

Received February 23, 2021, accepted March 24, 2021, date of publication April 5, 2021, date of current version April 13, 2021.

Digital Object Identifier 10.1109/ACCESS.2021.3071163

BSUV-Net 2.0: Spatio-Temporal Data Augmentations for Video-Agnostic Supervised Background Subtraction

M. OZAN TEZCAN¹, (Member, IEEE), PRAKASH ISHWAR¹, (Senior Member, IEEE),
AND JANUSZ KONRAD¹, (Fellow, IEEE)

Department of Electrical and Computer Engineering, Boston University, Boston, MA 02215, USA

Corresponding author: M. Ozan Tezcan (mtezcan@bu.edu)

This work was supported in part by the Advanced Research Projects Agency-Energy (ARPA-E) under Agreement DEAR0000944, and in part by the donation of Titan Graphical Processing Units (GPUs) from NVIDIA Corporation.

ABSTRACT Background subtraction (BGS) is a fundamental video processing task which is a key component of many applications. Deep learning-based supervised algorithms achieve very good performance in BGS, however, most of these algorithms are optimized for either a specific video or a group of videos, and their performance decreases dramatically when applied to unseen videos. Recently, several papers addressed this problem and proposed video-agnostic supervised BGS algorithms. However, nearly all of the data augmentations used in these algorithms are limited to the spatial domain and do not account for temporal variations that naturally occur in video data. In this work, we introduce spatio-temporal data augmentations and apply them to one of the leading video-agnostic BGS algorithms, BSUV-Net. We also introduce a new cross-validation training and evaluation strategy for the CDNet-2014 dataset that makes it possible to fairly and easily compare the performance of various video-agnostic supervised BGS algorithms. Our new model trained using the proposed data augmentations, named BSUV-Net 2.0, significantly outperforms state-of-the-art algorithms evaluated on unseen videos of CDNet-2014. We also evaluate the cross-dataset generalization capacity of BSUV-Net 2.0 by training it solely on CDNet-2014 videos and evaluating its performance on LASIESTA dataset. Overall, BSUV-Net 2.0 provides a ~5% improvement in the F-score over state-of-the-art methods on unseen videos of CDNet-2014 and LASIESTA datasets. Furthermore, we develop a real-time variant of our model, that we call Fast BSUV-Net 2.0, whose performance is close to the state of the art.

INDEX TERMS Background subtraction, foreground detection, scene independent, scene agnostic, deep learning, data augmentation.

I. INTRODUCTION

Background subtraction (BGS) is one of fundamental video processing blocks frequently used in applications such as advanced video surveillance, human activity recognition, autonomous navigation, etc. [1], [2]. BGS can be defined as binary video segmentation aiming to extract foreground and background regions in each video frame.

End-to-end BGS algorithms can be loosely grouped into three categories: (i) Unsupervised algorithms, (ii) *video- or video-group-optimized* supervised algorithms and (iii) *video-agnostic* supervised algorithms. Unsupervised

algorithms attempt to mathematically model the background and extract the foreground pixels accordingly. Several model-based approaches in this category, such as PAWCS [3] and WiseNetMD [4], achieve very competitive performance, however they are currently outperformed by deep learning-based supervised algorithms.

Several recent papers introduced *video- or video-group-optimized* algorithms [5]–[9] and *video-agnostic* [10], [11] supervised BGS algorithms. The first category report results for methods that have been trained and tested on the same set of videos and their performance on unseen videos is not reported. On the other hand, *video-agnostic* algorithms report results on unseen videos by training and testing on disjoint sets of videos. One of the most successful *video-agnostic*

The associate editor coordinating the review of this manuscript and approving it for publication was Ze Ji¹.

BGS algorithms, BSUV-Net, uses spatial and semantic information from different time scales to improve performance on unseen videos. However, due to limited amount of labeled BGS data, BSUV-Net's performance on very challenging scenarios is still insufficient for real-world applications.

One of the most successful approaches for increasing the generalization capacity of computer vision algorithms trained with limited data is the use of data augmentation. Spatial data augmentations such as random crops, rotations, color changes, etc. have proved very successful in image-related tasks [12], [13]. A simple *spatio-temporal* data augmentation was introduced in BSUV-Net [10] to handle illumination differences between videos and resulted in some performance improvement. However, to the best of our knowledge, besides that work there are no other data augmentation attempts tailored to BGS that make use of both spatial and temporal information in a comprehensive manner. In this paper, we introduce a comprehensive suite of spatio-temporal data augmentation methods and adapt them to BSUV-Net. The proposed augmentations address some key BGS challenges, such as PTZ (pan-tilt-zoom) operation, camera jitter and presence of intermittently-static objects. We conduct *video-agnostic* performance analysis and show that these data augmentations significantly increase algorithm's performance for targeted categories without any significant loss of performance in other categories. We also show that a network trained on a combination of several spatio-temporal data augmentations outperforms state of the art (SOTA) methods on unseen videos by $\sim 5\%$ in terms of F-score. Furthermore, we demonstrate a cross-dataset generalization capacity of BSUV-Net 2.0 by training it solely on CDNet-2014 and testing it on LASIESTA [14], a completely unseen dataset. Here too, BSUV-Net 2.0 outperforms SOTA methods by $\sim 5\%$ (F-score). The main contributions of this work are as follows:

- 1) **Spatio-Temporal Data Augmentation:** We introduce spatio-temporal data augmentation methods for BSUV-Net to mimic challenging BGS scenarios, such as PTZ operation, camera jitter and presence of intermittently-static objects (e.g., cars stopped at a streetlight). Our experimental results show that these augmentations significantly improve the performance on unseen videos of corresponding categories.
- 2) **Fair evaluation strategy for CDNet-2014:** Although CDNet-2014 is an extensive BGS dataset, it lacks a training/testing split for use by supervised learning approaches. We introduce a split of CDNet-2014 videos into 4 groups to be used for cross-validation. In this way, we can easily evaluate any supervised BGS algorithm on all CDNet-2014 videos in a video-agnostic manner. This will simplify algorithm performance comparisons in the future.
- 3) **State-of-the-art and real-time results:** Our proposed algorithm outperforms SOTA on CDNet-2014 and LASIESTA for unseen videos by $\sim 5\%$. We also

introduce a real-time variant of BSUV-Net 2.0 which runs at ~ 29 FPS and performs on-par with SOTA.

II. RELATED WORK

A. UNSUPERVISED BGS ALGORITHMS

The early attempts at BGS have relied on probabilistic background models such as Gaussian Mixture Model (GMM) [15] and Kernel Density Estimation (KDE) [16]. Following the idea of BGS based on background modeling, more sophisticated algorithms were introduced (e.g., SubSENSE [17], PAWCS [3], SWCD [18] and WisenetMD). Recently, VTD-FastICA [19] has applied independent component analysis to multiple frames, whereas Giraldo and Bouwmans have introduced a graph-based algorithm that considers the instances in a video as nodes of a graph and computes BGS predictions by minimizing the total variation [20], [21]. Finally, RT-SBS [22] and RTSS [23] have combined unsupervised BGS algorithms with deep learning-based semantic segmentation algorithms, such as PSPNet [24], to improve BGS predictions.

B. VIDEO- OR VIDEO-GROUP-OPTIMIZED SUPERVISED BGS ALGORITHMS

The early attempts of deep learning at BGS have focused on *video-* or *video-group-optimized* algorithms, which are tested on the same videos that they are trained on. Usually, their performance on unseen videos is not reported. *Video-optimized* algorithms train a new set of weights for each video using some of the labeled frames from this very test video, while *video-group-optimized* ones train a single network for the whole dataset by using some labeled frames from the whole dataset. They all achieve near-perfect results [6]–[8], [25]. Although these algorithms might be very useful for speeding up the labeling process of new videos, their performance drops significantly when they are applied to unseen videos [10], [11], [26]. Clearly, they are not suitable for real-world applications.

C. VIDEO-AGNOSTIC SUPERVISED BGS ALGORITHMS

Recently, several supervised-BGS algorithms for unseen videos have been introduced. ChangeDet [27], 3DFR [11], and 3DCD [26] proposed end-to-end convolutional neural networks for BGS that use both spatial and temporal information based on previous frames and a simple median-based background model. Similarly, Kim and Ha [28] introduced a U-Net-based [29] neural network that uses a concatenation of the current frame and several background models generated at different time scales as the input. During evaluation, these methods divided videos of a popular BGS dataset, CDNet-2014 [30], into a training set and a testing set, and reported results for the test videos, unseen by the algorithm during training. Although all three algorithms outperform unsupervised algorithms on their own test sets, their true performance is unknown since no results were reported for the full dataset. Furthermore, these algorithms cannot be

compared with each other since each used a different train/test split. Recently, BSUV-Net [10] also used U-Net architecture but added a novel augmentation step to handle illumination variations, thus improving performance on unseen videos. The authors proposed an evaluation approach on the whole CDNet-2014 dataset by using 18 training/test sets and demonstrated one of the top performances on the dataset's evaluation server. However, their approach with 18 training/test sets is complicated and makes any comparison of future algorithms against BSUV-Net difficult.

