

CNN BASED BACKDROP PURGING.

Ashrith Appani
ashrithappani4@gmail.com

Abstract

Backdrop Purging is a common pre-processing step in computer vision and video processing for object tracking, people recognition, and other tasks. Several successful background-subtraction algorithms have recently been proposed, however nearly all of the best-performing ones are supervised. The availability of some annotated frames of the test video during training is critical to their performance. As a result, there is no literature on their performance on completely "unseen" videos. We provide a new supervised background-subtraction technique for unseen films (BSUV-Net) based on a fully-convolutional neural network in this paper. The current frame and two background frames collected at various time scales, along with their semantic segmentation maps, are fed into our network. We also offer a new data-augmentation strategy that mitigates the influence of illumination differences between the background frames and the current frame in order to limit the risk of overfitting. In terms of F-measure, recall, and precision, BSUV-Net beats state-of-the-art algorithms assessed on unseen videos on the CDNet-2014 dataset.

1. Introduction

Backdrop Purging is a foundational, lowlevel task in computer vision and video processing. The aim of BACKDROP PURGING is to segment an input video frame into regions corresponding to either foreground or background. It is frequently used as a pre-processing step for higher-level tasks such as object tracking, people and motor-vehicle recognition, human activity recognition, etc. Since BACKDROP PURGING is often the first pre-processing step, the accuracy of its output has an overwhelming impact on the overall performance of subsequent steps. Therefore, it is critical that BACKDROP PURGING produce as accurate a foreground/background segmentation as possible.

Traditional BACKDROP PURGING algorithms are unsupervised and rely on a background model to predict foreground regions. PAWCS, SWCD and WisenetMD are considered to be state-of-the-art unsupervised BACKDROP PURGING algorithms. However, since they rely on the accuracy of the background model, they encounter difficulties when

applied to complex scenes. Recently, ensemble methods and a method leveraging semantic segmentation have been proposed which significantly outperform traditional algorithms.

The success of deep learning in computer vision did not bypass BACKDROP PURGING research. A number of *supervised* deeplearning BACKDROP PURGING algorithms have been developed with performance easily surpassing that of traditional methods. However, most of these algorithms have been tuned to either one specific video or to a group of similar videos, and their performance on unseen videos has not been evaluated. For example, FgSegNet uses 200 frames from a test video for training and the remaining frames from the *same* video for evaluation. If applied to an unseen video, its performance drops significantly.

In this paper, we introduce *Backdrop Purging for Unseen Videos* (BSUV-Net), a fully-convolutional neural network for predicting foreground of an *unseen* video. A key feature of our approach is that the training and test sets are composed of frames originating from different videos. This guarantees that no ground-truth data from the test videos have been shown to the network in the training phase. By employing two reference backgrounds at different time scales, BSUV-Net addresses two challenges often encountered in BACKDROP PURGING: varying scene illumination and intermittently-static objects that tend to get absorbed into the background. We also propose novel data augmentation which further improves our method's performance under varying illumination. Furthermore, motivated by recent work on the use of semantic segmentation in BACKDROP PURGING, we improve our method's accuracy by inputting semantic information along with the reference backgrounds and current frame. The main contributions of our work are as follows:

1. **Supervised BACKDROP PURGING for Unseen Videos:** Although supervised algorithms, especially neural networks, have significantly improved BACKDROP PURGING performance, they are tuned to a specific video and thus their performance on *unseen* videos deteriorates dramatically. To the best of our knowledge, BSUV-Net is the first supervised BACKDROP PURGING algorithm that is truly generalizable to *unseen* videos.
2. **Data Augmentation for Increased Resilience to Varying Illumination:** Changes in scene illumination pose a major challenge to BACKDROP PURGING algorithms. To mitigate this, we develop a simple, yet effective, data augmentation technique. Using a simple additive model we vary the illumination of the current frame and the reference background frames that are fed into BSUV-Net during training. This enables us to effectively tackle various illumination change scenarios that may be present in test videos.
3. **Leveraging Semantic and Multiple Time-Scale Information:** BSUV-Net improves foreground-boundary segmentation accuracy by accepting semantic information as one of its inputs. This is unlike in an earlier BACKDROP PURGING method which used semantic information as a *post-processing* step. The other network inputs are the current frame (to be segmented) and a two-frame background model with data from different time scales. While one background frame, based on *distant* history, helps with the discovery of intermittently-static objects, the other frame, based on *recent* history, is key for handling dynamic factors such as illumination changes.

Based on our extensive experiments on the CDNet-2014 dataset, BSUV-Net outperforms state-of-the-art BACKDROP PURGING algorithms evaluated on *unseen* videos

2. Related Work

A wide range of BACKDROP PURGING algorithms have been developed in the past, each having some advantages and disadvantages over others. Since this is not a survey paper, we will not cover all BACKDROP PURGING variants. Instead, we will focus only on recent top-performing methods. We divide these algorithms into 3 categories: (i) BACKDROP PURGING by (unsupervised) background modeling, (ii) supervised BACKDROP PURGING tuned to a single video or a group of videos, (iii) Improving BACKDROP PURGING algorithms by post-processing.

