

Using a Supervised Method Without Supervision for Foreground Segmentation

Levi Kassel

The Hebrew University of Jerusalem
Jerusalem, Israel

levi.kassel@mail.huji.ac.il

Michael Werman

The Hebrew University of Jerusalem
Jerusalem, Israel

michael.werman@mail.huji.ac.il

Abstract

Neural networks are a powerful framework for foreground segmentation in video, robustly segmenting moving objects from the background in various challenging scenarios. The premier methods are those based on supervision requiring a final training stage on a database of tens to hundreds of manually segmented images from the specific static camera.

In this work, we propose a method to automatically create a scene specific "artificial" database that is sufficient for training the supervised methods so that it performs better than current unsupervised methods. It is based on combining an unsupervised foreground segmenter to extract objects from the training images and randomly inserting these objects back in their correct location into a background image.

Test results are shown on test sequences in CDnet.¹

1. Introduction

Foreground segmentation/background subtraction is one of the central tasks in the field of computer vision due to its numerous applications including surveillance of human activities in public spaces, traffic monitoring, and industrial machine vision.

Deep learning has become the principal methodology in foreground segmentation where each pixel in an image is classified as foreground or background.

One of the great difficulties with supervised learning, especially in the field of foreground segmentation, is the labeling of data that usually requires extensive human hands-on work.

One of the known methods designed to eliminate the need to create data manually is to use synthetic data. The problem with synthetic data in computer vision is the ability to create data that is similar enough to the real data so that the learned system can be generalized to the original task

with the real data.

This paper proposes substituting the hand-labeled data needed for the premier (supervised) method with an automatically created "artificial" database. It is based on an unsupervised weak foreground segmenter that extracts "good" objects from the training images and then to create the fine tuning data we randomly chose from these objects and insert them in their original position into a background image. It is especially pertinent to static cameras as the objects found are automatically in the "right" location, being of the "right" size, color and shape, and in the "right" lighting conditions. In this way, we create an "scene specific" system in a unsupervised way. This is applicable to a gigantic number of cameras, as cited in CNBC ², "One billion surveillance cameras will be watching around the world in 2021".

2. Related work

Deep learning methods are currently the premier methods in video foreground segmentation.

Many foreground segmentation methods work with one image at a time without taking into consideration temporal aspects between the frames. For example Lin et al. [10] use the Fully Convolutional Semantic Networks (FCSN) [6] for foreground segmentation. The authors concatenate the current frame with the background image channel-wise and feed the 6-channel image to a FCN. Although their model achieves good results in simple videos, it fails in "noisy" environments and in detecting small objects.

In 2019, Lim and Keles presented FgSegNetv2[5], a modified VGG 16 network is used as an encoder for the network, obtaining high-resolution feature maps, which are input for the Feature Pooling Module (FPM) and consequently as input for the decoder, working with two Global Average Pooling (GAP) modules. This method used up to 200 annotated training frames of each scene and achieves state of the art performance on the CDnet 2014 dataset[14]. Note that the performance of FgSegNet v2 drops dramati-

¹This research was supported by the DFG.

²<https://www.cnbc.com/2019/12/06/one-billion-surveillance-cameras-will-be-watching-globally-in-2021.html>

BSUV-Net v2						
Category	25f	50f	100f	200f	500f	1000f
baseline	0.935	0.955	0.980	0.984	0.985	0.983
shadow	0.915	0.939	0.949	0.960	0.973	0.970
dynamic	0.900	0.931	0.952	0.956	0.954	0.950

Table 1. F-measure as a function of the number of images in augmented training set constructed from BSUV-Net v2 objects

cally when applied to unseen videos as shown by Tezcan et al. [13].

Tezcan et al. [11] presented the BSUV-Net 2.0 that mainly focused on data augmentation to enhance jitter, pan-tilt and intermittently-static objects especially trying to account for temporal variations that naturally occur in video achieving state of the art results in the unsupervised case for various scenarios.

Synthetic data has been employed for data augmentation. For example, Richter et al.[9] and Ros et al. [3] develop a virtual reality tool from the world of computer games for this purpose. A big advantage of this method is that you can produce endless samples of tagged data and you can create situations as you wish.

Dwivedi et al.[4] used the power of synthetic data for instance detection. They used a 'cut paste and learn' paradigm to cut objects and paste them on various background images to create a dataset for instance detection. They used Poisson blending and Gaussian smoothing on the pasted image to avoid boundary artifacts.

In order to benefit from using synthetic data it is necessary to use foreground elements of the right categories; humans, vehicles, etc., with the the correct style of the elements of the foreground in a specific video. And of course, the synthetic objects should be placed in the correct parts of the image with the right size and projection parameters. This entails a serious intervention for every scene and still there is a big difference between the appearance of synthetic data and reality.

