Image editing-based data augmentation for illumination-insensitive background subtraction

Dimitrios Sakkos, Edmond S. L. Ho, Hubert P. H. Shum and Garry Elvin

Department of Computer and Information Sciences, Northumbria University, Newcastle upon Tyne, UK

Abstract

Purpose – A core challenge in background subtraction (BGS) is handling videos with sudden illumination changes in consecutive frames. In our pilot study published in, Sakkos:SKIMA 2019, we tackle the problem from a data point-of-view using data augmentation. Our method performs data augmentation that not only creates endless data on the fly but also features semantic transformations of illumination which enhance the generalisation of the model.

Design/methodology/approach – In our pilot study published in SKIMA 2019, the proposed framework successfully simulates flashes and shadows by applying the Euclidean distance transform over a binary mask generated randomly. In this paper, we further enhance the data augmentation framework by proposing new variations in image appearance both locally and globally.

Findings – Experimental results demonstrate the contribution of the synthetics in the ability of the models to perform BGS even when significant illumination changes take place.

Originality/value – Such data augmentation allows us to effectively train an illumination-invariant deep learning model for BGS. We further propose a post-processing method that removes noise from the output binary map of segmentation, resulting in a cleaner, more accurate segmentation map that can generalise to multiple scenes of different conditions. We show that it is possible to train deep learning models even with very limited training samples. The source code of the project is made publicly available at https://github.com/dksakkos/illumination_augmentation

Keywords Background subtraction, Convolutional neural networks, Synthetics, Data augmentation, Illumination-invariant

Paper type Research paper

1. Introduction

The challenge in background subtraction (BGS) is to identify the pixels belonging to the background, which comprises the static areas in image such as the sky and roads, from the foreground, the areas that move against the background such as cars and humans (Barnich and Van Droogenbroeck, 2011). A large number of real-world applications, such as person re-identification (Bazzani *et al.*, 2013), object tracking (Shen *et al.*, 2016), gesture recognition (Yeo *et al.*, 2013), vehicle tracking (Sirikuntamat *et al.*, 2015), video recognition (Li and Yue, 2019), action recognition (Ho *et al.*, 2016), crowd analysis (Wang *et al.*, 2017a) and even use cases from the medical domain (Them *et al.*, 2016; McCay *et al.*, 2019, 2020), depend on accurate and robust BGS as a first step in their pipelines.

Sudden illumination changes provide a particularly difficult challenge, since they cannot be captured by a background model. Such changes in lighting conditions can be caused either by weather conditions or electric lights and result in colour changes involving a significant number of pixels. Due to the difference in visual appearance in consecutive frames,

BGS becomes inaccurate. The timing of these changes could be short, such as switching a light on/off or a piece of cloud blocking the sun, making it tough for the system to adjust to the new condition in a timely manner. In addition, a scene and the objects that appear in it will drastically transform during the night. It is necessary for an algorithm to be able to adjust in these kinds of conditions.

State-of-the-art deep learning algorithms allow adaptation to sudden illumination changes if a huge amount of training data is provided. However, obtaining labelled data is very costly and there are only limited datasets available in the community. As a solution, data augmentation methods are proposed to perform image-based operations on the data, such as mirroring or cropping, to synthesize a larger dataset. However, simple image tricks cannot effectively generate images with realistic illumination changes. Another solution is adding a small amount of noise to create a new, synthetic image that is similar than the original in context but different in colour distribution. A major advantage is that each synthetic image will be unique, due to the added noise being random. However, the downside is that the added noise does not have any semantic meaning. Therefore, although the synthetic images do increase the generalisation of the model simply by obscuring pixels in the original image, they do not offer any additional knowledge regarding differing lighting conditions in the same scene. So the synthetic images only slightly increase the generalisation power of the model.

To overcome this challenge, we propose a new data augmentation technique which synthesises the light-based effects of different degrees of brightness. Such effects include shadows and halos of different size, placed in random locations of the input image. In addition, global illumination changes are also included, in order to increase the generalisation abilities of the model to scenes filmed at various times of the day and night. Such augmented data allows us to provide extra semantic information to the BGS model in terms of illumination for better generalisation performance. The pilot study, published in (Sakkos *et al.*, 2019a), demonstrated the effectiveness and feasibility of such an approach. In this paper, we extend the work by introducing new data augmentation techniques to handle additional variations to the input image locally and globally. A wide range of new experiments are conducted to evaluate the effectiveness of the new framework. The results show that the proposed technique is superior to regular augmentation methods and can significantly boost the segmentation results even in scenes that feature illumination conditions unseen to the model.

We further propose a post-processing method that can successfully remove noise from the output binary map of segmentation. The method is based on the fact that contiguous frames have minimal changes between them and thus, the potential areas of the output that include foreground objects can be limited. Our experiments indicate that the proposed method improves the BGS results in our quantitative and qualitative evaluations on the benchmark dataset SABS (Brutzer *et al.*, 2011).

The main contributions of this work can be summarized as follows:

(1) A novel synthetic image generation method which focuses on local and global image effects for robust BGS under challenging illumination conditions.

(2) An illumination-invariant deep neural network for BGS.

(3) A post-processing technique based on temporal coherence for the refinement of the segmentation results.

The rest of the paper is organised as follows. First, we review related work on BGS and particularly focus on illumination-aware systems in Section 2. Second, we present how we synthesize images by including local, global and combined illumination changes and present our new post-processing technique, in Section 3. Third, we present the dataset we created and

explain how we train the network in four. Next we present the experimental results and discuss the performance of our proposed methods in Section 5. Finally, we conclude the paper and discuss future directions in Section 6.

2. Related work
In this section, we review works that are related to this paper. We first review traditional approaches of BGS, which involves statistical models like Gaussian Mixture Models and Principal Component Analysis (PCA). We then review BGS method that utilizes deep-learning, with a particular focus on supervised methods that yield promising performance. We finally look at the problem of illumination in BGS and existing solutions, which motivate our research.

2.1 Traditional approaches
BGS in video is a popular research topic, and the manipulation of illumination to improve accuracy has many researches. Siva *et al.* (2015) demonstrated that the pixel intensity values affected by sudden local illumination change can be modelled by combining a GMM with a conditional probabilistic function based on an extension of Zivkovic *et al.* (Zivkovic and Heijden, 2006). Boulmerka and Allili (Boulmerka and Said Allili, 2018) combine a GMM with inter-frame correlation analysis and histogram matching. Chen *et al.* (2018a) use a number of GMMs to construct spanning trees for hierarchical superpixel segmentation. They report that extending their model with optical flow for modelling temporal information increases the segmentation accuracy. Shen *et al.* (2016) propose an efficient approach to BGS by reducing the dimensionality of the input data with a random projection matrix. Finally, they apply a GMM on the projected data.

