{\tau=1}^{t-1} \left(\mu \left(\mathcal{DT}_{3,\tau}(\mathbf{l}_t^{[i]}) \right)^2 \right) \quad (22)$$

where $\mathcal{DT}_{3,\tau}(\mathbf{l}_t^{[i]})$ is the Directional Chamfer Distance for the particle $\mathbf{l}_t^{[i]}$ computed over the DCD tensor extracted from the real image at time step τ. Combining Eq. 22 and Eq. 19 we obtain:

$$p(\mathbf{z}_{i,t} | \mathbf{l}_t^{[i]}, \bar{\mathbf{a}}_t) \propto e^{-\left(\Upsilon_{\mathcal{DT}_3}(t-1) - \mu \left(\mathcal{DT}(\mathbf{z}_{i,t}, \mathbf{l}_t^{[i]}, \bar{\mathbf{a}}_t) \right)^2 \right)} \quad (23)$$

which provides the weight of the particle $\mathbf{l}_t^{[i]}$ considering the real observations up to time step $t-1$ *plus* the new

11

(synthetic) observation $\mathbf{z}_{i,t}$.

We finally multiply the likelihood of Eq. 23 by a regularization term γ that penalizes already considered views:

$$p(\mathbf{z}_{i,t}|\mathbf{l}_t^{[i]},\bar{\mathbf{a}}_t) \propto \gamma \cdot e^{-\left(\Upsilon_{\mathcal{DT}_3}(t-1)\right)-\mu\left(\mathcal{DT}(\mathbf{z}_{i,t},\mathbf{l}_t^{[i]},\bar{\mathbf{a}}_t)\right)^2} \quad (24)$$

γ is defined as the ratio between the number of object template points seen (i.e., projected into the image) up to time t, including the synthetic observation $\mathbf{z}_{i,t}$, and the total number of points that compose the object template. In other words, the regularization term γ rewards the disocclusions. The pseudo code of the proposed next-best-view strategy is reported in Algorithm 2. At each iteration, Algorithm 2 takes a decision about the next view (line 13), it uses the new image to further refine the objects position (i.e., the particles) and to refine the current belief (lines 15-18) and finally it provides a new generation of particles exploiting an importance resampling algorithm [54] (line 19). The importance resampling step removes particles with low weights with a high probability while condensing more particles around areas where particles have high weights. An example of the proposed strategy is shown in Fig. 8.
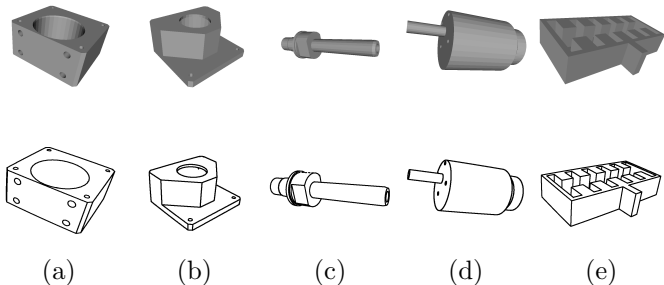
6. Experiments

Figure 9: [First row]: Some objects used in the experiments. (a),(b) and (c) are metal objects with high-reflectance surfaces, while (d) and (e) are black, plastic objects, and present a strong visible-light absorption. [Second row]: The edge templates extracted from the CAD models.

We present four different experimental validations. The first experiment aims to show that FDCM [7], that is the backbone of $\mathrm{D}^2\mathrm{CO}$, outperforms in our datasets another recent state-of-the-art object detection algorithm. In a second experiments, we compare $\mathrm{D}^2\mathrm{CO}$ to other registration techniques, showing state-of-the-art performances with a gain in speed up to a factor 10. Then we compare our registration algorithm in both the single-view and multi view-settings, showing for the latter a remarkable performance improvement. In the last experiment we show the effectiveness of the proposed model-based next-best-view strategy, comparing our method with two baseline active perception strategies.

We collected three datasets using the objects used in the RoCKIn@Work benchmarks [55]. The first dataset (**Dataset 1**) is composed by 60 1024x768 grey level images of different scenes that contain up to 5 untextured objects (see Fig. 9). The second dataset (**Dataset 2**) is composed by 90 1920x1200 grey level images of scenes that contain up to 20 untextured objects. The images of **Dataset 2** are divided in groups of three images, where each group contains images of the same scene captured from different viewpoints: this dataset is used to test the multi-view extension of $\mathrm{D}^2\mathrm{CO}$. In both datasets, objects are disposed in arbitrary 3D positions, often mutually occluded. In many images of **Dataset 1** and **Dataset 2** we added a background board with multiple patterns in order to simulate a (visual) cluttered background (e.g., see Fig. 13) The last dataset (**Dataset 3**) is composed by 1800 1920x1200 grey level images of scenes that include a box with many different objects of interest (e.g. Fig. 5). The images have been taken with a mobile manipulator, moving both the robot base and the camera mounted on the robot arm end effector (Fig. 1) in different viewpoints, sampling the camera positions in a hemisphere that encloses the working area. We recovered the robot base position using a simple marker-based localization strategy. **Dataset 3** has been used for the validation of the next-best-view strategy.

