

Deep Extreme Cut: From Extreme Points to Object Segmentation

K.-K. Maninis* S. Caelles* J. Pont-Tuset L. Van Gool
Computer Vision Lab, ETH Zürich, Switzerland




Figure 1. **Example results of DEXTR:** The user provides the extreme clicks for an object, and the CNN produces the segmented masks.

Abstract

This paper explores the use of extreme points in an object (left-most, right-most, top, bottom pixels) as input to obtain precise object segmentation for images and videos. We do so by adding an extra channel to the image in the input of a convolutional neural network (CNN), which contains a Gaussian centered in each of the extreme points. The CNN learns to transform this information into a segmentation of an object that matches those extreme points.

We demonstrate the usefulness of this approach for guided segmentation (grabcut-style), interactive segmentation, video object segmentation, and dense segmentation annotation. We show that we obtain the most precise results to date, also with less user input, in an extensive and varied selection of benchmarks and datasets. All our models and code will be made publicly available.

1. Introduction

Deep learning techniques have revolutionized the field of computer vision since their explosive appearance in the ImageNet competition [33], where the task is to classify images into predefined categories, that is, algorithms produce one label for each input image. Image and video segmentation, on the other hand, generate dense predictions where each pixel receives a (potentially different) output classification. Deep learning algorithms, especially Convolutional Neural Networks (CNNs), were adapted to this scenario by removing the final fully connected layers to produce dense predictions.

*First two authors contributed equally

Supervised techniques, those that train from manually-annotated results, are currently the best performing in many public benchmark and challenges [42, 12]. In the case of image and video segmentation, the supervision is in the form of dense annotations, *i.e.* each pixel has to be annotated in an expensive and cumbersome process. Weakly-supervised techniques, which train from incomplete but easier-to-obtain annotations, are still significantly behind the state of the art. Semi-automatic techniques, which need a human in the loop to produce results, are another way of circumventing the expensive training annotations but need interaction at test time, which usually comes in the form of a bounding box [7] or scribbles [20] around the object of interest. How to incorporate this information at test time without introducing unacceptable lag, is also a challenge.

This paper tackles all these scenarios in a unified way and shows state-of-the-art results in all of them in a variety of benchmarks and setups. We present Deep Extreme Cut (DEXTR), that obtains an object segmentation from its four extreme points: the left-most, right-most, top, and bottom pixels. Figure 1 shows an example result of our technique.

In the context of semi-automatic object segmentation, we show that information from extreme clicking results in significantly more accurate segmentations than the ones obtained from bounding-boxes (PASCAL, COCO, Grabcut) in a Grabcut-like formulation. DEXTR outperforms other methods using extreme points or object proposals (PASCAL), and provides a better input to video object segmentation (DAVIS-2016, DAVIS-2017). DEXTR can also incorporate more points beyond the extreme ones, which further refines the quality (PASCAL).

DEXTR can also be used to obtain dense annotations to train supervised techniques. We show that we obtain very accurate annotations with respect to the ground truth, but more importantly, that algorithms trained on the annotations obtained by our algorithm perform as good as when trained from the ground-truth ones. If we add the cost to obtain such annotations into the equation, then training using DEXTR is significantly more efficient than training from the ground truth for a given target quality.

We perform an extensive and comprehensive set of experiments on COCO, PASCAL, Grabcut, and DAVIS, to demonstrate the effectiveness of our approach. We will make all code, pre-trained models and annotation interfaces used in this project publicly available.

2. Related Work

Weakly Supervised Signals for Segmentation: Numerous alternatives to expensive pixel-level segmentation have been proposed and used in the literature. Image-level labels [26], noisy web labels [1, 16] and scribble-level labels [20] are some of the supervisory signal that have been used to guide segmentation methods. Closer to our approach, [3] employs point-level supervision in the form of a single click to train a CNN for semantic segmentation and [25] use central points of an imaginary bounding box to weakly supervise object detection. Also related to our approach, [7, 17] train semantic segmentation methods from box supervision. Recently, Papadopoulos et al. proposed a novel method for annotating objects by extreme clicks [24]. They show that extreme clicks provide additional information to a bounding box, which they use to enhance GrabCut-like object segmentation from bounding boxes. Different than these approaches, we use extreme clicking as a form of guidance for deep architectures, and show how this additional information can be used to further boost accuracy of segmentation networks, and help various applications.

Instance Segmentation: Several works have tackled the task of grouping pixels by object instances. Popular grouping methods provide instance segmentation in the form of automatically segmented object proposals [14, 30]. Other variants provide instance-level segmentation from a weak guiding signal in the form of a bounding box [32]. Accuracy for both groups of methods has increased by recent approaches that employ deep architectures trained on large dataset with strong supervisory signals, to learn how to produce class-agnostic masks from patches [28, 29], or from bounding boxes [40]. Our approach relates to the second group, since we utilize information from extreme clicks to group pixels of the same instance, with higher accuracy.