PCA-related techniques are used for modelling the background of a video with an eigenspace. Since PCA retains the most significant eigenvectors, the foreground of the input image cannot be represented by the background model, as long as it is not static. The foreground can then be recovered with a difference image between the output of the model and the input frame (Oliver *et al.*, 2000). Candès *et al.* (2011) developed an efficient algorithm (RPCA) for decomposing the data into a low-rank matrix and a sparse matrix, which are representing the background and foreground in the BGS scenario, respectively. Recently, Ibadi and Isquierdo (Erfanian Ebadi and Izquierdo, 2018) extended RPCA by using a tree-structured sparse matrix to represent the input images. Although their method performs well on standard datasets, it fails in videos with sudden illumination changes like the *Light Switch* sequence of the *SABS* dataset.

The major weakness of GMM-based illumination-aware methods is that it cannot model significant illumination changes across consecutive frames as mentioned in Andrews and Antoine (2014), which happens frequently in real-world environments such as switching on and off the light. At the same time, while PCA methods have better robustness in modelling illumination changes, it lacks the semantic knowledge of the scene, resulting in a sub-optimal performance.

2.2 Deep-learning based approaches
Deep learning has improved system performance significantly in many areas. Here, we review how deep learning techniques and how they have been used in BGS in the past.

Deep learning approaches use variants of the fully convolutional network (FCN) proposed by Long *et al.* (2015). This is a special kind of convolutional neural networks with no fully connected layers, specifically designed for dense prediction tasks such as image segmentation. Most BGS methods follow the trend of recent generic image segmentation networks (He *et al.*, 2017; Chen *et al.*, 2018b; Zhao *et al.*, 2017; Yu *et al.*, 2017; Wang *et al.*, 2018)

and treat videos as a collection of images while disregarding the temporal information. Following the success of earlier approaches in object detection (Cai *et al.*, 2016), image dehazing (Ren *et al.*, 2016), segmentation (Roy *et al.*, 2016) etc., Lim and Keles (2018) and Zeng and Zhu (2018) attempt to improve their binary maps by employing multi-scale feature aggregation. While Zeng and Zhu (2018) realise this idea simply by concatenating features from different layers, Lim and Keles (2018) employ multi-scale inputs, as previously done by Lu (2014). Wang *et al.* (2017b) also adopt the same preprocessing, but they refine the original CNN output by feeding it into another CNN.

A 3D convolution-based approached is proposed by Sakkos *et al.* Sakkos *et al.* (2018) to exploit the relationship between a block of ten-frame for BGS tasks. In Berjón *et al.* (2018), the background model of the Kernel Density Estimation-based system is updated using information from previous frames. Group property information is exploited in both spatial and temporal domains in the sparse signal recovery-based approached proposed by Liu *et al.* (2018). A recent work (Javed *et al.*, 2018) further demonstrated incorporating spatio-temporal constraints to improve (Erfanian Ebadi and Izquierdo, 2018) results in better performance. We also employ deep-learning to perform BGS. However, we particularly focus on the problem of poor performance in different illumination conditions and propose methods to tackle it.

2.3 Dealing with illumination

A major challenge for BGS is the illumination conditions, in which pixels belonging to the same object may look different under different illuminations. At the same time, deep-learning requires a huge amount of training data and annotating the lighting condition could be difficult and time-consuming. While existing deep-learning networks may automatically predict illumination information (Gardner *et al.*, 2017) and generate the illumination map (Song and Funkhouser, 2019; Gardner *et al.*, 2019), such approaches are typically used indoor with a limited number of light sources.

To solve this problem, Sakkos *et al.* Sakkos *et al.* (2019b) trained a multi-task generative adversarial network that modelled pairwise dark and bright images. The core idea is to apply gamma correction (Poynton, 1993) to synthesize images under different illuminations, with a generative adversarial network (Goodfellow *et al.*, 2014) attempting to learn the data probability distribution for generating such images. A major advantage is that the method can synthesize pairwise images of different lighting conditions for effective training in both indoor and outdoor scene. However, such a method does not explicitly model the light sources that consist of different properties (Goesele *et al.*, 2003), and therefore generate images that may not necessarily realistic. In this work, we solve the problem with a data augmentation approach, in which we perform image-based modelling of the light sources to synthesize different lighting conditions. Such a method create more realistic training images such that even a simple VGG16 (Simonyan and Zisserman, 2015) network can perform highly accurate BGS.

3. Methodology

In this section, we introduce the proposed data augmentation approaches and explain how a wide range of local (Section 3.1) and global (Section 3.2) illumination change effects can be synthesised using the proposed system. Our local illumination methods focus on editing a local region to simulate lamp-post and shadow effects caused by light sources from different angles (e.g. Plate 1 and Plate 2 (a) and (b)). For global illumination change effects, we propose editing the entire image by altering the contrast, sharpness and image saturation (for example, Plate 2 (c) and (d), and Plate 3). In addition, we present a post-processing output refinement method (Section 3.4) that takes into account temporal information to further

enhance segmentation results. Finally, we incorporate the data augmentation and output refinement approaches into a unified deep neural network architecture 3.5 to perform BGS. For the sake of clarity, all the symbols used in this paper are listed in Table 1.

3.1 Local changes

Light sources can roughly be divided into *direct* and *indirect* lights. A direct light refers to light falling on a specific area or the surface of an object. In this research, we define this as a local illumination change since only a small region in the image is being affected. Our pilot study (Sakkos *et al.*, 2019a) demonstrated the effectiveness of an efficient approach to synthesize images with local illumination change. In this section, we first review the technique proposed in Sakkos *et al.* (2019a). Next, we present a new local illumination augmentation approach to simulate a direct light source from different angles, including lamp-post and shadow effects.

3.1.1 Lamp-post and shadow effects. The shape of different light sources can vary significantly. However, circular shapes, such as street lights and light bulbs are common. In Sakkos *et al.* (2019a), we proposed to edit the intensity of a circular region on an image to simulate a "lamppost" light source (i.e. by increasing the pixel value) or a shadow effect (i.e. by decreasing the pixel value). Firstly, our method generates the centre i of the circular region randomly:

$$i = I(w, h), w \in W, h \in H, I = W \times H \tag{1}$$

where W, H the width and height of the input image I respectively.