Each image of the three datasets is provided with the ground truth positions of each object, obtained with an externally guided procedure. All the experiments were performed running the algorithms on a standard laptop with an Intel Core i7-3820QM CPU with 2.70GHz, using a single core. All the compared algorithms has been implemented in C++ without any strong optimization. When possible, they share the same codebase and the same parameters, enabling an objective performance and timing comparison. An open-source C++ implementation of the object detection and localization algorithms, along with some of the used datasets, are freely available for download at:

http://www.dis.uniroma1.it/~labrococo/D2CO[7].

6.1. Object Detection

In the first experiment, we compare our object matching approach, that is strongly related with FDCM presented in [7], with a state-of-the-art object detection approach (called in the plots LINE-2D) described in [6] (since we are using only images, we use the LINE-2D version). In our experiments we don't perform any memory linearization as in [6], since we don't assume that the object projection remain the same every $x - y$ translation. Anyhow, this modification does not affect the performance of LINE-2D: it just runs slower.

We perform object matching starting from a set of 1260 pose candidates for each one of the 60 scenes of our dataset. Each of these candidates are acquired by sampling a large

---

[7]Case-sensitive URL.

cubic scene region containing the target objects. Fig. 10 shows how FDCM outperforms LINE-2D in terms of correct detection rate against the number of false positives.

6.2. Object Registration

We compared our approach (D^{2}CO) to DC-ICP [7], LM-ICP [28], C-ICP (ICP that exploits the Chamfer distance) and Direct (a simple coarse-to-fine object registration strategy that uses a Gaussian pyramid of gradient magnitudes images). All the tested algorithms share the same inner loop's stopping criteria parameters. We set the number of external ICP iterations to 50: we verified that this is a good trade-off in order to reach reliable results in our dataset. For each image and for each object in the image, we sampled 100 random positions around the ground truth position (e.g., the blue templates reported in Fig. 11). We used these positions as initial guesses for the registration algorithms we are testing. The final estimated position (e.g., the red templates reported in Fig. 11) is checked against the ground truth: if the total angular error was less than 0.1 radians, and the total error of translation was less than 5 mm, the registration was considered correct.

The proposed D^{2}CO algorithm outperforms the other methods in almost all tests (Fig. 12), while getting a gain in speed of a factor 10 compared to the second most competitive approach (see Table 1).

6.3. Multi-view Object Registration

We tested the multi-view object registration strategy presented in Sec. 4.3 employing images from **Dataset 2**. We applied D^{2}CO for both the single and multi-view settings. In the first case, only the first captured image has been exploited for registration; in the second case the first image plus N_v further images has been exploited, we set here $N_v = 2$. The camera positions for the second and third view has been randomly chosen. Results are reported in Fig. 14: in almost all cases the multi-view algorithm outperforms the single-view algorithm by a big margin, achieving very good results also in presence of strong occlusions (e.g., Fig. 13).

6.4. Next Best View

We compared our next-best-view strategy (called MI-MAX in the plots) with other two approaches: i) A random algorithm (called Random in the plots) which at every iteration moves the camera in a new random position, chosen from the positions not yet selected; ii) A maximum disocclusion approach (called DIS in the plots). The DIS algorithm maximizes the visibility gain for each object candidate $c \in C$, with C the set of all the detected object candidates, taking into account their confidence Ψ_c (i.e., their score). To this aim, the following objective function has to be maximized:

$$I(a) = \sum_{c \in C} \Psi_c * G(c, a) \qquad (25)$$

where a is the camera pose. As visibility gain measure $G(c, a)$, we consider the number of new visible points of the candidate c seen from a. Given the set A of all the possible views (i.e. camera poses), the next best view a^* is then computed by:

$$a^* = \underset{a \in A}{\mathrm{argmax}} \, I(a) \qquad (26)$$

We compared the correct localization rate and the average number of false positives provided by each algorithm (Fig. 15) using **Dataset 3**. In all the experiments, we used 40 object candidates, 200 particles, 1000 object combinations, and 32 possible next views. In order to sample the objects combinations (see Sec. 5.2) we used one (Fig. 15(a,b,e,f)) or three different objects types (Fig. 15 (c,d,g,h)). In **Dataset 3**, we split the objects by their level of occlusion: Fig. 15 (a,c,e,g) report experiments for no or slightly occluded objects, Fig. 15 (b,d,f,h) report experiments for highly occluded objects.

As reported in the plots of Fig. 15, MI-MAX outperforms the other algorithms in all experiments, showing good performances even if the object combinations have been generated using only one object type (Fig. 15 (a,b,e,f)). Exploiting more objects (Fig. 15 (c,d,g,h)) the combinations better approximate the scene thus the next-best-view strategy is more effective. MI-MAX takes on average 7 seconds to compute the next view.

