# MVTec D2S: Densely Segmented Supermarket Dataset

Patrick Follmann, Tobias Böttger, Philipp Härtinger, Rebecca König, Markus Ulrich

MVTec Software GmbH, www.mvtec.com {follmann,boettger,haertinger,koenig,ulrich}@mvtec.com



Abstract. We introduce the Denselv Segmented Supermarket (D2S) dataset, a novel benchmark for instance-aware semantic segmentation in an industrial domain. It contains 21000 high-resolution images with pixel-wise labels of all object instances. The objects comprise groceries and everyday products from 60 categories. The benchmark is designed such that it resembles the real-world setting of an automatic checkout, inventory, or warehouse system. The training images only contain objects of a single class on a homogeneous background, while the validation and test sets are much more complex and diverse. To further benchmark the robustness of instance segmentation methods, the scenes are acquired with different lightings, rotations, and backgrounds. We ensure that there are no ambiguities in the labels and that every instance is labeled comprehensively. The annotations are pixel-precise and allow using crops of single instances for articial data augmentation. The dataset covers several challenges highly relevant in the field, such as a limited amount of training data and a high diversity in the test and validation sets. The evaluation of state-of-the-art object detection and instance segmentation methods on D2S reveals significant room for improvement.

**Keywords:** instance segmentation, segmentation dataset, industrial application

### 1 Introduction

The task of *instance-aware semantic segmentation* (*instance segmentation* for short) can be interpreted as the combination of *semantic segmentation* and *object detection*. While *semantic segmentation* methods predict a semantic category for each pixel [20], *object detection* focuses on generating bounding boxes for all object instances within an image [25]. As a combination of both, *instance* 

segmentation methods generate pixel-precise masks for all object instances in an image. While solving this task was considered a distant dream a few years ago, the recent advances in computer vision have made instance segmentation a key focus of current research [20,8,16]. This is especially due to the progress in deep convolutional networks [14] and the development of strong baseline frameworks such as Faster R-CNN [25] and Fully Convolutional Networks (FCN) [20].

All top-performing methods in common instance segmentation challenges are based on deep learning and require a large amount of annotated training data. Accordingly, the availability of large-scale datasets, such as *ADE20K* [32], *Cityscapes* [2], *ImageNet* [27], *KITTI* [5], *COCO* [19], *Mapillary Vistas* [23], *VOC* [4], *Places* [31], *The Plant Phenotyping Datasets* [22], or *Youtube-8M* [1], is of paramount importance.

Most of the above datasets focus on everyday photography or urban street scenes, which makes them of limited use for many industrial applications. Furthermore, the amount and diversity of labeled training data is usually much lower in industrial settings. To train a visual warehouse system, for instance, the user typically only has a handful of images of each product in a fixed setting. Nevertheless, at runtime, the products need to be robustly detected in very diverse settings. Only few datasets focus on industry-relevant challenges in the context of warehouses. The Freiburg Groceries Dataset [12], SOIL-47 [13], and the Supermarket Produce Dataset [26] contain images of supermarket products, but only provide class annotations on image level, and hence no segmentation. The Grocery Products Dataset [6] and GroZi-120 [21] include bounding box annotations that can be used for object detection. However, not all object instances in the images are labeled separately. To the best of our knowledge, none of the existing industrial datasets provides pixel-wise annotations on instance level. In this paper, we introduce the Densely Segmented Supermarket (D2S) dataset, which satisfies the industrial requirements described above. The training, validation, and test sets are explicitly designed to resemble the real-world applications of automatic checkout, inventory, or warehouse systems.

Contributions. We present a novel instance segmentation dataset with highresolution images in a real-world, industrial setting. The annotations for the 60 different object categories were obtained in a meticulous labeling process and are of very high quality. Specific care was taken to ensure that every occurring instance is labeled comprehensively. We show that the high-quality region annotations of the training set can easily be used for artificial data augmentation. Using both the original training data and the augmented data leads to a significant improvement of the average precision (AP) on the test set by about 30 percentage points. In contrast to existing datasets, our setup and the choice of the objects ensures that there is no ambiguity in the labels and an AP of 100% is achievable by an algorithm that performs flawlessly. To evaluate the generalizability of methods, the training set is considerably smaller than the validation and test sets and contains mainly images that show instances of a single category on a homogeneous background. Overall, the dataset serves as a de-



**Fig. 1.** The *D2S* image acquisition setup. Each scene was rotated ten times using a turntable. For each rotation, three images under different illuminations were acquired

manding benchmark and resembles real-world applications and their challenges. The dataset will be made publicly available<sup>1</sup>.

## 2 The Densely Segmented Supermarket Dataset

The overall target of the dataset is to realistically cover the real-world applications of an automatic checkout, inventory, or warehouse system. For example, existing automatic checkout systems in supermarkets identify isolated products that are conveyed on a belt through a scanning tunnel [11,3]. Even though such systems often provide a semi-controlled environment, external influences (e.g. lighting changes) cannot be completely avoided. Furthermore, the system's efficiency is higher if non-isolated products can be identified as well. Consequently, methods should be able to segment also partly occluded objects. Also, the background behind the products is not constant in many applications because of different types of storage racks in a warehouse system or because of dirt on the conveyer belt of a checkout system in the supermarket, for example.

For the D2S dataset, we acquired a total of 21 000 images in 700 different scenes with various backgrounds, clutter objects, and occlusion levels. In order to obtain systematic test settings and to reduce the amount of manual work, a part of the image acquisition process was automated. Therefore, each scene was rotated ten times with a fixed angle step and acquired under three different illuminations.

Setup. The setup used for the image acquisition is depicted in Fig. 1. A high-resolution industrial color camera with  $1920 \times 1440$  pixels was mounted above a turntable. The camera was intentionally mounted off-centered with respect to the rotation center of the turntable to introduce more variations in the rotated images.

