



HAL
open science

Object segmentation in depth maps with one user click and a synthetically trained fully convolutional network

Matthieu Grard, Romain Brégier, Florian Sella, Emmanuel Dellandréa,
Liming Chen

► To cite this version:

Matthieu Grard, Romain Brégier, Florian Sella, Emmanuel Dellandréa, Liming Chen. Object segmentation in depth maps with one user click and a synthetically trained fully convolutional network. 2018. hal-01674511v1

HAL Id: hal-01674511

<https://hal.science/hal-01674511v1>

Preprint submitted on 3 Jan 2018 (v1), last revised 24 Sep 2018 (v2)

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

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

Object segmentation in depth maps with one user click and a synthetically trained fully convolutional network

Matthieu Grard^{1,2}, Romain Brégier^{1,3},
Florian Sella¹, Emmanuel Dellandréa², and Liming Chen²

¹Siléane, 17 rue Descartes, F-42000 St Étienne, France
m.grard@sileane.com

²Université de Lyon, CNRS, École Centrale de Lyon
LIRIS UMR5205, F-69134 Lyon, France

³Univ. Grenoble Alpes, Inria, CNRS
Grenoble INP, LIG, F-38000 Grenoble, France

Abstract. With more and more household objects built on planned obsolescence and consumed by a fast-growing population, hazardous waste recycling has become a critical challenge. Given the large variability of household waste, current recycling platforms mostly rely on human operators to analyze the scene, typically composed of many object instances piled up in bulk. Helping them by robotizing the unitary extraction is a key challenge to speed up this tedious process. Whereas supervised deep learning has proven very efficient for such object-level scene understanding, *e.g.*, generic object detection and segmentation in everyday scenes, it however requires large sets of per-pixel labeled images, that are hardly available for numerous application contexts, including industrial robotics.

We thus propose a step towards a practical interactive application for generating an object-oriented robotic grasp, requiring as inputs only one depth map of the scene and one user click on the next object to extract. More precisely, we address in this paper the middle issue of object segmentation in top views of piles of bulk objects given a pixel location, namely seed, provided interactively by a human operator. We propose a two-fold framework for generating edge-driven instance segments. First, we repurpose a state-of-the-art fully convolutional object contour detector for seed-based instance segmentation by introducing the notion of *edge-mask duality* with a novel patch-free and contour-oriented loss function. Second, we train one model using only synthetic scenes, instead of manually labeled training data. Our experimental results show that considering edge-mask duality for training an encoder-decoder network, as we suggest, outperforms a state-of-the-art patch-based network in the present application context.

Keywords: Interactive instance segmentation, supervised pixel-wise learning, synthetic training images

1 Introduction

Waste recycling is one of the main challenges lying ahead, as the ever-growing amount of obsolete household objects and industrial materials outstrips the deployed human resources. Current state legislations and policies thus entice industrialists to develop robotized waste sorting, which stands out as a competitive approach to protect humans from hazardous waste collection sites (toxic emissions, radioactive radiations, *etc.*). However, given the large variability of elements to sort, and the cluttered nature of waste, human expertise is hardly dispensable despite the great advances in machine learning. Towards a fast and accurate processing, enhancing human operators decisions with computer vision-guided robotic grasping seems therefore one of the current most suitable strategies. Numerous works on supervised learning for object-wise scene understanding in images of everyday scenes have indeed demonstrated remarkable results for real-time applications: *generic object detection*, *i.e.*, generating category-free object proposals, either as bounding box proposals [12, 24], or as binary masks, namely segment proposals [10, 14, 22, 23]; *semantic segmentation*, *i.e.*, assigning an instance-agnostic object category to each pixel [16, 18]; *object contour detection*, *i.e.*, assigning to each pixel a probability of being an object boundary [5, 30].




Fig. 1: A test application of our approach on a real robotic setup: (1) A user clicks on the object to extract. (2) The proposed model (FCN-DOL) delineates the selected object in a single forward pass. (3) An off-the-shelf grasp search engine can thus generate a relevant object-oriented robotic grasp.

In this work, we aim at extracting one by one each instance of a dense pile of objects in random poses. As illustrated by Fig. 1, we consider the following scenario: (1) a depth-augmented image of the scene is captured; (2) a human operator clicks on the next object to extract; (3) the selected object is automatically delineated; (4) robotic grasps on the segmented object are automatically detected. This paper describes the third step, *i.e.*, the segmentation of the se-

lected object. As the ultimate goal is to successfully grasp each item of the scene, object instances must be correctly distinguished despite potential occlusion and entanglement, which is hardly achievable without an explicit notion of object. In the context of waste sorting, prior knowledge is not available: the target objects to be extracted, as well as their corresponding model, are unknown. Differently from industrial bin-picking of many instances of a known manufactured object, traditional pose estimation techniques [4] cannot be applied. We thus adapt object segmentation methods used for everyday image analysis to the context of industrial robotic picking. Indeed, state-of-the-art techniques [10, 22] first perform a coarse localization by extracting a set of image patches assumed to contain one instance. While this approach is relevant for the typical foreground/background paradigm inherent to most indoor and outdoor scenes, it can be ambiguous when the scene is only composed of objects densely stacked on top of each other, as a patch may be equally shared by many objects (cf. Fig. 2). The proposed approach alleviates this problem by considering a pixel included in the object, instead of a bounding box. Interestingly, one can notice in the present application context that inter-instance boundaries are complementary to instance segments, since the background is generally fully hidden and all objects are targets. This suggests besides that object contour detection may be used as a backbone for object segmentation. In the following, we refer to this property as *edge-mask duality*.