Interactive Segmentation from points: Interactive segmentation methods have been proposed in order to reduce

annotation time. In this context, the user is asked to gradually refine a method by providing additional labels to the data. Grabcut [32] is one of the pioneering works for the task, segmenting from bounding boxes by gradually updating an appearance model. Our method relates with interactive segmentation using points as the supervisory signal. Click Carving [15] interactively updates the result of video object segmentation by user-defined clicks. Recent methods use these ideas in the pipeline of deep architectures. iFCN [41] guides a CNN from positive and negative points, automatically acquired from the ground-truth masks. RIS-Net [10] build on iFCN to improve the result by adding local context. Our method significantly improves the results by using 4 points as the supervisory signal: the extreme points.

3. Method

3.1. Extreme points

One of the most common ways to perform weakly supervised segmentation is drawing a bounding box around the object of interest [4, 38, 32, 18]. However, in order to draw the corners of a bounding box, the user has to click points outside the object, drag the box diagonally, and adjust it several times to obtain a tight, accurate bounding box. This process is cognitively demanding, with increased error rates and labelling times [24].

Recently, Papadopoulos et al. [24] have shown a much more efficient way of obtaining a bounding box using extreme clicks, spending on average 7.2 seconds instead of 34.5 seconds required for drawing a bounding box around an object [35]. They show that extreme clicking leads to high quality bounding boxes that are on par with the ones obtained by traditional methods. These extreme points belong to the top, bottom, left-most and right-most parts of the object. Extreme-clicking annotations by definition provide more information than a bounding box; they contain four points that are on the boundary of the object, from which one can easily obtain the bounding-box. We use extreme points for object segmentation leveraging their two main outcomes: the points and their inferred bounding box.

3.2. Segmentation from Extreme Points

The overview of our method is shown in Figure 2. The annotated extreme points are given as a guiding signal to the input of the network. To this end, we create a heatmap with activations in the regions of extreme points. We center a 2D Gaussian around each of the points, in order to create a single heatmap. The heatmap is concatenated with the RGB channels of the input image, to form a 4-channel input for the CNN. In order to focus on the object of interest, the input is cropped by the bounding box, formed from the extreme point annotations. To include context on the resulting crop, we relax the tight bounding box by several pixels.

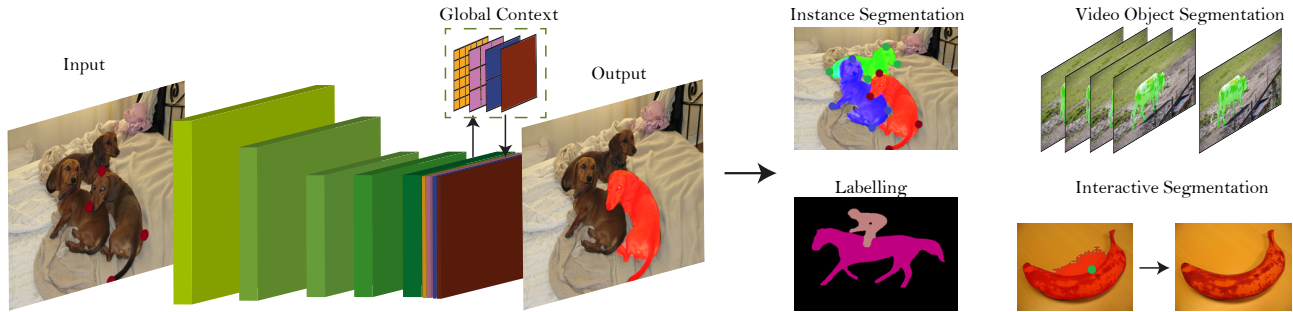


Figure 2. **Architecture of DEXTR:** Both the RGB image and the labeled extreme points are processed by the CNN to produce the segmented mask. The applicability of this method is illustrated for various tasks: Instance, Semantic, Video, and interactive segmentation.

After the pre-processing step that comes exclusively from the extreme clicks, the input consists of an RGB crop including an object, plus its extreme points.

We choose *ResNet-101* [13] as the backbone of our architecture, as it has been proven successful in a variety of segmentation methods [6, 12]. We remove the fully convolutional layers as well as the max pooling layers in the last two stages to preserve acceptable output resolution for dense prediction, and we introduce atrous convolutions in the two last stages to maintain the same receptive field. After the last *ResNet-101* stage, we introduce a pyramid scene parsing module [42] to aggregate global context to the final feature map. Initializing the weights of the network from pre-training on ImageNet has been proven beneficial for various tasks [22, 39, 12]. For most experiments, we use the provided Deeplab-v2 model pre-trained on ImageNet, and fine-tuned on PASCAL for semantic segmentation.

The output of the CNN is a probability map representing whether a pixel belongs to the object that we want to segment or not. The CNN is trained to minimize the standard cross entropy loss, which takes into account that different classes occur with different frequency in a dataset:

$$\mathcal{L} = \sum_{j \in Y} w_{y_j} C(y_j, \hat{y}_j), \quad j \in 1, \dots, |Y| \quad (1)$$

where w_{y_j} depends on the label y_j of pixel j . In our case we define w_{y_j} with $y_j \in \{0, 1\}$ as the inverse normalized frequency of labels inside the minibatch. $C(\cdot)$ indicates the standard cross-entropy loss between the label and the prediction \hat{y}_j . The balanced loss has proven to perform very well in boundary detection [39, 23], where the majority of the samples belong to the background class. We note that our method is trained from strong mask-level supervision, on publicly available datasets, using the extreme points as a guiding signal to the network.