2.1. BACKDROP PURGING by Background Modeling

Nearly all traditional BACKDROP PURGING algorithms first compute a background model, and then use it to predict the foreground. While a simple model based on the mean or median of a subset of preceding frames offers only a single background value per pixel, a probabilistic Gaussian Mixture Model (GMM) allows a range of background values. This idea was improved by creating an online procedure for the update of GMM parameters in a pixel-wise manner. Kernel Density Estimation (KDE) was introduced into BACKDROP PURGING as a non-parametric alternative to GMMs and was subsequently improved. The probabilistic methods achieve better performance compared to single-value models for dynamic scenes and scenes with small background changes.

In, Barnich and Droogenbroeck introduced a sample-based background model. Instead of implementing a probabilistic model, they modeled the background by a set of sample values per pixel and used a distance-based model to decide whether a pixel should be classified as background or foreground. Since color information alone is not sufficient for complex cases, such as illumination changes, Bilodeau *et al.* introduced Local Binary Similarity Patterns (LBSP) to compare the current frame and background using spatiotemporal features instead of color. St-Charles *et al.* combined color and texture information, and introduced a word-based approach, PAWCS. They considered pixels as background words and updated each word's reliability by its persistence. Similarly, SuBSENSE by St-Charles *et al.* combines LBSP and color features, and employs pixel-level feedback to improve the background model.

Recently, Isik *et al.* introduced SWCD, a pixel-wise, sliding-window approach leveraging a dynamic control system to update the background model, while Lee *et al.* introduced WisenetMD, a multi-step algorithm to eliminate false positives in dynamic backgrounds. In, Sultana *et al.* introduced an unsupervised background estimation method based on a generative adversarial network (GAN). They use optical flow to create a motion mask and then in-paint covered regions with background values estimated by a GAN. The foreground is then computed by subtracting the estimated background from the current frame followed by morphological operations.

They, however, do not achieve state-of-the-art results. Zeng *et al.* introduced RTSS which uses deep learning-based semantic segmentation predictions to improve the background model used inSubSENSE.

2.2. Supervised BACKDROP PURGING

Although Backdrop Purging has been extensively studied in the past, the definition of a supervised BACKDROP PURGING algorithm is still vague. Generally speaking, the aim of a supervised BACKDROP PURGING algorithm is to learn the parameters (e.g., neuralnetwork weights) of a complex function in order to minimize a loss function of the labeled training frames. Then, the performance of the algorithm is evaluated on a separate set of test frames. In this section we divide the supervised BACKDROP PURGING algorithms into three groups namely, *video-optimized*, *videogroup-optimized* and *video-agnostic* depending on which frames and videos they use during training and testing.

Several algorithms use some frames from a test video for training and all the frames of the *same* video for evaluating performance on that video. In such algorithms, parameter values are optimized separately for each video. We will refer to this class of algorithms as *video-optimized* BACKDROP PURGING algorithms. In another family of algorithms, randomly-selected frames from a *group* of test videos are used for training and all the frames of the same videos are used for testing. Since some frames from *all* test videos are used for training, we will refer this class of algorithms as *video-groupoptimized* algorithms. Note that, in both of these scenarios the algorithms are neither optimized for nor evaluated on *unseen* videos and to the best of our knowledge all of the top-performing supervised BACKDROP PURGING algorithms to-date are either *video-optimized* or *video-group-optimized*. In this paper, we introduce a new category of supervised BACKDROP PURGING algorithms, called *video-agnostic* algorithms, that can be applied to unseen videos with no or little loss of performance. To learn parameters, a *video-agnostic* algorithm uses frames from a set of training videos but for performance evaluation it uses a completely different set of videos.

In recent years, supervised learning algorithms based on convolutional neural networks (CNNs) have been widely applied to BACKDROP PURGING. The first CNN-based BACKDROP PURGING algorithm was introduced in. This is a *video-optimized* algorithm which produces a single foreground probability for the center of each 27×27 patch of pixels. A method proposed in uses a similar approach, but with a modified CNN which operates on patches of size 31×31 pixels.

