

RESEARCH ARTICLE

A Cloud-Based Deep Learning Framework for Early Detection of Pushing at Crowded Event Entrances

AHMED ALIA^{1,2,3}, MOHAMMED MAREE⁴, MOHCINE CHRAIBI¹,
ANAS TOMA⁵, AND ARMIN SEYFRIED^{1,2}

¹Institute for Advanced Simulation, Forschungszentrum Jülich, 52425 Jülich, Germany

²Department of Computer Simulation for Fire Protection and Pedestrian Traffic, University of Wuppertal, 42285 Wuppertal, Germany

³Department of Information Technology, An-Najah National University, Nablus, Palestine

⁴Department of Information Technology, Arab American University, 00970 Jenin, Palestine

⁵Department of Electrical and Computer Engineering, An-Najah National University, Nablus, Palestine

Corresponding authors: Ahmed Alia (a.alia@fz-juelich.de), Mohammed Maree (mohammed.maree@aaup.edu), and Armin Seyfried (a.seyfried@fz-juelich.de)

This work was funded by the German Federal Ministry of Education and Research (BMBF: funding number 01DH16027) within the Palestinian-German Science Bridge project framework and by the Deutsche Forschungsgemeinschaft (DFG, German Research Foundation)—491111487.

This work involved human subjects or animals in its research. Approval of all ethical and experimental procedures and protocols was granted by the Ethics Board at the University of Wuppertal, Germany, and performed in line with the Declaration of Helsinki.

ABSTRACT Crowding at the entrances of large events may lead to critical and life-threatening situations, particularly when people start pushing each other to reach the event faster. Automatic and timely identification of pushing behavior would help organizers and security forces to intervene early and mitigate dangerous situations. In this paper, we propose a cloud-based deep learning framework for automatic early detection of pushing in crowded event entrances. The proposed framework initially modifies and trains the EfficientNetV2B0 Convolutional Neural Network model. Subsequently, it integrates the adapted model with an accurate and fast pre-trained deep optical flow model with the color wheel method to analyze video streams and identify pushing patches in real-time. Moreover, the framework uses live capturing technology and a cloud-based environment to collect video streams of crowds in real-time and provide early-stage results. A novel dataset is generated based on five real-world experiments and their associated ground truth data to train the adapted EfficientNetV2B0 model. The experimental setups simulated a crowded event entrance, while the ground truths for each video experiment were generated manually by social psychologists. Several experiments on the videos and the generated dataset are carried out to evaluate the accuracy and annotation delay time of the proposed framework. The experimental results show that the proposed framework identified pushing behaviors with an accuracy rate of 87% within a reasonable delay time.

INDEX TERMS Artificial intelligence, computer vision, convolutional neural network, deep learning, image classification, intelligent system, machine learning, pushing behavior detection.

I. INTRODUCTION

The entrances of large-scale events such as sport venues, concerts, and religious gatherings are organized as bottlenecks for access control, ticket validation, or security check [1].

The associate editor coordinating the review of this manuscript and approving it for publication was Okyay Kaynak¹.

In these scenarios, some pedestrians might start pushing each other to gain faster access to the event. According to Üsten et al. [2], pushing for forward motion is defined as “a behavior that can involve using arms, shoulders, or elbows; or simply the upper body, in which one person actively applies force to another person (or people) to overtake, while shifting their direction to the side or back, or force them to

move forward more quickly.” Additionally, using gaps in the crowd is considered as a strategy of pushing because it is a form of overtaking [2]. Indeed, such behavior increases the crowd’s density over time [1], [3], resulting in the lack of comfort zones and, more importantly, can lead to dangerous situations [4], [5]. In such cases, early pushing detection is essential, as it can provide valuable information to the organizers and the security team for better crowd management, thereby ensuring a smoother flow at entrances with higher safety [6]. Since manual identification of pushing behavior in the early stages can be complex or impossible, developing an automatic detection framework in real-time or near real-time is crucial. However, automatic pushing detection is still a challenging task due to the highly-dense crowds, the diversity of pushing behavior strategies, and the varying features for pushing behavior representation, which still requires further investigation and identification [7].