In order to segment an object, our method uses an object-centered crop, therefore there is a much higher number of samples belonging to the foreground than to the background and the use of a balanced loss proves to be beneficial.

Alternatives for each of the components used in our final model have been studied, and a detail comparison can be found in Section 4.1.

3.3. Use cases for DEXTR

Class-agnostic Instance Segmentation: One application of DEXTR is class-agnostic instance segmentation. In this task, we click on the extreme points of an object in an image, and we obtain a mask prediction for it. The selected object can be of any class, as our method is class agnostic.

In Section 4.2, we compare our method with the state of the art in two different datasets, PASCAL and Grabcut, where we improve current results. We also analyse the generalization of our method to other datasets and to unseen categories. We conclude positive results in both experiments, the performance drop for testing in a different dataset than the one used for training is very small and the result achieved is the same whether the class has been seen during training or not.

Annotation: The common annotation pipeline for segmentation can also be assisted by DEXTR. In this framework, instead of detailed polygon labels, the workload of the annotator is reduced to only providing the extreme points of an object, and DEXTR produces the desired segmentation. In this pipeline, the labelling cost is reduced by a factor of 10 (from 79s needed for a mask, to 7.2 seconds needed for the extreme clicks) [24].

In Section 4.3, the quality of the produced masks are validated for two tasks: instance segmentation, and semantic segmentation. We show that our method produces very accurate masks, on par with the ground-truth annotation in terms of quality, with much less effort.

Video Object Segmentation: DEXTR can also improve the pipeline of video object segmentation. We focus on the semi-supervised setting where methods use one or more masks as inputs to produce the segmentation of the whole video. Our aim is to replace the costly per pixel annotation masks by the masks produced by our algorithm after the user has selected the extreme points of a certain object, and re-train strongly supervised state-of-the-art video segmentation architectures.

In Section 4.4, we provide results on two different dataset, DAVIS-2016 and DAVIS-2017. We conclude that

state-of-the-art results can be achieved reducing the annotation time by a factor of 5. Moreover, for almost any specific annotation budget, better results can also be obtained using a higher number of masks produced by our algorithm rather than expensive per-pixel annotated masks.

Interactive User Input: The pipeline of DEXTR can also be used in the frame of interactive segmentation from points [41, 40]. We work on the case where the user labels the extreme points of an object, but is nevertheless not satisfied with the obtained results. The natural thing to do in such case is to annotate an extra point (not extreme) in the region that segmentation fails, and expect for a refined result. Given the nature of extreme points, we expect that the extra point also lies in the boundary of the object.

To simulate such behaviour, we first train DEXTR on a first split of a training set of images, using the 4 extreme points as input. For the extra point, we infer on an image of the second split of the training set, and compute the accuracy of its segmentation. If the segmentation is accurate (eg. $IoU \geq 0.8$), the image is excluded from further processing. In the opposite case ($IoU < 0.8$), we select a fifth point in the erroneous area. To simulate human behaviour, we perturbate its location and we train the network with 5 points as input. Results presented in Section 4.5 indicate that it is possible to recover performance on the difficult examples, by using such interactive user input.

3.4. Implementation Details

Automatic Generation of Extreme Points: In [24], extreme points were obtained by crowd-sourcing. Instead, we train our network by taking advantage of already existing segmentation datasets [8, 21], by recovering the extreme points from the masks of already annotated objects. In order to simulate human behaviour, the recovered points are jittered by 10 pixels. Automatically obtained extreme points provide us with sufficient amount of data ($\sim 25k$ and $\sim 600k$ objects for PASCAL and COCO, respectively). The extreme points are provided as an extra channel to the input of the CNN, in the form of Gaussians with standard deviation of 10 pixels. To focus on a particular object, we crop the input from the extreme points, and relax by 50 pixels on each side, for increasing context. In Section 4.1 we show that DEXTR segments accurately with both manually and automatically obtained extreme points.

Training and testing details: DEXTR is trained on PASCAL for 100 epochs. For several experiments where we need to train on a different dataset, we train on COCO training set for 10 epochs. The learning rate is set to 0.0025, with momentum of 0.9 and weight decay of $5 * 10^{-4}$. A mini-batch of 5 objects is used for PASCAL, whereas for COCO, due to the large size of the database we train on 4 GPUs with an effective batch size of 20. Training on PAS-

CAL takes approximately 20 hours on a Titan-X GPU, and 5 days on COCO. Testing the network is fast, requiring 80 ms for a forward pass.

4. Experimental Validation

Our method is extensively validated in five publicly available databases: PASCAL [8], COCO [21], DAVIS-2016 [27], DAVIS-2017 [31], and Grabcut [32], for various experimental setups that show its applicability and generalization capabilities. We use DEXTR trained on PASCAL (augmented by the labels of SBD [11] following the common practice - 10582 images), unless indicated differently. We perform an ablation study to separately validate all components of our method in Section 4.1. Class-agnostic instance segmentation experiments from extreme points are presented in Section 4.2, whereas 4.3 and 4.4 are dedicated to how DEXTR contributes to labelling and video object segmentation pipelines, respectively. Section 4.5 presents our method as an interactive segmenter from points.