Next, the diameter d of the circular region is again determined randomly in order to create different illumination change effects. The range of the diameter size is defined as follow:

$$d = k \times \min(W, H), k \in \left(\frac{1}{5}, \frac{1}{2}\right). \tag{2}$$

where k is the kernel size of the binary mask M_1.

To replicate a realistic fading out effect near the edge of the circular region, we gradually decrease the brightening/darkening effect from the centre to the edge (i.e. the attenuation of light).

| The mask M_1 | The mask M_2 | Original image | After effect |
| (a) | (b) | (c) | (d) |

| Ellipsoid-shaped shadow | Ellipsoid-shaped flash of light | Low saturation | High saturation |
| (a) | (b) | (c) | (d) |

Specifically, we first calculate the binary mask M_1 of the pixels to be altered using the following formula:

$$M_1(x,y) = 1 \Leftrightarrow (x - w)^2 + (y - h)^2 \leq d^2 \tag{3}$$

Consequently, all pixels inside the circular region on the mask image will have the value of 1 and zero everywhere else. The light attenuation effect can then be simulated by using the Euclidean Distance Transform (EDT). Given a binary mask B, EDT is defined as:

$$EDT_x(B) = \min_b \left(\|x - b\|_{L_2} \right), \quad \forall b \in B, \tag{4}$$

where L_2 is the Euclidean norm. By applying the EDT on M_1, the mask for local changes M_2 can be calculated as follows:

$$M_2 = \text{EDT}(M_1) \tag{5}$$

Having created the new mask for the local change, the intensity of the pixel inside the circular region (i.e. the masking region) on the image will be edited using the formula:

$$I_s = I \pm (M_2 \times z), \quad z \in [120, 160], \tag{6}$$

where I and I_s are the original and new synthetic image, respectively. z a random integer to further produce a wider variety of synthetic images, and $\pm$ is either pixel-wise addition or subtraction, chosen with equal probability. In summary, a lamp-post effect can be simulated by using the addition operation in Eq. (6). On the other hand, shadow effects can be created when the subtraction operation is used. An example of the circular masks M_1, M_2 and the image editing effect are illustrated in Plate 1.

3.1.2 Illumination angle effects. Although the method described above is very effective in creating a very realistic "lamp-post" lighting effect, it does not cover all variations of flashes and shadows. In reality, many lights often shine at an angle, resulting in a lit area that is not perfectly circular. As a result, the circular-shaped lighting effect has to be distorted and transformed in order to simulate the real-world effect. In computer graphics, such an effect can be created by estimating the 3D position and orientation of the light source in the virtual world, as well as the location and 3D shape of the objects in the scene. A realistically lit image can be generated by rendering the scene using techniques such as ray-tracing (Glassner, 1989). However, the aforementioned approach is computationally costly and requires detailed 3D information about the scene which is not available in 2D video and images. On the other hand, from our observation, the light effect closely resembles an ellipse rather than a circle in most cases. This motivates us to improve the realism of the lighting effect by using different mask shapes directly in the image space.

In particular, we propose using an ellipse-shaped mask in our simulation to cover these cases. Such a light-weight approach can effectively improve the realism of the lighting effect while minimizing the additional computation required for the framework in the training stage. More specifically, the following formula is used instead:

$$M_1(x,y) = 1 \Leftrightarrow \frac{(x - w)^2}{d^2} + \frac{(y - h)^2}{d_0^2} \leq 1 \tag{7}$$

where $d_0 = a_d * d, \quad a_d \in [0.3, 0.8]$. We also randomly rotate the ellipse along its axis.

3.2 Global changes

In some cases, global illumination changes can occur. For example, lightning during a storm may instantly increase brightness, and once the rain is over global illumination will change

again. In order to model such illumination changes, we need to alter the pixels across the whole image, rather than in a small patch.

We synthesize global illumination changes as:

$$I_s = I \pm z, \quad z \in [40, 80], \tag{8}$$

where I, z and $\pm$ are as previously defined. In this case the illumination noise z needs to be slightly diminished, since the whole image is affected.

3.2.1 Contrast augmentation. In addition to image brightness augmentation, we further enhance the variance of our global augmentation settings with contrast changes. The contrast of an image plays an important role in highlighting different objects in the scene. Low contrast images usually look softer and flatter, as well as lacking shadows and highlights. In reality, various occurrences can result in low contrast images. One of the common situations is lens flare in the image, where a bright light source scatters the light directly into the lens. Inspired by this observation, we propose a new data augmentation approach that varies the contrast of the image to improve the robustness of the framework. Specifically, we alter the contrast of the original image by applying the following formula:

$$I_s = 128 + c^*(I - 128) \tag{9}$$

where c is the contrast factor. For $c < 1$ the contrast is decreased; conversely we can increase it by setting $c > 1$. In our experiments, we let $c \in [0.2, 2]$. Example images are shown in Figure 2.

3.2.2 Sharpness augmentation. Finding the sharp borders between different objects in the image will provide a clear separation between foreground and background and will certainly contribute positively toward the BGS task. In contrast, the blurring effects caused by low illumination as well as motion blur will have a negative impact. Inspired by this observation, we propose incorporating sharpness/blurring in the data augmentation algorithm.

Blurring can be easily achieved by convolving the original image with a $n \times n$ low-pass filter, which is averaging the neighbouring pixels of the input image. Specifically, we use $F_b = \frac{1}{25} \times [1]^{5x5}$.

Sharpening an image can be done in the same manner, with the use of the filter

$$F_s = \frac{1}{9} \begin{bmatrix} 0 & -1 & 0 \\ -1 & 5 & -1 \\ 0 & -1 & 0 \end{bmatrix} \tag{10}$$

An example of the edited images is illustrated in Plate 3c and 3d.

| convolution | Pooling | Upsampling | Decoder Block | Sigmoid | ⊕ Concatenation |

Overview
(a)

| Decoder Block | Convolution | Batch Norm | ReLu |

Decoder block
(b)

Figure 1.
The CNN was used for the experiments. The encoder is initialised from VGG16 (Simonyan and Zisserman, 2015) and is keep fixed during training. ReLu layers are used after every convolution and are omitted from Figure (a) for clarity

Figure 2.
F-Measure values on
different thresholds for
each model

3.2.3 Colour saturation. The colour saturation of an image refers to the intensity of the colour. The higher the saturation, the more colourful the image is. On the other hand, an image resembles a grey-scale one when the saturation is very low. Such a difference in image appearance is similar to the situation when illumination changes significantly. Here, we propose to further incorporate colour saturation in the data augmentation process. This improves the framework by enhancing its robustness against significant changes in the colour of the image caused by illumination changes.