7. Conclusion and Future Work

This paper introduces a novel model-based active detection and localization framework well suited for textureless objects placed in highly cluttered environments. We propose a registration strategy (D^{2}CO) that leverages the Directional Chamfer Distance tensor in a direct and efficient way. D^{2}CO enables to speed-up the registration process while preserving the wide basin of convergence provided by the DCD tensor.

Our system plans the sensing process by selecting the most informative next views that maximize the mutual information between the current state and the next observations. We propose to generate the future observations by efficiently sampling combinations of object hypothesis extracted by the object detector, while modeling the state by means of a set of particles, re-generated at each iteration with an importance resampling algorithm. The proposed active strategy takes into account the uncertainty in both the detection and the object localization estimates, while dealing with the objects mutual occlusions in a probabilistic way.

We reported several experiments performed on different challenging datasets acquired using a mobile manipulator: these datasets include many images of untextured objects often in presence of occlusions and cluttered background.
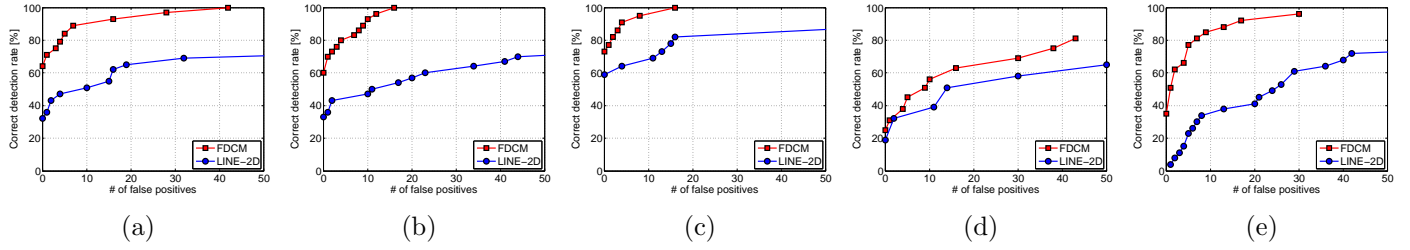
Figure 10: True positives rate plotted against the number of false positives (the letters (a),..(e) refer to the objects of Fig. 9).
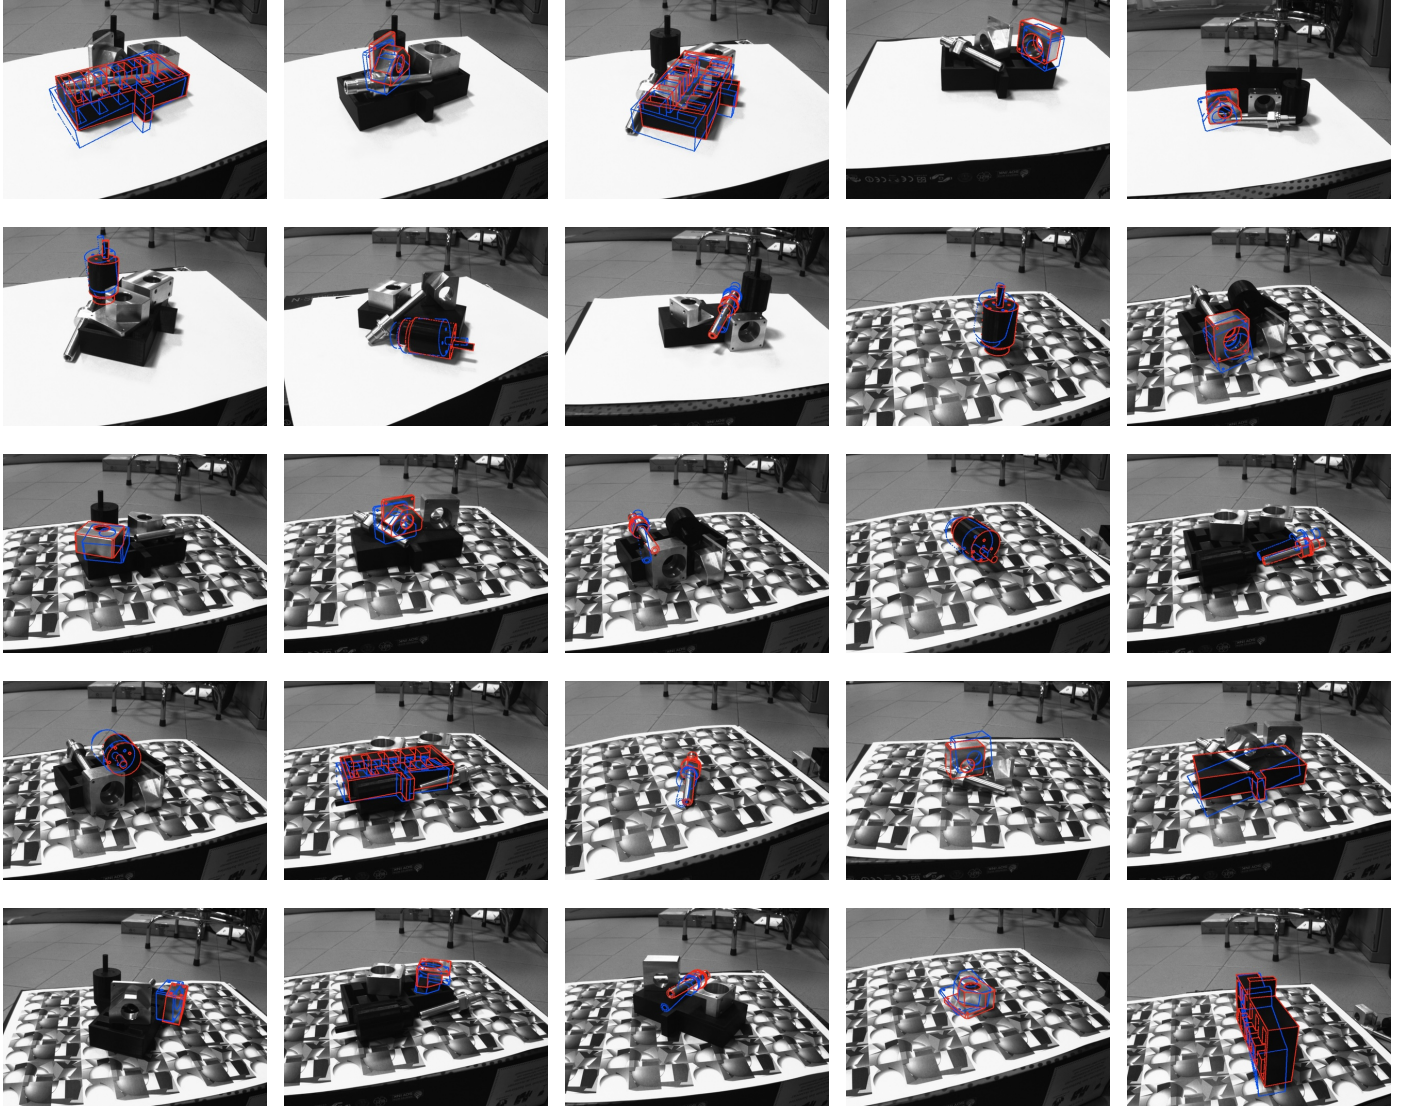