4.1. Ablation Study

The following sections show a number of ablation experiments in the context of class agnostic instance segmentation to quantify the importance of each of the components of our algorithm and to justify various design choices. Section 4.1 summarizes these results.

Architecture: We use *ResNet-101* as the backbone architecture, and compare two different alternatives. The first one is a straightforward fully convolutional architecture (Deeplab-v2 [6]) where the fully connected and the two last max pooling layers are removed; and the last two stages are substituted with dilated (or atrous) convolutions. This keeps the size of the prediction in reasonable limits ($8 \times$ lower than the input). We also tested a region-based architecture, similar to Mask R-CNN [12], with a re-implementation of the ResNet-101-C4 variant [12], which uses the fifth stage (C5) for regressing a mask from the Region of Interest (RoI), together with the re-implementation of the RoI-Align layer. For more details please refer to [12]. In the first architecture, the input is a patch around the object of interest, whereas in the latter the input is the full image, and cropping is applied at the RoI-Align stage. Deeplab-v2 performs +3.9% better. We conclude that the output resolution of ResNet-101-C4 (28×28) is inadequate for the level of detail that we target.

Bounding boxes vs. extreme points: We study the performance of Deeplab-v2 as a foreground-background classifier given a bounding box compared to the extreme points. In the first case, the input of the network is the cropped image around the bounding box plus a margin of 50 pixels to include some context. In the second case, the extreme points are fed together in a fourth channel of the input to

guide the segmentation. Including extreme points to the input increases performance by +12.9%, which suggest that they are a source of very valuable information that the network uses additionally to guide its output.

Loss: For the task of class-agnostic instance segmentation, we compare two binary losses, i.e. the standard cross-entropy and a class-balanced version of it, where the loss for each class in the batch is weighted by its inverse frequency. Class-balancing the loss gives more importance to the less frequent classes, and has been successful in various tasks [39, 5]. DEXTR also performs better when the loss is balanced, leading to a performance boost of +3.3%.

Full image vs. crops: Having the extreme points annotated allows for focusing on specific regions in an image, cropped by the limits specified by them. In this experiment, we compare how beneficial it is to focus on the region of interest, rather than processing the entire image. To this end, we crop the region surrounded by the extreme points, relaxing it by 50 pixel for increased context and compare it against the full image case. We notice that cropping increases performance by +7.9%, and is especially beneficial for the small objects of the database. Similar findings have been reported for video object segmentation by [19].

Atrous spatial pyramid (ASPP) vs. pyramid scene parsing (PSP) module: Pyramid Scene Parsing Network [42] steps on the Deeplab-v2 [6] architecture to further improve results on Semantic Segmentation. Their main contribution was a global context module (PSP) that employs global features together with the local features for dense prediction. We compare the two network heads, the original ASPP [6], and the recent PSP module [42] for our task. The increased results of the PSP module (+2.3%) indicate that the PSP module builds a global context that is also useful in our case.

Manually vs. automatically generated extreme points: The authors of [24] perform a large-scale crowd sourcing to obtain extreme point annotations. In our work, instead, we obtain the extreme points automatically by adding some random noise around the true extreme points obtained from the ground truth. We analyse the differences, by comparing our automatic sources of extreme points to the manual ones provided by the authors of [24]. Since [24] annotated the training set of PASCAL, we use DEXTR trained on COCO and evaluate on the same labelled set of objects. Table 1 shows that the manual extreme points lead to -5% performance, which we interpret as them being less accurate with respect to the ground-truth segmentation. To back up this reasoning and put the numbers in context, the last column shows the best result an oracle could get by segmenting within the bounding box defined by the manual extreme

	Ours Manual	Ours Automatic	Oracle Manual
IoU	80.1%	85.1%	94.4%

Table 1. **Manually [24] vs. automatically annotated extreme points in PASCAL:** IoU for manually and ground truth generated extreme points.

Method	\mathcal{J}	\mathcal{F}	$\mathcal{J}\&\mathcal{F}$
Manual extreme points	80.9	81.9	81.4
Automatic extreme points	79.5	82.9	81.2

Table 2. **Manually vs. automatically generated extreme points:** Comparison in all the frames of DAVIS-2016. Metrics are region accuracy (\mathcal{J}), contour accuracy (\mathcal{F}), and their mean.

points. The upper bound is -5.6% below the perfect result, which aligns well with the drop DEXTR suffers.

In a second experiment in this direction, different annotators marked all video sequences of DAVIS-2016 [27] with extreme points. Annotation time (average of all 1376 frames of the validation set) was 7.5 seconds per frame, in line with [24] (7.2 s. per image). Table 2 shows the results of DEXTR when tested on DAVIS-2016 from manual vs. automatic extreme points. DEXTR results in very high quality results and no significant differences are observed between the two types of extreme points. Please note that no DAVIS image was used for training the models.