Instead of using a patch-wise algorithm, Zeng and Zhu introduced the Multiscale Fully-Convolutional Neural Network (MFCN) which can predict the foreground of the entire input image frame in one step. Lim and Keles proposed a triplet CNN which uses siamese networks to create features at three different scales and combines these features within a transposed CNN. In a follow-up work, they removed the triplet networks and used dilated convolutions to capture the multiscale information. In, Bakkay *et al.* used generative adversarial networks for BACKDROP PURGING. The generator performs the BACKDROP PURGING task, whereas the discriminator tries to classify the BACKDROP PURGING map as real or fake. Although all these algorithms perform very well on various BACKDROP PURGING datasets, it is important to note that they are all *video-optimized*, thus they will suffer a performance loss when tested on unseen videos. In, Babae *et al.* designed a *video-group-optimized* CNN for BACKDROP PURGING. They randomly selected 5% of CDNet-2014 frames as a training set and developed a single network for all of the videos in this dataset. In, Sakkos *et al.* used a 3DCNN to capture the temporal information in addition to the color information. Similarly to, they trained a single algorithm using 70% of frames in CDNet-2014 and then used it to predict the foreground in all videos of the dataset. Note that even these approaches do not generalize to other videos since some ground truth data from each video exists in the training set. Table 1 compares and summarizes the landscape of supervised BACKDROP PURGING algorithms and the methodology used for training and evaluation.

As discussed above, none of the CNN-based BACKDROP PURGING algorithms to-date have been designed for or tested on unseen videos with no ground truth at all. This limits their practical utility since it is not possible to label some frames in *every* new video. Since the publication of our preprint, we have learned about a recent BACKDROP PURGING algorithm named 3DFR, which uses 3D spatiotemporal convolution blocks in an encoder-decoder architecture to predict background in an unseen video. However, only reports evaluation results on 10 out of the 53 videos of CDNet2014.

2.3. Improving BACKDROP PURGING Algorithms by Post-Processing

Over the last few years, many deep-learning-based algorithms were developed for the problem of semantic segmentation and they achieved state-of-the-art performance. In, Braham and Droogenbroeck introduced a post-processing step for BACKDROP PURGING algorithms based on semantic segmentation predictions. Given an input frame, they predicted a segmentation map using PSPNet and obtained pixel-wise probability predictions for semantic labels such as person, car, animal, house etc. Then, they manually grouped these labels into two sets foreground and background labels, and used this information to improve any BACKDROP PURGING algorithm’s output in a post-processing step. They obtained very competitive results by using SubSENSE as the BACKDROP PURGING algorithm.

Bianco *et al.* introduced an algorithm called IUTIS which combines the results produced by several BACKDROP PURGING algorithms. They used genetic programming to determine how to combine several BACKDROP PURGING algorithms’ outputs using a sequence of basic binary operations, such as logical “and/or”, majority voting and median filtering. Their best result was achieved by using 5 top-performing BACKDROP PURGING algorithms on the CDNet2014 dataset at the time of publication. Zeng *et al.* followed the same idea, but instead of genetic programming, used a fully-convolutional neural network to fuse several BACKDROP PURGING results into a single output, and outperformed IUTIS on CDNet-2014.

3. Proposed Algorithm: BSUV-Net

3.1. Inputs to BSUV-Net

Segmenting an unseen video frame into foreground and background regions without using any information about the background would be an ill-defined problem. In BSUV-Net, we use two reference frames to characterize the background. One frame is an “empty” background frame, with no people or other objects of interest, which can typically be extracted from the beginning of a video e.g., *via* median filtering over a large number of *initial* frames. This provides an accurate reference that is very helpful for segmenting intermittently static objects in the foreground. However, due to dynamic factors, such as illumination variations, this reference may

Table 1. Training/evaluation methodologies of supervised BACKDROP PURGING algorithms on CDNet-2014.

Algorithm	Are some frames from test videos used in training?		Training and evaluation methodology
Braham-CNN-BACKDROP PURGING	Yes	First half of the labeled frames of the test video	<i>video-optimized</i>
MFCNN	Yes	Randomly selected 200 frames from the first 3000 labeled frames of the test video	<i>video-optimized</i>
Wang-CNN-BACKDROP PURGING FGSegNet BScGAN	Yes	Hand picked 200 labeled frames of the test video	<i>video-optimized</i>
Babae-CNN-BACKDROP PURGING	Yes	5% of the labeled frames of all videos	<i>video-group-optimized</i>
3D-CNN-BACKDROP PURGING	Yes	70% of the labeled frames of all videos	<i>video-group-optimized</i>
BSUV-Net (proposed)	No	No frame from test videos is used in training	<i>video-agnostic</i>

not be valid after some time. To counteract this, we use another reference frame that characterizes *recent* background, for example by computing median of 100 frames preceding the frame being processed. However, this frame might not be as accurate as the first reference frame since we cannot guarantee that there will be no foreground objects in it (if such objects are present for less than 50 frames, the temporal median will suppress them). By using two reference frames captured at different time scales, we aim to leverage benefits of each frame type.