We edit the saturation of the whole image by converting it from the RGB to the HSV colour space and directly changing the saturation attribute. Specifically, we scale the second dimension of the HSV space which corresponds to saturation using a parameter $s \in [0, 2]$. For $s < 1$, colour is diminished; conversely for $s > 1$ the colours become more saturated. An example of the edited image is illustrated in Plate 2c and 2d.

3.3 Combined changes

To capture both local and global illumination changes in the scene, we combine Eq. (6) and Eq. (8) into the following:

$$I_s = z_1 \pm (I \pm (M_2 \times z_2)), z_1 \in [40, 80], z_2 \in [120, 160] \qquad (11)$$

Sample images synthesised from our system can be found in Plate 4. Since both the positioning and the intensity of the masks are random, this method can effectively cover all kinds of illumination changes. Additionally, hundreds of different synthetic images can be generated from a single frame. Therefore, given a small video, we can generate enough unique synthetic images to train a very deep network.

3.4 Output refinement via temporal coherence

While the proposed augmentation method works for still images and videos alike, in the latter case we can exploit motion information to refine the segmentation results as follows:


Low contrast	High contrast	Blurring	Sharpening
(a)	**(b)**	**(c)**	**(d)**

Plate 3.
Global changes include contrast augmentation, blur and sharpness. Examples of after decreasing and increasing the effects are shown above

| $G_b L_b$	$G_b L_d$	$G_d L_b$	$G_d L_d$
(a)	**(b)**	**(c)**	**(d)**

Plate 4.
Combination of global and local illumination changes. The subfigures (a) and (b) depict a combination of a brightening global filter with a bright and dark local filter respectively. On the other hand, subfigures (c) and (d) implement the darkening global filter

Lemma 3.1. Let $o_t = (i, j)$ be a pixel of an object at time t. Then, the corresponding $o_{t+1} \in \{(i \pm \delta i, j \pm \delta j)\}$, where δ is a small integer.

Based on this, we can create a refining map to highlight the areas of the input image that are likely to contain pixels of the foreground in the next frame. The map acts as a weight matrix that refines the probabilities of each pixel of the model output. Essentially, this refining map needs to be designed in a way that the predicted foreground probability of the next frame is not scaled down. This is a desired property since the change between two contiguous frames is minimal and most foreground pixels remain in the same class. Secondly, those pixels of the refining mask that are adjacent to foreground pixels need to be assigned with a probability value very close to 1, as it is highly possible for the foreground object to move into this area. As the distance becomes larger, the values need to be gradually scaled down. Eventually, the pixels that are furthest away from the foreground have the smallest probability.

Given each timestamp t and a video frame F_t, we can construct the refining mask r_t in the following way: First, we obtain the model output p_t, the pixels of which represent the probability of them belonging to the foreground class. Then, r_t can be generated by applying the Euclidean distance transform on p_t:

$$r_t = \text{EDT}(p_t), \tag{12}$$

where EDT is as defined in Eq. (4).

While this is a valid approach for existing moving objects, we need to account for new objects entering the scene at any moment. As a result, we set the values of r_t located around the border to 1. Therefore, the mask does not penalise new objects entering the frame. The end result of the refining mask r_t is depicted in Plate 5.

Having defined the process of creating the probability map, the refinement is performed in a post-processing manner. During testing, we obtain the model output of the current frame and calculate the probability map, which is used to filter the model output on the next frame. Thus, p_{t+1} can be refined by scaling its probability values according to r_t as follows:

$$p_{t+1}' = p_{t+1} \times r_t, \tag{13}$$

where $\times$ is the pixel-wise multiplication operator and p_{t+1}' denotes the refined segmentation result.

### 3.5 Illumination-invariant deep networks

In this work, we incorporate the proposed local and global data augmentation techniques into a deep learning framework for BGS. To keep all other variables fixed, the same network architecture is used for all experiments. We use a VGG16 backbone (Simonyan and Zisserman, 2015), which is transformed to a fully convolutional network by removing the fully connected layers and appending a decoder. As seen in Figure 1, the decoder mirrors the

Plate 5.
The (a) probability mask that is used for refining the output, created from the (b) segmentation mask of the previous frame. Bright colours indicate high probability, whereas dark colours represent low probability values

Segmentation mask of the previous frame

Probabilistic mask

encoder and has the same number of channels. Additionally, the encoder features are concatenated with those of the decoder which are of the same size to enable information flow. With the proposed illumination-focused data augmentation techniques, the proposed model is more robust on inputs with significant illumination changes. The experimental results are presented in Section 5.

4. Experiment settings

4.1 Dataset

In this work, a wide range of illumination-focus data augmentation techniques is proposed. In order to evaluate the performance of the new methods on BGS tasks, a dataset with significant illumination changes is needed. In particular, the Stuttgart Artificial Background Subtraction dataset (SABS) (Brutzer *et al.*, 2011) fulfil all the requirements and it is used in other illumination-aware BGS approaches in the literature (Sakkos *et al.*, 2019a, 2019b).

The SABS dataset (Brutzer *et al.*, 2011) contains videos with challenging illumination conditions and makes the BGS difficult. There is a wide range of environmental lighting conditions (such as day-time and night scenes), as well as other light sources (such as switching on/off the lights inside the shops in the street scenes). Following our pilot study (Sakkos *et al.*, 2019a), the sequence *Darkening* is used for training our models in all of our experiments. The *Light Switch* video is then used as the unseen testing sequence. In Plate 7, a number of sample training and testing frames are illustrated. Note that we did not use the

Plate 6.
Default augmentation techniques (from left to right): image mirroring, centre cropping and adding noise


Plate 7.
The SABS dataset used for evaluating the models. The first row depicts the training sequence *Darkening*, while the second row shows the testing video *LightSwitch*. The columns show frames from the start, middle and ending parts of the video. Note that in the middle of the *LightSwitch* sequence the store light switches off, causing major changes to the background

rest of the SABS dataset (Brutzer *et al.*, 2011) since those videos do not contain any significant illumination changes and thus not suitable for evaluating our method (see Table 1).