Fig. 2: A use case showing that it is less ambiguous to consider the task “Segment the object including this pixel” (c) rather than “Segment the object inside this box” (b) as a patch may be equally shared by several objects.

Moreover, state-of-the-art learning-based edge detection and segmentation tasks require training datasets annotated on a per-pixel basis. Despite many efforts [7], notably for real indoor scenes featuring household objects [15, 29], and table-top sparse clutters [25], none matches the camera point of view, neither the typical scene setup of encountered waste sorting applications. Although collecting adequate data is certainly not a problem given the recent spread of affordable RGB-D sensors, annotating reliably new images remains tedious, as current methods consist in time-consuming manual labeling. Real image datasets

are therefore limited in terms of content, and costly to extend. This makes the availability of appropriate training data a factor hindering the widespread deployment of pixel-wise learning-based techniques in industrial environments.

In this paper, towards an interactive application for robotized waste sorting, we address the middle issue of learning object segmentation in depth maps, given one user click per object to segment. We aim at demonstrating that leveraging edge-mask duality with synthetic depth maps for end-to-end instance segmentation enables higher boundary precision than a patch-based approach, while ensuring a framework fully pluggable in an interactive interface. Starting from the observations that (1) realistic RGB-D images of synthetic scenes can be automatically generated by simulation, (2) learning from synthetic depth-augmented images has already shown promising results for semantic segmentation of indoor scenes [8] and disparity estimation [19], (3) fully convolutional networks trained end to end for object contour detection has proven very effective [30], our contribution is two-fold:

- **A synthetic expandable dataset for industrial robotic picking.** With custom scripts on top of Blender [2], we produce realistic rendered RGB-D images of object instances stacked in bulk, thus ensuring time-saving and error-free annotations used as pixel-wise ground truth during training.
- **A scheme to train a fully convolutional edge detector for seed-based instance segmentation.** First, we show that object contour detectors can be trained on synthetic depth maps while performing almost equivalently on real data at inference time. Second, we introduce edge-mask duality with a novel dual-objective loss function to train a fully convolutional encoder-decoder in segmenting objects clicked interactively by a human operator. We show that the proposed approach, more suited for a practical human-robot interface, outperforms a state-of-the-art patch-based network when objects are piled up in bulk.

Our paper is organized as follows. After reviewing the state of the art in Section 2, we describe the proposed training scheme in Section 3 and our training data production pipeline in Section 4. Then, Section 5 details our experimental protocol. Results are discussed in Section 6.

2 Related Work

2.1 Synthetic training data

To our knowledge, learning from synthetic RGB-D images for pixel-wise object segmentation tasks has yet received little attention. In the context of generic object detection and recognition, a first kind of approach consists in generating synthetic images by projecting 3D models onto natural “background” images [20, 28]. However, these synthetic images contain by construction at most a few object instances that are implicitly assumed to be distinguishable from their environment, thus excluding many situations like the present case of dense

stack of object instances. A second type of approach consists in simulating scenes from scratch. Mainly dedicated to semantic segmentation, synthetic datasets of annotated urban [27] and indoor [9] scenes have recently emerged. Following the creation of these datasets, it has been shown that depth-based semantic segmentation of indoor scenes with a convolutional neural network (CNN) can be boosted by using synthetic images as training data and ten times less real images only for fine-tuning [8]. Also, in the footsteps of FlowNet [6] for optical flow estimation using synthetic images, training a CNN only on sufficiently realistic rendered image patches has again recently proved efficient for the purposes of disparity estimation on outdoor scenes with DispNet [19]. In continuity with the latter simulation-based approaches, we explore the context of robotic picking when the target objects are undefined. We show that real depth maps, whose costly pixel-wise ground-truth annotations are unsustainable in industrial environments, may be fully replaced by rendered depth maps when feeding structured random forests or fully convolutional networks for object contour detection during training.