Figure 11: Some registration results obtained using the single-view D²CO algorithm (**Dataset 1**), the initial guess is reported in blue, while the final position estimate is reported in red.

Table 1: Average object registration time (milliseconds).

Algorithm	D²CO	DC-ICP	LM-ICP	C-ICP	Direct
Time (msec)	56.49	659.09	68.43	601.10	65.83
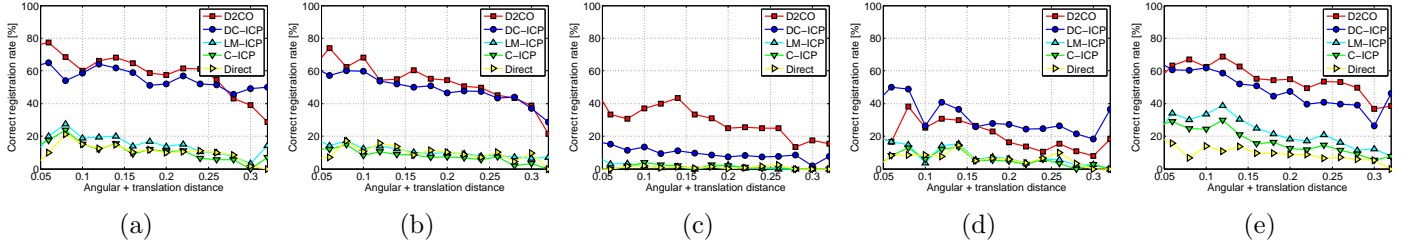
Figure 12: Correct registrations rate plotted against the distance (angle + translation) of the initial guess from the ground truth position (the letters (a),..(e) refer to the objects of Fig. 9).
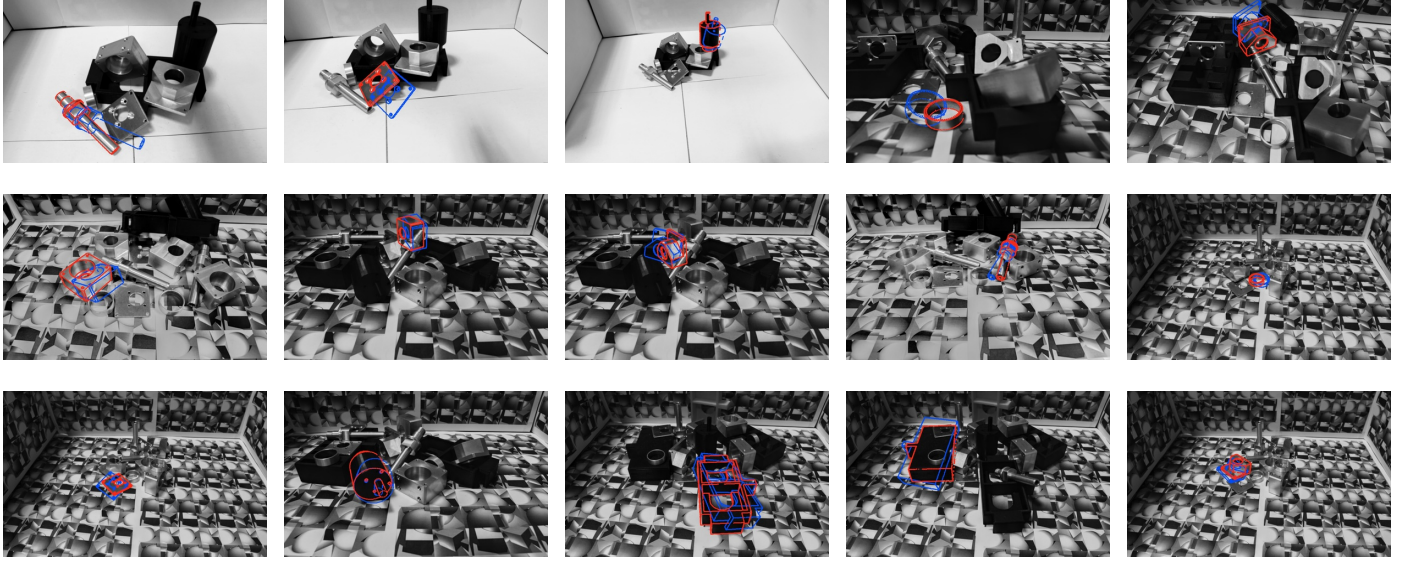