These decisions and interventions are bypassed in our method because the foreground elements are automatically extracted from the real images in the video sequence. Therefore, our method is much more specific and appropriate to the environment in which the videos are shot.

3. Building the "supervised" training dataset

The supervised foreground segmentation methods require a "supervised" labelled set of images. Here we describe how we automatically generate a video sequence dependent, training set which is sufficient for the training of supervised based background subtraction architectures.

We use the foreground segmentation results of an unsupervised method to extract objects from the video. Subsets of these extracted objects are randomly chosen and inserted into a background image to produce a training set.

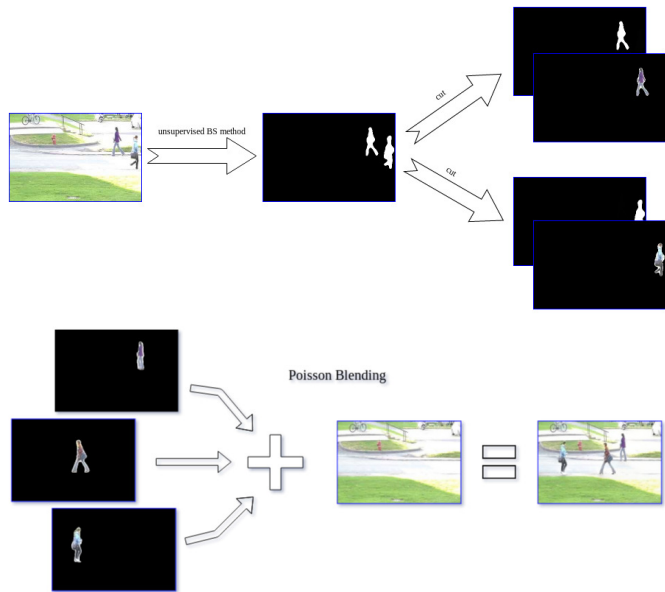


Figure 1. (top) construction of the foreground object database. (bottom) annotated image generation process




Figure 2. Examples from our 'artificial' database based on videos from the CDnet2014 dataset, (top) generated images with the foreground objects, and (bottom) corresponding foreground masks.

As these objects are spatially positioned exactly where they were found, cars are on the right roads, pedestrians on the traversed paths, and everything is where it really should be and with the correct projection and lighting, without having to analyze the scenario.

3.1. Choosing the objects

As we mentioned before, the foreground objects are extracted with unsupervised methods in a 'cut & paste' fashion such that a foreground object and its mask are extracted to create a foreground database, see Fig 1. In this work we use the state of the art unsupervised method on the CDnet dataset [14], BSUV-Net v2 [12].

3.2. Building the dataset

The construction of the database is as follows:

Given the objects \mathcal{O} , from the previous step and a background image \mathcal{I} , randomly choose objects from \mathcal{O} and paste them in their original spatial locations onto \mathcal{I} , see Fig 1. To



Figure 3. (top) randomly placed objects in 'highway', (bottom) objects of 'PETS2006' in 'pedestrian' scene

avoid boundary artifacts, Poisson editing [8], as Dwibedi et al. suggested, was used. It turned out that a Gaussian smoothing [4] step was counterproductive. These images comprise the dataset required for the supervised method, examples are in Fig 2.

4. Experiments

We tested our method with FgSegNet V2 by Lim et al. [5], which is currently the top-ranked method in the Change Detection 2014 Challenge[14], SBI2015[7], and UCSD Background Subtraction[1]. FgSegNet V2 requires a final training stage using 25-200 labeled images from the

Category	baseline	shadow	dynamic
BSUV-Net v2			
Reference	0.962	0.956	0.905
FgSegNet V2			
BSUV	0.964	0.950	0.946
Objects from different scene & random location	0.739	0.673	0.648
Same scene objects & random location	0.923	0.893	0.854
Same scene objects & specific location	0.984	0.960	0.956

Table 2. Different object selection and placement strategies, The BSUV results are followed by using FgSegNet V2 trained using four different options, the annotated frames from BSUV, pasting objects from other scenes in random locations, pasting objects from BSUV on the same video in random locations, and pasting objects from BSUV on the same video in their correct locations

video being evaluated and is thus a supervised method. Our contribution is to use automatic object detection on the images used for fine-tuning converting the method to being completely unsupervised but still scenario dependant.

Our method is especially relevant for static cameras so we tested three of the datasets from CDnet: baseline, shadow, and dynamic background. The background images are pixel-wise medians of a sequence of 50 frames. The frames used to extract objects were the same 200 used by FgSegNet V2[5].