D. BGS VERSUS VIDEO-OBJECT SEGMENTATION

Several video-object segmentation datasets and challenges have been introduced in the last few years [31]–[35]. That body of work can be easily confused with BGS. However, there exist important differences that make video-object segmentation algorithms not suitable for BGS. For example, the unsupervised multi-object segmentation dataset in the DAVIS Challenge [35] includes annotations only for objects that capture human attention the most. No annotations are provided for foreground objects that are not in the focus of attention. Since one of the main applications of BGS is video surveillance, the detection of every moving object is critical (e.g., a person on whom the video is not focused may still act suspiciously). Furthermore, the semi-supervised multi-object segmentation dataset in the DAVIS Challenge assumes that semantic segmentation of the first frame of each video is available. Youtube-video object/instance segmentation datasets [31], [32] include annotations for objects from a set of predefined categories. The annotations for a single category might include both stationary background objects and moving foreground objects (e.g., there is no distinction between a parked motorbike [background object] and a moving motorbike [foreground object]). Thus, a comparison of object/instance segmentation and BGS is not very meaningful. In this paper, we do not investigate algorithms designed for video-object segmentation since, despite similarities, they focus on a different task.

In this paper, we improve the performance of BSUV-Net by introducing multiple spatio-temporal data augmentations designed to attack the most common challenges in BGS. We name our improved algorithm BSUV-Net 2.0 and show that it significantly outperforms state-of-the-art BGS algorithms on unseen videos. We also introduce a real-time version of BSUV-Net 2.0 and call it Fast BSUV-Net 2.0. Finally, we propose a 4-fold cross-validation strategy to facilitate fair and streamlined comparison of unsupervised and *video-agnostic* supervised algorithms, which should prove useful for future BGS algorithm comparisons on CDNet-2014.

III. SUMMARY OF BSUV-NET

BSUV-Net is among the top-performing BGS algorithms designed for unseen videos. We briefly summarize it below.

BSUV-Net is a U-Net-based [29] CNN which takes a concatenation of 3 images as input and produces a probabilistic foreground estimate. The input consists of two

background models captured at different time scales and the current frame. One background model, called “empty”, is a manually-selected static frame void of moving objects, whereas the other model, called “recent”, is the median of previous 100 frames. All three input images consist of 4 channels: R, G, B color channels and a foreground probability map (FPM). FPM is an initial foreground estimate for each input image computed by DeepLabv3 [36], a semantic segmentation algorithm that does not use any temporal information. For more details on network architecture and the FPM channel of BSUV-Net, please refer to the original paper [10]. BSUV-Net uses relaxed Jaccard index as the loss function:

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

where $\hat{\mathbf{Y}} \in [0, 1]^{w \times h}$ is the predicted foreground probability map, $\mathbf{Y} \in \{0, 1\}^{w \times h}$ is the ground-truth foreground label, T is a smoothing parameter and m, n are spatial locations.

The authors of BSUV-Net also proposed a novel data-augmentation method for video that addresses illumination differences (ID) often present between video frames. In an ablation study, they demonstrated a significant impact of this augmentation on the overall performance. In this paper, we expand on this idea and introduce a new category of data augmentations designed specifically for spatio-temporal video data.

IV. SPATIO-TEMPORAL DATA AUGMENTATIONS

In this section, we first introduce mathematical notation. Then, we propose new spatio-temporal augmentations and also describe the illumination-difference augmentation which was proposed in BSUV-Net. Fig. 1 shows one example of each of the proposed augmentations.

A. NOTATION

Let us consider an input-label pair of BSUV-Net. The input consists of $\mathbf{I}_E, \mathbf{I}_R, \mathbf{I}_C \in \mathbb{R}^{w \times h \times 4}$ an empty background, a recent background and the current frame, respectively, where w, h are the width and height of each image. Each image has 4 channels: three colors (R, G, B) plus FPM discussed in Section III. Similarly, let $\mathbf{I}_{FG} \in \{0, 1\}^{w \times h}$ be the corresponding foreground label field where 0 represents the background and 1 – the foreground.

Although the resolution of input images varies from video to video, it is beneficial to use a single resolution during training in order to leverage parallel processing of GPUs. Therefore, the first augmentation step we propose is *spatio-temporal cropping* that maps each video to the same spatial resolution. In the second step, we propose two additional augmentations that modify video content but not size.

In our two-step process, in the first step we use different cropping functions to compute $\tilde{\mathbf{I}}_E, \tilde{\mathbf{I}}_R, \tilde{\mathbf{I}}_C \in \mathbb{R}^{\tilde{w} \times \tilde{h} \times 4}$ and $\tilde{\mathbf{I}}_{FG} \in \{0, 1\}^{\tilde{w} \times \tilde{h}}$ from $\mathbf{I}_E, \mathbf{I}_R, \mathbf{I}_C$ and \mathbf{I}_{FG} where \tilde{w}, \tilde{h} are the desired width and height after cropping. In the second step,

we apply post-crop augmentations to compute $\widehat{\mathbf{I}}_{\mathbf{E}}, \widehat{\mathbf{I}}_{\mathbf{R}}, \widehat{\mathbf{I}}_{\mathbf{C}} \in \mathbb{R}^{\widetilde{w} \times \widetilde{h} \times 4}$ and $\widehat{\mathbf{I}}_{\mathbf{FG}} \in \{0, 1\}^{\widetilde{w} \times \widetilde{h}}$ from $\widetilde{\mathbf{I}}_{\mathbf{E}}, \widetilde{\mathbf{I}}_{\mathbf{R}}, \widetilde{\mathbf{I}}_{\mathbf{C}}$ and $\widetilde{\mathbf{I}}_{\mathbf{FG}}$. Below, we explain these two steps in detail.

B. SPATIO-TEMPORAL CROP

Here we describe 3 augmentation techniques to compute $\widetilde{\mathbf{I}}_{\mathbf{E}}, \widetilde{\mathbf{I}}_{\mathbf{R}}, \widetilde{\mathbf{I}}_{\mathbf{C}}, \widetilde{\mathbf{I}}_{\mathbf{FG}}$ from $\mathbf{I}_{\mathbf{E}}, \mathbf{I}_{\mathbf{R}}, \mathbf{I}_{\mathbf{C}}, \mathbf{I}_{\mathbf{FG}}$, each addressing a different BGS challenge. We begin by defining a cropping function, to be used in this section, as follows:

$$\mathcal{C}(\mathbf{I}, i, j, h, w) = \mathbf{I} \left[\left[i - \frac{h}{2} \right] : \left[i + \frac{h}{2} \right], \right. \\ \left. \times \left[j - \frac{w}{2} \right] : \left[j + \frac{w}{2} \right], 1 : 4 \right]$$

where i, j are the center coordinates, h, w are height and width of the crop, $\lceil \cdot \rceil$ denotes the ceiling function and $a : b$ denotes the range of integer indices $a, a + 1, \dots, b - 1$.

1) SPATIALLY-ALIGNED CROP

This is an extension of the widely-used spatial cropping for individual images. Although this is straightforward, we provide a precise definition in order to clearly define steps in subsequent sections.

The output of a spatially-aligned crop is defined follows:

$$\widetilde{\mathbf{I}}_k = \mathcal{C}(\mathbf{I}_k, i, j, \widetilde{h}, \widetilde{w}) \quad \text{for all } k \in \{\mathbf{E}, \mathbf{R}, \mathbf{C}, \mathbf{FG}\},$$

where i, j are randomly-selected spatial indices of the center of the crop. This formulation allows us to obtain a fixed-size, spatially-aligned crop from the input-label pair.

2) RANDOMLY-SHIFTED CROP

One of the most challenging scenarios for BGS algorithms is camera jitter which results in random spatial shifts between consecutive video frames. However, since the variety of such videos is limited in public datasets, it is not trivial to learn the behavior of camera jitter using a data-driven algorithm. In order to address this, we introduce a new data augmentation method by simulating camera jitter. As a result, spatially-aligned inputs become randomly shifted. This is formulated as follows:

$$\widetilde{\mathbf{I}}_k = \mathcal{C}(\mathbf{I}_k, i_k, j_k, \widetilde{h}, \widetilde{w}) \quad \text{for all } k \in \{\mathbf{E}, \mathbf{R}, \mathbf{C}, \mathbf{FG}\},$$

where i_k, j_k are randomly-selected, but such that $i_{\mathbf{C}} = i_{\mathbf{FG}}$ and $j_{\mathbf{C}} = j_{\mathbf{FG}}$ to make sure that the current frame and foreground labels are aligned. By using different center spatial indices for background images and the current frame, we emulate camera jitter effect in the input.