Figure 13: Some registration results obtained using the multi-view D^{2}CO algorithm (**Dataset 2**), the initial guess is reported in blue, final position estimate is reported in red.
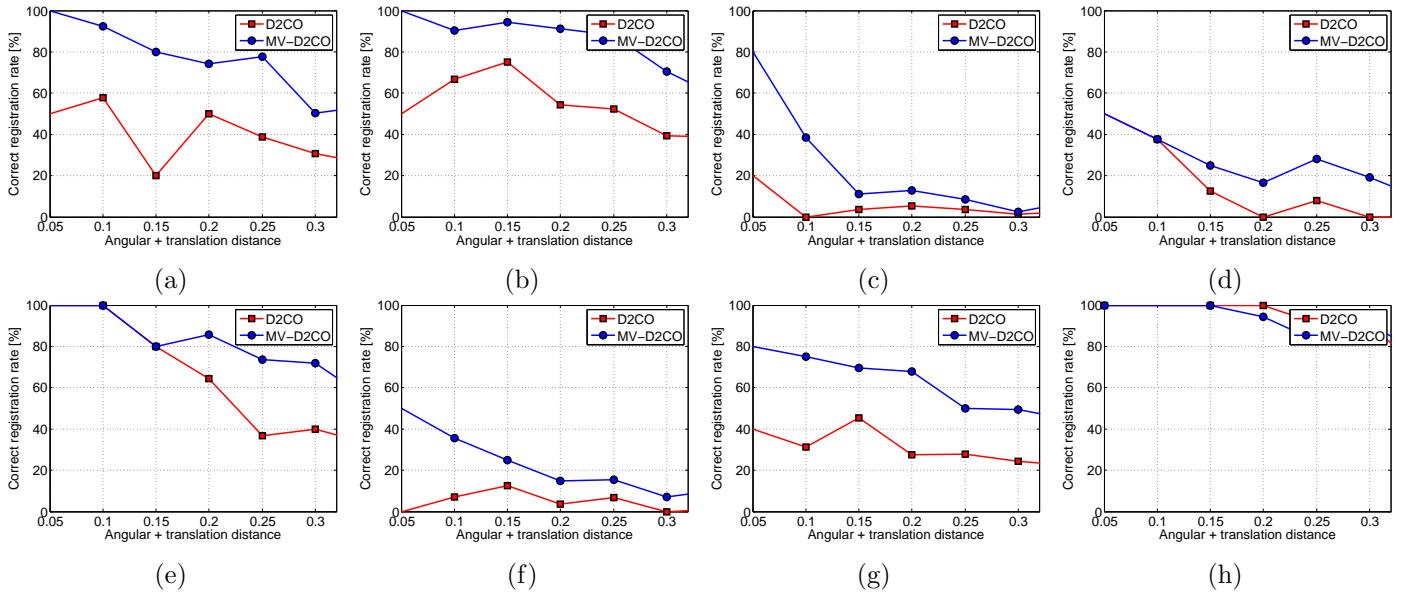