<sup>&</sup>lt;sup>1</sup> https://www.mvtec.com/research



Fig. 2. Each scene was acquired at ten different rotations in steps of  $36^{\circ}$ . The camera was mounted slightly off-centered in order to introduce more variation in the images

*Objects.* An overview of the 60 different classes is shown in Fig. 6. The object categories cover a selection of common, everyday products such as fruits, vegetables, cereal packets, pasta, and bottles. They are embedded into a class hierarchy tree. The branches of the tree split the classes into groups of different packaging. This results in neighboring leafs being visually very similar, while distant nodes are visually more different, even if they are semantically similar products, e.g. single apples in comparison to a bundle of apples in a cardboard tray. The class hierarchy can be used, for instance, for advanced training and evaluation strategies similar to those used by YOLO9000 [24]. However, it is not used in the scope of this paper.

Rotations. Each scene was rotated ten times in increments of  $36^{\circ}$ . The turntable allowed to automate this process and to ensure that the rotation angles are precise. An example of the ten rotations for a scene from the training set is displayed in Fig. 2.

*Lighting.* To evaluate the robustness of methods to illumination changes and different amounts of reflection, each scene and rotation was acquired under three different lighting settings. For this purpose an LED ring light was attached to the camera. The illumination was set to span a large spectrum of possible lightings. Hence, as displayed in Fig. 3 (top), the dark images were underexposed and the bright images overexposed.

*Background.* The validation and test scenes have a variety of different backgrounds that are shown in Fig. 3 (*bottom*). This allows to evaluate the generalizability of approaches. In contrast, the training set is restricted to images with a single homogeneous background. It is kept constant to imitate the settings of a warehouse system, where new products are mostly imaged within a fixed environment and not in the test scenario.

*Occlusion and Clutter.* As indicated in Fig. 4, occlusions in the dataset may arise from objects of the same class, objects of a different class, or from clutter objects. Clutter objects have a category that is not present in the training images. They



**Fig. 3.** (*Top*) Each scene was acquired under three different lightings. (*Bottom*) As opposed to the training set (where a single uniform background is used), the test and validation sets include three additional backgrounds. This allows for a detailed evaluation of the robustness of the methods

were added explicitly to the validation and test images to evaluate the robustness of methods to novel objects. Examples of the selected clutter objects are shown in Fig. 5.



Fig. 4. Objects appear with different amounts of occlusion. These may either be caused by objects of the same class, objects of a different class or by clutter objects not within the training set



Fig. 5. To test the robustness of approaches to unseen clutter objects, objects not within the training set were added to the validation and test sets (e.g., a mouse pad and a black foam block)

### 3 Dataset Splitting

In contrast to existing datasets for instance-aware semantic segmentation, such as VOC [4] and COCO [19], the D2S training set has a different distribution with respect to image and class statistics than the validation and test sets. The complexity of the captured scenes as well as the average number of objects per image are substantially higher in the validation and test sets (see Table 1). The motivation for this choice of split is to follow common industrial requirements, such as: low labelling effort, low complexity of training set acquisition for easy replicability, and the possibility to easily add new classes to the system.



Fig. 6. Overview of the 60 different classes within the D2S dataset

The split is performed on a per-scene basis: all 30 images of a scene, i.e. all combinations of the ten rotations and three lightings, are included in either the training, the validation, or the test set. In the following, we describe the rules for generating the splits.

Training Split. To meet the industrial requirements mentioned above, the training scenes are selected to be as simple as possible. Hence, the scenes have a homogeneous background and mostly contain only one object. Furthermore, the amount of occlusions is reduced to a minimum. To summarize, we add scenes to the training split that

- contain only objects of one category<sup>2</sup>,
- provide new views of an object,
- only contain objects with no or marginal overlap,
- have no clutter and a homogeneous background.

The total number of scenes in the training set is 147, resulting in 4380 images of 6900 objects. The rather small training set should encourage work towards the generation of augmented or synthetic training data, for instance using generative adversarial networks [7,15,10,29,30].

*Validation and Test Splits.* The remaining scenes are split between the validation and the test set. They consist of scenes with

 $<sup>^2</sup>$  In order to provide similar views of each object class as they are visible in the validation and test set, four scenes were added to the training set that contain two distinct classes.

split	all	train	val	test
# scenes	700	146	120	434
# images	21000	4380	3600	13020
# objects	72447	6900	15654	49893
# objects/image	3.45	1.58	4.35	3.83
# scenes w. occlusion	393	10	84	299
# scenes w. clutter	86	0	18	68
rotations		$\checkmark$	$\checkmark$	$\checkmark$
lighting variation		$\checkmark$	$\checkmark$	$\checkmark$
background variation			$\checkmark$	$\checkmark$
clutter			$\checkmark$	$\checkmark$

**Table 1. Split statistics.** Due to our splitting strategy, the number of images and the number of instances per image is significantly lower for the training set. The complexity of validation and test scenes is approximately the same

- single or multiple objects of different classes,

- touching or occluded objects,
- clutter objects and
- varying background.

These scenes were chosen such that the generalization capabilities of approaches can be evaluated. Additionally, current methods struggle with heavy occlusion and novelty detection. These issues are addressed by this choice of splits as well. The split between validation and test set was performed on subgroups of images containing the same number of total and occluded objects. This ensures that both sets have approximately the same distribution. The ratio of the number of scenes in the validation and test set is chosen to be 1:4. The reasons for this decision are twofold: First, the evaluation of the model on a small validation set is faster. Second, we do not want to encourage training on the validation set, but stimulate work on approaches that require little training data or use augmentation techniques. The statistics of the number of images and objects in the splits are visualized in Table 1.

### 4 Statistics & Comparison

In this section, we compare our dataset to VOC [4] and COCO [19]. These datasets have encouraged many researchers to work on instance segmentation and are frequently used to benchmark state-of-the-art methods.