2.2 Instance segmentation

Due to a lack of large real-world RGB-D datasets, instance segmentation algorithms have been mostly performed on RGB-only images of everyday scenes, augmented with crowd-sourced pixel-wise annotations [17]. Concurrently with the emergence of end-to-end deep learning for instance segmentation, images were typically represented as a graph of superpixels, often built from boundary-related cues. To this end, popular edge detectors, using structured random forests to classify patches [31] or lately end-to-end fully convolutional "encoder-decoder" networks [30], were employed to give edge-preserving starting cues. Unsupervised instance proposals generation then typically consisted in multi-scale combinatorial grouping [23] based upon generic hand-crafted features, *e.g.*, instance candidate's area or perimeter. Later, learning to propose instance candidates was notably introduced with conditional random fields (CRFs) [14], that are jointly trained for binary graph cuts while maximizing global recall. Lately, Facebook AI Research released a network that predicts scored instance proposals from image patches [21], trained end to end on the MS COCO dataset [17], and outperforming previous state-of-the-art approaches. A top-down refinement scheme was further proposed [22] to refine the boundaries of the predicted masks, since without explicit notion of object contour during training, feed-forward deep architectures generally fail to encode accurate boundaries due to progressive down-sampling. Within the context of semantic instance segmentation, recurrent neural networks have been employed to sequentially segment each instance belonging to one inferred category [26], by heading fully convolutional networks for semantic segmentation with long short-term memory (LSTM) layers, so as to enable a recurrent spatial inhibition. In contrast to previous approaches, we leverage the duality between inter-instance boundaries and instance masks (described in Section 3) directly within a fast feed-forward architecture by introducing an explicit object contour learning during training for instance segmentation, so as to

obtain masks with sharp and consistent boundaries in the case of objects piled up in bulk.

3 Leveraging edge-mask duality

Motivated by the many recent works on feed-forward fully convolutional architectures that demonstrated strong learning capabilities and fast inference time to be part of a responsive interactive interface, we follow the idea that when the scene is full of target objects, typically in waste sorting applications, the inter-instance boundaries and the set of instance masks are dual representations, what we call *edge-mask duality*. This interestingly suggests that feature encoding-decoding for object contour detection may be used as backbone for instance segmentation. One may indeed expect a better object pose invariance and a finer segmentation when learning local inter-instance boundaries instead of the spatial continuity inside the masks. We thus propose to repurpose a state-of-the-art encoder-decoder object contour detector to predict instance segments by appending during training to the input and the output respectively a seed pixel and the instance mask including the seed. State-of-the-art approaches [10, 22] link a patch, either randomly generated or predicted, to one instance, which re-

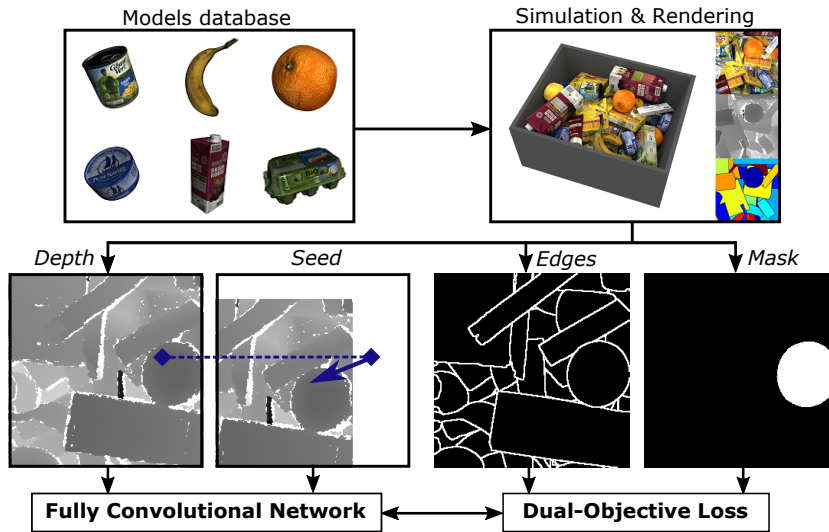


Fig. 3: Our training scheme to generate instance segments by introducing the notion of *edge-mask duality* with a novel loss function. A fully convolutional encoder-decoder network is fed with one image and its translation centering the user click on one instance, namely seed, to predict the overall inter-instance boundary map and the instance mask including the input seed. The proposed dual-objective loss drives the network to explicitly learn segmentation from boundaries.

quires size priors and may be ambiguous (cf. Fig. 2). We use instead a translation centering a pixel location assumed to be included in one instance, namely seed, here provided by an external user, as illustrated by Fig. 3.