Figure 14: Single-view D^{2}CO VS multi-view D^{2}CO (MV-D^{2}CO in the plots) correct registrations rate (**Dataset 2**, the letters (a),..,(e) refer to the objects of Fig. 9, the letters (f), (g), (h) refer to other objects in our dataset).
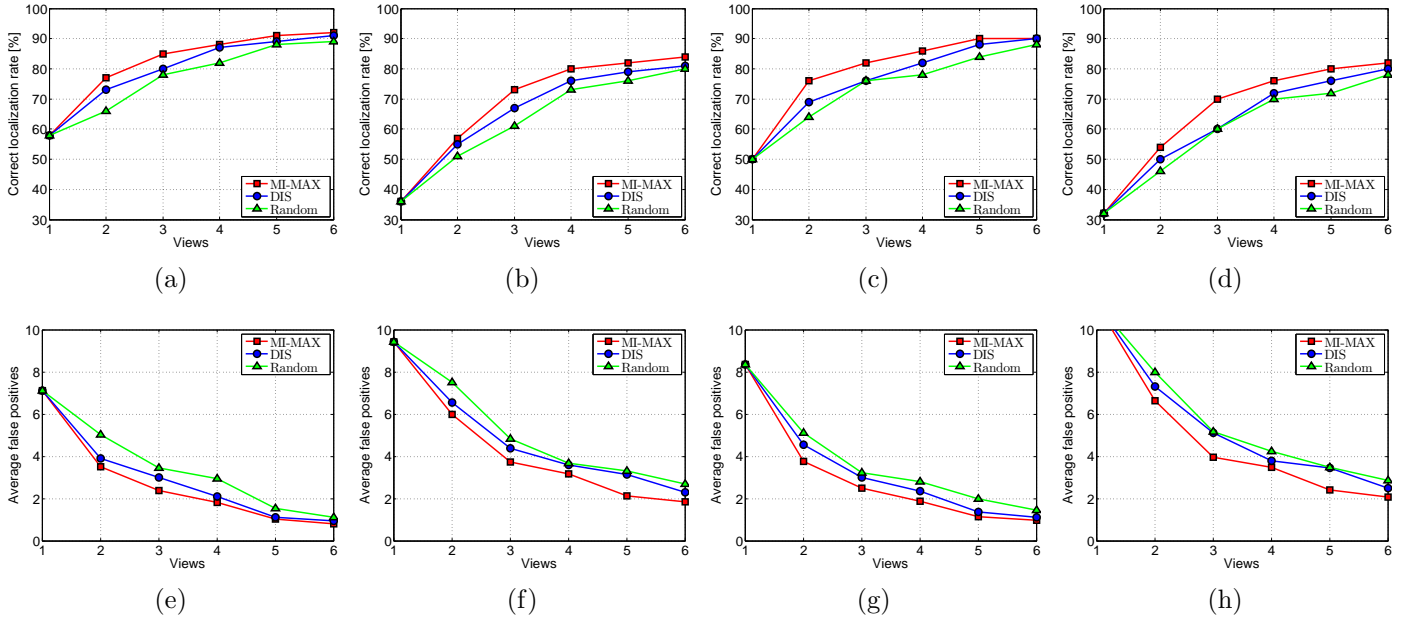
Figure 15: (a,b,c,d): Correct localization rate plotted against the numbers of views. (e,f,g,h): Average number of false positives plotted against the numbers of views. (a,c,e,g): Objects are no or slightly occluded (occlusions are less that 30% of the object surface) in the first view. (b,d,f,h): Objects are highly occluded (occlusions are greater than 30% of the object surface). (a,b,e,f): One object type (i.e., the object class) has been used to generate the object combinations. (c,d,g,h): Three different objects types have been used to generate the object combinations.

The results confirm the effectiveness of the proposed methods.

We are currently improving the system by integrating 3D information provided by a depth sensor in both the object detection and the objects combinations sampling algorithm. We also plan to enable the system to perform a new object detection step at each iteration, in order to discover previously unseen objects to be also used to improve the observations generation process. A further future work is to develop a massively parallel implementation of the proposed next-best-view strategy, in order to enable a fast exploration of the scene.

References

[1] S. Savarese, F. Li, 3D generic object categorization, localization and pose estimation, in: Proc. of the IEEE International Conference on Computer Vision (ICCV), 2007, pp. 1–8.

[2] T. Yeh, J. J. Lee, T. Darrell, Fast concurrent object localization and recognition, in: Proc. of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR), 2009, pp. 280–287.

[3] Ø. Skotheim, M. Lind, P. Ystgaard, S. A. Fjerdingen, A flexible 3D object localization system for industrial part handling, in: Proc. of the IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS), 2012, pp. 332 –3333.

[4] D. Fischinger, M. Vincze, Empty the basket - a shape based learning approach for grasping piles of unknown objects, in: Proc. of the IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS), 2012, pp. 2051–2057.

[5] D. Holz, A. Topalidou-Kyniazopoulou, J. Stückler, S. Behnke, Real-time object detection, localization and verification for fast robotic depalletizing, in: Proc. of the IEEE/RSJ International Conference on Intelligent Robots and Systems IROS, 2015, pp. 1459–1466.

[6] S. Hinterstoisser, C. Cagniart, S. Ilic, P. F. Sturm, N. Navab, P. Fua, V. Lepetit, Gradient response maps for real-time detection of textureless objects, IEEE Transactions on Pattern Analysis and Machine Intelligence 34 (5) (2012) 876–888.