Dataset Statistics. As summarized in Table 2, D2S contains significantly more object instances than VOC, but fewer than COCO. Specifically, although the D2S training set is larger than that of VOC, the number of training objects is

Table 2. Dataset statistics. Number of images and objects per split, average number of objects per image and number of classes for *D2S* (ours), *VOC 2012*, and *COCO*. \*For *VOC 2012* and *COCO*, the object numbers are only available for the training and validation set

Dataset		VOC	COCO	D2S
# images	all	4369	163957	21000
	train	1464	118287	4380
	val	1449	5000	3600
	test	1456	40670	13020
# objects	all	-	-	72447
	train	3507	849941	6900
	val	3422	36335	15654
	test	-	-	49893
# obj/img		$2.38^{*}$	$7.19^{*}$	3.45
# classes		20	80	60

less than 1% of those in COCO. This choice was made intentionally, as in many industrial applications it is desired to use as few training images as possible. In contrast, the proportion of validation images is significantly larger for D2S in order to enable a thorough evaluation of the generalization capabilities. On average, there are half as many objects per image in D2S as in COCO.

Class Statistics. Since the images of COCO and VOC were taken from flickr<sup>3</sup>, the distribution of object classes is not uniform. In both datasets, the class *person* dominates, as visualized in Fig. 7: 31% and 25% of all objects belong to this class for COCO and VOC, respectively. Moreover, 10% of the classes with

<sup>3</sup> https://www.flickr.com



**Fig. 7.** Ratio of objects per class for *D2S* (*orange*), *VOC* (*green*) and *COCO* (*violet*). In *COCO* and *VOC*, the class *person* is dominant and some classes are underrepresented. In *D2S*, the number of objects per class is uniformly distributed. Note that for *COCO* and *VOC* the diagram was calculated based on train and val splits



Fig. 8. Number of images per class and split sorted by the total number of images per class for *D2S*. The number of images per class is almost uniformly distributed

the highest number of objects are represented by 51% and 33% of all objects, while only 5.4% and 13.5% of the objects belong to the 25% of classes with the lowest number of objects. This class imbalance is valid since both *COCO* and *VOC* represent the real world where some classes naturally appear more often than others. In the evaluation all classes are weighted uniformly. Therefore, the class imbalance inherently poses a challenge to learn all classes equally well, independent from the number of training samples. For example, the *COCO 2017* validation set contains nine instances of the class *toaster*, but 10777 instances of *person*. Nevertheless, both categories are equally weighted in the calculation of the mean average precision, which is the metric used for ranking the methods in the *COCO* segmentation challenge.

There is no such class imbalance in D2S. In the controlled environment of the supermarket scenario, all classes have the same probability to appear in an image. The class with the highest number of objects is represented by only 2.7% of all objects. Only 14% of the objects represent the 10% of classes with the highest number of objects, while 19% of the objects are from the 25% of classes with the lowest number of objects. The class distribution of D2S is visualized in Fig. 8, where the number of images per class is shown in total and for each split. As mentioned in the previous section, the number of images for each class is rather low in the training split. This is especially the case for classes that have a similar appearance for different views, such as *kiwi* and *orange\_single*. Note that, although the split choice between validation and test set is not made on the class level, each class is well represented in both sets. The key challenge of the D2S dataset is thus not the handling of underrepresented classes, but the low amount of training data.

Label Consistency. It is difficult to ensure that all object instances in large realworld datasets are labeled consistently. On the one hand, it is hard to establish a reliable review process for the labeling of large datasets, e.g. to avoid unlabeled objects. On the other hand, some labels are ambiguous by nature, for instance



**Fig. 9.** Large real-world datasets are extremely difficult to label consistently. In the examples from *ADE20K*, *VOC* and *COCO*, some labels are missing (*from left to right*): a window, the sofa, some donuts, and the painting of a person

a painting of a person. Fig. 9 shows examples for label inconsistencies from ADE20K [32], VOC and COCO.

In D2S, the object classes are unambiguous and have been labeled by six expert annotators. We ensured that all present objects are annotated with high quality labels. A perfect algorithm, which flawlessly detects and segments every object in all images of the D2S dataset, will achieve an AP of 100%. This is not the case for COCO, VOC, and ADE20K. In these datasets, if an algorithm correctly detects one of the objects that is not labeled, the missing ground truth leads to a false positive. Furthermore, if such an object is not found by an algorithm, the resulting false negative is not accounted for. As algorithms improve, this might prevent better algorithms from obtaining higher scores in the benchmarks. It should be noted that in COCO, this problem is addressed using crowd annotations, i.e. regions containing many objects of the same class that are ignored in the evaluation. However, crowd annotations are not present in all cases.

# 5 Benchmark

In this section, we provide first benchmark results for our dataset. We evaluate the performance of state-of-the-art methods for object detection [25,18] and instance segmentation [8,16]. We experiment with various training sets, which differ in the number of rotations and the availability of under- and overexposed images. Furthermore, we evaluate a simple approach for augmenting the training data artificially.

#### 5.1 Evaluated Methods

Object Detection. For the object detection task, we evaluate the performance of Faster R-CNN [25] and RetinaNet [18]. We use the official implementations of both methods, which are provided in the Detectron<sup>4</sup> framework. Both methods use a ResNet-101 [9] backbone with Feature Pyramid Network [17].

Instance Segmentation. For the instance segmentation task, we evaluate the performance of Mask R-CNN [8] and FCIS [16]. We use the official implementation of Mask R-CNN in the Detectron framework and the official implementation of

<sup>&</sup>lt;sup>4</sup> https://github.com/facebookresearch/Detectron

11

FCIS provided by the authors<sup>5</sup>. Mask R-CNN uses a ResNet-101 with Feature Pyramid Network as backbone, while FCIS uses a plain ResNet-101. Since both methods output boxes in addition to the segmentation masks, we also include them in the object detection evaluation.