3) PTZ CAMERA CROP

Another challenging BGS scenario is PTZ camera operation. While such videos are very common in surveillance, they form only a small fraction of public datasets. Therefore, we introduce another data augmentation technique specific to this challenge.

Since PTZ videos do not have a static empty background frame, BSUV-Net [10] handles them differently than other categories. Instead of empty and recent backgrounds,

the authors suggest to use recent and more recent background, 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 such frames. To simulate this kind of behavior, we introduce two types of PTZ camera crops: (i) zooming camera crop, (ii) moving camera crop.

The zooming camera crop is defined as follows:

$$\widetilde{\mathbf{I}}_k = \mathcal{C}(\mathbf{I}_k, i, j, \widetilde{h}, \widetilde{w}) \quad \text{for all } k \in \{\mathbf{C}, \mathbf{FG}\}, \\ \widetilde{\mathbf{I}}_k = \frac{1}{N^z} \sum_{n=0}^{N^z-1} \widetilde{\mathbf{I}}_k^n \quad \text{for all } k \in \{\mathbf{E}, \mathbf{R}\}, \text{ where} \\ \widetilde{\mathbf{I}}_k^n = \mathcal{R} \left(\mathcal{C}(\mathbf{I}_k, i, j, \widetilde{h}(1 + nz_k), \widetilde{w}(1 + nz_k)), \widetilde{h}, \widetilde{w} \right)$$

where $z_{\mathbf{E}}, z_{\mathbf{R}}$ represent zoom factors for empty and recent backgrounds and N^z represents the number of zoomed in/out frames to use in averaging. In our experiments, we use $-0.1 < z_{\mathbf{E}}, z_{\mathbf{R}} < 0.1$ and $5 < N^z < 15$ to simulate real-world camera zooming. $\mathcal{R}(\mathbf{I}, \widetilde{h}, \widetilde{w})$ is an image resizing function that changes the resolution of \mathbf{I} to $(\widetilde{w}, \widetilde{h})$ using bilinear interpolation. Note, that using positive values for z_k simulates zooming in whereas using negative values simulates zooming out. Fig. 1(d) shows an example of zoom-in.

Similarly, the moving camera crop is defined as follows:

$$\widetilde{\mathbf{I}}_k = \mathcal{C}(\mathbf{I}_k, i, j, \widetilde{h}, \widetilde{w}) \quad \text{for all } k \in \{\mathbf{C}, \mathbf{FG}\}, \\ \widetilde{\mathbf{I}}_k = \frac{1}{N_k^m} \sum_{n=0}^{N_k^m-1} \widetilde{\mathbf{I}}_k^n \quad \text{for all } k \in \{\mathbf{E}, \mathbf{R}\}, \text{ where} \\ \widetilde{\mathbf{I}}_k^n = \mathcal{C}(\mathbf{I}_k, i + np, j + nq, \widetilde{h}, \widetilde{w})$$

where p, q are the vertical and horizontal shift amounts per frame and $N_{\mathbf{E}}^m, N_{\mathbf{R}}^m$ represent the number of empty and recent moving background crops to use for averaging. This simulates camera pan and tilt. In our experiments, we use $-5 < p, q < 5$ and $5 < N_{\mathbf{E}}^m, N_{\mathbf{R}}^m < 15$ to simulate real-world camera movements.

C. POST-CROP AUGMENTATIONS

In this section, we define several content-modifying augmentation techniques to compute $\widehat{\mathbf{I}}_{\mathbf{E}}, \widehat{\mathbf{I}}_{\mathbf{R}}, \widehat{\mathbf{I}}_{\mathbf{C}}, \widehat{\mathbf{I}}_{\mathbf{FG}}$ from $\widetilde{\mathbf{I}}_{\mathbf{E}}, \widetilde{\mathbf{I}}_{\mathbf{R}}, \widetilde{\mathbf{I}}_{\mathbf{C}}, \widetilde{\mathbf{I}}_{\mathbf{FG}}$. These augmentations can be applied after any one of the spatio-temporal crop augmentations.

1) ILLUMINATION DIFFERENCE

Illumination variations are common, especially in long videos, for example due to changes in natural light or lights being turned on/off. A temporal data augmentation technique to handle illumination changes was introduced in BSUV-Net [10] with the goal of increasing the network's generalization capacity for unseen videos. We use this augmentation here as well, formulated in our notation as follows:

$$\widehat{\mathbf{I}}_k[i, j, c] = \widetilde{\mathbf{I}}_k[i, j, c] + \mathbf{d}_k[c] \\ \times \text{for } k \in \{\mathbf{E}, \mathbf{R}, \mathbf{C}\}, c = 1, 2, 3$$

where $\mathbf{d}_{\mathbf{E}}, \mathbf{d}_{\mathbf{R}}, \mathbf{d}_{\mathbf{C}} \in \mathbb{R}^3$ represent illumination offsets applied to RGB channels of the input images.

2) INTERMITTENT-OBJECT ADDITION

Another challenge for BGS are scenarios when objects enter a scene but then stop and remain static for a long time. Even very successful BGS algorithms, after some time, predict these objects as part of the background for they rely on recent frames to estimate the background model. BSUV-Net overcomes this challenge by using inputs from multiple time scales, however it still underperforms on videos with intermittently-static objects. To address this, we introduce another spatio-temporal data augmentation specific to this challenge.

We use a masking-based approach for intermittently-static objects as follows. In addition to the cropped inputs $\tilde{\mathbf{I}}_E, \tilde{\mathbf{I}}_R, \tilde{\mathbf{I}}_C, \tilde{\mathbf{I}}_{FG}$, we also use cropped inputs from videos with intermittently-static objects defined as $\tilde{\mathbf{I}}_E^{IO}, \tilde{\mathbf{I}}_R^{IO}, \tilde{\mathbf{I}}_C^{IO} \in \mathbb{R}^{\tilde{w} \times \tilde{h} \times 4}$ and $\tilde{\mathbf{I}}_{FG}^{IO} \in \{0, 1\}^{\tilde{w} \times \tilde{h}}$. We copy foreground pixels from the intermittently-static input and paste them into the original input to synthetically create an intermittent object. This can be formulated as follows:

$$\begin{aligned} \hat{\mathbf{I}}_E &= \tilde{\mathbf{I}}_E \\ \hat{\mathbf{I}}_k &= \tilde{\mathbf{I}}_{FG}^{IO} \odot \tilde{\mathbf{I}}_k^{IO} + (1 - \tilde{\mathbf{I}}_{FG}^{IO}) \odot \tilde{\mathbf{I}}_k \text{ for } k \in \{C, R\}, \\ \hat{\mathbf{I}}_{FG} &= \tilde{\mathbf{I}}_{FG}^{IO} + (1 - \tilde{\mathbf{I}}_{FG}^{IO}) \odot \tilde{\mathbf{I}}_{FG} \end{aligned}$$

where \odot denotes Hadamard (element-wise) product. Fig. 1(f) shows an example of intermittent object addition. Note, that this augmentation requires prior knowledge of examples with intermittently-static objects which can be found in some public datasets.

D. COMBINING SPATIO-TEMPORAL AUGMENTATIONS

While the augmentations defined above can all be used by themselves to improve the BGS performance on related categories, combining multiple or even all of them might result in a better algorithm for a general unseen video of which the category is unknown. However, combining the crop algorithms is not trivial since it is not practical to apply more than one crop function to a single input. Thus, we use online augmentation, where we randomly augment every input while forming mini-batches. The augmentation steps are as follows:

- 1) randomly select one of the spatial crop augmentations and apply it to the input,
- 2) apply the illumination change augmentation using randomized illumination values,
- 3) apply intermittent object addition to $p\%$ of the inputs.

Clearly, a different combination of augmentations will be applied to the same input in different epochs. We hope this will significantly increase the generalization capacity of our network.

V. VIDEO-AGNOSTIC EVALUATION STRATEGY FOR SUPERVISED ALGORITHMS

The most commonly used BGS datasets with a variety of scenarios and pixel-wise ground-truth annotations are CDNet-2014 [30], LASIESTA [14] and SBMI2015 [37]. Among these 3 datasets, only CDNet-2014 has a well-maintained




FIGURE 1. Image augmentation examples. Each row shows an example for one of the augmentations: (a) original input, (b) spatially-aligned crop, (c) randomly-shifted crop, (d) PTZ camera crop, (e) illumination difference, (f) intermittent-object addition.

evaluation server, that keeps a cumulative performance record of the uploaded algorithms. Moreover, it has been the most widely-used dataset for BGS in recent years with publicly-available evaluation results for nearly all of the published BGS algorithms.