Formally, suppose we have a set of annotated depth maps. For one depth map, let \mathcal{G} be the ground-truth instance label map and \mathcal{E} the ground-truth inter-instance contour map, that is, for each pixel \mathbf{p} ,

$$\mathcal{E}_{\mathbf{p}} = \begin{cases} 1 & \text{if } \exists(\mathbf{q}, \mathbf{q}') \in \mathcal{N}_{\mathbf{p}}^2 \text{ such that } \mathcal{G}_{\mathbf{q}} \neq \mathcal{G}_{\mathbf{q}'} \\ 0 & \text{otherwise} \end{cases} \quad (1)$$

where $\mathcal{N}_{\mathbf{p}}$ denotes the 8-connected neighborhood of pixel \mathbf{p} . In other words, $\mathcal{E}_{\mathbf{p}}$ is the expected probability that \mathbf{p} is an instance boundary. Let $\{\mathcal{M}_k\}_k$ be the set of the corresponding ground-truth instance binary masks. One can simply define the instance mask $\mathcal{M}_{\mathbf{s}}$ at seed pixel \mathbf{s} as a two-operation process on the contour map:

$$\mathcal{M}_{\mathbf{s}} = \mathcal{C}(\bar{\mathcal{E}}, \mathbf{s}) \quad (2)$$

where $\mathcal{C}(\mathcal{X}, \mathbf{s})$ is the connected subset of \mathcal{X} including \mathbf{s} , and $\bar{\mathcal{X}}$ the complementary of \mathcal{X} . $\mathcal{M}_{\mathbf{s}}$ thus defines the expected probability for each pixel \mathbf{p} to belong to the instance including the seed \mathbf{s} . Reciprocally, the instance contour map can be trivially derived from the set of instance masks.

Feed-forward networks require one-to-one or many-to-one relationships between the training inputs and expected outputs, however in the context of instance segmentation, the corresponding set of instances is of varying cardinality. Circumventing this issue thus requires each expected instance mask to be associated with one different input. State-of-the-art approaches [22] implement one-to-one relationships by defining a set of patches, assuming that one patch contains one instance. In this work, we instead consider a set \mathcal{S} of seed pixels and propose the following mapping:

$$\begin{aligned} \mathcal{S} &\rightarrow \{\mathcal{M}_k\}_k \\ \mathbf{s} &\mapsto \mathcal{M}_{\mathbf{s}} \end{aligned} \quad (3)$$

so that each seed pixel maps to the instance it belongs to. Exploiting the translational invariance of convolutional networks and Equation 3, we can then define a pixel-wise training set to use encoder-decoder networks for instance segmentation, by mapping one depth map and one seed \mathbf{s} to one contour map \mathcal{E} and one instance mask $\mathcal{M}_{\mathbf{s}}$, as depicted by Figure 3. At training time, we select randomly a few seeds inside each ground-truth instance. At inference time, the user provides a seed by clicking on the instance.

To introduce the edge-mask duality in the learning, we train the encoder-decoder network end to end by minimizing the following sum of pixel-wise logistic loss functions

$$\mathcal{L}(\theta) = \frac{1}{N} \sum_{n=1}^N \left\{ \ell(\lambda_e, \mathcal{E}^{(n)}, \hat{\mathcal{E}}^{(n)}) + \sum_{\mathbf{s} \in \mathcal{S}} \ell(\lambda_m, \mathcal{M}_{\mathbf{s}}^{(n)}, \hat{\mathcal{M}}_{\mathbf{s}}^{(n)}) \right\} \quad (4)$$

where θ denotes the network’s parameters, N the number of training samples, ℓ has the following form:

$$\ell(\lambda, \mathcal{Y}, \hat{\mathcal{Y}}) = - \sum_{\mathbf{p}} \left\{ (1 - \mathcal{Y}_{\mathbf{p}}) \log(1 - \sigma(\hat{\mathcal{Y}}_{\mathbf{p}})) + \lambda \mathcal{Y}_{\mathbf{p}} \log(\sigma(\hat{\mathcal{Y}}_{\mathbf{p}})) \right\} \quad (5)$$

$\mathcal{Y}_{\mathbf{p}} \in \{0, 1\}$ and $\hat{\mathcal{Y}}_{\mathbf{p}} \in \mathbb{R}$ are respectively the expected output and the encoder-decoder network’s response at pixel \mathbf{p} , λ is a trade-off parameter between ”object contour” or ”instance mask” pixels and inactive ones, and σ is the sigmoid function. While the term $\ell(\lambda_e, \mathcal{E}, \hat{\mathcal{E}})$ focuses the network’s attention on boundaries, the term $\ell(\lambda_m, \mathcal{M}_s, \hat{\mathcal{M}}_s)$ enables to learn per-seed connectivity, so that encoding-decoding for predicting instance contours on the whole image and instance masks is jointly learned, as suggested by Equation 2. In practice, we set $\lambda_e = 10$ and $\lambda_m = 1$. By definition, \mathcal{L} is fully differentiable as a linear combination of logistic losses. With the proposed loss function, instance segmentation is performed in one forward pass while learning explicitly driven by inter-instance boundaries.

4 Synthesizing training images

To train the network, we generate a synthetic dataset using Blender [2] with custom code to simulate scenes of objects in bulk and render the corresponding RGB-D top views.