Training. All methods are trained end-to-end. The network weights are initialized with the COCO-pretrained models provided by the respective authors. The input images are resized to have a shorter side of 800 pixels (600 pixels for FCIS, respectively). All methods use horizonal flipping of the images at training time. FCIS uses online hard example mining [28] during training.

#### 5.2 Evaluation Metric

The standard metric used for object detection and instance segmentation is mean average precision (mAP) [4]. It is used, for instance, for the ranking of state-of-the-art methods in the *COCO* segmentation challenge [19]. We compute the mAP exactly as in the official COCO evaluation tool<sup>6</sup> and give its value in percentage points. The basic average precision (AP) is the area under the precision-recall curve, computed for a specific intersection over union (IoU) threshold. In order to reward algorithms with better localization, the AP is usually averaged over multiple IoU thresholds, typically the interval [0.5, 0.95] in steps of 0.05. The mAP is the mean over APs of all classes in the dataset. In the following, we just use the abbreviation AP for the value averaged over IoUs and classes. When referring to class-averaged AP for a specific IoU threshold, e.g. 0.5, we write AP<sub>50</sub>.

#### 5.3 Data Augmentation

In order to keep the labeling effort low and still achieve good results, it is crucial to artificially augment the existing training set such that it can be used to train deep neural networks. Hence, we experiment with a simple data augmentation technique, which serves as baseline for more sophisticated approaches. In particular, we simulate the distribution of validation and test set using only the annotations of the training set. For this purpose, we assemble 10 000 new artificial images that contain one to fifteen objects randomly picked from the training split. We denote the augmented data as **aug** in Table 3. For each generated image, we randomly sample the lighting and number of object instances. For each instance, we randomly sample its class, the orientation, and the location in the image. The background of these images is the plain turntable. We make sure that the instances' region centers lie on the turntable and that occluded objects have a visible area larger than 5000 pixels. Fig. 10 shows example images of the artificially augmented dataset for all three different lightings. Due to the high-quality annotations without margins around the object border, the artificially

<sup>&</sup>lt;sup>5</sup> https://github.com/msracver/FCIS

<sup>&</sup>lt;sup>6</sup> https://github.com/cocodataset/cocoapi

assembled images have an appearance that is very similar to the original test and validation images.



Fig. 10. The artificial augmented training set is generated by randomly assembling objects from the basic training set

#### 5.4 Results

When trained on the full training set train and evaluated on the test set, the instance segmentation methods provide solid baseline APs of 49.5% (Mask R-CNN) and 45.6% (FCIS). The object detection results are on a similar level, with APs of 46.5% (Mask R-CNN), 44.0% (FCIS), 46.1% (Faster R-CNN), and 51.0% (RetinaNet). Tables 3 and 4 show the results in full detail.

Ablation Study. As aforementioned, the D2S splits are based on scenes, i.e. all rotations and lightings for one placement of objects are included in the same split. In order to evaluate the importance of these variations and the ability of the instance segmentation methods to learn invariance with respect to rotations and illumination, we perform an ablation study. For this purpose, we create three subsets of the full training set train. The train\_rot0 set contains all three lightings, but only the first rotation of each scene. The train\_light0 set contains only the default lighting, but all ten rotations of each scene. The train\_rot0\_light0 set contains only the default lighting and the first rotation for each scene.

The resulting AP values of the instance segmentation methods Mask R-CNN and FCIS are summarized in Table 3 (top). As expected, we obtain the best results when training on the full train set. Training only on the first rotation reduced the AP on the test set by 15.7% and 9.1% for Mask R-CNN and FCIS, respectively. Training only with default lighting reduced the AP slightly by 3.4% for Mask R-CNN and increased the AP by a neglible 0.4% for FCIS. Training on train\_rot0\_light0 reduced the AP by 13.2% and 12.9%, respectively. Overall, the results indicate that the models are more invariant to changes in lighting than to rotations of the objects.

Data Augmentation. As shown in Table 3, training on the augmented dataset aug boosts the AP on the test set to 76.1% and 69.8% for Mask R-CNN and FCIS, respectively. This is significantly higher than the 49.5% and 45.6% achieved by training on the original train set. Combining the sets train and aug to train+aug further improves the AP by 8.3% and 2.7%, respectively.

**Table 3. Instance segmentation benchmark results on the test set.** Mean average precision values for models trained on different training sets. (*Top*) Training on different subsets of the train set. (*Bottom*) Training on augmented data yields the highest AP values

	M	ask R-C	NN	FCIS			
	AP	$AP_{50}$	$AP_{75}$	AP	$AP_{50}$	$AP_{75}$	
train	49.5	57.6	51.3	45.6	58.3	51.3	
train_rot0	33.8	41.6	35.6	36.5	47.5	41.8	
train_light0	46.1	54.8	48.0	46.0	59.3	52.0	
train_rot0_light0	36.3	45.1	38.6	32.7	43.4	38.1	
aug	71.6	86.9	81.7	69.8	87.6	82.4	
train+aug	79.9	89.1	85.3	72.5	88.1	83.5	

 Table 4. Object detection benchmark results on the test set.
 Mean average

 precision values for models trained on different training sets

	Ma	sk R-C	CNN	FCIS			Faster R-CNN			RetinaNet		
	AP	$AP_{50}$	$\mathrm{AP}_{75}$	AP	$AP_{50}$	$AP_{75}$	AP	$AP_{50}$	$\mathrm{AP}_{75}$	AP	$AP_{50}$	$AP_{75}$
train	46.5	58.3	53.5	44.0	59.4	51.7	46.1	55.2	49.7	51.0	61.0	52.8
train_rot0	34.1	42.5	38.3	34.6	48.2	41.3	36.7	46.9	41.5	32.9	39.8	34.5
train_light0	45.5	55.7	49.5	44.0	60.3	51.9	43.7	53.9	47.8	51.7	62.0	53.6
train_rot0_light0	35.7	46.0	40.5	29.9	43.9	35.4	34.3	44.3	39.0	31.6	38.9	33.2
aug	72.9	87.9	82.0	69.9	88.1	80.7	73.5	88.4	82.2	74.2	86.9	81.4
train+aug	78.3	89.8	84.9	68.3	88.5	80.9	78.0	90.3	84.8	80.1	89.6	84.5