Distance-map vs. fixed points: Recent works [41, 40, 10] that focus on segmentation from (not-extreme) points use the distance transform of positive and negative annotations as an input to the network, in order to guide the segmentation. We compare with their approach by substituting the fixed Gaussians to the distance transform of the extreme points. We notice a performance drop of -1.3%, suggesting that using fixed Gaussians centered on the points is a better representation when coupled with extreme points. In Section 4.2 we compare to such approaches, showing that extreme points provide a much richer representation than arbitrary points on the foreground and the background of an object.

Summary Table 3 summarizes the main ablated results that have been discussed above, analyzing all components. Table 4 illustrates the building blocks that lead to the best performing variant for our method. All in all, we start by a Deeplab-v2 base model working on bounding boxes. We add the PSP module (+2.3%), the extreme points in the input of the network (+12.9%), and more annotated data from SBD (+1%) to reach maximum accuracy. The improvement comes mostly because of the guidance from extreme points, which highlights their importance for the task.

Component #1	Component #2	Gain in IoU
Region-based	Deeplab-v2	+3.9%
Bounding Boxes	Extreme Points	+12.9%
Cross Entropy	Balanced BCE	+3.3%
Full Image	Crop on Object	+7.9%
ASPP	PSP	+2.3%
Fixed Points	Distance Map	-1.3%

Table 3. **Ablation study:** Comparison between choices in various components of our system. Mean IoU over all objects in PASCAL VOC 2012 val set.

Method	IoU (%)
Deeplab-v2	75.3
+ PSP	77.6
+ Extreme Points	90.5
+ SBD data (Ours)	91.5

Table 4. **Ablation study:** Building performance in PASCAL VOC 2012 val set.

Method	IoU Aut.	IoU Man.
Sharpmask [29] bounding box	73.1%	xx.x%
Sharpmask [29] upper bound	78.0%	78.0%
Extreme clicks [24]	-	73.6%
Ours	85.1%	80.1%

Table 5. **Comparison in PASCAL:** IoU when using automatic and manually generated extreme points against class-agnostic instance segmentation methods.

4.2. Class-agnostic Instance Segmentation

Comparison to the State of the Art in PASCAL: We compare our method against state-of-the-art class-agnostic instance segmentation methods in Table 5. When we compare to extreme clicks [24] using their manual extreme points, DEXTR gets a boost of +6.5%. We then compare to two other baselines using SharpMask [29], the state-of-the-art object proposal technique. In the first row, we evaluate the proposal (out of 1000) whose bounding box best overlaps with the ground-truth bounding box, mimicking a naive algorithm to segment boxes from proposals. The second row shows the upper bound of SharpMask, that is, the best proposal against the ground truth, selected by an oracle. Both approaches are well below our result (-12.0% and -7.1%). Figure 3 illustrates some results obtained by our method on PASCAL.

Comparison to the State of the Art in Grabcut: We use our best PASCAL model and we test it in the Grabcut [32] dataset. This dataset contains 50 images each one with one annotated object from various categories, some of them not belonging to any of the PASCAL ones (banana, scissors, kangaroo, etc.). The evaluation metric is the error rate: the percentage of misclassified pixels within the bounding boxes provided by [18]. Our method achieves the lowest error rate by 1.1% (32% relative improvement).

Method	Error Rate (%)
GrabCut [32]	8.1
KernelCut [36]	7.1
OneCut [37]	6.7
Extreme clicks [24]	5.5
BoxPrior [18]	3.7
MILCut [38]	3.6
DeepGC [40]	3.4
Ours	2.3

Table 6. **Comparison in Grabcut:** Error rates compared to the state of the art techniques.

	Train	Test	IoU
Unseen categories	PASCAL	COCO MVal w/o PASCAL classes	80.3%
	PASCAL	COCO MVal only PASCAL classes	79.9%
Dataset generalization	PASCAL	COCO MVal	80.1%
	COCO	COCO MVal	82.1%
	COCO	PASCAL	87.8%
	PASCAL	PASCAL	89.8%

Table 7. **Generalization to unseen classes and across datasets:** Intersection over union results of training in one setup and testing on another one. MVal stands for mini-val.

Generalization to unseen categories and across datasets:

Table 7 shows our results when trained on a certain dataset/categories (first column) and tested in another one (second column). In order to make a fair comparison, all the models are pre-trained only on Imagenet [33] for image labeling and trained on the specified dataset for category-agnostic instance segmentation. The first two rows show that our technique is indeed class agnostic, since the model trained on PASCAL achieves roughly the same performance in COCO mini-val regardless of the categories tested. The remaining rows shows that DEXTR also generalizes very well across datasets, since differences are around only 2% of performance drop.

4.3. Annotation

As seen in the previous section, DEXTR is able to generate high-quality class-agnostic masks given only extreme points as input. The resulting masks can in turn be used to train other deep architectures for other tasks or datasets, that is, we use extreme points as a way to annotate dataset with object segmentations. In this experiment we use the resulting masks to train for two different tasks: (a) class-agnostic instance segmentation, and (b) semantic segmentation. We compare to the full ground-truth supervision given a certain annotation budget or a certain number of annotated images. We work on PASCAL, and since we need to generate masks for the training set, we use DEXTR trained on COCO.