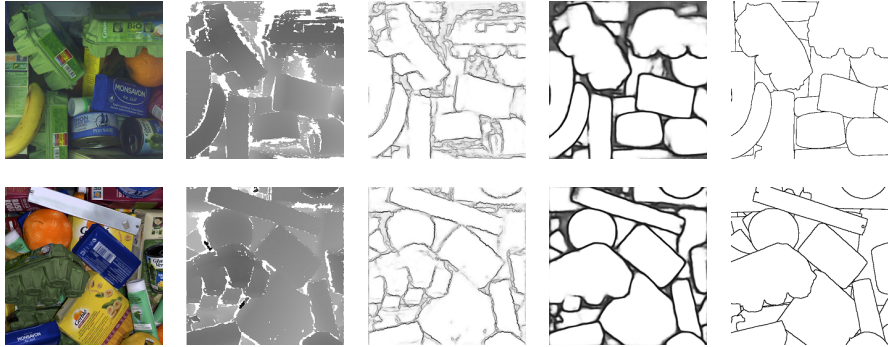


Fig. 4: Detected object contours on real (first row) and synthetic (second row) test images, both using the same model trained only on our synthetic depth maps. From left to right: RGB, depth, contours using a structured random forest (RF) [5], contours using a fully convolutional network (FCN) [30], ground truth.

Simulation We model a static bin, an RGB-D sensor and some objects successively dropped above the bin in random pose using Blender’s physics engine. We add some intraclass geometric variability if needed, by applying isotropic and anisotropic random scaling. In both our real robotic setup and Blender, depth is recovered using an active binocular stereoscopic system composed of two RGB cameras and a pseudo-random pattern projector to add artificial texture to the scene. Calibration, rectification and matching are performed using standard off-the-shelf algorithms.

Rendering Once all instances have been dropped, we render each camera view using Cycles render engine and process the pair of rendered images in the exact same way as a pair of real images. The depth sensor’s noise model is thus implicitly embedded during stereo matching, which gives realistic depth maps without additional post-processing, contrary to [8]. Figure 4 shows a qualitative comparison between real and synthetic images from our dataset.

5 Experimental setup

We conduct two sets of experiments: (1) we train two state-of-the-art object contour detectors – a structured random forest (RF) [5] and a fully convolutional network (FCN) [30] – on synthetic depth maps to achieve equivalent performance on real data at inference time; (2) we show that learning instance segmentation by training a fully convolutional encoder-decoder network with our novel dual-objective loss function (FCN-DOL) outperforms a patch-based network (DeepMask) [22] in terms of boundary precision, while requiring only seed pixels instead of patches. In all our experiments, we focus on synthetic depth as rendered RGB is less realistic and depth conveys more generalizable information.

Data preparation To evaluate our synthetic training, we capture 25 RGB-D top views of real-world scenes featuring 15 household objects. Four expert human annotators manually label each image with the ground-truth contours. Using an off-the-shelf 3D scanner, we scan each of the 15 objects, and generate 1,000 synthetic images of multi-object scenes, *i.e.*, scenes with many instances of many objects, and 1,500 mono-object scenes (100 per object), *i.e.*, scenes with many instances of one object. The multi-object scenes are randomly divided into 3 subsets with the following distribution: 400 for training, 200 for validation, and 400 for test. All mono-object scenes are used only for test. In addition, 5 synthetic and 5 real images are manually labeled separately, each one by 3 different annotators, to evaluate the quality of manual annotation and the impact of ground truth on contour detection performances. We limit the manual annotation to a few images due to the burden of the task, that required 68 hours in total.

Performance metrics For contour detection, we use standard metrics from [1]: the best F-score, and the corresponding recall and precision, with a fixed threshold on dataset scale. We add the *synthetic gap*, *i.e.* the absolute difference be-

tween scores on real (R) and synthetic (S) test images. For instance segmentation, we compute two quantities: the average intersection over union ratio (IoU) over the best ground-truth-matching proposal set; and the boundary precision on instance scale using the same metric as for contour detection. Given one image, a proposal best matches a ground-truth instance if it has the highest IoU among all proposals for one of the ground-truth instances, and if the corresponding IoU is higher than 50%.

Settings for evaluating the use of synthetic images We train two object contour detectors, state-of-the-art in their kind: a structured random forest (RF) from [5] of 4 trees and a fully convolutional network (FCN) from [30] with a VGG-16-based encoder-decoder architecture – on synthetic depth maps to perform equivalently on real data at inference time. For RF, the trees are trained on 10^6 patches extracted from 100 synthetic multi-object images (10^4 patches – as many positive as negative ones – per training image). Multi-scale option and non-maximal suppression are not enabled at test time. Other parameters are set to their default value. For FCN, we perform 180 epochs on 400 synthetic images, using Caffe [13]. Depth is introduced as a new input layer replacing the BGR inputs, therefore both encoder and decoder’s weights are learned. Weights are initialized like in [30], except that two channels of the input layer are removed. We set a constant learning rate of 10^{-4} and a weight decay of 10^{-4} . For both RF and FCN, tests are performed on the 25 manually annotated real depth maps (R-tests), and on 25 synthetic depth maps (S-tests) of multi-object scenes. Human annotators are evaluated on 5 real and 5 synthetic images.