Surveillance cameras have recently been widely integrated with computer vision techniques to automatically identify abnormal behaviors from crowds [8], [9]. Within the realm of computer vision, pushing behavior can be classified as abnormal behavior. Machine learning algorithms, particularly Convolutional Neural Network (CNN) architectures, have remarkably succeeded in several computer vision tasks; among these is abnormal behavior detection in crowds [10]. One of the critical reasons for this success is that CNN can learn the relevant features [11], [12] and classification automatically from data without human intervention [13], [14]. Although CNN architectures are powerful for modeling human behaviors, building an accurate model requires a large training dataset [15], [16], which is often unavailable. Researchers have developed hybrid-based approaches that integrate CNN with handcrafted feature descriptors to address this limitation [17], [18]. These approaches employ descriptors to obtain useful data, which is subsequently used by CNN to learn and identify abnormal behavior automatically. Due to the limited availability of labeled data for pushing behavior, hybrid-based approaches may be more appropriate for automatically identifying pushing behavior. For example, Alia et al. [7] proposed a hybrid deep learning and visualization framework for pushing behavior detection in video recordings of crowded event entrances. Unfortunately, this framework does not cope with early detection requirements because it is slow and can not work with the live camera stream. To the best of our knowledge, despite the numerous computer vision and machine learning approaches reported in the literature, none of them can detect pushing behavior in real-time or near real-time from crowds.

In order to address the above limitations, this article introduces a novel cloud-based deep learning framework for pushing patch detection in live video streams acquired from crowded event entrances. In this framework, we propose: 1) Integrating a robust deep optical flow model (GPU-based pre-trained Recurrent All-pairs Field Transforms (RAFT) [19]) with the color wheel method [20], [21] to accurately and rapidly extract the visual motion information from

the crowd. 2) Adapting and training EfficientNetV2B0-based CNN [22] using visual motion information to detect pushing patches accurately. 3) Using live camera technology and a cloud environment to provide more powerful computational resources and help to collect and annotate the video stream of the crowd in real-time.

The main contributions of this article are summarized as follows:

- 1) To the best of our knowledge, we propose the first real-time or near real-time automatic framework dedicated to early identifying pushing behavior in human crowds.
- 2) We introduce a new video analysis and pushing detection approach based on integrating an adapted version of EfficientNetV2B0, GPU-based pre-trained RAFT model, and color wheel method.
- 3) We create a novel dataset for pushing behavior, using five real-world experiments with their associated ground truths. This dataset is not only used as a training and evaluation resource for our adapted EfficientNetV2B0, but can also be a valuable asset for future research in this area.
- 4) We perform a thorough performance comparison of fifteen CNN architectures for pushing detection using the generated dataset.

The rest of the paper is organized as follows. Section II reviews the related studies of video-based abnormal human behavior detection. The proposed framework is presented in Section III. Section IV discusses the evaluation process and experimental results. Finally, the conclusion and future work are summarized in Section V.

II. RELATED WORK

Generally, identifying pushing behavior in videos falls under the field of computer vision, specifically in the task of abnormal behavior detection. CNNs have played a crucial role in significant advancements in this area [23]. Consequently, in this section, our objective is to examine several abnormal behavior detection approaches that have been developed using CNNs.

A customized CNN-based method to identify abnormal activities in videos was presented by Tay et al. [24]. The authors trained a customized CNN for feature extraction and labeling using normal and abnormal samples. In another study, Alafif et al. [18] proposed two methods of identifying abnormal behaviors in small and large-scale crowd videos. The first method employs a combination of a CNN model and a random forest classifier to detect anomaly behaviors at the object level in a small-scale crowd. In contrast, the second method utilizes two classifiers to recognize abnormal behaviors in a large-scale crowd. The initial model, finds the frames containing abnormal behaviors, while the second classifier, You Only Look Once (version 2), processes those frames to identify abnormal behaviors exhibited by individuals. The effectiveness of these techniques relies heavily on utilizing

CNNs to learn features from labeled datasets containing both normal and abnormal behaviors. A large training dataset of normal and abnormal behaviors is necessary to create an accurate and adaptable CNN model. However, obtaining such a dataset is often unattainable for various abnormal behaviors, including pushing behavior.