Since the images are generated on-the-fly during training and are not saved on disk, the training dataset is slightly different for each experiment. However, when the same parameters are used, those differences become non-significant due to a large amount of generated data. The list of experiments and their hyper-parameters is given in Table 2

Symbol	Description
I	Input image
W	Width of I
H	Height of I
i	Centre of the local mask
d	Diameter of the local mask
M_1	Initial mask for local changes
K	Kernel size of the mask M_1
Z	Illumination intensity in terms of pixel values
M_2	Adjusted mask for local changes
c	Contrast factor
F_b	Blurring filter
F_s	Sharpening filter
S	Saturation parameter
C	Contrast parameter
R	Post-processing mask
P	Model prediction
p'	Refined model prediction
N	Number of pixels of all input frames

Table 1. Table of symbols

Name	Description	Threshold
Baseline	No augmentation	0.8
Default	Common augmentation: Mirror, crop and noise	0.7
L_a	Local changes with $z \in (80, 120), k \in (1/2, 2/3) \times G$	0.7
L_b	Local changes with $z \in (80, 120), k \in (1/5, 1/2) \times G$	0.7
L_c	Local changes with $z \in (120, 160), k \in (1/5, 1/2) \times G$	0.6
G_{low}	Global, low intensity changes with $z \in (20, 60)$	0.9
G_{med}	Global, medium intensity changes with $z \in (40, 80)$	0.6
G_{high}	Global, high intensity changes with $z \in (60, 100)$	0.8
GL	Global and local changes with $z_{global} \in (40, 80)$ and $z_{local} \in (120, 160)$	0.7
GL_{refine}	The GL model, after applying the post-processing method	0.6
GL_{sb}	GL plus sharpening and blurring augmentation	0.3
GL_s	GL plus colour saturation augmentation	0.3
GL_c	GL plus contrast augmentation	0.1
GL_e	GL with ellipsoid-shaped masks for local changes	0.3
GL_{all}	GL plus all the above	0.3
GL_{AD}	GL plus all the above plus default augmentation	0.2

Note(s): Parameters k, z and G denote the kernel size of the mask M_1, the illumination intensity in terms of pixel values and the resolution of the smallest dimension of the input image respectively. The last column shows the threshold that maximised the F-Measure of the segmentation mask

Table 2. The different augmentation settings that were tested in our experiments

4.2 Training parameters
In this section, the parameters used in training the deep learning models are explained. Firstly, a mini-batch approach with the batch size set to 1 is used. Next, we used the Adam optimiser (Diederik, 2014) with betas $b_1 = 0.9$ and $b_2 = 0.999$. Thirdly, to avoid overfitting, the training process is terminated if there is no improvement after 5 epochs. The initial learning rate is $lr = 0.001$ and is reduced by a factor of 0.1 if the model does not improve for 2 epochs.

As another measure against overfitting, we freeze the encoder of our network. Specifically, the first five convolutional blocks of VGG16 are fixed and we only train the decoder. Therefore, although the total number of parameters is around 19M, we only train 4.3M.

The optimal ratio between the unaltered, original images and the new, augmented training samples is not the same for every problem. It can depend on the size of the training set and the variance of the images' appearance. In our case, we find that augmenting two-thirds of our training set yields the best results.

We observe that most frames contain more pixels of the background than the foreground – some frames might not even depict any moving objects at all. Considering this, we conclude that the loss function needs to balance the classes as to not allow the model to be biased towards the background class. As a solution, we use the weighted cross-entropy loss, which is formally defined as follows:

$$G_s = wt[-\log\sigma(x)] + (1 - t)[-\log(1 - \sigma(x))], \tag{14}$$

where w is the weight coefficient, x is the predicted label, t is the target label and $\sigma(x) = \frac{1}{1+e^{-x}}$ is a sigmoid function. The weight w is calculated according to the ground truth frames with the following formula:

$$w = \frac{N}{2 \times [N_b, N_f]}, \tag{15}$$

where N denotes the number of pixels of all input frames and N_b, N_f are those pixels that belong to the background and foreground respectively.

4.3 Implementation details
We use the Keras library (Chollet *et al.*, 2015) for training our models. Furthermore, for the quick deployment of the proposed model, the *Segmentation models* (Yakubovskiy, 2019) library is used. The full code is uploaded on GitHub [1]. The Graphics Processing Unit (GPU) that was used in all our experiments is a GeForce GTX TITAN *X*.

4.4 Evaluation metric
For evaluating our experiments, we use the following metrics: *F-Measure (FM), Intersection over Union (IoU), Matthews correlation (MC)*. We provide the formal definitions below:

$$\text{Precision} = \frac{\text{TP}}{\text{TP} + \text{FP}} \tag{16}$$

$$\text{Recall} = \frac{\text{TP}}{\text{TP} + \text{FN}} \tag{17}$$

$$\text{FM} = \frac{2 \times \text{Precision} \times \text{Recall}}{\text{Precision} + \text{Recall}} \tag{18}$$

$$\text{IoU} = \frac{\text{TP}}{\text{TP} + \text{FP} + \text{FN}} \tag{19}$$

$$MC = \frac{TP \times TN - FP \times FN}{\sqrt{(TP + FP)(TP + FN)(TN + FP)(TN + FN)}} \tag{20}$$

where TP, TN, FP, FN denote the true positive, true negative, false positive and false negative pixels respectively.

5. Results

We perform extensive evaluations on the proposed method. In particular, a wide range of different augmentation settings (Table 2) were evaluated. We also compare against the regular augmentation techniques. We implement a "default" augmenter which performs the following image transformations: *horizontal flipping, random cropping* and *noise addition*, as depicted in Plate 6. The *cropping* operation performs centre cropping with random image sizes, whereas the *noise* option adds salt and pepper noise drawn from a Gaussian distribution. The amount of noise is fixed to 0.05. To maximise the variance of the augmented data's appearance, all operations have a 50% probability of taking place.

In the following section, we will evaluate the proposed method quantitatively to determine the optimal settings.

5.1 Quantitative evaluations

To quantitatively evaluate the performance of the proposed data augmentation techniques as well as the new post-processing method, we follow the commonly used metrics as in the previous work on the SABS dataset. The details are stated in section 4.4.