Class-agnostic Instance Segmentation: We train DEXTR on COCO and infer the model on the objects of PASCAL train set. We train a different instance of our method on either the generated annotations or on




Figure 3. **Qualitative results by DEXTR on PASCAL:** Each instance with the input extreme points and the resulting mask overlaid.

the original ground-truth masks. Figure 4.3 compares the performance on PASCAL validation as a function of annotation budget (time to label the training data, *i.e.* 79s for segmentation masks and 7.2 seconds for extreme points) (left), and the number of annotated objects instances (right).

The results trained on our masks are significantly better than training from the ground truth on the same budget (e.g. 70% IoU with 9-minute annotation vs. 1h50, or 82% IoU for 1 hour annotation vs. 62.5%). Surprisingly, the results trained from our masks are slightly better even when training with the same number of images. We hypothesize that the more intricate details of the ground-truth annotations, which are lost on our masks, might make the training harder.




Figure 4. **Quality vs. annotation budget:** IoU for instance segmentation trained on our masks or the ground truth, as a function of annotation budget (left) and number of training images (right).

Semantic Segmentation: In order to evaluate the usefulness of our produced masks for other tasks, we move to semantic segmentation. We sort the results by annotation budget and number of images used for training. We use Deeplab [6] with a ResNet-101 backbone, and the PSP [42] head as the semantic segmentation network. To keep training time manageable, we do not use multi-scale training/testing. We report results on the PASCAL VOC 2012 val set, and measure performance by the standard mIoU measure (IoU per-category and averaged over categories).

Figure 4.3 shows the results, with respect to annotation

budget and number of images as before. For completeness, we also report the results of PSPNet [42] (●) by evaluating the model provided by the authors (pre-trained on COCO, with multi-scale training and testing). As before, DEXTR’s annotations are significantly more effective given a certain budget, and reach practically the same performance than ground truth when given the same number of annotations.




Figure 5. **Quality vs. annotation budget:** mIoU for semantic segmentation trained on our masks or the input, as a function of annotation budget (left) and the number of annotated images (right).

4.4. Video Object Segmentation

We test DEXTR also for Video Object Segmentation on the DAVIS datasets [27, 31]. We focus on the semi-supervised setting *i.e.* the mask in one or more frames of the object that we want to segment is given as input to the algorithm, and as before we will compare the results obtained from the masks obtained by DEXTR or the ground truth having a certain annotation budget. We assume that the annotation time of the DAVIS masks is the same than that of COCO [21] (79 seconds per instance), despite the former are significantly more accurate.

We use OSVOS [5], as a state-of-the-art semi-supervised video object segmentation technique, which heavily relies on the appearance of the annotated frame, and their code is publicly available. Figure 6 (left) shows the performance of OSVOS in DAVIS-2016 [27] trained on the ground truth mask (—) or the masks generated by DEXTR from extreme points (—). We reach the same performance as using one ground-truth annotated mask with an annotation

Trained on	4 points	4 points-all	5 points	5 points + OHEM
IoU	59.6%	69.0%	69.2%	73.2%

Table 8. **Interactive Object Segmentation Evaluation:** Average IoU on difficult cases of PASCAL VOC 2012 validation dataset.

budget 5 times smaller. Once we train with more than one ground-truth annotated mask, however, even though we can generate roughly ten times more masks, we cannot achieve the same accuracy. We believe this is so because DAVIS-2016 sequences have more than one semantic instance per mask while we only annotate a global set of extreme points, which confuses DEXTR.

To corroborate this intuition, we perform the same experiment in DAVIS-2017 [31], where every mask contains only one semantic instance. Figure 6 (right) shows that the performance gap with respect to using the full ground-truth mask is much smaller than in DAVIS-2016. Overall, we conclude that DEXTR is also very efficient to reduce annotation time in video object segmentation.




Figure 6. **Quality vs. annotation budget in Video Object Segmentation:** OSVOS performance when trained from our masks or the ground truth, on DAVIS-2016 (left) and on DAVIS-2017 (right).

4.5. Interactive Object Segmentation

DEXTR for Interactive Segmentation: We experiment on PASCAL VOC 2012 for interactive object segmentation. We split the training dataset into two equal splits. Initially, we train DEXTR on the first split and test on the second. We then focus on the objects with inaccurate segmentations, i.e. $\text{IoU} < 0.8$, to simulate the ones on which a human would mark a fifth point on the boundary with the highest error. From the perspective of network training, this can be interpreted as Online Hard Example Mining (OHEM) [34], where one only needs to back-propagate gradients for the training examples that lead to the highest losses. Results are presented in Table 8.

We evaluate performance on PASCAL validation set. We first select the objects that lead to poor performance ($\text{IoU} < 0.8$) when applying the network trained on the first split. We report the average IoU on them (338 objects - 59.6%).

Using the network trained further on the hard examples, with a fifth boundary point, performance increases to 73.2% (“5 points + OHEM”).