Braham *et al* have shown that leveraging results of semantic segmentation significantly improves the performance of a BACKDROP PURGING algorithm, for example by using semantic segmentation results in a post-processing step. In BSUV-Net, we follow a different idea and use semantic information as an additional input channel to our neural network. In this way, we let our network learn how to use this information. To extract semantic segmentation information, we used a state-of-the-art CNN called DeepLabv3 trained on ADE20K, an extensive semantic-segmentation dataset with 150 different class labels and more than 20,000 images with dense annotations. Let us denote the set of object classes in ADE20K as $C = \{c_0, c_1, \dots, c_{149}\}$. Following the same procedure as in, we divided these classes into two sets: foreground and background objects. As foreground objects, we used person, car, cushion, box, book, boat, bus, truck, bottle, van, bag and bicycle. The rest of the classes are used as background objects. The softmax layer of DeepLabv3 provides pixel-wise class probabilities p_{c_j} for $c_j \in C$. Let $I[m, n]$ be an input frame at spatial location m, n and

$I[m, n]$. We compute a foreground probability map (FPM)

$$S[m, n] = \sum_{c_j \in F} p_{c_j} I[m, n], \text{ where } F \text{ is the set of fore-}$$

We use the current, recent and empty frames in color, each along with its FPM, as the input to BSUV-Net. Clearly, the number of channels in BSUV-Net's input layer loss function, defined as follows: is 12 for each frame consists of 4 channels (R,G,B,FPM).

3.2. Network Architecture and Loss Function

We use a UNET-type [19] fully-convolutional neural network (FCNN) with residual connections. The architecture of where $Y \in \{0, 1\}^{w \times h}$ is the ground truth of X , T is a smoothing parameter and m, n are the spatial locations.

Let $\{p_c, [m, n]\}^{149}$ be the predicted probability distribution of BSUV-Net has two parts: encoder and decoder. In the encoder network, we use 2×2 maxpooling operators to decrease the spatial dimensions. In the decoder network, we use up-convolutional layers (transposed convolution with a stride of 2) to increase the dimensions back to those of the input. In all convolutional and upconvolutional layers, we use 3×3 convolutions as in VGG. The residual connections from the encoder to the decoder help the network combine low-level visual information gained in the initial layers with high-level visual information gained in the deeper layers. Since our aim is to increase the performance on unseen videos, we use strong batch normalization (BN) and spatial dropout (SD) layers to increase the generalization capacity. Specifically, we use a BN layer after each convolutional and up-convolutional layer, and an SD layer before each max-pooling layer. Since our task can be viewed as a binary segmentation, we use a sigmoid layer as the last layer in BSUV-Net. The operation of the overall network can be defined as a nonlinear map

$G(W) : X \rightarrow Y$ where $X \in \mathbb{R}^{w \times h \times 12}$ is a 12-channel input, w and h are its spatial dimensions, W represents the parameters of neural network G , and $Y \in [0, 1]^{w \times h}$ is a pixel-wise foreground probability prediction. Note that

since this is a fully-convolutional neural network, it does not require fixed input size; any frame size can be used, but some padding may be needed to account for max-pooling operations.

In most BACKDROP PURGING datasets, the number of background pixels is much larger than the number of foreground pixels. This class imbalance creates significant problems for the commonly used loss functions, such as cross-entropy and mean-squared error. A good alternative for unbalanced binary datasets is

the Jaccard index. Since the network output is a probability map, we opted for a relaxed form of the Jaccard index as the

$$T + \mathbf{m}, (Y[m, n]Y[m, n])$$

$$JR(Y, Y) = \frac{Y[m, n] + Y[m, n] - Y[m, n]Y[m, n]}{T + \mathbf{m}, n}$$

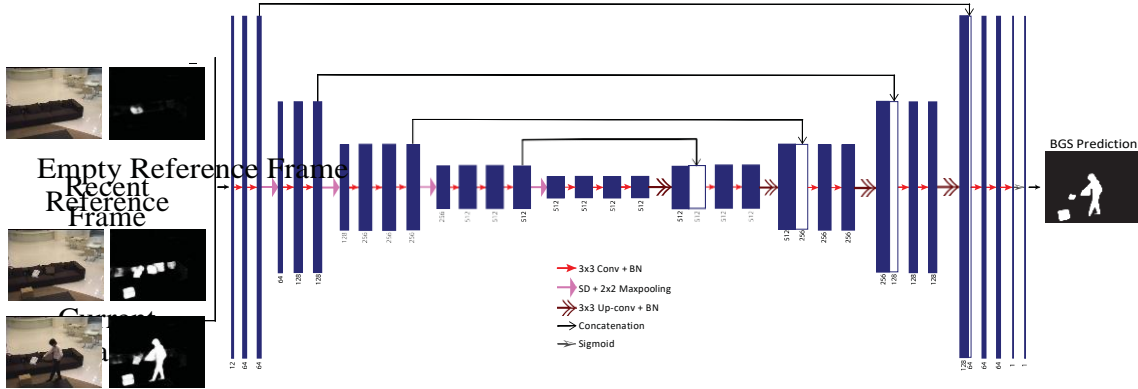


Figure 1. Network architecture of BSUV-Net. BN stands for batch normalization and SD stands for spatial dropout. Grayscale images at the network input show foreground probability maps (FPM) of the corresponding RGB frames.