Settings for evaluating the use of edge-mask duality We train two feed-forward deep architectures only on synthetic depth maps of multi-object scenes: (i) the state-of-the-art patch-based network referred to as DeepMask [22], built on a 50-layer residual encoder [11] and taking patches as input; (ii) the previous encoder-decoder network referred to as FCN [30] but trained using our dual-objective loss function (FCN-DOL). For these trainings, we removed all partially occluded instances from ground truth, so that each model is trained only on instances whose mask is connected and of reasonably large area, *i.e.*, graspable instances. DeepMask is initialized with the weights pretrained on MS COCO, and fine-tuned by performing 100 epochs on our 400 synthetic training images, with a constant learning rate of 10^{-4} and data augmentation (shift and scale jittering). FCN-DOL is initialized with our FCN model for contour detection. We add one input layer, corresponding to the translated image that centers the seed pixel, and one output layer, corresponding the object mask including the seed. We then perform 120 epochs driven by our dual-objective loss function, with a constant learning rate of 10^{-4} , using 2 random pixel seeds per ground-truth instance and data augmentation (4 rotations and horizontal mirror). At test time, we define a constant regular grid of seeds over the whole image. For each image, we obtain a set of object mask proposals by forwarding each seed of the grid, and binarizing the continuous output mask with a threshold of 0.8. For both DeepMask and

FCN-DOL, tests are performed on 400 synthetic depth maps of multi-object scenes and 1,500 synthetic images of mono-object scenes.

6 Results


Synthetic over real training images Using simulation enables significant time-savings: annotating a real image takes 40 min in average while generating a synthetic view about 5 min (with a quad-core 3.5GHz Intel Xeon E5 for physics simulation, and a Nvidia Quadro M4000 for GPU rendering), thus 8 times faster. Interestingly, both edge detectors trained on synthetic depth maps still achieve good results in practice when tested on our real data, producing contour maps visually similar to those obtained with a synthetic input, as illustrated in Fig. 4. However, Tab. 1 shows synthetic gaps of 16 and 18 points in F-score inducing that our synthetic data is not realistic enough, here in terms of depth noise modeling. Nevertheless, this gap can be partially relativised, considering the non-zero gap (12 points) achieved by humans. S-tests are indeed conducted against perfect ground truth while R-tests against error-prone annotations, as humans miss hardly detectable edges in noisy areas. Bridging the gap with more realism should be part of a future work to limit any overtraining in the synthetic domain.

		Test									
		Recall			Precision			F-score			
	Modality	Train	25R	25S	Gap	25R	25S	Gap	25R	25S	Gap
RF [5]	Depth	100S	.54	.70	.16	.53	.69	.16	.53	.69	.16
FCN [30]	Depth	400S	.43	.60	.17	.78	.93	.15	.55	.73	.18
Humans	RGB	-	.48	.65	.17	.98	.92	.06	.64	.76	.12

Table 1: Comparative results for instance contour detection on real (R) and synthetic (S) images, using a random forest (RF) [5] and a fully convolutional network (FCN) [30] both trained only on our synthetic depth maps. The synthetic gap (Gap) is the absolute difference between R and S-scores. Humans are given RGB since they are unable to distinguish instances only from depth.

		400Mu		1500Mo	
	Train	Average Best IoU	Boundary Precision	Average Best IoU	Boundary Precision
DeepMask [22]	400Mu	.83	.51	.88	.33
FCN-DOL (Ours)	400Mu	.89	.90	.88	.90

Table 2: Intersection over Union (IoU) and boundary precision of the best matching instances on synthetic multi-object (Mu) and mono-object (Mo) scenes.



(a) Training sample (b) DeepMask [22] (c) FCN-DOL (Ours) (d) Ground truth

Fig. 5: Best matching instance proposals on synthetic test depth maps. With one depth-only training on multi-object scenes (a), our edge-driven model (c) gives more edge-preserving instances and better generalizes to unknown arrangements, in contrast to patch-based networks (b).

Seed-based edge-driven over patch-based instance segments As shown by Fig. 5, our dual-objective loss for instance segmentation by learning explicitly inter-object boundaries leads to sharper mask boundaries than using a state-of-the-art patch-based network, when the scene is a pile of bulk objects as it is the case in a waste sorting application. These results are corroborated by Tab. 2, with a gain of 39 points in boundary precision. In addition, our model performs equally on mono-object setups though it was trained only on multi-object scenes, whereas the patch-based network’s performance drops by 18 points. This suggests that introducing the edge-mask duality at training time enables a better generalization on new arrangements. In a real interactive interface, the human-guided network helps to relevantly reduce the grasp search space as illustrated by Fig. 1. Although the proposed model implies one forward pass per seed, this is not problematic in our application context as we trust the user to perform a relevant click. Our experimental dataset should nevertheless be extended to a much larger

database of 3D models such as ShapeNet [3] to investigate the generalization on new objects, as our current images lack variability in terms of object category.