Object Detection. We conduct the same ablation study for the task of object detection. The resulting AP values for all training splits of the methods Faster R-CNN and RetinaNet, as well as the results of instance segmentation methods Mask R-CNN and FCIS evaluated on bounding box level, are summarized in Table 4. It is interesting to note, that these AP values are not always better than the AP values obtained for the more difficult task of instance segmentation. For all methods the overall performance is very similar. Reducing the training set to only one rotation or only one lighting per scene results in worse performance. Analogously, augmenting the dataset by generating artificial training images results in a strong improvement.

Qualitative results. We show qualitative results of the best-performing method Mask R-CNN in Fig. 11. Furthermore, Fig. 12 shows typical failure cases we observed for Mask R-CNN and FCIS on the D2S dataset. More qualitative results are provided in the supplementary material.



Fig. 11. (*Top*) Ground truth annotations from the *D2S* val and test sets. (*Bottom*) Results of Mask R-CNN trained on the train set. The classes are indicated by different colors



**Fig. 12.** Typical failure cases of of Mask R-CNN and FCIS on *D2S.* (*From left to right*) (1) Nearby objects are detected as a single instance. (2) Segmentation mask spans to neighboring objects. (3 and 4) Background is falsely detected as object

# 6 Conclusion

We have introduced D2S, a novel dataset for instance-aware semantic segmentation that focuses on real-world industrial applications. The dataset addresses several challenges highly relevant in the field, such as dealing with very limited training data. The training set is intentionally small and simple, while the validation and test sets are much more complex and diverse. As opposed to existing datasets, D2S has a very uniform distribution of the samples per class. Furthermore, the fixed acquisition setup prevents ambiguities in the labels, which in turn allows flawless algorithms to achieve an AP of 100%. We further showed how the high-quality annotations can easily be utilized for artificial data augmentation to significantly boost the performance of the evaluated instance segmentation methods from an AP of 49.5% and 45.6% to 79.9% and 72.5%, respectively. Overall, the benchmark results indicate a significant room for improvement of the current state-of-the-art. We believe the dataset will help to boost research on instance-aware segmentation and leverage new approaches for artificial data augmentation.

# Appendix

We provide the following supplementary material:

- Qualitative results of the instance segmentation models
- Per-class results
- The class tree
- A video showing the quality of annotations and the variation of the dataset<sup>7</sup>

# A Qualitative Results

Figures 13 and 14 show qualitative instance segmentation results of FCIS and Mask R-CNN, respectively. The models have been trained on different splits of the D2S dataset. Note that the model trained on train+aug performs significantly better than the model trained only on train, especially for objects with occlusion. All models struggle with textured background and reflections.

# **B** Per-Class Results

Tables 5 and 6 show the per-class AP values for FCIS and Mask R-CNN, respectively, for all 60 classes and all training splits. In general, the highest APs are achieved for classes with large objects and no reflections, such as *lettuce, salad\_iceberg* or *oranges*. The data augmentation boosts particularly the worst-performing classes, such as *banana\_single, cucumber, zucchini* and all *gepa\_bio\_und\_fair\_\**-classes.

# C Data Augmentation

The aug split consists of 10 000 artificially generated images. In order to evaluate the influence of the number of training images, we used subsets of 2 000, 4 000, 6 000, and 8 000 images of aug, respectively, to train FCIS pretrained on ImageNet. The resulting AP values for both validation and test set are shown in Fig. 15. It can be seen that using more images results in better performance. However, already with 2 000 generated images the AP increases by over 20 percentage points compared to using train.

# D Class Tree

The class tree assigns superclasses to classes according to their appearance or packaging in a hierarchical manner. It is visualized in Fig. 16.

<sup>&</sup>lt;sup>7</sup> https://www.mvtec.com/research



Fig. 13. FCIS qualitative results. A collection of exemplary results of FCIS trained on different splits. Note that complex backgrounds often lead to false detections. The object classes are color-coded



**Fig. 14. Mask R-CNN qualitative results.** A collection of exemplary results of Mask R-CNN trained on different splits. Note that complex backgrounds often lead to false detections. The object classes are color-coded

Table 5. FCIS AP values per class. The AP values of the FCIS models trained on different training sets are calculated on the test set. The mean AP is calculated over all 60 classes. The highest class AP values are highlighted in bold.