3.3. Resilience to Illumination Change by Data Augmentation

Since neural networks have millions of parameters, they are very prone to overfitting. A widely-used method for reducing overfitting in computer-vision problems is to enlarge the dataset by applying several data augmentations such as random crops, rotations and noise addition. Since we are dealing with videos in this paper, we can also add augmentation in the temporal domain.

In real-life BACKDROP PURGING problems, there might be a significant illumination difference between an empty background frame acquired at an earlier time and the current frame. However, only a small portion of videos in CDnet-2014 capture significant illumination changes which limits BSUVNet’s generalization performance. Therefore, we introduce a new data-augmentation technique to account for global

illumination changes between the empty reference frame and the current frame. Suppose that $\mathbf{R}_E \in \mathbb{R}^{w,h,3}$ represents the RGB channels of an empty reference frame. Then, an augmented version of \mathbf{R}_E can be computed as $\mathbf{R}_E[m, n, c] = \mathbf{R}_E[m, n, c] + d[c]$ for $c = 1, 2, 3$, where $d \in \mathbb{R}^3$ represents a frame-specific global RGB change in our illumination model. By choosing d randomly for each

example during training (see Section 4.2 for details), we can make the network resilient to illumination variations.

considered hard-shadow pixels as background. Our treatment of hard shadows is consistent with what is done in CDNet-2014 for the change-detection task.

In CDNet-2014, the author propose seven binary performance metrics to cover a wide range of BACKDROP PURGING cases: recall (Re), specificity (Sp), false positive rate (FPR), false negative rate (FNR), percentage of wrong classification (PWC), precision (Pr) and F-measure (F_1). They also introduced two ranking-based metrics namely “average ranking” (R) and “average ranking across categories” (R_{cat}) which combine all 7 metrics into ranking scores. The details of these rankings can be found at “changedetection.net”.

4.2. Training and Evaluation Details

As discussed in Section 2.2, we use a *video-agnostic* evaluation methodology in all experiments. This allows us to measure an algorithm’s performance on real-world-like tasks when no ground-truth labels are available. To evaluate BSUV-Net performance on all videos in CDNet-2014, we used 18 different combinations of training/test video sets. The splits are structured in such a manner that every video appears in the test set of exactly one split, but when it does so, it does not appear in the training set for that split. Detailed information about these sets can be found in the supplementary material.

Let us denote the m -th combination, by $(V_{train}^m, V_{test}^m)$. Then, $\cup_{m=1}^{18} V_{test}^m$ is equal to the set

4. Experimental Results

4.1. Dataset and Evaluation Metrics

In order to evaluate the performance of BSUV-Net, we used CDNet-2014, the largest BACKDROP PURGING dataset with 53 natural videos from 11 categories including challenging scenarios such as shadows, night videos, dynamic background, etc. The spatial resolution of videos varies from 320×240 to 720×526 pixels. Each video has a region of interest labelled as either 1) foreground, 2) background, 3) hard shadow or 4) unknown motion. When measuring an algorithm’s performance, we ignored pixels with unknown motion label and all 53 videos in CDNet-2014. During training, we used 200 frames suggested in for each video in V^m .

When training on different sets V^m , we used *exactly the same* hyperparameter values across all sets to make sure that we are not tuning our network to specific videos. In all of our experiments, we used ADAM optimizer with a learning rate of 10^{-4} , $\beta_1 = 0.9$, and $\beta_2 = 0.99$. The minibatch size was 8 and we trained for 50 epochs. As the empty background frame, we used the median of all foreground-free frames within the first 100 frames. In a few videos containing highway traffic, the first 100 frames did not contain a single foreground-free frame. For these videos, we hand-picked empty frames (e.g., in groups) and used their median as the empty reference. Although this may seem like a limitation, in practice one can randomly sample several hundred frames at the same time of the day across several days (similar illumination) and median filter them to obtain an empty background frame (due to random selection, a moving object is unlikely to occupy the same location in more than 50% of frames). Since there is no single empty background frame in videos from the pantiltzoom (PTZ) category, we slightly changed the inputs. Instead of “empty background + recent background” pair we used “recent background + more recent background” pair, where the recent background is computed as the median of 100 preceding frames and the more recent background is computed as the median of 30 preceding frames.