Since one of our aims is to compare the performance of BSUV-Net 2.0 with SOTA *video-agnostic* BGS algorithms on unseen videos, the availability of public results for these algorithms is critical. Therefore, we use CDNet-2014 as our evaluation dataset. CDNet-2014 is a comprehensive dataset that provides some ground-truth frames from all 53 videos to the public, but keeps others internally for algorithm comparison. Since it does not include any videos with no labeled frames, it is not directly suitable for testing of *video-agnostic* supervised algorithms. Consequently, most of the leading algorithms are either *video-optimized* or *video-group-optimized* and achieve near-perfect results by over-fitting the training data. However, these results are not generalizable to unseen videos [10], [26], [27]. Several researchers addressed this problem by designing generalizable networks and evaluating their algorithms on unseen videos by using different videos in training and testing [10], [11], [26]–[28]. Yet, there is no common strategy for testing the performance of supervised

BGS algorithms on CDNet-2014 for unseen videos. Some of the recent papers divide the dataset into two folds, train their algorithm on one of the folds and test on the other one. Since they only report the results on the test videos that they selected, their results might be biased towards their test set and not directly comparable with unsupervised algorithms. While BSUV-Net [10] proposes a *video-agnostic* evaluation strategy for the full CDNet-2014 dataset by using 18 training/testing video sets, this strategy might also be biased and is computationally expensive.

In this paper, we introduce a simple and intuitive 4-fold cross-validation strategy for CDNet-2014. We grouped all videos in the dataset and each category into 4 folds as evenly as possible (Table 1). The proposed *video-agnostic* evaluation strategy is to train any supervised BGS algorithm on three of the folds and test on the remaining fold and replicate the same process for all 4 combinations. This approach will provide results on the full CDNet-2014 dataset which can be uploaded to the evaluation server to compare against SOTA. We believe this cross-validation strategy will be very beneficial for the evaluation of future BGS algorithms.

VI. EXPERIMENTAL RESULTS

A. DATASET AND EVALUATION DETAILS

We evaluate the performance of our algorithm on CDNet-2014 [30] using the evaluation strategy described in Section V. In CDNet-2014, the spatial resolution of videos varies from 320×240 to 720×526 pixels. The videos are labeled pixel-wise as follows: 1) foreground, 2) background, 3) hard shadow or 4) unknown motion. As suggested in [30], during evaluation we ignore pixels with unknown motion label and consider hard-shadow pixels as background.

In performance evaluation, we use metrics reported by CDNet-2014, namely recall (Re), specificity (Sp), false positive rate (FPR), false negative rate (FNR), percentage of wrong classifications (PWC), precision (Pr) and F-score (F_1). We also report two ranking-based metrics, “average ranking” (R) and “average ranking across categories” (R_{cat}), which combine all 7 metrics into ranking scores. A detailed description of these rankings can be found in [38].

In order to better understand the performance of BSUV-Net 2.0 on unseen videos, we also performed a *cross-dataset* evaluation by training our model on CDNet-2014 and testing it on a completely different dataset, LASIESTA [14]. LASIESTA is an extensive BGS dataset which includes 24 different videos from various indoor and outdoor scenarios. It includes a “Simulated Motion” category that is comprised of fixed-camera videos that are post-processed to mimic camera pan, tilt and jitter [14].

B. TRAINING DETAILS

In order to train BSUV-Net 2.0, we use similar parameters to the ones used for BSUV-Net. The same parameters are used for each of the four cross-validation folds. We used ADAM optimizer with a learning rate of 10^{-4} , $\beta_1 = 0.9$,

TABLE 1. Sets used in 4-fold cross-validation on CDNet-2014.

category	video	S ₁	S ₂	S ₃	S ₄
baseline	highway	✓			
	pedestrians		✓		
	office			✓	
	PETS2006				✓
bad weather	blizzard	✓			
	skating		✓		
	wetSnow			✓	
	snowFall				✓
intermittent object motion	sofa	✓			
	winterDriveway		✓		
	parking			✓	
	abandonedBox			✓	
	streetLight				✓
low framerate	port 0.17fps	✓			
	tramCrossroad 1fps		✓		
	tunnelExit 0.35fps			✓	
	turnpike 0.5fps				✓
PTZ	continuousPan	✓			
	intermittentPan		✓		
	zoomInZoomOut			✓	
	twoPositionPTZCam				✓
thermal	corridor	✓			
	lakeSide	✓			
	library		✓		
	diningRoom			✓	
	park				✓
camera jitter	badminton	✓			
	traffic		✓		
	boulevard			✓	
	sidewalk				✓
shadow	copyMachine	✓			
	busStation		✓		
	cubicle			✓	
	peopleInShade			✓	
	bungalows				✓
dynamic background	backdoor				✓
	overpass	✓			
	fountain02	✓			
	fountain01		✓		
	boats		✓		
	canoe			✓	
night videos	fall				✓
	bridgeEntry	✓			
	busyBoulevard	✓			
	tramStation		✓		
	winterStreet		✓		
	fluidHighway			✓	
turbulence	streetCornerAtNight				✓
	turbulence0	✓			
	turbulence1		✓		
	turbulence2			✓	
	turbulence3				✓

and $\beta_2 = 0.99$. The mini-batch size was 8 and the number of epochs was 200. As the empty background frame, we used manually-selected frames introduced in [10]. We used the median of preceding 100 frames as the recent background.

In terms of spatio-temporal data augmentations, we use an online approach to randomly change the parameters under the following constraints. The random pixel shift between inputs is sampled from $\mathcal{U}(0, 5)$ where $\mathcal{U}(a, b)$ denotes uniform random variable between a to b . The zooming-in ratios are sampled from $\mathcal{U}(0, 0.02)$ and $\mathcal{U}(0, 0.04)$ for recent and empty backgrounds, respectively, while the zooming-out ratios are sampled from

TABLE 2. Comparison of different spatio-temporal augmentations on CDNet-2014 based on F-score. SAC: spatially-aligned crop, RSC: randomly-shifted crop, PTZ: PTZ camera crop, ID: illumination difference, IOA: intermittent object addition. The values in boldface font show the best performance for each category.

SAC	RSC	PTZ	ID	IOA	Bad weather	Low framerate	Night	PTZ	Thermal	Shadow	Int. obj. motion	Camera jitter	Dynamic backgr.	Base-line	Turbulence	Overall
✓					0.9442	0.7886	0.6982	0.6564	0.8960	0.9848	0.7732	0.8237	0.8517	0.9878	0.7285	0.8303
✓	✓				0.9439	0.8217	0.6941	0.6304	0.8854	0.9818	0.7620	0.9043	0.8745	0.9865	0.7354	0.8382
✓		✓			0.9315	0.7961	0.6557	0.6815	0.8905	0.9795	0.7458	0.8999	0.8674	0.9838	0.7409	0.8339
✓			✓		0.9489	0.7606	0.7605	0.6579	0.9024	0.9855	0.7503	0.8270	0.8364	0.9874	0.7341	0.8319
✓				✓	0.9456	0.7550	0.7233	0.6383	0.8997	0.9836	0.9312	0.8359	0.8709	0.9883	0.7023	0.8431
✓	✓	✓	✓	✓	0.9272	0.8114	0.6841	0.6725	0.8960	0.9811	0.8489	0.9163	0.8848	0.9834	0.8056	0.8556

$\mathcal{U}(-0.02, 0)$ and $\mathcal{U}(-0.04, 0)$. We use $N^z = 10$. The horizontal pixel shift for moving camera augmentation is sampled from $\mathcal{U}(0, 5)$ with $N_E^m = 20$ and $N_R^m = 10$. We perform no vertical-shift augmentation since CDNet-2014 does not include any videos with vertical camera movement. For illumination change, assuming $[0, 1]$ as the range of pixel values, we use $\mathbf{d}_R[k] = \mathbf{d}_C[k] = I + I_k$ where $I \sim \mathcal{N}(0, 0.1^2)$ and $I_k \sim \mathcal{N}(0, 0.04^2)$ for $k \in \{1, 2, 3\}$. Similarly, $\mathbf{d}_E[k] = \mathbf{d}_C[k] + I^E + I_k^E$ where $I^E \sim \mathcal{N}(0, 0.1^2)$ and $I_k^E \sim \mathcal{N}(0, 0.04^2)$ for $k \in \{1, 2, 3\}$. Lastly, for intermittent object addition, we always use the “intermittent object motion” inputs from the current training set and apply this augmentation to $p = 10\%$ of the inputs only. Details of the training and evaluation implementation, with all of the defined augmentations, will be made publicly available upon publication. During inference, binary maps are obtained by thresholding network output at $\theta = 0.5$.