In order to overcome the shortage of large datasets comprising normal and abnormal behaviors, some researchers have utilized one-class classifiers with datasets consisting only of normal behaviors. It is easier to obtain or create a dataset that contains only normal behavior than a dataset that includes both normal and abnormal behaviors [25], [26]. The fundamental concept behind the one-class classifier is to exclusively learn from normal behaviors, thereby establishing a class boundary between normal and undefined (abnormal) classes. For example, Sabokrou et al. [25] employed a pre-trained CNN for extracting motion and appearance information from crowded scenes. Subsequently, they utilized a one-class Gaussian distribution to construct the classifier using datasets comprised of normal behavior. Similarly, in [26] and [27], the authors developed one-class classifiers by utilizing a dataset of normal samples. In [26], Xu et al. employed a convolutional variational autoencoder to extract features, followed by the use of multiple Gaussian models to detect abnormal behavior. Meanwhile, in [27], a pre-trained CNN model was utilized for feature extraction, with one-class support vector machines being used to identify abnormal behavior. Another study by Ilyas et al. [28] utilized a pre-trained CNN and a gradient sum of the frame difference to extract significant features. Following this, three support vector machines were trained on normal behavior to detect abnormal behaviors. Generally, the one-class classifier is commonly used when the target behavior class or abnormal behavior is infrequent or poorly defined [29]. However, pushing behavior is well-defined and not rare, particularly in high-density and competitive situations. Furthermore, this type of classifier regards new normal behavior as abnormal.

To overcome the limitations of CNN-based and one-class classifier approaches, several studies have combined multi-class CNN with one or more handcrafted feature descriptors [10], [28]. As an example, Duman and Erdem [17] utilized the traditional Farnebäck optical flow approach in conjunction with CNN to detect anomalous behavior. They extracted direction and speed information using Farnebäck and CNN, and then utilized a convolutional long short-term memory network to construct the classifier. Similarly, Hu [30] employed a combination of the histogram of gradient and CNN for feature extraction, while a least-squares support vector was used for classification. Almazroey and Jaraya [31] focused on utilizing the Lucas-Kanade optical flow method, pre-trained CNN, and feature selection method (neighborhood component analysis) to extract relevant features. They then used a support vector machine to generate a trained classifier. In a different study [32], Zhou et al. introduced a CNN-based method to identify and locate abnormal activities. This approach integrated optical flow with

CNN for feature extraction and utilized a CNN for classification. Direkoglu [10] utilized the Lucas-Kanade optical flow method and CNN to extract relevant features and identify “escape and panic behaviors”.

Most of the hybrid-based approaches for abnormal behavior detection that were reviewed have limited efficiency in detecting pushing since 1) The descriptors used in these approaches can only extract limited essential data from high-density crowds to represent pushing behavior. 2) Some CNN architectures commonly utilized in these approaches may not be effective in dealing with the increased variations within pushing behavior (intra-class variance) and the substantial resemblance between pushing and non-pushing behaviors (high inter-class similarity), which can potentially result in misclassification. To benefit from the power of hybrid-based approaches on a small dataset, Alia et al. [7] introduced a hybrid framework for pushing patch detection in video recordings of crowds. The authors utilized a robust handcrafted feature descriptor and efficient CNN architecture in this framework. In more details, the framework used a deep optical flow technique to extract the motion information from the crowds. This information is then analyzed using an EfficientNetB0-based CNN and false reduction algorithms to identify and label pushing patches in the video. However, this framework does not cope with early detection requirements due to three reasons. First, it can only handle offline-recorded videos. Second, The deep optical flow technique employed in motion extraction is slow because it was performed on the CPU. Third, it needs to identify pushing patches for the whole video before producing the output. Moreover, as reported by the authors, the accuracy of the framework decreases with complex scenarios of pushing.

To sum up, the reviewed methods have limitations regarding early pushing detection in crowded human environments. On the one hand, approaches that rely solely on CNNs for feature extraction require a large dataset containing normal and abnormal behaviors, which is typically unavailable for pushing scenarios. On the other hand, one-class classifiers are often used for infrequent or poorly defined target behavior or abnormal behavior. However, pushing behavior is well-defined and common, particularly in high-density and competitive scenarios. Additionally, this type of classifier may misclassify new normal behavior as abnormal. Although hybrid-based approaches may be more suitable for pushing behavior, existing methods do not meet the requirements for early pushing detection in human crowds. To overcome these limitations, this article proposes a novel framework that adapts the EfficientNetV2B0 model and integrates it with GPU-based RAFT, wheel color method and live camera technology on a cloud platform. The following section provides a detailed discussion of the framework.