Although BSUV-Net can accept frames of any spatial dimension, we used a *fixed* size of 224×224 pixels (randomly cropped from the input frame) so as to leverage parallel GPU processing in the training process. We applied random data augmentation at the beginning of each epoch. For illumination resilience, we used the data augmentation method of Section 3.3 with $d[c] = I + I_c$, where I is the same for all c and sampled from $N(0, 0.1^2)$, while I_c is independently sampled from $N(0, 0.04^2)$ for each c . We also added random Gaussian noise from $N(0, 0.01^2)$ to each pixel in each color channel. For pixel values, we used double precision numbers that lie between 0 and 1.

In the evaluation step, we did not apply any scaling or cropping to the inputs. To obtain binary maps, we applied thresholding with threshold $\theta = 0.5$ to the output of the sigmoid layer of BSUV-Net.

4.3. Quantitative Results

Table 2 compares BSUV-Net against state-of-the-art BACKDROP PURGING algorithms in terms of the seven metrics and two rankings discussed in Section 4.1. All quantitative results shown in this paper are computed by “changedetection.net” evaluation servers to reflect the real performance on test data. Since BSUV-Net is *video-agnostic*, comparing it with *video optimized* or *video-group-optimized* algorithms would not be fair and we omit them. Instead, we compare BSUV-Net with state-of-the-art unsupervised algorithms, namely SWCD, WisenetMD and PAWCS, which, by definition, are *video-agnostic*. We exclude RTSS and 3DFR in Table 2 since their results on the test frames are not available on “changedetection.net”. We include the results of IUTIS-5 and SemanticBACKDROP PURGING, but we list them separately because these are *post-processing* algorithms. Note that, both IUTIS-5 and SemanticBACKDROP PURGING can be applied to any BACKDROP PURGING algorithm from Table 2, including BSUV-Net, to improve its performance. To show this, we also report the result of BSUV-Net post-processed by SemanticBACKDROP PURGING. In the *self-contained algorithms* category, we also list FgSegNet v2 since it is currently the best performing algorithm on CDNet-2014. However, since FgSegNet v2’s performance reported on “changedetection.net” has been obtained in a *video-optimized* manner, we trained it anew in a *video-agnostic* manner using the same methodology that we used for BSUV-Net. As expected, this caused a huge performance decrease of FgSegNet v2 compared to its *video-optimized* training. As is clear from Table 2, BSUV-Net outperforms its competitors on almost all of the metrics. The F-measure performance demonstrates that BSUV-Net achieves excellent results without compromising either recall or precision. Table 2 also shows that the performance of BSUV-Net can be improved even further by combining it with SemanticBACKDROP PURGING. The combined algorithm outperforms all of the *video-agnostic*

algorithms that are available on “changedetection.net”.

Table 3 compares the per-category F-measure performance of BSUV-Net against state-of-the-art BACKDROP PURGING algorithms. For RTSS, the values of performance metrics shown in Table 3 are as reported in their paper. Columns 2-12 report the F-measure for each of the 11 categories from “changedetection.net”, while the last column reports the mean F-measure across all categories. Similarly to Table 2, we divided this table into *post-processing* and *self-contained algorithms*. It can be observed that BSUV-Net achieves the best performance in 5 out of 11 categories. It has a striking performance advantage in the “night” category. All videos in this category are traffic-related and many cars have headlights turned on at night which causes significant local illumination variations in time. BSUV-Net’s excellent performance in this category demonstrates that the proposed model is indeed largely illumination-invariant.

BSUV-Net performs poorer than other algorithms in “camera jitter” and “dynamic background” categories. We believe this is related to the empty and recent background frames we are using as input. The median operation used to compute background frames creates very blurry images for these categories since the background is not static. Thus, BSUV-Net predicts some pixels in the background as foreground and increases the number of false positives.

4.4. Visual Results

A visual comparison of BSUV-Net with SWCD and WisenetMD is shown in Figure 2. Each column shows a sample frame from one of the videos in one of the 8 categories. It can be observed that BSUV-Net produces visually the best results for almost all categories.

In the “night” category, SWCD and WisenetMD produce many false positives because of local illumination changes. BSUV-Net produces better results since it is designed to be illumination-invariant. In the “shadow” category, BSUV-Net performs much better in the shadow regions. Results in the “intermittent object motion” and “baseline” categories show that BSUV-Net can successfully detect intermittently-static

Table 2. Comparison of methods for unseen videos from CDNet-2014. For fairness, we separated the post-processing and self-contained algorithms.