The results of our experiments are presented in Table 3 and 4. We first review the performance of the basic local and global data augmentation approaches presented in our pilot study (Sakkos *et al.*, 2019a). In Table 3, the common augmentation approach achieved a 7% improvement in the main evaluation metric F-Measure (FM) over the baseline which does not employ any data augmentation, which indicates that general augmentation methods can improve results significantly. We further presented the results obtained by our combined (global and local) data augmentation, namely *GL*,

Table 3.
Comparison between no augmentation, common augmentation and method proposed in our pilot study (Sakkos *et al.*, 2019a) which covers global and local illumination changes

Settings	Recall ↑	Sp ↑	FPR ↓	FNR ↓	PWC ↓	FM ↑	Precision ↑	IoU ↑	Matthews ↑
No augm	0.4606	0.9933	0.0067	0.5394	1.9172	0.5288	0.6207	0.3594	0.5253
Common augm	0.5440	0.9937	0.0063	0.4560	1.6767	0.6025	0.6750	0.4311	0.5976
GL (Sakkos *et al.*, 2019a)	0.7687	0.9941	0.0059	0.2313	1.1189	0.7624	0.7562	0.6161	0.7567

Settings	Recall ↑	Sp ↑	FPR ↓	FNR ↓	PWC ↓	FM ↑	Precision ↑	IoU ↑	Matthews ↑
Ellipsoid masks	0.7806	0.9942	0.0058	0.2194	1.0836	0.7709	0.7614	0.6272	0.7654
Sharp/Blur	0.7925	0.9941	0.0059	0.2075	1.0587	0.7776	0.7633	0.6361	0.7723
Saturation	0.7813	0.9945	0.0055	0.2187	1.0515	0.7763	0.7714	0.6344	0.7710
Contrast	0.7752	0.9955	0.0045	0.2248	0.9601	0.7904	0.8062	0.6535	0.7857
All	0.8097	0.9948	0.0052	0.1903	0.9478	0.7996	0.7898	0.6661	0.7948
All + Default	0.8191	0.9949	0.0051	0.1809	0.9211	0.8060	0.7932	0.6750	0.8014

Table 4.
Individual contributions of new augmentation methods

proposed in our pilot study (Sakkos *et al.*, 2019a). Our method achieved an outstanding 76.24%, which is significantly higher than the common augmentation approach by 16% and the baseline by 23%. Our method also outperformed other methods in every single metric. These results highlight the importance of targeted, task-specific data augmentation and demonstrate the superiority of the proposed method against illumination-agnostic augmentation.

Next, we focus on the new data augmentation approaches proposed in this work. The results, which are depicted in Table 4, show that all four of the newly proposed methods further improve the model performance in segmenting the foreground objects. The F-Measure values are ranging from 77.09% to 79.04%, a very good performance which is further increased to 79.96% when we combine all new methods. Finally, when the proposed augmentation methods are combined with the default ones, the model accuracy reaches an excellent 80.6%. That is an improvement of 4.36% FM score than *GL* proposed in our pilot study (Sakkos *et al.*, 2019a). Therefore, we can deduce that each new method introduces modifications to the training data that offer improvements in different areas, and also all methods complement each other.

We further evaluate the performance of the proposed post-processing method and the results are presented in Table 5. Evidently, our method improves the result of all experiments. This improvement fluctuates between 0.1 and 0.84% with an average of 0.65%. Our best model reaches an F-Measure score of 81.27%. This shows that the post-processing can further boost the performance even the new augmentation techniques have already achieved a high accuracy level.

5.2 Qualitative evaluations

In this section, we visualize the segmentation results to evaluate the proposed augmentation methods qualitatively. We picked three representative frames from the start, middle and end in the testing video sequence for a fair comparison. The selected frames have different illumination conditions which allow us to evaluate the performance of the methods in all situations. The BGS results depicted in Plates 8 and 9 show that the proposed augmentation approaches generated higher quality segmentation masks. In particular, local illumination augmentation leads to masks with much fewer false positives, effectively suppressing noise. On the other hand, global augmentation offers significant improvement on the true positives, while discarding some noise as well. However, the results are further improved when combining global and local augmentation, with the model predictions showing minimal noise and being accurate even when the foreground is very dark. In this case, the contrast and colour/brightness difference between the foreground and background objects is very low, therefore our contrast and blurring augmentation helps the model segment the foreground more accurately.

Method	FM	Post-processed
Baseline (GL) (Sakkos *et al.*, 2019a)	0.7624	0.7697
Sharp/Blur	0.7776	0.78527
Saturation	0.7763	0.7846
Contrast	0.7904	0.7988
Ellipsoid masks	0.7709	0.7719
All	0.7996	0.8055
All + Default	0.8060	0.8127

Table 5.
Accuracy improvements of our post-processing method. Numbers represent FM scores

input frames

no augmentation

default augmentation

local changes

global changes

global and local

ground truths

Plate 8.
Comparison between
different augmentation
techniques

input frames

Ellipsoid masks

Saturation

Sharpening/Blurring

Contrast

All

All+Default

5.3 Ablation studies
In this section, we justify the selection of hyperparameters used in this work. A comprehensive list of the experiments and the corresponding settings are presented in Table 2. In general, we investigate the kernel size used for local augmentation, in addition to the range of the pixel intensity change. Furthermore, we provide the optimal threshold which maximises the model performance. The quantitative results of ablation testing can be found in Table 6, while the performance of each model under different threshold is depicted in Figure 2.

6. Conclusion and future work
In this paper, we have proposed a data augmentation framework to generate structured training data for BGS on videos with significant lighting changes. Specifically, to improve the realism of image synthesis, we have proposed to separate the illumination changes into local and global components, and we proposed some novel designs to effectively model the respective illumination effects. We have further proposed a post-process technique to refine the background mask for generating more accurate results. Our framework has improved the training process and generalisation of a deep neural network for BGS, as it can provide an unlimited amount of training data that represents challenging illumination conditions. Experimental results have shown our method outperforming existing work, and achieving the highest score of 81.27% in all comparative setups.

There are several interesting future directions for this research.

Firstly, the geometric information of the objects in the scenes directly affects the illuminations. While we currently model the effects by adjusting the intensity of the pixels, the results can potentially be improved if we explicitly model the geometry and deduce the occlusion information between objects. This will allow us to generate more realistic lighting and shadowing effects.

Secondly, the location of applying the local effects is randomly selected in our current design. This has the advantage of improving the robustness of the system. However, this also has the disadvantage of creating scenes that may be semantically incorrect. We would like to learn the correlations between illumination effects and the locations to apply them from real-world images.

Finally, under the current design, the data augmentation is used as a pre-processing method for creating more training data. On one hand, this is advantageous as the computation cost of training the deep network becomes independent of that of the data augmentation process. Therefore, the augmented data can be pre-computed and reused repeatedly. However, it is likely that the generated training data might not be the most effective samples for training the network and maximising its generalisation capabilities. From this point of view, coupling image generation with training and optimising both tasks makes sense. Hence, in the future, we would like to explore methods like adversarial training and measure their impact on data augmentation.