III. THE PROPOSED FRAMEWORK

In this section, we describe the proposed framework for early detection of pushing within the live camera stream of crowded event entrances, where the camera is fixed

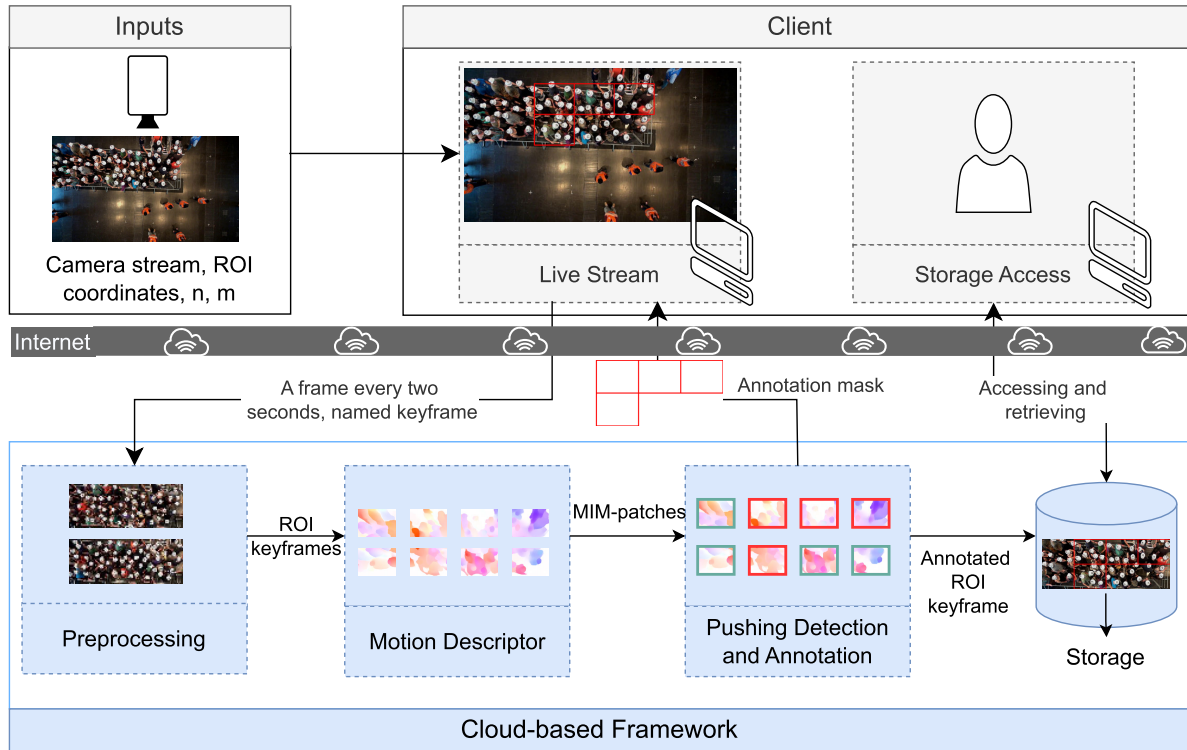


FIGURE 1. The proposed framework architecture. ROI refers to the entrance area (Region Of Interest). User-defined row (n) and column (m) are used to split Motion Information Map (MIM) into $n \times m$ patches.

and top-view. Fig. 1 shows the architecture of our framework which comprises three major components: preprocessing; motion descriptor; and pushing detection and annotation. The first component aims to collect and process the live camera stream, as well as display the stream on the web client in real-time. Simultaneously, the second component, the motion descriptor, employs the GPU-based RAFT model and color wheel method [20], [21] to extract the visual motion information from the crowd. Finally, the pushing detection and annotation component utilizes the adapted and trained EfficientNetV2B0 model to analyze the visual motion information and detect pushing patches. Notably, it directly annotates the regions that contain pushing behavior on the live stream on the web client. The following sections provide a more detailed discussion of the three components.