Limitations of the proposed model and future work The proposed model relies on a state-of-the-art encoder-decoder structure with 130M parameters against less than 60M parameters for DeepMask. This, and the lack of residual blocks, hinder the training, which explains why we performed a two-phase training with three times more fine-tuning epochs than DeepMask in total to reach a stable accuracy. Building a new encoder-decoder architecture with residual blocks and less filters for the layer connecting the encoder and the decoder – currently 79% of the parameters to train – could alleviate this issue. Also, thanks to the translation invariance of fully convolutional networks, any pixel belonging to the object may be selected, and in practice, we can expect a reasonable seed, as a human user instinctively clicks near the object’s centroid. As depicted in Fig. 6, a click on a boundary produces a low-contrasted output as the selected seed doesn’t belong to any instance. A click near a boundary returns a continuous segmentation map more sensitive to the binarization threshold. The robustness on this point could be increased with a smarter seed selection at training time.




Fig. 6: Examples of clicks (the white disks) using the proposed model. From left to right: RGB, predicted edges, predicted continuous mask, binarized masks with respectively a threshold of 0.2, 0.5, 0.8. From top to bottom: a click near the instance’s centroid; a click on a boundary; a click near a boundary.

7 Conclusion

We presented a simulation-based framework towards an interactive application for robotized waste sorting. Realistic synthetic RGB-D images are generated

to train an encoder-decoder network for edge-driven instance segmentation in depth maps. Assuming edge-mask duality in the images, we introduce a novel dual-objective loss function to explicitly focus the learning on inter-instance boundaries. Given seeds provided interactively by a human operator, the network thus produces instance masks with sharper boundaries than a state-of-the-art patch-based approach, in a single forward pass, and better generalizes to unknown object arrangements. Its deep fully convolutional architecture enables strong learning capabilities into a responsive human-robot interface to speed up the overall process of selecting, extracting and sorting objects.

References

1. P. Arbeláez, M. Maire, C. Fowlkes, and J. Malik. Contour detection and hierarchical image segmentation. *IEEE Transactions on Pattern Analysis and Machine Intelligence (TPAMI)*, 33(5):898–916, 2011.
2. Blender Online Community. *Blender - a 3D modelling and rendering package*. Blender Foundation, Blender Institute, Amsterdam, 2016.
3. A. X. Chang, T. A. Funkhouser, L. J. Guibas, P. Hanrahan, Q.-X. Huang, Z. Li, S. Savarese, M. Savva, S. Song, H. Su, J. Xiao, L. Yi, and F. Yu. ShapeNet: An Information-Rich 3D Model Repository. *CoRR*, abs/1512.03012, 2015.
4. C. Choi, Y. Taguchi, O. Tuzel, M.-Y. Liu, and S. Ramalingam. Voting-based pose estimation for robotic assembly using a 3D sensor. In *ICRA*, pages 1724–1731. IEEE, 2012.
5. P. Dollár and C. L. Zitnick. Fast Edge Detection Using Structured Forests. *CoRR*, abs/1406.5549, 2014.
6. A. Dosovitskiy, P. Fischer, E. Ilg, P. Husser, C. Hazirbas, V. Golkov, P. van der Smagt, D. Cremers, and T. Brox. FlowNet: Learning Optical Flow with Convolutional Networks. In *ICCV*, pages 2758–2766. IEEE Computer Society, 2015.
7. M. Firman. RGBD Datasets: Past, Present and Future. In *The IEEE Conference on Computer Vision and Pattern Recognition (CVPR) Workshops*, June 2016.
8. A. Handa, V. Patraucean, V. Badrinarayanan, S. Stent, and R. Cipolla. Understanding RealWorld Indoor Scenes with Synthetic Data. In *CVPR*, pages 4077–4085. IEEE Computer Society, 2016.
9. A. Handa, V. Patraucean, S. Stent, and R. Cipolla. SceneNet: An Annotated Model Generator for Indoor Scene Understanding. In D. Kragic, A. Bicchi, and A. D. Luca, editors, *ICRA*, pages 5737–5743. IEEE, 2016.
10. K. He, G. Gkioxari, P. Dollár, and R. Girshick. Mask-RCNN. In *ICCV*, 2017.
11. K. He, X. Zhang, S. Ren, and J. Sun. Deep Residual Learning for Image Recognition. In *CVPR*, pages 770–778. IEEE Computer Society, 2016.
12. J. H. Hosang, R. Benenson, P. Dollr, and B. Schiele. What makes for effective detection proposals? *CoRR*, abs/1502.05082, 2015.
13. Y. Jia, E. Shelhamer, J. Donahue, S. Karayev, J. Long, R. Girshick, S. Guadarrama, and T. Darrell. Caffe: Convolutional Architecture for Fast Feature Embedding. *arXiv preprint arXiv:1408.5093*, 2014.
14. P. Krahenbuhl and V. Koltun. Learning to Propose Objects. In *CVPR*, pages 1574–1582. IEEE, 2015.
15. K. Lai, L. Bo, X. Ren, and D. Fox. A Large-Scale Hierarchical Multi-View RGB-D Object Dataset. In *ICRA*, pages 1817–1824. IEEE, 2011.