Settings	Recall ↑	Sp ↑	FPR ↓	FNR ↓	PWC ↓	FM ↑	Precision ↑	IoU ↑	Matthews ↑
			(a) Ablation studies for global changes						
L_a	0.5467	0.9962	0.0038	0.4533	1.4290	0.6412	0.7752	0.4719	0.6442
L_b	0.5958	0.9951	0.0049	0.4042	1.4219	0.6619	0.7444	0.4946	0.6589
L_c	0.6294	0.9954	0.0046	0.3706	1.3189	0.6903	0.7643	0.5271	0.6870
			(b) Ablation studies for local and global changes						
G_{low}	0.7103	0.9927	0.0073	0.2897	1.3877	0.7051	0.6999	0.5445	0.6980
G_{med}	0.7082	0.9942	0.0058	0.2918	1.2464	0.7263	0.7454	0.5703	0.7202
G_{high}	0.6679	0.9952	0.0048	0.3321	1.2405	0.7155	0.7704	0.5570	0.7111

Table 6.
Ablation studies for local changes

Note

1 https://github.com/dksakkos/illumination_augmentation

References

Andrews, S. and Antoine, V. (2014), "A comprehensive review of background subtraction algorithms evaluated with synthetic and real videos", *Computer Vision and Image Understanding*, Vol. 122, pp. 4-21.

Barnich, O. and Van Droogenbroeck, M. (2011), "ViBe: a universal background subtraction algorithm for video sequences", *IEEE Transactions on Image Processing*, Vol. 20 No. 6, pp. 1709-1724.

Bazzani, L., Cristani, M. and Murino, V. (2013), "Symmetry-driven accumulation of local features for human characterization and re-identification", *Computer Vision and Image Understanding*, Vol. 117 No. 2, pp. 130-144.

Berjón, D., Cuevas, C., Morán, F. and García, N. (2018), "Real-time nonparametric background subtraction with tracking-based foreground update", *Pattern Recognition*, Vol. 74, pp. 156-170.

Boulmerka, A. and Said Allili, M. (2018), "Foreground segmentation in videos combining general Gaussian mixture modeling and spatial information", *IEEE Transactions on Circuits and Systems for Video Technology*, Vol. 28 No. 6, pp. 1330-1345.

Brutzer, S., Hoferlin, B. and Gunther, H. (2011), "Evaluation of background subtraction techniques for video surveillance", *CVPR 2011, Providence, RI*, pp. 1937-1944, doi: 10.1109/CVPR.2011.5995508.

Cai, Z., Fan, Q., Feris, R.S. and Vasconcelos, N. (2016), "A unified multi-scale deep convolutional neural network for fast object detection", in Leibe, B., Matas, J., Sebe, N. and Welling, M. (Eds), *Computer Vision – ECCV*, Springer, Cham, pp. 354-370.

Candès, E.J., Li, X., Ma, Y. and Wright, J. (2011), "Robust principal component analysis?", *Journal of the ACM*, Vol. 58 No. 3, pp. 1-37.

Chen, M., Xing, W., Yang, Q., Li, Q., Wang, G. and Yang, M.-H. (2018), "Spatiotemporal GMM for background subtraction with superpixel hierarchy", *IEEE Transactions on Pattern Analysis and Machine Intelligence*, Vol. 40 No. 6, pp. 1518-1525.

Chen, L.-C., Papandreou, G., Kokkinos, I., Murphy, K. and Yuille, A.L. (2018), "Symmetry-driven accumulation of local features for human characterization and re-identification", *IEEE Transactions on Pattern Analysis and Machine Intelligence*, Vol. 40 No. 4, pp. 834-848.

Chollet, F., *et al.* (2015), "Keras", available at: https://keras.io.

Diederik, P. (2014), "Kingma and Jimmy Ba. Adam: a method for Stochastic optimization", *arXiv e-prints*, arXiv:1412.6980.

Erfanian Ebadi, S. and Izquierdo, E. (2018), "Foreground segmentation with tree-structured sparse RPCA", *IEEE Transactions on Pattern Analysis and Machine Intelligence*, Vol. 40, pp. 2273-2280.

Gardner, M.-A., Sunkavalli, K., Yumer, E., Shen, X., Gambaretto, E., Gagné, C. and Lalonde, J.-F. (2017), "Learning to predict indoor illumination from a single image", *ACM Transactions on Graphics*, Vol. 36 No. 6, article no. 176, doi: 10.1145/3130800.3130891.

Gardner, M.-A., Hold-Geoffroy, Y., Sunkavalli, K., Gagne, C. and Lalonde, J.-F. (2019), "Deep parametric indoor lighting estimation", *The IEEE International Conference on Computer Vision (ICCV)*.

Glassner, A.S., (1989), (Ed.), *An Introduction to Ray Tracing*, Academic Press, London, GBR.

Goesele, M., Granier, X., Heidrich, W. and Seidel, H.P. (2003), "Accurate light source acquisition and rendering", *Proceeding of SIGGRAPH '03 (Special Issue of ACM Transactions on Graphics)*, pp. 621-630.

Goodfellow, I, Pouget-Abadie, J., Mirza, M., Xu, Bing, Warde-Farley, D., Ozair, S., Courville, A. and Bengio, Y. (2014), "Generative adversarial nets", in Ghahramani, Z., Welling, M., Cortes, C., Lawrence, N.D. and Weinberger, K.Q. (Eds), *Advances in Neural Information Processing Systems*, Curran Associates, Vol. 27, pp. 2672-2680.

He, K., Gkioxari, G., Dollár, P. and Ross, B. (2017), "Girshick. Mask r-cnn", *IEEE International Conference on Computer Vision (ICCV)*, pp. 2980-2988.

Ho, E.S.L., Chan, J.C.P., Chan, D.C.K. and Shum, H.P.H., Yiu, M.C. and Yuen, P.C. (2016), "Improving posture classification accuracy for depth sensor-based human activity monitoring in smart environments", *Computer Vision and Image Understanding*, Vol. 148, pp. 97-110, doi: 10.1016/j.cviu.2015.12.011.

Javed, S., Mahmood, A., Al-Maadeed, S., Bouwmans, T. and Ki Jung, S.(2018), "Moving object detection in complex scene using spatiotemporal structured-sparse RPCA", *IEEE Transactions on Image Processing*, Vol. 28 No. 2, pp. 1007-1022.

Li, C. and Yue, M. (2019), "Three-stream convolution networks after background subtraction for action recognition", *Artificial Intelligence and Soft Computing*, pp. 12-24.