A. PREPROCESSING

In order to reduce the computational time of the framework without sacrificing performance, the preprocessing component directly displays the client camera stream on the web client. At the same time, it collects only the data required for detection purposes from the live stream. Let $\{f^t\}$ represents the live camera stream, where t is the time of the frame f in the stream. Firstly, this component displays the live stream on the web client in real-time without uploading it to the cloud. Then, a frame f^t is collected from the stream every two seconds, hereafter referred to as keyframe (examples in Fig. 3a). After that, this component utilizes the

user-defined coordinates in pixel units to crop the entrance area f^t (ROI keyframe) from its corresponding keyframe f^t . Finally, f^t is submitted as an input to the second component. For the brevity, we name the ROI keyframe sequence $\{\tilde{f}^t, \tilde{f}^{t+2}, \tilde{f}^{t+4}, \dots\}$ as $\{\tilde{f}_i | i = 1, 2, 3, \dots\}$, where i is the order of the ROI keyframe in the stream, and t is the time in seconds. Fig. 3b displays two examples of \tilde{f}_i .

B. MOTION DESCRIPTOR

Using this component, we aim to extract the crowd's motion characteristics at the patch level. More specifically, this component estimates the motion direction, magnitude, and associated spatio-temporal information from the crowds, and accordingly visualizes this information. The displayed information includes relevant features that are important for representing the pushing behavior. As shown in Fig. 2, the component uses GPU-based pre-trained RAFT model and color wheel method to achieve its purpose. Unlike the majority of the already used optical flow methods [34], [35], a GPU-based pre-trained RAFT model performs well in terms of speed, accuracy, and generality for dense crowds [7], [19]. This model was created by training an ensemble of CNN and recurrent neural networks on the Sintel dataset to calculate the optical flow between two images. For further details about the model, we refer the reader to [19]. Firstly, the component uses the pre-trained model to calculate the displacement of each pixel (x, y) between each pair of \tilde{f}_i and \tilde{f}_{i+1} , generating the dense displacement field d_i . Each pixel location (x, y) in

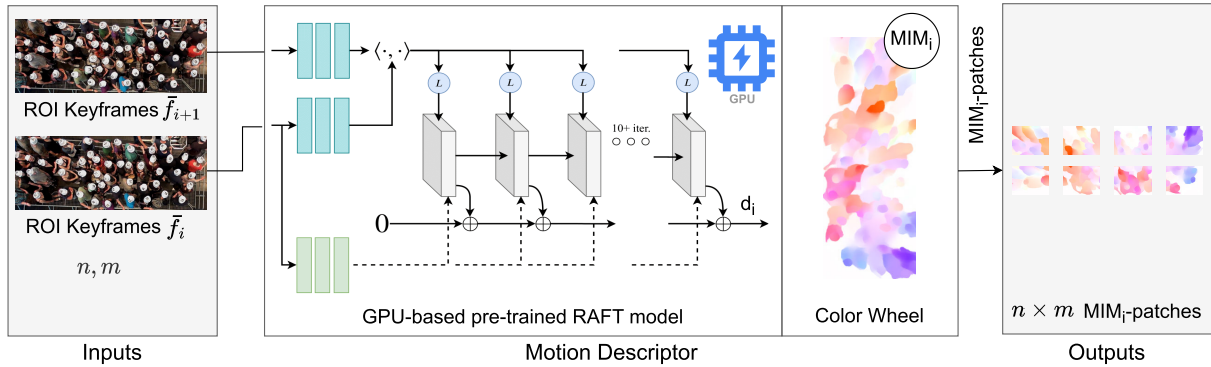


FIGURE 2. Motion descriptor component pipeline. i is the order of the ROI keyframe in the stream. d refers to a dense displacement field. MIM represents motion information map.