Method	<i>R</i>	<i>R_{cat}</i>	<i>Re</i>	<i>Sp</i>	<i>FPR</i>	<i>FNR</i>	<i>PWC</i>	<i>Pr</i>	<i>F₁</i>
<i>Post-processing algorithms</i>									
BSUV-net + SemanticBACKDRO P PURGING*	9.00	14.00	0.817 9	0.9944	0.005 6	0.182 1	1.132 6	0.831 9	0.798 6
IUTIS-5* + SemanticBACKDRO PURGING*	9.43	11.45	0.789 0	0.996 1	0.003 9	0.211 0	1.072 2	0.830 5	0.789 2
IUTIS-5*	11.43	10.36	0.784 9	0.9948	0.005 2	0.215 1	1.198 6	0.808 7	0.771 7
<i>Self-contained algorithms</i>									
BSUV- net	9.29	13.1 8	0.820 3	0.9946	0.005 4	0.179 7	1.140 2	0.811 3	0.786 8
SWCD	15.43	19.0 0	0.7839	0.9930	0.007 0	0.216 1	1.341 4	0.752 7	0.758 3
WisenetM D	16.29	15.1 8	0.8179	0.9904	0.009 6	0.182 1	1.613 6	0.753 5	0.766 8
PAWCS	14.00	15.4 5	0.7718	0.994 9	0.005 1	0.228 2	1.199 2	0.785 7	0.740 3
FgSegNet v2	44.57	44.0 9	0.5119	0.9411	0.058 9	0.488 1	7.350 7	0.485 9	0.371 5

Table 3. Comparison of methods according to the per-category F-measure for unseen videos from CDNet-2014.

Method	Bad weather	Low framerate	Night	PTZ	Thermal	Shadow	Int. obj. motion	Camera jitter	Dynamic background	Baseline	Turbulence	Overall
<i>Post-processing algorithms</i>												
BSUV-net + SemanticBACKDR OP PURGING*	0.8730	0.6788	0.6815	0.6562	0.8455	0.9664	0.7601	0.7788	0.8176	0.9640	0.7631	0.7986
IUTIS-5* + SemanticBACKDR OP PURGING*	0.8260	0.7888	0.5014	0.5673	0.8219	0.9478	0.7878	0.8388	0.9489	0.9604	0.6921	0.7892
IUTIS-5*	0.8248	0.7743	0.5290	0.4282	0.8303	0.9084	0.7296	0.8332	0.8902	0.9567	0.7836	0.7717
<i>Self-contained algorithms</i>												
BSUV-net	0.813	0.6197	0.6987	0.6282	0.8581	0.9233	0.7499	0.7743	0.7967	0.9693	0.7051	0.7868
RTSS	0.8662	0.6771	0.5295	0.5489	0.8510	0.9551	0.7864	0.8396	0.9325	0.9597	0.7630	0.7917
SWCD	0.8233	0.7374	0.5807	0.4545	0.8581	0.8779	0.7092	0.7411	0.8645	0.9214	0.7735	0.7583
Wisenet MD	0.8616	0.6404	0.5701	0.3367	0.8152	0.8984	0.7264	0.8228	0.8376	0.9487	0.8304	0.7535
PAWCS	0.8152	0.6588	0.4152	0.4615	0.8324	0.8913	0.7764	0.8137	0.8938	0.9397	0.6450	0.7403
FgSegNet v2	0.3277	0.2482	0.2800	0.3503	0.6038	0.5295	0.2002	0.4266	0.3634	0.6926	0.0643	0.3715

objects. It is safe to say that BSUV-Net is capable of simultaneously handling the discovery of intermittently-static objects and also the dynamic factors such as illumination change.

An inspection of results in the “dynamic background” category shows that BSUV-Net has detected most of the foreground pixels but failed to detect the background pixels around the foreground objects. We believe this is due to the blurring effect of the median operation that we used in the computation of background frames. Using more advanced background models as an input to BSUV-Net might improve the performance in this category.

4.5. Ablation Study

One of the contributions of BSUV-Net is its multi-channel input composed of two background frames from different time scales and a foreground probability map (FPM). Another contribution is temporal data augmentation tailored to handling illumination changes. In Table 4, we explore their impact on precision, recall and F-measure. Each column on the left represents one characteristic and each row represents a different combination of these characteristics. RGB channels of the current frame are used in all of the combinations. “Empty BG” and “Recent BG” refer to the use of empty and/or recent background frames, respectively, in addition to the current frame. “Data aug.” refers to temporal data augmentation described in Section 3.3 and “FPM” refers to the use of semantic FPM channel in addition to the RGB channels for all input frames. It is clear that all these characteristics have a significant impact on the overall performance. Using only the current frame as input results in very poor metrics. The introduction of empty or/and recent background frames leads to a significant improvement. Adding temporal data augmentation or/and FPM channels further improves the performance and the final network achieves state-of-the-art results.

Table 4. Impact of background frames, data augmentation for temporal illumination change and FPM on BSUV-Net performance.