[7] M. Liu, O. Tuzel, A. Veeraraghavan, Y. Taguchi, T. K. Marks, R. Chellappa, Fast object localization and pose estimation in heavy clutter for robotic bin picking, International Journal of Robotic Research 31 (8) (2012) 951–973.

[8] M. Imperoli, A. Pretto, D^2CO: Fast and robust registration of 3D textureless objects using the Directional Chamfer Distance, in: Proc. of the International Conference on Computer Vision Systems (ICVS), 2015, pp. 316–328.

[9] P. J. Besl, N. D. McKay, A method for registration of 3-D shapes, IEEE Transaction on Pattern Analysis and Machine Intelligence 14 (2) (1992) 239–256.

[10] R. Eidenberger, J. Scharinger, Active perception and scene modeling by planning with probabilistic 6D object poses, in: Proc. of the IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS), 2010, pp. 1036–1043.

[11] D. Holz, M. Nieuwenhuisen, D. Droeschel, J. Stückler, A. Berner, J. Li, R. Klein, S. Behnke, Active recognition and manipulation for mobile robot bin picking, in: Gearing Up and Accelerating Crossfertilization between Academic and Industrial Robotics Research in Europe:: Technology Transfer Experiments from the ECHORD Project, Springer International Publishing, 2014, pp. 133–153.

[12] A. Andreopoulos, J. K. Tsotsos, A theory of active object localization, in: Proc. of the IEEE International Conference on Computer Vision (ICCV), 2009, pp. 903–910.

[13] J. Denzler, C. M. Brown, Information theoretic sensor data selection for active object recognition and state estimation, IEEE Transaction on Pattern Analysis and Machine Intelligence 24 (2) (2002) 145–157.

[14] D. Marr, H. K. Nishihara, Representation and recognition of the spatial organization of three-dimensional shapes, Proc. of the Royal Society of London. Series B, Biological Sciences

200 (1140) (1978) 269–294.

[15] D. G. Lowe, Three-dimensional object recognition from single two-dimensional images, Artificial Intelligence 31 (3) (1987) 355 – 395.

[16] P. Viola, M. Jones, Rapid object detection using a boosted cascade of simple features, in: Proc. of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR), 2001, pp. 511–518.

[17] D. G. Lowe, Distinctive image features from scale-invariant keypoints, International Journal of Computer Vision 60 (2) (2004) 91–110.

[18] N. Dalal, B. Triggs, Histograms of oriented gradients for human detection, in: Proc. of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR), Vol. 1, 2005, pp. 886–893.

[19] P. F. Felzenszwalb, R. B. Girshick, D. McAllester, D. Ramanan, Object detection with discriminatively trained part-based models, IEEE Transactions on Pattern Analysis and Machine Intelligence 32 (9) (2010) 1627–1645.

[20] F. Viksten, P.-E. Forssén, B. Johansson, A. Moe, Comparison of local image descriptors for full 6 degree-of-freedom pose estimation, in: Proc. of the IEEE international conference on Robotics and Automation (ICRA), 2009.

[21] R. B. Girshick, J. Donahue, T. Darrell, J. Malik, Region-based convolutional networks for accurate object detection and segmentation, IEEE Transactions on Pattern Analysis and Machine Intelligence 38 (1) (2016) 142–158.

[22] M. Everingham, L. Van Gool, C. K. Williams, J. Winn, A. Zisserman, The pascal visual object classes (VOC) challenge, International Journal of Computer Vision 88 (2) (2010) 303–338.

[23] J. J. Lim, H. Pirsiavash, A. Torralba, Parsing IKEA objects: Fine pose estimation, in: Proc. of the IEEE International Conference on Computer Vision (ICCV), 2013, pp. 2992–2999.

[24] M. Aubry, D. Maturana, A. A. Efros, B. C. Russell, J. Sivic, Seeing 3D chairs: Exemplar part-based 2D-3D alignment using a large dataset of cad models, in: Proc. of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR), 2014, pp. 3762–3769.

[25] C. B. Choy, M. Stark, S. Corbett-Davies, S. Savarese, Enriching object detection with 2D-3D registration and continuous viewpoint estimation, in: Proc. of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR), 2015, pp. 2512–2520.

[26] D. Damen, P. Bunnun, A. Calway, W. W. Mayol-Cuevas, Real-time learning and detection of 3D texture-less objects: A scalable approach, in: Proc. of the British Machine Vision Conference (BMVC), 2012, pp. 1–12.

[27] A. Pretto, S. Tonello, E. Menegatti, Flexible 3D localization of planar objects for industrial bin-picking with monocamera vision system, 2013, pp. 168–175.

[28] A. W. Fitzgibbon, Robust registration of 2D and 3D point sets, in: Proc. of the British Machine Vision Conference (BMVC), 2001, pp. 662–670.

[29] B. Jian, B. C. Vemuri, Robust point set registration using gaussian mixture models, IEEE Transaction on Pattern Analysis and Machine Intelligence 33 (8) (2011) 1633–1645.