The average F-measure dependence on the number of augmented frames of the training set is shown in Table 1. Surprisingly more images is not always better, maybe due to inserting too many false positives.

To understand the power of our copy and correct location paste method in a static camera environment, we carried out two more experiments. In one we used the objects that were taken from the same scene but placed in random positions and thus do not take into account the natural location where they should be and in the other we pasted objects from other scenes, see Fig 3 for an illustration. In both cases there was a decrease in the quality of the results, see Table 2.

To see that our method works with different unsupervised methods we carried out the same test based on the semanticBGS[2] algorithm and saw similar improvements, see Table 3.

SemanticBGS						
Category	25f	50f	100f	200f	500f	1000f
baseline	0.943	0.953	0.976	0.985	0.985	0.980
shadow	0.917	0.943	0.959	0.963	0.970	0.968
dynamic	0.895	0.933	0.952	0.953	0.942	0.906

Table 3. F-measure as a function of the number of images in augmented training set constructed from SemanticBGS objects

5. Summary

We showed how a supervised foreground segmentation algorithm can be made unsupervised by replacing the supervision with augmented frames based on objects extracted with an unsupervised algorithm. This leads to a state of the art unsupervised foreground segmentation.

References

- [1] Ucsd anomaly detection dataset. In [//www.svcl.ucsd.edu/projects/anomaly/dataset.htm](http://www.svcl.ucsd.edu/projects/anomaly/dataset.htm), 2010. 3
- [2] Marc Braham, Sébastien Piérard, and Marc Van Droogenbroeck. Semantic background subtraction. In *2017 IEEE International Conference on Image Processing (ICIP)*, pages 4552–4556. IEEE, 2017. 3
- [3] Alexey Dosovitskiy, German Ros, Felipe Codevilla, Antonio Lopez, and Vladlen Koltun. Carla: An open urban driving simulator. *arXiv preprint arXiv:1711.03938*, 2017. 2
- [4] Debidatta Dwibedi, Ishan Misra, and Martial Hebert. Cut, paste and learn: Surprisingly easy synthesis for instance detection. In *Proceedings of the IEEE International Conference on Computer Vision*, pages 1301–1310, 2017. 2, 3
- [5] Long Ang Lim and Hacer Yalim Keles. Learning multi-scale features for foreground segmentation. *Pattern Analysis and Applications*, 23(3):1369–1380, 2020. 1, 3
- [6] Jonathan Long, Evan Shelhamer, and Trevor Darrell. Fully convolutional networks for semantic segmentation. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 3431–3440, 2015. 1
- [7] Lucia Maddalena and Alfredo Petrosino. Towards benchmarking scene background initialization. In *International conference on image analysis and processing*, pages 469–476. Springer, 2015. 3
- [8] Patrick Pérez, Michel Gangnet, and Andrew Blake. Poisson image editing. In *ACM SIGGRAPH 2003 Papers*, pages 313–318. 2003. 3
- [9] Stephan R Richter, Vibhav Vineet, Stefan Roth, and Vladlen Koltun. Playing for data: Ground truth from computer games. In *European conference on computer vision*, pages 102–118. Springer, 2016. 2
- [10] Dimitrios Sakkos, Heng Liu, Jungong Han, and Ling Shao. End-to-end video background subtraction with 3d convolutional neural networks. *Multimedia Tools and Applications*, 77(17):23023–23041, 2018. 1
- [11] M. Ozan Tezcan, Prakash Ishwar, and Janusz Konrad. Bsuv-net 2.0: Spatio-temporal data augmentations for video-agnostic supervised background subtraction. *IEEE Access*, 9:53849–53860, 2021. 2
- [12] M Ozan Tezcan, Prakash Ishwar, and Janusz Konrad. Bsuv-net 2.0: Spatio-temporal data augmentations for video-agnostic supervised background subtraction. *IEEE Access*, 9:53849–53860, 2021. 2
- [13] Ozan Tezcan, Prakash Ishwar, and Janusz Konrad. Bsuv-net: A fully-convolutional neural network for background subtraction of unseen videos. In *Proceedings of the IEEE/CVF Winter Conference on Applications of Computer Vision*, pages 2774–2783, 2020. 2
- [14] Yi Wang, Pierre-Marc Jodoin, Fatih Porikli, Janusz Konrad, Yannick Benezeth, and Prakash Ishwar. Cdnet 2014: An expanded change detection benchmark dataset. In *Proceedings of the IEEE conference on computer vision and pattern recognition workshops*, pages 387–394, 2014. 1, 2, 3