Since the increased performance is partially due to the

Method	Number of Clicks		IoU @ 4 clicks	
	PASCAL @85%	Grabcut@90%	PASCAL	Grabcut
GraphCut [4]	> 20	> 20	41.1	59.3
Geodesic matting [2]	> 20	> 20	45.9	55.6
Random Walker [9]	16.1	15	55.1	56.9
iFCN [41]	8.7	7.5	75.2	84.0
RIS-Net [10]	5.7	6.0	80.7	85.0
Ours	4.0	4.0	91.5	94.4

Table 9. **PASCAL and Grabcut evaluation:** Comparison to interactive segmentation methods in terms of number of clicks to reach a certain quality and in terms of quality at 4 click.

increased amount of training data (first split + hard examples of the second split), we need to disentangle the two sources of performance gain. To this end, we train DEXTR on 4 points, by appending the hard examples of the second split to the first split of our training set (“4 points-all”).

Results suggest that DEXTR learns to handle more input information given interactively in the form of boundary clicks, to improve results of poorly segmented difficult examples (+4.2%).

Interestingly, OHEM is a crucial component for improving performance: without it the network does not focus on the difficult examples (only 11% of objects of the second training split are hard examples), and fails to improve on the erroneous region indicated by the fifth boundary point (“5 points”).

Comparison to the State of the Art: To illustrate the importance of features generated from extreme points, we compare against the state-of-the-art methods for interactive segmentation, by considering extreme points as 4 clicks. Table 9 shows the number of clicks that each method needs to reach the performance of DEXTR, as well as their performance when the input is 4 clicks, in PASCAL and Grabcut. DEXTR reaches about 10% higher performance at 4 clicks than the best competing method. This further demonstrates the enhanced performance of the CNN, when guided by extreme points.

5. Conclusions

We have presented DEXTR, a CNN architecture for semi-automatic segmentation that turns extreme clicking annotations into accurate object masks. In an ablated study, we show how information from extreme clicking is an crucial ingredient for accurate segmentation, and that we can employ already existing datasets to automatically acquire extreme points. The applicability of our method is illustrated from a series of experiments regarding semantic, instance, video, and interactive segmentation in 5 different datasets. DEXTR can also be used as an accurate and efficient object mask annotation tool, reducing the labelling costs by a factor of 10.

Acknowledgements: Research funded by the EU Framework Programme for Research and Innovation Horizon 2020 (Grant No. 645331, EurEyeCase), and the Swiss Commission for Technology and Innovation (CTI, Grant No. 19015.1 PFES-ES, NeGeVA). We gratefully acknowledge support by armasuisse, thank NVidia Corporation for donating the GPUs used in this project, and Dim P. Papadopoulos for sharing the resources of [24] with us.