C. ABLATION STUDY

We assess the impact of each spatio-temporal data augmentation method defined in Section IV. As the baseline network, we use BSUV-Net with only spatially-aligned crop augmentation and random Gaussian noise sampled from $\mathcal{N}(0, 0.01^2)$. We evaluate the proposed spatio-temporal augmentations against this baseline by including the spatially-aligned crop among spatial crop augmentations, as explained in Section IV-D. In PTZ camera crop, for each input, we randomly select one of the following: zooming in, zooming out, moving right or moving left. Table 2 shows F-scores for each category of CDNet-2014 computed locally¹ for frames with publicly-available ground truth. We report the median of results for every 5th epoch between 150th and 200th epochs to disregard small fluctuations in the learning process. We perform this across all four splits proposed in Table 1.

Fig. 2 shows some visual results for these algorithms for 5 videos. It can be observed that each augmentation type significantly improves the performance on related categories (randomly shifted crop – on “Camera jitter”, PTZ camera crop – on “PTZ”, illumination difference – on “Shadow”, intermittent object addition – on “Intermittent object motion”), but combining all augmentations decreases

¹We provide only locally-computed results because if the results of the ablation study were uploaded to the CDNet-2014 evaluation server, they would have not been made public since they all come from the same algorithm. Moreover, this simplifies corroboration of our results by independent parties by not requiring uploads to the evaluation server.

TABLE 3. Efficiency vs performance trade-off for BSUV-Net 2.0 on CDNet-2014. FPS is calculated using PyTorch 1.3 implementation on a node with single Nvidia Tesla P100 GPU.

	Re	Pr	F ₁	FPS	
				320 × 240	640 × 480
BSUV-Net 2.0	0.85	0.89	0.86	~ 6	~ 2.5
Fast BSUV-Net 2.0	0.84	0.84	0.81	~ 29	~ 13

the performance significantly on some categories (e.g., night and intermittent object motion). We believe this is due to trade-offs between the effects of different augmentations. For example, when a static background object starts moving it should be labeled as foreground, but a network trained with a randomly-shifted crop augmentation can confuse this input with an input from the “Camera jitter” category and continue labeling the object as background. Still, the overall performance (last column in Table 2) of BSUV-Net 2.0 that uses all augmentations handily outperforms the overall performance for individual augmentations.

Since BGS is often applied as a pre-processing step in real-time video processing applications, computation speed is critical. As discussed in our previous work [10], one of the main bottlenecks of BSUV-Net is the computation of FPM for each channel – it decreases the overall computation speed significantly. On the other hand, either removing the FPM channel or predicting BGS by thresholding the FPM channel alone decreases the performance to values that are lower than that of some unsupervised algorithms [10]. In this work, we show that the performance of our model, even without the FPM channel but with augmentations, is better than the current state-of-the-art. We call this version of BSUV-Net 2.0, which uses 9 instead of 12 channels on input, Fast BSUV-Net 2.0. Table 3 shows a speed and performance comparison of the two versions. Clearly, while Fast BSUV-Net 2.0 has lower performance, it can be used in real-time applications at 320 × 240 spatial resolution, which is very similar to the resolution used in training. For higher-resolution videos, one can easily feed decimated frames into Fast BSUV-Net 2.0 and interpolate the resulting BGS predictions to the original resolution.

D. COMPARISON WITH STATE OF THE ART

Table 4 shows the performance of BSUV-Net 2.0 and Fast BSUV-Net 2.0 compared to state-of-the-art BGS algorithms that are designed for and tested on unseen videos. We did not include the results of *video-* or *video-group-optimized* algorithms since it is not fair to compare them against

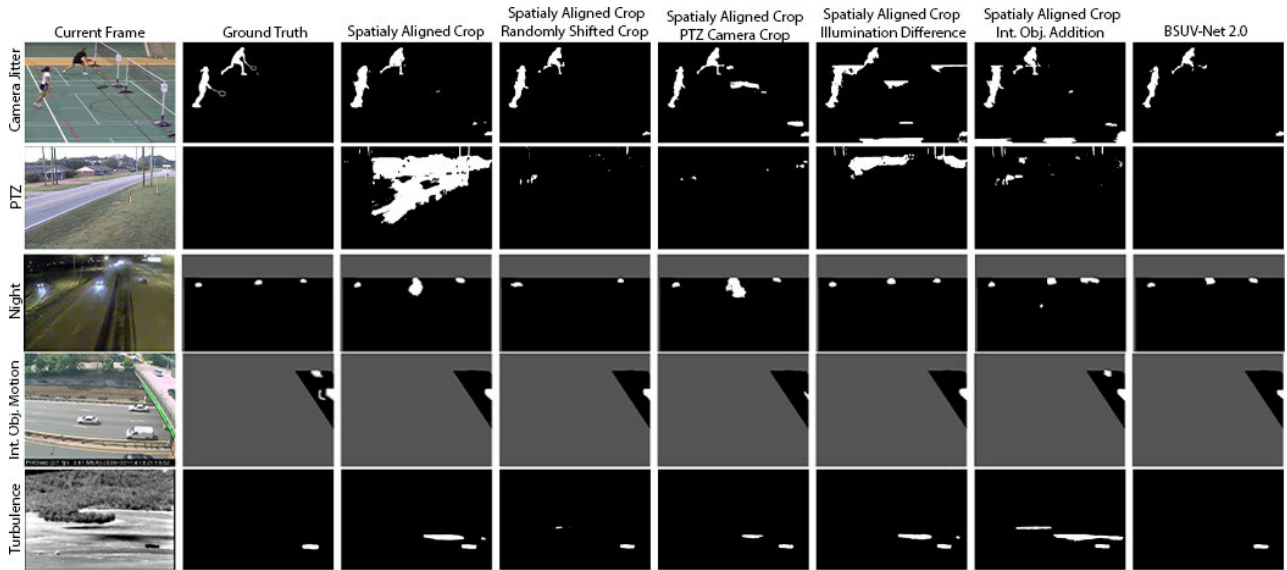


FIGURE 2. Visual comparison of different spatio-temporal augmentations on sample frames. Columns 1 and 2 show the current frame and its ground truth. The subsequent columns show predictions of our network trained with augmentations listed at the top of each column. The last column shows predictions of BSUV-Net 2.0 trained using a combination of all proposed augmentations. The dark-gray areas represent pixels outside of CDNet-2014 regions of interest.

TABLE 4. Official comparison of top BGS algorithms evaluated on unseen videos from CDNet-2014.

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>
BSUV-Net 2.0	7.71	7.73	0.8136	0.9979	0.0021	0.1864	0.7614	0.9011	0.8387
Fast BSUV-Net 2.0	7.86	10.36	0.8181	0.9956	0.0044	0.1819	0.9054	0.8425	0.8039
BSUV-Net + SemanticBGS [10]	9.57	14.27	0.8179	0.9944	0.0056	0.1821	1.1326	0.8319	0.7986
IUTIS-5 + SemanticBGS [39]	9.71	11.91	0.7890	0.9961	0.0039	0.2110	1.0722	0.8305	0.7892
IUTIS-5 [40]	12.14	10.91	0.7849	0.9948	0.0052	0.2151	1.1986	0.8087	0.7717
BSUV-Net [10]	9.71	14.00	0.8203	0.9946	0.0054	0.1797	1.1402	0.8113	0.7868
WisenetMD [4]	17.29	15.82	0.8179	0.9904	0.0096	0.1821	1.6136	0.7535	0.7668
FgSegNet v2 [41]	45.57	45.09	0.5119	0.9411	0.0589	0.4881	7.3507	0.4859	0.3715

TABLE 5. Official comparison of top BGS algorithms according to the per-category F-score on unseen videos from CDNet-2014.

Method	Bad weather	Low framerate	Night	PTZ	Thermal	Shadow	Int. obj. motion	Camera jitter	Dynamic backgr.	Base-line	Turbulence	Overall
BSUV-Net 2.0	0.8844	0.7902	0.5857	0.7037	0.8932	0.9562	0.8263	0.9004	0.9057	0.9620	0.8174	0.8387
Fast BSUV-Net 2.0	0.8909	0.7824	0.6551	0.5014	0.8379	0.8890	0.9016	0.8828	0.7320	0.9694	0.7998	0.8039
BSUV-Net + SemanticBGS	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 + SemanticBGS	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
BSUV-Net	0.8713	0.6797	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
WisenetMD	0.8616	0.6404	0.5701	0.3367	0.8152	0.8984	0.7264	0.8228	0.8376	0.9487	0.8304	0.7535
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

video-agnostic algorithms. This table shows official results computed by CDNet-2014 evaluation server,² so the results of our models slightly differ from those in Tables 2 and 3 (different ground-truth frames). We compare BSUV-Net 2.0 with some of the top-performing video-agnostic algorithms reported by this server. RTSS [23], 3DCD [26], 3DFR [11], ChangeDet [27] and Kim and Ha [28] are not included in this table since their results are not reported. Video-agnostic results of FgSegNet v2 are taken from [10]. BSUV-Net 2.0 outperforms all SOTA algorithms by at least ~5% in terms

²Full results can be accessed from jacarini.dinf.usherbrooke.ca/results2014/870/ and jacarini.dinf.usherbrooke.ca/results2014/871/

TABLE 6. F-score comparison of BSUV-Net 2.0 with video-agnostic supervised BGS algorithms that are not reported in changedetection.net. Each column shows test performance of the algorithm by using the training/testing split provided in respective paper.