[30] H. G. Barrow, J. M. Tenenbaum, R. C. Bolles, H. C. Wolf, Parametric correspondence and chamfer matching: Two new techniques for image matching, in: Proc. of International Joint Conference on Artificial Intelligence (IJCAI), 1977, pp. 659–663.

[31] G. Borgefors, Hierarchical chamfer matching: A parametric edge matching algorithm, IEEE Transaction on Pattern Analysis and Machine Intelligence 10 (6) (1988) 849–865.

[32] C. Choi, H. I. Christensen, 3D textureless object detection and tracking: An edge-based approach, in: Proc. of the IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS), 2012, pp. 3877–3884.

[33] J. Shotton, A. Blake, R. Cipolla, Multiscale categorical object recognition using contour fragments, IEEE Transactions on Pattern Analysis and Machine Intelligence 30 (7) (2008) 1270–1281.

[34] H. Cai, T. Werner, J. Matas, Fast detection of multiple textureless 3-D objects, in: Proc. of the International Conference

Computer Vision Systems (ICVS), 2013, pp. 103–112.

[35] R. Bajcsy, Active perception, Proc. of the IEEE 76 (8) (1988) 996–1005.

[36] L. E. Wixson, D. H. Ballard, Using intermediate objects to improve the efficiency of visual search, International Journal of Computer Vision 12 (2-3) (1994) 209–230.

[37] K. Shubina, J. K. Tsotsos, Visual search for an object in a 3D environment using a mobile robot, Computer Vision and Image Understanding 114 (5) (2010) 535–547.

[38] A. Aydemir, A. Pronobis, M. Gobelbecker, P. Jensfelt, Active visual object search in unknown environments using uncertain semantics, IEEE Transactions on Robotics 29 (4) (2013) 986–1002.

[39] C. Stachniss, G. Grisetti, W. Burgard, Information gain-based exploration using rao-blackwellized particle filters, in: Proceedings of Robotics: Science and Systems (RSS), 2005, pp. 65–72.

[40] S. Chen, Y. Li, Vision sensor planning for 3-D model acquisition, IEEE Transactions on Systems, Man, and Cybernetics, Part B: Cybernetics 35 (5) (2005) 894–904.

[41] Y. Li, Z. Liu, Information entropy-based viewpoint planning for 3-D object reconstruction, IEEE Transactions on Robotics 21 (3) (2005) 324–337.

[42] M. Krainin, B. Curless, D. Fox, Autonomous generation of complete 3D object models using next best view manipulation planning, in: Proc. of the IEEE International Conference on Robotics and Automation ICRA, 2011, pp. 5031–5037.

[43] S. Chen, Y. Li, N. M. Kwok, Active vision in robotic systems: A survey of recent developments, The International Journal of Robotics Research 30 (11) (2011) 1343–1377.

[44] C. Connolly, The determination of next best views, in: In Proc. of the IEEE International Conference on Robotics and Automation. Proceedings, 1985.

[45] D. Wilkes, J. Tsotsos, Active object recognition, in: Proc. of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR), 1992, pp. 136–141.

[46] B. Schiele, J. L. Crowley, Transinformation for active object recognition, in: In Proc. of the IEEE International Conference on Computer Vision (ICCV), 1998.

[47] H. Borotschnig, L. Paletta, M. Prantl, A. Pinz, Appearance-based active object recognition, Image and Vision Computing 18 (9) (1998) 715–727.

[48] F. Farshidi, S. Sirouspour, T. Kirubarajan, Robust sequential view planning for object recognition using multiple cameras, Image and Vision Computing 27 (8) (2009) 1072–1082.

[49] N. Atanasov, B. Sankaran, J. L. Ny, G. J. Pappas, K. Daniilidis, Nonmyopic view planning for active object detection, IEEE Transactions on Robotics 30 (5) (2014) 1078–1090.

[50] Z. Wu, S. Song, A. Khosla, F. Yu, L. Zhang, X. Tang, J. Xiao, 3D shapenets: A deep representation for volumetric shapes, in: Proc. of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR), 2015.

[51] R. Grompone Von Gioi, J. Jakubowicz, J.-M. Morel, G. Randall, Lsd: A fast line segment detector with a false detection control, IEEE Transaction on Pattern Analysis and Machine Intelligence 32 (4) (2010) 722–732.

[52] Y. Ma, S. Soatto, J. Kosecka, S. Sastry, An invitation to 3D vision, from images to models, Springer Verlag, 2003.

[53] T. M. Cover, J. A. Thomas, Elements of Information Theory, Wiley-Interscience, New York, NY, USA, 1991.

[54] S. Thrun, W. Burgard, D. Fox, Probabilistic Robotics, The MIT Press, 2005.

[55] F. Amigoni, E. Bastianelli, J. Berghofer, A. Bonarini, G. Fontana, N. Hochgeschwender, L. Iocchi, G. K. Kraetzschmar, P. U. Lima, M. Matteucci, P. Miraldo, D. Nardi, V. Schiaffonati, Competitions for benchmarking: Task and functionality scoring complete performance assessment, IEEE Robotics and Automation Magazine 22 (3) (2015) 53–61.