Lim, L.A. and Keles, H.Y. (2018), "Foreground segmentation using convolutional neural networks for multiscale feature encoding", *Pattern Recognition Letters*, Vol. 112, pp. 256-262.

Liu, X., Yao, J., Hong, X., Huang, X., Zhou, Z., Qi, C. and Zhao, G. (2018), "Background subtraction using spatio-temporal group sparsity recovery", *IEEE Transactions on Circuits and Systems for Video Technology*, Vol. 28 No. 8, pp. 1737-1751.

Long, J., Shelhamer, E. and Trevor Darrell (2015), "Fully convolutional networks for semantic segmentation", *IEEE Conference on Computer Vision and Pattern Recognition*, 2015.

Lu, X. (2014), "A multiscale spatio-temporal background model for motion detection", *IEEE International Conference on Image Processing (ICIP)*, 2014.

McCay, K.D., Ho, E.S.L., Marcroft, C. and Embleton, N.D. (2019), "Establishing pose based features using histograms for the detection of abnormal infant movements", *41st Annual International Conference of the IEEE Engineering in Medicine and Biology Society (EMBC)*, pp. 5469-5472, July 2019.

McCay, K.D., Ho, E.S.L., Shum, H.P.H., Fehringer, G., Marcroft, C. and Embleton, N.D. (2020), "Abnormal infant movements classification with deep learning on pose-based features", *IEEE Access*, Vol. 8, pp. 51582-51592.

Oliver, N., Rosario, B. and Pentland, A. (2000), "A Bayesian computer vision system for modeling human interactions", *IEEE Transactions on Pattern Analysis and Machine Intelligence*, Vol. 22 No. 8, pp. 831-843.

Poynton, C.A. (1993), "Gamma and its disguises: the nonlinear mappings of intensity in perception, CRTS, film, and video", *SMPTE*, Vol. 102, pp. 1099-1108.

Ren, W., Liu, S., Zhang, H., Pan, J., Cao, X. and Yang, M.H. (2016), "Single image dehazing via multi-scale convolutional neural networks", in Leibe, B., Matas, J., Sebe, N. and Welling, M. (Eds), *Computer Vision – ECCV 2016*, SpringerCham, pp. 154-169.

Roy, A. and Todorovic, S. (2016), "A multi-scale cnn for affordance segmentation in RGB images", in Leibe, B., Matas, J., Sebe, N. and Welling, M. (Eds), *Computer Vision – ECCV 2016*, Springer, Cham, pp. 186-201.

Sakkos, D., Liu, H., Han, J. and Shao, L. (2018), "End-to-end video background subtraction with 3d convolutional neural networks", *Multimedia Tools and Applications*, Vol. 77 No. 17, pp. 23023-23041.

Sakkos, D., Ho, E.S.L. and Shum, H.P.H. (2019), "Illumination-based data augmentation for robust background subtraction", *International Conference on Software, Knowledge, Information Management and Applications (SKIMA)*.

Sakkos, D., Edmond, S., Ho, L. and Shum, H.P.H. (2019), "Illumination-aware multi-task GANs for foreground segmentation", *IEEE Access*, Vol. 7 No. 1, pp. 10976-10986.

Shen, Y., Hu, W., Yang, M., Liu, J., Wei, B., Lucey, S. and Chou, C.T. (2016), "Real-time and robust compressive background subtraction for embedded camera networks", *IEEE Transactions on Mobile Computing*, Vol. 15 No. 2, pp. 406-418.

Simonyan, K. and Zisserman, A. (2015), "Very deep convolutional networks for large-scale image recognition", *International Conference on Learning Representations (ICRL)*, pp. 1-14.

Sirikuntamat, N., Satoh, S. and Chalidabhongse, T.H. (2015), "Vehicle tracking in low hue contrast based on CAMshift and background subtraction", *12th International Joint Conference on Computer Science and Software Engineering (JCSSE)*, pp. 58-62, July 2015.

Siva, P., Javad Shafiee, M., Li, F. and Wong, A. (2015), "PIRM: fast background subtraction under sudden, local illumination changes via probabilistic illumination range modelling", *IEEE International Conference on Image Processing (ICIP)*, 2015.

Song, S. and Funkhouser, T. (2019), "Neural illumination: lighting prediction for indoor environments", *The IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*.

Them, K., Kaul, M.G., Jung, C., Hofmann, M., Mummert, T., Werner, F. and Knopp, T. (2016), "Sensitivity enhancement in magnetic particle imaging by background subtraction", *IEEE Transactions on Medical Imaging*, Vol. 35 No. 3, pp. 893-900.

Wang, X., Lu, C., Jia, J. and Li, H. (2017), "Regularized stationary-time estimation for crowd analysis", *EEE Transactions on Pattern Analysis and Machine Intelligence*, Vol. 39 No. 5, pp. 981-994.

Wang, Y., Luo, Z. and Jodoin, P.M. (2017), "Interactive deep learning method for segmenting moving objects", *Pattern Recognition Letters*, Vol. 96, pp. 66-75.

Wang, P., Chen, P., Yuan, Y., Ding, L., Huang, Z., Hou, X. and Cottrell, G.W. (2018), "Understanding convolution for semantic segmentation", *IEEE Winter Conference on Applications of Computer Vision (WACV)*, pp. 1451-1460, 2018.

Yakubovskiy, P. (2019), "Segmentation models", available at: https://github.com/qubvel/segmentation_models.

Yeo, H.-S., Lee, B.-G. and Lim, H. (2013), "Hand tracking and gesture recognition system for human-computer interaction using low-cost hardware", *Multimedia Tools and Applications*, Vol. 74 No. 8, pp. 2687-2715, ISSN 1380-7501, doi: 10.1007/s11042-013-1501-1.

Yu, F., Koltun, V. and Funkhouser, T. (2017), "Dilated residual networks", *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pp. 636-644, July 2017.

Zeng, D. and Zhu, M. (2018), "Background subtraction using multiscale fully convolutional network", *IEEE Access*, Vol. 6, pp. 16010-16021.

Zhao, H., Shi, J., Qi, X., Wang, X. and Jia, J. (2017), "Pyramid scene parsing network", *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pp. 6230-6239, July 2017.

Zivkovic, Z. and Heijden, F. (2006), "Efficient adaptive density estimation per image pixel for the task of background subtraction", *Pattern Recognition Letters*, Vol. 27, pp. 773-780.

Corresponding author

Edmond S. L. Ho can be contacted at: e.ho@northumbria.ac.uk