Method	Training/Testing Split			
	3DCD	3DFR	ChangeDet	Kim
BSUV-Net 2.0	0.89	0.89	0.90	0.88
Fast BSUV-Net 2.0	0.84	0.84	0.86	0.88
3DCD [26]	0.86	-	-	-
3DFR [11]	-	0.86	-	-
ChangeDet [27]	-	-	0.84	-
Kim [28]	-	-	-	0.86

of F-score (0.8387 versus 0.7986 in Tables 4, 5). Fast BSUV-Net 2.0 also outperforms all state-of-the-art algorithms while being ~5 times faster than BSUV-Net 2.0 during inference

TABLE 7. Per-category F-score comparison of the cross-dataset performance of BSUV-Net 2.0 with the top-performing unsupervised BGS algorithms on LASIESTA.

Method	ISI	ICA	IOC	IIL	IMB	IBS	OCL	ORA	OSN	OSU	Overall
BSUV-Net 2.0	0.92	0.68	0.96	0.88	0.81	0.77	0.93	0.94	0.84	0.79	0.85
Cuevas [42]	0.88	0.84	0.78	0.65	0.93	0.66	0.93	0.87	0.78	0.72	0.81
Haines [43]	0.89	0.89	0.92	0.85	0.84	0.68	0.83	0.89	0.17	0.86	0.78
Maddalena2 [44]	0.95	0.86	0.95	0.21	0.91	0.40	0.97	0.90	0.81	0.88	0.78
Maddalena1 [45]	0.87	0.85	0.91	0.61	0.76	0.42	0.88	0.84	0.58	0.80	0.75

TABLE 8. Per-video F-score comparison of BSUV-Net 2.0 trained on two different configurations with the top-performing supervised BGS algorithms on unseen videos from LASIESTA. BSUV-Net 2.0 is trained only on CDNet-2014, as explained in Sec. VI-E. BSUV-Net 2.0* is trained on videos from LASIESTA that are not from the test set as described in [26].

Method	ISI-2	ICA-2	IOC-2	IIL-2	IMB-2	IBS-2	OCL-2	ORA-2	OSN-2	OSU-2	Overall
BSUV-Net 2.0	0.89	0.60	0.95	0.89	0.76	0.69	0.89	0.93	0.70	0.91	0.82
BSUV-Net 2.0*	0.98	0.99	0.97	0.97	0.88	0.95	0.97	0.97	0.55	0.95	0.92
3DCD [26]	0.86	0.49	0.93	0.85	0.79	0.87	0.87	0.87	0.49	0.83	0.79
FgSegNet v2 [7]	0.53	0.58	0.25	0.41	0.63	0.25	0.54	0.54	0.05	0.29	0.41
FgSegNet-M [41]	0.56	0.55	0.65	0.42	0.56	0.19	0.28	0.18	0.01	0.33	0.37

(Table 3). Table 5 shows the comparison of F_1 results for each category. This table includes RTSS using results reported in the paper [23]. In 7 out of 11 categories, either BSUV-Net 2.0 or Fast BSUV-Net 2.0 achieve the best performance, including most of the categories that we designed the augmentations for (an exception is the “Night” category). However, note that the best-performing algorithm in the “Night” category is BSUV-Net which uses only the illumination-difference augmentation. Thus, it focuses on videos with illumination differences such as night videos.

Fig. 3 qualitatively compares the performance of BSUV-Net 2.0 with state-of-the-art video-agnostic BGS algorithms on a few example videos from CDNet-2014. BSUV-Net 2.0 clearly produces the best visual results in a variety of scenarios. Results for *Camera jitter* and *PTZ* categories show the effectiveness of BSUV-Net 2.0 in removing false positives resulting from camera motion. In the example from *Intermittent object motion* category, the car on the left is starting to back-up from the driveway and most of the algorithms produce false positives at the location where the car was parked whereas BSUV-Net 2.0 successfully eliminates these false positives. Results for *Dynamic background* show that BSUV-Net 2.0 is very effective in accurately delineating the boundary between foreground objects and the background.

As discussed in Section II, 3DCD [26], 3DFR [11], ChangeDet [27] and Kim and Ha [28] are also among the best *video-agnostic* supervised algorithms, however each reports performance on a different subset of CDNet-2014, with the algorithm trained on the remaining videos. Table 6 shows the comparison of BSUV-Net 2.0 with these algorithms using the training/testing splits provided in respective papers in each column. BSUV-Net 2.0 clearly outperforms all four competitors, while Fast BSUV-Net 2.0 beats 2 out of 4, and does so with real-time performance.

E. CROSS-DATASET EVALUATION

In this section, we perform a cross-dataset evaluation to show the generalization capacity of BSUV-Net 2.0. We train

TABLE 9. F-score comparison of cross-dataset performance of different spatio-temporal augmentations on *moving camera* and *simulated motion* videos of LASIESTA. SAC: spatially-aligned crop, RSC: randomly-shifted crop, PTZ: PTZ camera crop.

SAC	RSC	PTZ	Indoor pan & tilt	Outdoor pan & tilt	Indoor jitter	Outdoor jitter
✓			0.48	0.56	0.81	0.75
✓	✓		0.52	0.42	0.88	0.85
✓		✓	0.58	0.58	0.84	0.50

BSUV-Net 2.0 using CDNet-2014 videos from S_2, S_3, S_4 sets shown in Table 1 and use S_1 as a validation set to select the best performing epoch. Then, we evaluate the results on a completely different dataset, LASIESTA [14].³ Table 7 shows the comparison of BSUV-Net 2.0 with top-performing unsupervised algorithms reported in [14]. Since the authors reported results only for categories of LASIESTA recorded with static cameras, we report results only on these categories. Clearly, BSUV-Net 2.0 outperforms its competitors on a completely unseen dataset by a significant margin.

In [26], Mandal *et al.* performed a *video-agnostic* evaluation of some supervised learning algorithms by training with 10 of the LASIESTA videos and evaluating on 10 unseen videos from LASIESTA. Table 8 shows a comparison of BSUV-Net 2.0 with unseen video performance of the algorithms reported in [26]. We show the results of BSUV-Net 2.0 trained with two different datasets. BSUV-Net 2.0 row shows the results of cross-dataset training whereas BSUV-Net 2.0* row shows the results of using the same training set that is used in [26] for a fair comparison. BSUV-Net 2.0 achieves significantly better results than state of the art even if the training set does not include any videos from LASIESTA. Since we train BSUV-Net 2.0* with videos from LASIESTA, it performs even better than BSUV-Net 2.0. This shows that the proposed spatio-temporal data augmentations are not specific to CDNet-2014 and can be very effective on other datasets as well. Note that the performance of BSUV-Net 2.0 is significantly better than that of BSUV-Net 2.0* on OSN-2,

³The empty backgrounds of LASIESTA videos are computed automatically as the median of all frames in the video. The recent backgrounds are computed similarly to CDNet-2014, as the median of previous 100 frames.