Current frame	Empty BG	Recent BG	Data aug.	FPM	Pr	Re	F_1
✓					0.361	0.550	0.347
					5	9	6
✓	✓				0.699	0.768	0.681
					4	6	9

✓		✓			0.697	0.706	0.635
					6	4	1
✓	✓	✓			0.765	0.760	0.715
					8	6	6
✓	✓	✓	✓		0.757	0.815	0.744
					4	9	7
✓	✓	✓		✓	0.780	0.774	0.745
					7	7	0
✓	✓	✓	✓	✓	0.811	0.820	0.786
					3	3	8

In this paper, we proposed to add semantic FPM channel as input in order to improve our algorithm’s performance. However, if the background and foreground object categories are chosen carefully, FPM can be used as a BACKDROP PURGING algorithm by itself. This would require prior information about the video (to compute FPM) and, therefore, would not qualify as a *video-agnostic* method. In our algorithm, however, we combine FPM information with RGB channels and background frames. When applying DeepLabv3 to compute




Figure 2. Visual comparison of sample results from BSUV-Net, SWCD and WisenetMD on unseen videos from CDNet-2014.

FPM frames, we pre-defined global background and foreground class categories which might be wrong for some of the videos. To demonstrate that our algorithm is not replicating FPM but leverages its semantic information to boost performance, we compared BSUV-Net with thresholded FPM used as a BACKDROP PURGING result (Table 5). It is clear that FPM alone is not a powerful tool for BACKDROP PURGING as it is significantly outperformed by BSUV-Net.

We manually selected empty background frames from among the initial frames as explained in Section 4.2. In Table 6, we show the impact of this manual process by comparing the manual selection strategy with an automatic one, that is using the median of all frames in the test video as a the *empty background frame*. Clearly, the manual selection slightly improves

precision while significantly decreasing recall. We believe this is due to the increase of false negatives caused by the appearance of some of the foreground objects in the empty background. Since videos in CDNet2014 are rather short (at most 10 minutes), in some cases the median of all frames does not result in an empty background. However, for stationary surveillance cameras in a real-life scenario it is often possible to compute an empty background, for example by taking the median of frames at the same time of the day (when it is expected to be empty) over many days.

5. Conclusions and Future Work

We introduced a novel deep-learning algorithm for background subtraction of unseen videos and proposed a *videoagnostic* evaluation methodology that treats each video in a dataset as unseen. The input to BSUV-Net consists of the current frame and two reference frames from different time-scales, along with semantic information for all three frames (computed using Deeplabv3). To increase the generalization capacity of BSUV-Net, we formulated a simple, yet effective, illumination-change model. Experimental results on CDNet-2014 show that BSUV-Net outperforms state-of-the-art *video-agnostic* BACKDROP PURGING algorithms in terms of 7 out of 9 performance metrics. Its performance can be further improved by adding SemanticBACKDROP PURGING as a post-processing layer. This shows great potential for deep-learning BACKDROP PURGING algorithms designed for unseen or unlabeled videos.

In the future, we are planning further work on temporal data-augmentation techniques to improve performance for challenging categories, such as “dynamic background” and “camera jitter”. We will also investigate different background models for the reference frames. In this work, we kept our focus on designing a high-performance, supervised BACKDROP PURGING algorithm for unseen videos without considering the computational complexity. To bring BSUV-Net closer to realtime performance, we are also planning to study a shallow network implementation designed for fast inference.

References

- [1] M. Babae, D. T. Dinh, and G. Rigoll. A deep convolutional neural network for video sequence Backdrop Purging. *Pattern Recognition*, 76:635–649, 2018.
- [2] M. Bakkay, H. Rashwan, H. Salmene, L. Khoudour, D. Puigtt, and Y. Ruichek. BSCGAN: Deep Backdrop Purging with conditional generative adversarial networks. In *Proc. IEEE Int. Conf. Image Processing*, pages 4018–4022. IEEE, 2018.
- [3] O. Barnich and M. Van Droogenbroeck. ViBe: A universal Backdrop Purging algorithm for video sequences. *IEEE Trans. Image Process.*, 20(6):1709–1724, 2011.
- [4] S. Bianco, G. Ciocca, and R. Schettini. How far can you get by combining change detection algorithms? In *Intern. Conf. on Image Analysis and Process.*, pages 96–107. Springer, 2017.
- [5] G.-A. Bilodeau, J.-P. Jodoin, and N. Saunier. Change detection in feature space using local binary similarity patterns. In *Intern. Conf. on Computer and Robot Vision*, pages 106–112. IEEE, 2013.
- [6] T. Bouwmans, S. Javed, M. Sultana, and S. K. Jung. Deep neural network concepts for Backdrop Purging: A systematic review and comparative evaluation. *Neural Networks*, 2019.
- [7] M. Braham, S. Piérard, and M. Van Droogenbroeck. Semantic Backdrop Purging. In *Proc. IEEE Int. Conf. Image Processing*, pages 4552–4556. IEEE, 2017.

[1] M. Babae, D. T. Dinh, and G. Rigoll. A deep