class	train	rot0	light0	rot0 light0	aug	train+aug
adelholzener alpenquelle classic 075	40.1	23.9	28.1	17.4	42.9	51.8
adelholzener alpenquelle naturell 075	46.7	27.9	36.7	23.8	68.0	69.1
adelholzener classic bio apfelschorle 02	49.6	21.0 24.0	47.9	8.0	64.8	73 7
adelholzener classic naturell 02	47.6	18.0	/0.3	10.1	62.6	71.1
adelholzener geurmet mineralwasser 02	54.0	22.4	55.3	15.6	71.8	78.8
augustinor lagorbraou holl 05	40.0	22.4	47.0	17.1	52.8	60.2
augustiner maischier 05	49.9 91.4	17.1	20.0	7.0	22.0	50.5
augustilier_weissbier_05	45 1	24.5	20.0	1.4	55.0	57.0
acca_cola_05	49.1	11 4	24.0	10.2 9 1	55.0	60.8
coca_cola_light_05	43.0	11.4 40.6	54.2	0.1	69.1	70.1
suntory_gokuri_innonade	26.0	49.0	04.0 27.0	44.0	42.1	12.1
tegernseer_nen_05	50.9	0.0	57.2	1.1	42.3	04.8
corny_nussvoll	51.0	03.0	55.8 50.0	43.1	89.9	91.0
corny_nussvoil_single	59.0	22.9	59.6	16.5	08.0	72.0
corny_schoko_banane	44.1	21.4	40.4	17.3	83.2	85.1
corny_schoko_banane_single	17.4	13.7	23.1	11.0	54.0	54.6
dr_oetker_vitalis_knuspermuesli_klassisch	46.3	21.7	27.5	16.8	80.0	83.1
koelln_muesli_fruechte	47.0	29.7	38.5	24.6	82.2	85.5
koelln_muesli_schoko	46.2	11.0	30.1	12.1	76.8	80.1
caona_kakaohaltiges_getraenkepulver	57.6	63.3	62.4	67.9	85.0	85.0
cocoba_fruehstueckskakao_mit_honig	41.6	37.6	44.7	39.9	72.3	75.4
cafe_wunderbar_espresso	45.5	43.6	53.1	34.5	82.3	74.0
douwe_egberts_professional_kaffee_gemahlen	41.7	37.3	40.6	30.8	71.8	68.7
gepa_bio_caffe_crema	40.7	36.3	43.5	38.6	62.4	64.8
gepa_italienischer_bio_espresso	52.4	32.9	49.9	36.5	59.6	66.7
apple_braeburn_bundle	58.9	68.2	66.0	67.9	84.7	88.9
apple_golden_delicious	53.6	43.6	57.3	40.8	80.5	85.8
apple_granny_smith	58.4	54.8	64.0	52.9	78.8	82.6
apple_roter_boskoop	63.4	67.3	65.9	51.0	82.1	86.2
avocado	55.6	56.4	54.9	52.5	83.4	84.4
banana_bundle	62.9	66.0	73.1	62.3	79.0	80.8
banana_single	23.9	19.8	26.3	16.3	51.6	51.4
clementine	50.8	59.9	58.9	66.2	80.3	81.6
clementine_single	48.2	50.1	48.6	40.8	68.8	72.8
grapes_green_sugraone_seedless	52.8	61.7	57.4	59.5	76.8	78.3
grapes_sweet_celebration_seedless	55.7	55.5	63.3	58.4	76.8	80.2
kiwi	31.8	27.6	31.5	24.1	72.0	75.4
orange_single	58.1	71.2	61.4	69.4	76.7	82.8
oranges	62.5	65.8	65.0	66.5	80.8	83.5
pear	60.2	45.2	59.3	27.7	75.7	77.9
pasta reggia elicoidali	49.0	31.1	53.9	29.5	88.4	86.7
pasta reggia fusilli	49.4	30.6	54.4	30.9	77.5	53.2
pasta reggia spaghetti	47.3	36.0	47.5	32.8	82.3	83.4
franken tafelreiniger	47.3	37.3	40.7	22.1	48.3	62.7
pelikan tintenpatrone canon	22.3	18.5	22.4	9.2	66.5	65.6
ethiquable gruener tee cevlon	53.1	44 7	54.0	34.3	70.5	80.2
gepa bio und fair fencheltee	14.9	9.4	13.8	5.3	55 7	57.6
gepa bio und fair kamillentee	26.5	21 /	27.4	13.3	64.1	50.3
gepa bio und fair kraeuterteemischung	20.0	15.3	19.8	13.0	67.4	61.5
gopa bio und fair pfofformingtoo	36.5	18.9	22.4	11.6	68.8	64.0
gepa_bio_und_fair_prefierInniztee	44.0	20.2	44.5	21.4	70.0	75.9
lilimoniono too conlanoo	44.0	20.5	44.5	21.4	62.7	70.6
kiiimanjaro_tea_eari_grey	40.0 97.6	20.7	41.0	20.0	03.7 60.6	70.0
cucumber	27.0	40.7	29.3	26.0	54.6	50.0
	30.1 47 1	40.7	57.1	30.9	04.0 79.0	05.4 70.7
Ieldsalat	47.1	37.0	50.3	42.0	78.9	79.7
ieutuce	07.3	79.5	10.6	81.1	91.7	92.2
rispentomaten	39.0	45.3	49.4	51.7	71.4	70.9
roma_rispentomaten	41.1	35.3	46.1	27.9	00.7	68.4
rucola	45.4	28.8	52.3	40.3	68.7	70.9
salad_iceberg	53.9	63.7	61.7	69.5	81.4	83.5
zucchini	44.1	29.2	44.6	24.3	64.4	71.3

Table 6. Mask R-CNN AP values per class. The AP values of the Mask R-CNN models trained on different training sets are calculated on the test set. The mean AP is calculated over all 60 classes. The highest class AP values are highlighted in bold.