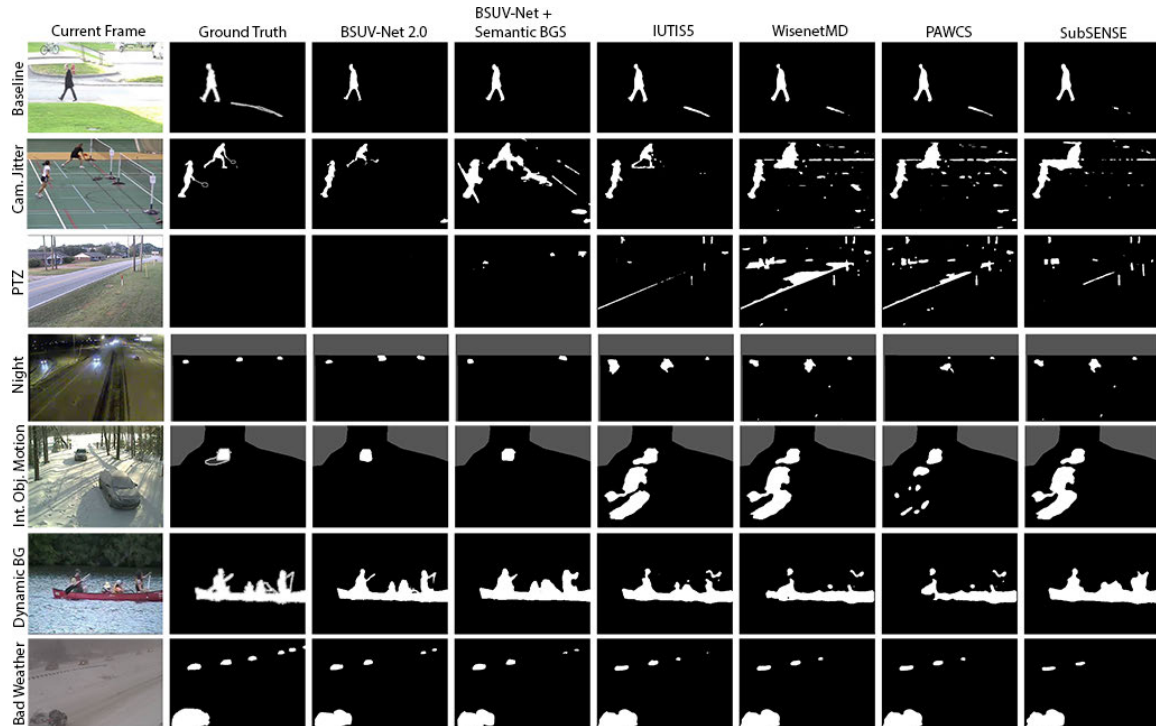


FIGURE 3. Qualitative comparison of top BGS algorithms on sample frames from different categories of CDNet-2014.

an outdoor video recorded in heavy snow. This is due to the fact that the training videos of LASIESTA do not include a heavy-snow video, however the training set of CDNet-2014 does. This also shows the importance of scene variety in the training dataset. Both Table 7 and 8 clearly show that BSUV-Net 2.0 is not specific to a dataset that it was trained on, but can successfully predict BGS of an unseen video.

In addition to the videos reported in Table 7, LASIESTA includes several BGS videos that are either recorded with a moving camera or post-processed to look like they were recorded with a moving camera. We group these videos under 4 categories:

- 1) *Indoor pan & tilt videos* (IMC-1, ISM-1, ISM-2, ISM-3),
- 2) *Outdoor pan & tilt videos* (OMC-1, OSM-1, OSM-2, OSM-3),
- 3) *Indoor jitter videos* (IMC-2, ISM-4, ..., ISM-12),
- 4) *Outdoor jitter videos* (OMC-2, OSM-4, ..., OSM-12).

Table 9 shows the F-score comparison of BSUV-Net 2.0 trained with different combinations of spatio-temporal data augmentations on these 4 categories. As expected, the randomly shifted crop augmentation achieves the best performance for videos with camera jitter whereas the PTZ augmentation achieves the best results for PTZ category. This further shows that the impact of spatio-temporal data augmentations is generalizable to different datasets.

VII. CONCLUSION

While background subtraction algorithms achieve remarkable performance today, they still often fail in challenging

scenarios such as shaking or panning/tilting/zooming cameras, or when moving objects stop for an extended period of time. In the case of supervised algorithms, this is largely due to the limited availability of labeled videos recorded in such scenarios – it is difficult to train end-to-end deep-learning algorithms for unseen videos. To address this, we introduced several spatio-temporal data augmentation methods to synthetically increase the number of inputs in such scenarios. Specifically, we introduced new augmentations for PTZ, camera jitter and intermittent object motion scenarios, and achieved significant performance improvements in these categories and, consequently, a better overall performance on CDNet-2014 dataset. We also introduced a real-time version of BSUV-Net 2.0 which still performs better than state-of-the-art methods and we proposed a 4-fold cross-validation data split for CDNet-2014 for easier comparison of future algorithms. Finally, we demonstrated a strong generalization capacity of BSUV-Net 2.0 using cross-dataset evaluation on LASIESTA in which the proposed model significantly outperforms the current state-of-the-art methods on a completely unseen dataset.

REFERENCES

- [1] T. Bouwmans, S. Javed, M. Sultana, and S. K. Jung, "Deep neural network concepts for background subtraction: A systematic review and comparative evaluation," *Neural Netw.*, vol. 117, pp. 8–66, Sep. 2019.
- [2] B. Garcia-Garcia, T. Bouwmans, and A. J. R. Silva, "Background subtraction in real applications: Challenges, current models and future directions," *Comput. Sci. Rev.*, vol. 35, Feb. 2020, Art. no. 100204.
- [3] P.-L. St-Charles, G.-A. Bilodeau, and R. Bergevin, "A self-adjusting approach to change detection based on background word consensus," in *Proc. IEEE Winter Conf. Appl. Comput. Vis.*, Jan. 2015, pp. 990–997.