References

- [1] E. Ahmed, S. Cohen, and B. Price. Semantic object selection. In *CVPR*, 2014. 2
- [2] X. Bai and G. Sapiro. Geodesic matting: A framework for fast interactive image and video segmentation and matting. *IJCV*, 82(2):113–132, 2009. 8
- [3] A. Bearman, O. Russakovsky, V. Ferrari, and L. Fei-Fei. What’s the Point: Semantic Segmentation with Point Supervision. *ECCV*, 2016. 2
- [4] Y. Y. Boykov and M.-P. Jolly. Interactive graph cuts for optimal boundary & region segmentation of objects in nd images. In *ICCV*, 2001. 2, 8
- [5] S. Caelles, K.-K. Maninis, J. Pont-Tuset, L. Leal-Taixé, D. Cremers, and L. Van Gool. One-shot video object segmentation. In *CVPR*, 2017. 5, 7
- [6] L.-C. Chen, G. Papandreou, I. Kokkinos, K. Murphy, and A. L. Yuille. Deeplab: Semantic image segmentation with deep convolutional nets, atrous convolution, and fully connected CRFs. *T-PAMI*, 2017. 3, 4, 5, 7
- [7] J. Dai, K. He, and J. Sun. Boxesup: Exploiting bounding boxes to supervise convolutional networks for semantic segmentation. In *ICCV*, 2015. 1, 2
- [8] M. Everingham, L. Van Gool, C. K. I. Williams, J. Winn, and A. Zisserman. The PASCAL Visual Object Classes Challenge 2012 (VOC2012) Results. 4
- [9] L. Grady. Random walks for image segmentation. *T-PAMI*, 28(11):1768–1783, 2006. 8
- [10] J. Hao Liew, Y. Wei, W. Xiong, S.-H. Ong, and J. Feng. Regional interactive image segmentation networks. In *ICCV*, 2017. 2, 5, 8
- [11] B. Hariharan, P. Arbeláez, L. Bourdev, S. Maji, and J. Malik. Semantic contours from inverse detectors. In *ICCV*, 2011. 4
- [12] K. He, G. Gkioxari, P. Dollár, and R. Girshick. Mask R-CNN. In *ICCV*, 2017. 1, 3, 4
- [13] K. He, X. Zhang, S. Ren, and J. Sun. Deep residual learning for image recognition. In *CVPR*, 2016. 3
- [14] J. Hosang, R. Benenson, P. Dollár, and B. Schiele. What makes for effective detection proposals? *T-PAMI*, 38(4):814–830, 2016. 2
- [15] S. D. Jain and K. Grauman. Click carving: Segmenting objects in video with point clicks. In *HCOMP*, 2016. 2
- [16] B. Jin, M. V. Ortiz Segovia, and S. Susstrunk. Webly supervised semantic segmentation. In *CVPR*, 2017. 2
- [17] A. Khoreva, R. Benenson, J. Hosang, M. Hein, and B. Schiele. Simple does it: Weakly supervised instance and semantic segmentation. In *CVPR*, 2017. 2
- [18] V. Lempitsky, P. Kohli, C. Rother, and T. Sharp. Image segmentation with a bounding box prior. In *ICCV*, pages 277–284, 2009. 2, 6
- [19] X. Li, Y. Qi, Z. Wang, K. Chen, Z. Liu, J. Shi, P. Luo, C. C. Loy, and X. Tang. Video object segmentation with re-identification. *The 2017 DAVIS Challenge on Video Object Segmentation - CVPR Workshops*, 2017. 5
- [20] D. Lin, J. Dai, J. Jia, K. He, and J. Sun. Scribble-sup: Scribble-supervised convolutional networks for semantic segmentation. In *CVPR*, 2016. 1, 2
- [21] 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 *ECCV*, 2014. 4, 7
- [22] J. Long, E. Shelhamer, and T. Darrell. Fully convolutional networks for semantic segmentation. In *CVPR*, 2015. 3
- [23] K. Maninis, J. Pont-Tuset, P. Arbeláez, and L. V. Gool. Convolutional oriented boundaries: From image segmentation to high-level tasks. *T-PAMI*, 2017. 3
- [24] D. P. Papadopoulos, J. R. Uijlings, F. Keller, and V. Ferrari. Extreme clicking for efficient object annotation. In *ICCV*, 2017. 2, 3, 4, 5, 6, 9
- [25] D. P. Papadopoulos, J. R. Uijlings, F. Keller, and V. Ferrari. Training object class detectors with click supervision. In *CVPR*, 2017. 2
- [26] D. Pathak, E. Shelhamer, J. Long, and T. Darrell. Fully convolutional multi-class multiple instance learning. In *ICLR*, 2015. 2
- [27] F. Perazzi, J. Pont-Tuset, B. McWilliams, L. Van Gool, M. Gross, and A. Sorkine-Hornung. A benchmark dataset and evaluation methodology for video object segmentation. In *CVPR*, 2016. 4, 5, 7
- [28] P. O. Pinheiro, R. Collobert, and P. Dollár. Learning to segment object candidates. In *NIPS*, 2015. 2
- [29] P. O. Pinheiro, T.-Y. Lin, R. Collobert, and P. Dollár. Learning to refine object segments. In *ECCV*, 2016. 2, 6
- [30] J. Pont-Tuset, P. Arbeláez, J. T. Barron, F. Marques, and J. Malik. Multiscale combinatorial grouping for image segmentation and object proposal generation. *T-PAMI*, 39(1):128–140, 2017. 2
- [31] J. Pont-Tuset, F. Perazzi, S. Caelles, P. Arbeláez, A. Sorkine-Hornung, and L. Van Gool. The 2017 DAVIS challenge on video object segmentation. *arXiv:1704.00675*, 2017. 4, 7, 8
- [32] C. Rother, V. Kolmogorov, and A. Blake. Grabcut: Interactive foreground extraction using iterated graph cuts. In *ACM TOG*, 2004. 2, 4, 6
- [33] O. Russakovsky, J. Deng, H. Su, J. Krause, S. Satheesh, S. Ma, Z. Huang, A. Karpathy, A. Khosla, M. Bernstein, A. C. Berg, and L. Fei-Fei. ImageNet Large Scale Visual Recognition Challenge. *IJCV*, 2015. 1, 6
- [34] A. Shrivastava, A. Gupta, and R. Girshick. Training region-based object detectors with online hard example mining. In *CVPR*, 2016. 8
- [35] H. Su, J. Deng, and L. Fei-Fei. Crowdsourcing annotations for visual object detection. In *AAAI Workshops*, 2012. 2
- [36] M. Tang, I. Ben Ayed, D. Marin, and Y. Boykov. Secrets of grabcut and kernel k-means. In *ICCV*, pages 1555–1563, 2015. 6

- [37] M. Tang, L. Gorelick, O. Veksler, and Y. Boykov. Grabcut in one cut. In *ICCV*, pages 1769–1776, 2013. 6
- [38] J. Wu, Y. Zhao, J.-Y. Zhu, S. Luo, and Z. Tu. MILcut: A sweeping line multiple instance learning paradigm for interactive image segmentation. In *CVPR*, 2014. 2, 6
- [39] S. Xie and Z. Tu. Holistically-nested edge detection. *IJCV*, pages 1–16, 2017. 3, 5
- [40] N. Xu, B. Price, S. Cohen, J. Yang, and T. Huang. Deep grabcut for object selection. In *BMVC*, 2017. 2, 4, 5, 6
- [41] N. Xu, B. Price, S. Cohen, J. Yang, and T. S. Huang. Deep interactive object selection. In *CVPR*, 2016. 2, 4, 5, 8
- [42] H. Zhao, J. Shi, X. Qi, X. Wang, and J. Jia. Pyramid scene parsing network. In *CVPR*, 2017. 1, 3, 5, 7