class	train	rot0	light0	rot0_light0	aug	train+aug
adelholzener_alpenquelle_classic_075	40.6	16.4	28.2	9.4	43.1	44.8
adelholzener_alpenquelle_naturell_075	53.8	23.8	37.1	23.1	72.1	76.3
adelholzener_classic_bio_apfelschorle_02	57.7	12.1	56.7	4.7	65.8	75.9
adelholzener_classic_naturell_02	48.5	11.3	59.3	36.3	68.4	79.0
adelholzener gourmet mineralwasser 02	69.9	36.9	61.4	45.7	76.2	86.8
augustiner lagerbraen hell 05	48.6	17.1	57.0	17.9	49.8	70.1
augustiner weissbier 05	15.6	12.7	21.9	13.0	41.2	50.5
coca cola 05	45.8	23.7	48.8	16.4	53.2	61.2
coca cola light 05	38.9	16.7	35.1	9.1	55.3	70.8
suntory gokuri limonade	67.1	51.6	67.2	54.0	59.4	83.0
togernsoor hell 03	44.8	11 7	38.5	6.0	47.7	67.8
comput pusquall	52.6	45.1	40.2	55.2	41.1	01.8
corpy_nussvoll_single	70.0	40.1	43.2 64.2	32.0	33.0 77.6	84.2
computed by barana	45.9	10.0	27.0	20.8	94.1	00.1
comy_schoko_banane	40.2	14.0	37.9	29.0	50.7	90.1
corny_schoko_banane_single	23.0	14.1	24.0	19.2	09.7	04.0
dr_oetker_vitalis_knuspermuesii_klassiscn	48.9	28.7	24.8	24.0	83.4	91.0
koelin_muesii_iruechte	51.1	24.3	33.8	20.7	88.8	93.3
koelin_muesli_schoko	57.8	4.4	38.6	8.1	85.3	90.2
caona_kakaohaltiges_getraenkepulver	56.9	45.3	44.9	56.0	82.0	88.0
cocoba_fruehstueckskakao_mit_honig	51.0	43.0	44.9	36.4	76.1	86.7
cafe_wunderbar_espresso	39.3	12.8	44.5	37.1	83.0	88.2
douwe_egberts_professional_kaffee_gemahlen	45.8	38.2	42.2	38.7	71.5	75.1
gepa_bio_caffe_crema	40.4	31.5	41.5	41.9	81.9	82.1
gepa_italienischer_bio_espresso	59.3	30.9	52.5	19.6	61.9	65.5
apple_braeburn_bundle	59.8	56.8	60.7	66.2	82.3	91.7
apple_golden_delicious	78.2	72.8	66.4	73.0	76.1	91.3
apple_granny_smith	72.5	71.4	65.9	66.4	75.5	87.4
apple_roter_boskoop	79.4	71.2	76.6	<b>79.8</b>	78.1	94.1
avocado	73.0	60.6	76.2	74.2	84.2	93.4
banana_bundle	51.4	47.5	55.8	46.9	76.8	86.6
banana_single	27.4	18.0	28.6	20.6	55.9	57.9
clementine	54.6	57.5	58.4	59.8	79.7	83.0
clementine_single	59.5	65.0	60.4	59.9	81.2	89.7
grapes_green_sugraone_seedless	44.0	29.4	27.8	34.2	68.9	77.2
grapes_sweet_celebration_seedless	61.5	54.9	51.4	46.8	75.7	80.7
kiwi	45.5	60.0	55.6	54.0	68.9	88.7
orange_single	69.6	70.5	66.5	72.7	81.5	91.2
oranges	61.1	51.1	61.2	65.8	81.3	86.3
Dear	56.4	40.7	74.8	70.5	76.5	85.1
pasta reggia elicoidali	43.1	23.3	48.2	38.8	88.5	90.5
pasta reggia fusilli	48.8	23.7	51.6	38.4	76.7	76.5
pasta reggia spaghetti	55.4	28.3	45.0	28.5	85.0	90.5
franken tafelreiniger	48.2	33.4	39.8	34 1	52.8	67.7
pelikan tintenpatrone canon	30.5	25.6	27.2	32.1	69.1	75.8
othiguable grueper too coulon	52.3	25.0	48.1	18.6	77.8	86.4
gene his und feir fencheltes	10.2	50.5	14.7	10.0	50.2	55.0
gepa bio und fair kamillantea	20.7	17.0	14.7	10.9	64.4	55.9 71.1
gepa_bio_und_fair_krannentee	27.0	10.1	17.8	0.4	59.1	72.6
gepa_bio_und_tain_kraeuterteennischung	27.0	50.1	24.5	9.4 7.0	60.2	72.0
gepa_bio_und_fair_prefierminztee	37.3	0.2 00.1	24.0	7.9	69.3 70.4	((.2
gepa_bio_und_fair_rooibostee	47.3	29.1	45.0	21.9	70.4	77.3
kilimanjaro_tea_earl_grey	46.1	31.6	38.4	25.4	64.1	75.6
cucumber	29.9	23.9	33.0	27.6	71.4	77.6
carrot	54.1	44.5	42.8	44.5	62.7	72.1
feldsalat	49.5	33.2	40.5	34.8	77.4	84.0
lettuce	51.6	55.0	34.7	38.4	93.4	96.2
rispentomaten	45.0	39.4	45.2	52.3	71.9	78.0
roma_rispentomaten	38.8	33.7	37.5	40.2	71.6	73.7
rucola	38.9	15.8	41.0	12.8	71.7	76.1
salad_iceberg	55.7	53.4	52.9	66.8	77.9	86.4
zucchini	49.4	23.9	51.0	24.8	68.7	84.0



Fig. 15. Data augmentation. Influence of number of artificially generated training images on AP for validation and test set



Fig. 16. Class Tree. Objects are sorted based mainly on their appearance and packaging rather than on their content. For example, different kinds of coffee are close to each other since they all have a textured, soft plastic packaging. But they are far from all sorts of tea, which is packaged in cardboard.