16. G. Lin, C. Shen, A. van den Hengel, and I. Reid. Efficient Piecewise Training of Deep Structured Models for Semantic Segmentation. In *CVPR*, pages 3194–3203. IEEE Computer Society, 2016.
17. T.-Y. Lin, M. Maire, S. Belongie, J. Hays, P. Perona, D. Ramanan, P. Dollár, and C. L. Zitnick. Microsoft COCO: Common Objects in Context. In D. J. Fleet, T. Pajdla, B. Schiele, and T. Tuytelaars, editors, *ECCV (5)*, volume 8693 of *Lecture Notes in Computer Science*, pages 740–755. Springer, 2014.
18. J. Long, E. Shelhamer, and T. Darrell. Fully Convolutional Networks for Semantic Segmentation. In *CVPR*, pages 3431–3440. IEEE, 2015.
19. N. Mayer, E. Ilg, P. Husser, P. Fischer, D. Cremers, A. Dosovitskiy, and T. Brox. A Large Dataset to Train Convolutional Networks for Disparity, Optical Flow, and Scene Flow Estimation. In *CVPR*, pages 4040–4048. IEEE Computer Society, 2016.
20. X. Peng, B. Sun, K. Ali, and K. Saenko. Learning Deep Object Detectors from 3D Models. In *ICCV*, pages 1278–1286. IEEE Computer Society, 2015.
21. P. H. O. Pinheiro, R. Collobert, and P. Dollár. Learning to Segment Object Candidates. In C. Cortes, N. D. Lawrence, D. D. Lee, M. Sugiyama, and R. Garnett, editors, *NIPS*, pages 1990–1998, 2015.
22. P. O. Pinheiro, T.-Y. Lin, R. Collobert, and P. Dollár. Learning to Refine Object Segments. In B. Leibe, J. Matas, N. Sebe, and M. Welling, editors, *ECCV (1)*, volume 9905 of *Lecture Notes in Computer Science*, pages 75–91. Springer, 2016.
23. J. Pont-Tuset, P. Arbelaez, J. T. Barron, F. Marqus, and J. Malik. Multiscale Combinatorial Grouping for Image Segmentation and Object Proposal Generation. *CoRR*, abs/1503.00848, 2015.
24. S. Ren, K. He, R. Girshick, and J. Sun. Faster R-CNN: Towards Real-Time Object Detection with Region Proposal Networks. In C. Cortes, N. D. Lawrence, D. D. Lee, M. Sugiyama, and R. Garnett, editors, *NIPS*, pages 91–99, 2015.
25. A. Richtsfeld, T. Morwald, J. Prankl, M. Zillich, and M. Vincze. Segmentation of Unknown Objects in Indoor Environments. In *IROS*, pages 4791–4796. IEEE, 2012.
26. B. Romera-Paredes and P. H. S. Torr. Recurrent instance segmentation. In B. Leibe, J. Matas, N. Sebe, and M. Welling, editors, *ECCV (6)*, volume 9910 of *Lecture Notes in Computer Science*, pages 312–329. Springer, 2016.
27. G. Ros, L. Sellart, J. Materzynska, D. Vázquez, and A. M. Lopez. The SYNTHIA Dataset: A Large Collection of Synthetic Images for Semantic Segmentation of Urban Scenes. In *CVPR*, pages 3234–3243. IEEE Computer Society, 2016.
28. A. Rozantsev, V. Lepetit, and P. Fua. On Rendering Synthetic Images for Training an Object Detector. *Computer Vision and Image Understanding*, 137:24–37, 2015.
29. N. Silberman, D. Hoiem, P. Kohli, and R. Fergus. Indoor Segmentation and Support Inference from RGBD Images. In A. W. Fitzgibbon, S. Lazebnik, P. Perona, Y. Sato, and C. Schmid, editors, *ECCV (5)*, volume 7576 of *Lecture Notes in Computer Science*, pages 746–760. Springer, 2012.
30. J. Yang, B. L. Price, S. Cohen, H. Lee, and M.-H. Yang. Object Contour Detection with a Fully Convolutional Encoder-Decoder Network. In *CVPR*, pages 193–202. IEEE Computer Society, 2016.
31. C. L. Zitnick and P. Dollár. Edge Boxes: Locating Object Proposals from Edges. In D. J. Fleet, T. Pajdla, B. Schiele, and T. Tuytelaars, editors, *ECCV (5)*, volume 8693 of *Lecture Notes in Computer Science*, pages 391–405. Springer, 2014.