- [4] S.-H. Lee, G.-C. Lee, J. Yoo, and S. Kwon, "WisenetMD: Motion detection using dynamic background region analysis," *Symmetry*, vol. 11, no. 5, p. 621, May 2019.
- [5] D. Zeng and M. Zhu, "Background subtraction using multiscale fully convolutional network," *IEEE Access*, vol. 6, pp. 16010–16021, 2018.
- [6] M. C. Bakkay, H. A. Rashwan, H. Salmame, L. Khoudour, D. Puig, and Y. Ruichek, "BSCGAN: Deep background subtraction with conditional generative adversarial networks," in *Proc. 25th IEEE Int. Conf. Image Process. (ICIP)*, Oct. 2018, pp. 4018–4022.
- [7] L. A. Lim and H. Y. Keles, "Learning multi-scale features for foreground segmentation," *Pattern Anal. Appl.*, vol. 23, no. 3, pp. 1369–1380, Aug. 2020.
- [8] M. Babae, D. T. Dinh, and G. Rigoll, "A deep convolutional neural network for video sequence background subtraction," *Pattern Recognit.*, vol. 76, pp. 635–649, Apr. 2018.
- [9] D. Sakkos, H. Liu, J. Han, and L. Shao, "End-to-end video background subtraction with 3D convolutional neural networks," *Multimedia Tools Appl.*, vol. 77, pp. 1–19, Dec. 2017.
- [10] M. O. Tezcan, P. Ishwar, and J. Konrad, "BSUV-Net: A fully-convolutional neural network for background subtraction of unseen videos," in *Proc. IEEE Winter Conf. Appl. Comput. Vis. (WACV)*, Mar. 2020, pp. 2774–2783.
- [11] M. Mandal, V. Dhar, A. Mishra, and S. K. Vipparthi, "3DFR: A swift 3D feature reduction framework for scene independent change detection," *IEEE Signal Process. Lett.*, vol. 26, no. 12, pp. 1882–1886, Dec. 2019.
- [12] L. Taylor and G. Nitschke, "Improving deep learning with generic data augmentation," in *Proc. IEEE Symp. Ser. Comput. Intell. (SSCI)*, Nov. 2018, pp. 1542–1547.
- [13] C. Shorten and T. M. Khoshgoftaar, "A survey on image data augmentation for deep learning," *J. Big Data*, vol. 6, no. 1, p. 60, Dec. 2019.
- [14] C. Cuevas, E. M. Yáñez, and N. García, "Labeled dataset for integral evaluation of moving object detection algorithms: LASIESTA," *Comput. Vis. Image Understand.*, vol. 152, pp. 103–117, Nov. 2016.
- [15] C. Stauffer and W. E. L. Grimson, "Adaptive background mixture models for real-time tracking," in *Proc. IEEE Comput. Soc. Conf. Comput. Vis. Pattern Recognit.*, vol. 2, Jun. 1999, pp. 246–252.
- [16] A. Elgammal, R. Duraiswami, D. Harwood, and L. S. Davis, "Background and foreground modeling using nonparametric kernel density estimation for visual surveillance," *Proc. IEEE*, vol. 90, no. 7, pp. 1151–1163, Jul. 2002.
- [17] P.-L. St-Charles, G.-A. Bilodeau, and R. Bergevin, "SuBSENSE: A universal change detection method with local adaptive sensitivity," *IEEE Trans. Image Process.*, vol. 24, no. 1, pp. 359–373, Jan. 2015.
- [18] Ş. Işık, K. Özkan, S. Günel, and Ö. N. Gerek, "SWCD: A sliding window and self-regulated learning-based background updating method for change detection in videos," *J. Electron. Imag.*, vol. 27, no. 2, p. 1, Mar. 2018.
- [19] Y. Huang, Q. Jiang, and Y. Qian, "A novel method for video moving object detection using improved independent component analysis," *IEEE Trans. Circuits Syst. Video Technol.*, early access, Sep. 10, 2020, doi: 10.1109/TCSVT.2020.3023175.
- [20] J. H. Giraldo and T. Bouwmans, "GraphBGS: Background subtraction via recovery of graph signals," 2020, *arXiv:2001.06404*. [Online]. Available: <http://arxiv.org/abs/2001.06404>
- [21] J. H. Giraldo and T. Bouwmans, "Semi-supervised background subtraction of unseen videos: Minimization of the total variation of graph signals," in *Proc. IEEE Int. Conf. Image Process. (ICIP)*, Oct. 2020, pp. 3224–3228.
- [22] A. Cioppa, M. V. Droogenbroeck, and M. Braham, "Real-time semantic background subtraction," in *Proc. IEEE Int. Conf. Image Process. (ICIP)*, Oct. 2020, pp. 3214–3218.
- [23] D. Zeng, X. Chen, M. Zhu, M. Goesele, and A. Kuijper, "Background subtraction with real-time semantic segmentation," *IEEE Access*, vol. 7, pp. 153869–153884, 2019.
- [24] H. Zhao, J. Shi, X. Qi, X. Wang, and J. Jia, "Pyramid scene parsing network," in *Proc. IEEE Conf. Comput. Vis. Pattern Recognit. (CVPR)*, Jul. 2017, pp. 2881–2890.
- [25] P. W. Patil, K. M. Biradar, A. Dudhane, and S. Murala, "An end-to-end edge aggregation network for moving object segmentation," in *Proc. IEEE/CVF Conf. Comput. Vis. Pattern Recognit. (CVPR)*, Jun. 2020, pp. 8149–8158.
- [26] M. Mandal, V. Dhar, A. Mishra, S. K. Vipparthi, and M. Abdel-Mottaleb, "3DCD: Scene independent end-to-end spatiotemporal feature learning framework for change detection in unseen videos," *IEEE Trans. Image Process.*, vol. 30, pp. 546–558, 2021.
- [27] M. Mandal and S. K. Vipparthi, "Scene independency matters: An empirical study of scene dependent and scene independent evaluation for CNN-based change detection," *IEEE Trans. Intell. Transp. Syst.*, early access, Oct. 23, 2020, doi: 10.1109/ITITS.2020.3030801.
- [28] J.-Y. Kim and J.-E. Ha, "Foreground objects detection using a fully convolutional network with a background model image and multiple original images," *IEEE Access*, vol. 8, pp. 159864–159878, 2020.
- [29] O. Ronneberger, P. Fischer, and T. Brox, "U-Net: Convolutional networks for biomedical image segmentation," in *Proc. Int. Conf. Med. Image Comput. Comput. Assist. Intervent.* Cham, Switzerland: Springer, 2015, pp. 234–241.
- [30] N. Goyette, P.-M. Jodoin, F. Porikli, J. Konrad, and P. Ishwar, "A novel video dataset for change detection benchmarking," *IEEE Trans. Image Process.*, vol. 23, no. 11, pp. 4663–4679, Nov. 2014.
- [31] N. Xu, L. Yang, Y. Fan, D. Yue, Y. Liang, J. Yang, and T. Huang, "YouTube-VOS: A large-scale video object segmentation benchmark," 2018, *arXiv:1809.03327*. [Online]. Available: <http://arxiv.org/abs/1809.03327>
- [32] L. Yang, Y. Fan, and N. Xu, "Video instance segmentation," in *Proc. IEEE/CVF Int. Conf. Comput. Vis. (ICCV)*, Oct. 2019, pp. 5188–5197.
- [33] 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 *Proc. IEEE Conf. Comput. Vis. Pattern Recognit. (CVPR)*, Jun. 2016, pp. 724–732.
- [34] J. Pont-Tuset, F. Perazzi, S. Caelles, P. Arbeláez, A. Sorkine-Hornung, and L. Van Gool, "The 2017 DAVIS challenge on video object segmentation," 2017, *arXiv:1704.00675*. [Online]. Available: <http://arxiv.org/abs/1704.00675>
- [35] S. Caelles, J. Pont-Tuset, F. Perazzi, A. Montes, K.-K. Maninis, and L. Van Gool, "The 2019 DAVIS challenge on VOS: Unsupervised multi-object segmentation," 2019, *arXiv:1905.00737*. [Online]. Available: <http://arxiv.org/abs/1905.00737>
- [36] L.-C. Chen, G. Papandreou, F. Schroff, and H. Adam, "Rethinking atrous convolution for semantic image segmentation," 2017, *arXiv:1706.05587*. [Online]. Available: <http://arxiv.org/abs/1706.05587>
- [37] L. Maddalena and A. Petrosino, "Towards benchmarking scene background initialization," in *Proc. Int. Conf. Image Anal. Process.* Cham, Switzerland: Springer, 2015, pp. 469–476.
- [38] N. Goyette, P.-M. Jodoin, F. Porikli, J. Konrad, and P. Ishwar, "Changetection.Net: A new change detection benchmark dataset," in *Proc. IEEE Comput. Soc. Conf. Comput. Vis. Pattern Recognit. Workshops*, Jun. 2012, pp. 1–8.
- [39] M. Braham, S. Pierard, and M. Van Droogenbroeck, "Semantic background subtraction," in *Proc. IEEE Int. Conf. Image Process. (ICIP)*, Sep. 2017, pp. 4552–4556.
- [40] S. Bianco, G. Ciocca, and R. Schettini, "How far can you get by combining change detection algorithms?" in *Proc. Int. Conf. Image Anal. Process.* Cham, Switzerland: Springer, 2017, pp. 96–107.
- [41] L. A. Lim and H. Y. Keles, "Foreground segmentation using convolutional neural networks for multiscale feature encoding," *Pattern Recognit. Lett.*, vol. 112, pp. 256–262, Sep. 2018.
- [42] D. Berjón, C. Cuevas, F. Morán, and N. García, "Real-time nonparametric background subtraction with tracking-based foreground update," *Pattern Recognit.*, vol. 74, pp. 156–170, Feb. 2018.
- [43] T. S. F. Haines and T. Xiang, "Background subtraction with Dirichlet process mixture models," *IEEE Trans. Pattern Anal. Mach. Intell.*, vol. 36, no. 4, pp. 670–683, Apr. 2014.
- [44] L. Maddalena and A. Petrosino, "The SOBS algorithm: What are the limits?" in *Proc. IEEE Conf. Comput. Vis. Pattern Recognit. Workshops*, Jun. 2012, pp. 21–26.
- [45] L. Maddalena and A. Petrosino, "A self-organizing approach to background subtraction for visual surveillance applications," *IEEE Trans. Image Process.*, vol. 17, no. 7, pp. 1168–1177, Jul. 2008.



M. OZAN TEZCAN (Member, IEEE) received the B.S. degree in electrical and electronics engineering from Koç University, Istanbul, Turkey, in 2016. He is currently pursuing the Ph.D. degree in electrical and computer engineering with Boston University, Boston, MA, USA. His research interests include computer vision and machine learning.



PRAKASH ISHWAR (Senior Member, IEEE) received the B.Tech. degree in electrical engineering from the Indian Institute of Technology Bombay, in 1996, and the M.S. and Ph.D. degrees in electrical and computer engineering from the University of Illinois Urbana–Champaign, in 1998 and 2002, respectively. After working two years as a Postdoctoral Researcher with the Department of Electrical Engineering and Computer Sciences, University of California, Berkeley,

he joined the Faculty of Boston University, where he is currently a Professor with the Department of Electrical and Computer Engineering and the Division of Systems Engineering and an Affiliate Professor with the Department of Computer Science, College of Arts and Sciences. He has published more than 100 peer-reviewed articles and received numerous grants. He served as an Elected Member for the IEEE Signal Processing Theory and Methods Technical Committee and the IEEE Image, Video, and Multidimensional Signal Processing Technical Committee. He was a recipient of the 2005 NSF CAREER Award, a co-recipient of the AVSS10 Best Paper Award, a co-winner of the ICPR10 Aerial View Activity Classification Challenge, and a coauthor of the Student Best Paper Award in the NIPS14 Workshop on Analysis of Rank Data. He has served as an Associate Editor for the IEEE TRANSACTIONS ON SIGNAL PROCESSING for two terms.



JANUSZ KONRAD (Fellow, IEEE) received the master's degree from the Technical University of Szczecin, Poland, and the Ph.D. degree from McGill University, Montreal, QC, Canada. He is currently a Professor with Boston University, Boston, MA, USA. His research interests include video processing and computer vision, visual sensor networks, stereoscopic and 3-D imaging and displays, and human–computer interfaces. He was a co-recipient of the 2001 IEEE Signal Processing

Magazine Award, the 2004–2005 EURASIP Image Communications Best Paper Award and the AVSS-2010 Best Paper Award. He was the Technical Program Co-Chair of ICIP-2000 and AVSS-2010 as well as the General Chair of AVSS-2013. He has been actively involved in the IEEE Signal Processing Society. He is currently a Senior Editor of the IEEE TRANSACTIONS ON IMAGE PROCESSING, and recently served as a Distinguished Lecturer and as Member-at-Large of its Conference Board. He is also an Area Editor of the EURASIP *Signal Processing: Image Communications* journal.

• • •