# References

- Abu-El-Haija, S., Kothari, N., Lee, J., Natsev, P., Toderici, G., Varadarajan, B., Vijayanarasimhan, S.: Youtube-8m: A large-scale video classification benchmark. CoRR abs/1609.08675 (2016) 2
- Cordts, M., Omran, M., Ramos, S., Rehfeld, T., Enzweiler, M., Benenson, R., Franke, U., Roth, S., Schiele, B.: The cityscapes dataset for semantic urban scene understanding. In: CVPR. pp. 3213–3223 (2016) 2
- 3. ECRS: RAPTOR. https://www.ecrs.com/products/point-of-sale-pos/ accelerated-checkout/,, accessed 2018-03-07 3
- Everingham, M., Eslami, S.M.A., Gool, L.J.V., Williams, C.K.I., Winn, J.M., Zisserman, A.: The pascal visual object classes challenge: A retrospective. International Journal of Computer Vision 111(1), 98–136 (2015) 2, 5, 7, 11
- 5. Geiger, A., Lenz, P., Stiller, C., Urtasun, R.: Vision meets robotics: The KITTI dataset. International Journal of Robotics Research 32(11), 1231–1237 (2013) 2

- George, M., Floerkemeier, C.: Recognizing products: A per-exemplar multi-label image classification approach. In: ECCV. Springer (2014) 2
- Gurumurthy, S., Kiran Sarvadevabhatla, R., Venkatesh Babu, R.: DeLiGAN: Generative adversarial networks for diverse and limited data. In: CVPR. pp. 166–174 (2017) 6
- He, K., Gkioxari, G., Dollar, P., Girshick, R.: Mask R-CNN. In: ICCV. pp. 1059– 1067 (2017) 2, 10
- He, K., Zhang, X., Ren, S., Sun, J.: Deep residual learning for image recognition. In: CVPR. pp. 770–778 (2016) 10
- Huang, X., Li, Y., Poursaeed, O., Hopcroft, J., Belongie, S.: Stacked generative adversarial networks. In: CVPR. pp. 5077–5086 (2017) 6
- 11. ITAB: HyperFLOW. https://itab.com/en/itab/checkout/self-checkouts/,, accessed 2018-03-07 3
- Jund, P., Abdo, N., Eitel, A., Burgard, W.: The freiburg groceries dataset. CoRR abs/1611.05799 (2016), https://arxiv.org/abs/1611.05799 2
- 13. Koubaroulis, D., Matas, J., Kittler, J.: Evaluating colour-based object recognition algorithms using the SOIL-47 database. In: ACCV. p. 2 (2002) 2
- LeCun, Y., Bengio, Y., Hinton, G.: Deep learning. Nature 521(7553), 436–444 (2015) 2
- Li, J., Liang, X., Wei, Y., Xu, T., Feng, J., Yan, S.: Perceptual generative adversarial networks for small object detection. In: CVPR. pp. 1222–1230 (2017) 6
- Li, Y., Qi, H., Da, J., Ji, X., Wei, Y.: Fully convolutional instance-aware semantic segmentation. In: CVPR. pp. 2359–2367 (2017) 2, 10
- Lin, T.Y., Dollár, P., Girshick, R., He, K., Hariharan, B., Belongie, S.: Feature pyramid networks for object detection. In: CVPR (2017) 10
- Lin, T.Y., Goyal, P., Girshick, R., He, K., Dollár, P.: Focal Loss for Dense Object Detection. In: ICCV (2017) 10
- Lin, T., Maire, M., Belongie, S.J., Hays, J., Perona, P., Ramanan, D., Dollár, P., Zitnick, C.L.: Microsoft COCO: common objects in context. In: ECCV. pp. 740– 755 (2014) 2, 5, 7, 11
- Long, J., Shelhamer, E., Darrell, T.: Fully convolutional networks for semantic segmentation. In: CVPR. pp. 3431–3440 (2015) 1, 2
- Merler, M., Galleguillos, C., Belongie, S.: Recognizing groceries in situ using in vitro training data. In: CVPR. pp. 1–8 (2007) 2
- Minervini, M., Fischbach, A., Scharr, H., Tsaftaris, S.A.: Finely-grained annotated datasets for image-based plant phenotyping. Pattern recognition letters 81, 80–89 (2016) 2
- Neuhold, G., Ollmann, T., Rota Bulo, S., Kontschieder, P.: The mapillary vistas dataset for semantic understanding of street scenes. In: ICCV. pp. 4990–4999 (2017) 2
- 24. Redmon, J., Farhadi, A.: Yolo9000: Better, faster, stronger. In: CVPR (2017) 4
- Ren, S., He, K., Girshick, R.B., Sun, J.: Faster R-CNN: Towards real-time object detection with region proposal networks. IEEE Transactions on Pattern Analysis and Machine Intelligence 39(6), 62–66 (2017) 1, 2, 10
- Rocha, A., Hauagge, D.C., Wainer, J., Goldenstein, S.: Automatic fruit and vegetable classification from images. Computers and Electronics in Agriculture 70(1) (2010) 2
- Russakovsky, O., Deng, J., Su, H., Krause, J., Satheesh, S., Ma, S., Huang, Z., Karpathy, A., Khosla, A., Bernstein, M.S., Berg, A.C., Li, F.: ImageNet large scale visual recognition challenge. International Journal of Computer Vision 115(3), 211– 252 (2015), https://doi.org/10.1007/s11263-015-0816-y 2

- 28. Shrivastava, A., Gupta, A., Girshick, R.: Training region-based object detectors with online hard example mining. In: CVPR. pp. 761–769 (2016) 11
- Shrivastava, A., Pfister, T., Tuzel, O., Susskind, J., Wang, W., Webb, R.: Learning from simulated and unsupervised images through adversarial training. In: CVPR. pp. 2107–2116 (2017) 6
- Zhang, H., Xu, T., Li, H., Zhang, S., Wang, X., Huang, X., Metaxas, D.N.: Stack-GAN: Text to photo-realistic image synthesis with stacked generative adversarial networks. In: ICCV. pp. 5907–5915 (2017) 6
- Zhou, B., Khosla, A., Lapedriza, À., Torralba, A., Oliva, A.: Places: An image database for deep scene understanding. CoRR abs/1610.02055 (2016), http:// arxiv.org/abs/1610.02055 2
- Zhou, B., Zhao, H., Puig, X., Fidler, S., Barriuso, A., Torralba, A.: Semantic understanding of scenes through the ADE20K dataset. CoRR abs/1608.05442 (2016), http://arxiv.org/abs/1608.05442 2, 10