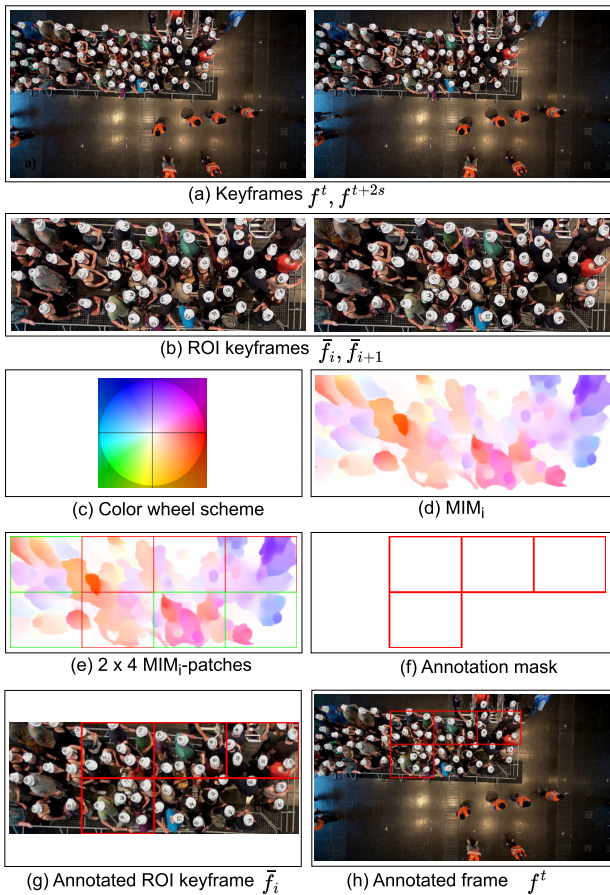


FIGURE 3. An illustration of two keyframes (experiment Entrance_2 [33], two ROI keyframes, color wheel schema [21], MIM, 2 x 4 MIM-patches, annotation mask, annotated ROI keyframe and annotated frame. t is the time of the frame f in the stream. i is the order of the ROI keyframe in the stream. s means second. The red boxes indicate pushing patches, while the green boxes mean non-pushing patches.

d_i is presented by a vector, given by

$$\langle u_{(x,y)}, v_{(x,y)} \rangle_{\bar{f}_i, \bar{f}_{i+1}} = RAFT(\langle x, y \rangle_{\bar{f}_i, \bar{f}_{i+1}}), \quad (1)$$

where u and v are horizontal and vertical displacements of a pixel at the (x, y) location between \bar{f}_i and \bar{f}_{i+1} , respectively.

This implies that d_i is a matrix of the vectors, as described in

$$d_i = \left\{ \langle u_{(x,y)}, v_{(x,y)} \rangle_{\bar{f}_i, \bar{f}_{i+1}} \right\}_{(x,y)=(1,1)}^{(w,h)}, \quad (2)$$

where w and h are the \bar{f}_i width and height, respectively.

After the estimation of d_i , the descriptor applies the color wheel method to deduce the visual motion information from d_i . It begins by calculating the direction θ and magnitude of each vector $\langle u_{(x,y)}, v_{(x,y)} \rangle$ in d_i using Eq. (3) and Eq.(4), respectively. The color wheel then visualizes the magnitude and direction information to generate MIM_i from the calculated information, where $MIM_i \in \mathbb{R}^{w \times h \times 3}$, and 3 is the number of channels in MIM_i. Fig. 3c is the color wheel scheme, and Fig. 3d is an example of MIM_i that is generated from the pair of \bar{f}_i and \bar{f}_{i+1} (Fig. 3b). According to the wheel schema, the color represents the motion direction, while the color intensity denotes the motion magnitude or speed.

$$\theta(\langle x, y \rangle_{\bar{f}_i, \bar{f}_{i+1}}) = \pi^{-1} \arctan\left(\frac{v_{(x,y)}}{u_{(x,y)}}\right). \quad (3)$$

$$mag(\langle x, y \rangle)_{c_i} = \sqrt{u_{(x,y)}^2 + v_{(x,y)}^2} \quad (4)$$

The motion descriptor component divides each MIM_i into $n \times m$ MIM_i-patches to help the framework localizing pushing in ROI. The MIM_i-patches can be expressed as $\{p_{i,k} \in \mathbb{R}^{(w/m) \times (h/n) \times 3} \mid k = 1, 2, \dots, n \times m\}$, where k is the order of the patch in MIM_i. For more clarity, MIM_i (Fig. 3d) is divided into 2 x 4 MIM-patches (Fig. 3e). It is worth noting that the patch should cover an area on the ground that can accommodate a group of pedestrians, as crowd characteristics are required for representing pushing behavior. To summarize, the MIM-patches represent the output of the motion descriptor component and the input of the next component.

C. PUSHING DETECTION AND ANNOTATION

The primary purpose of this component (Fig.4a) is to localize the pushing patches in the live stream, as well as blurring and storing the annotated ROI keyframes in the cloud storage. Labeling MIM-patches as pushing or non-pushing is the most important aspect of localizing pushing in the live stream. Therefore, we created an efficient binary classifier by

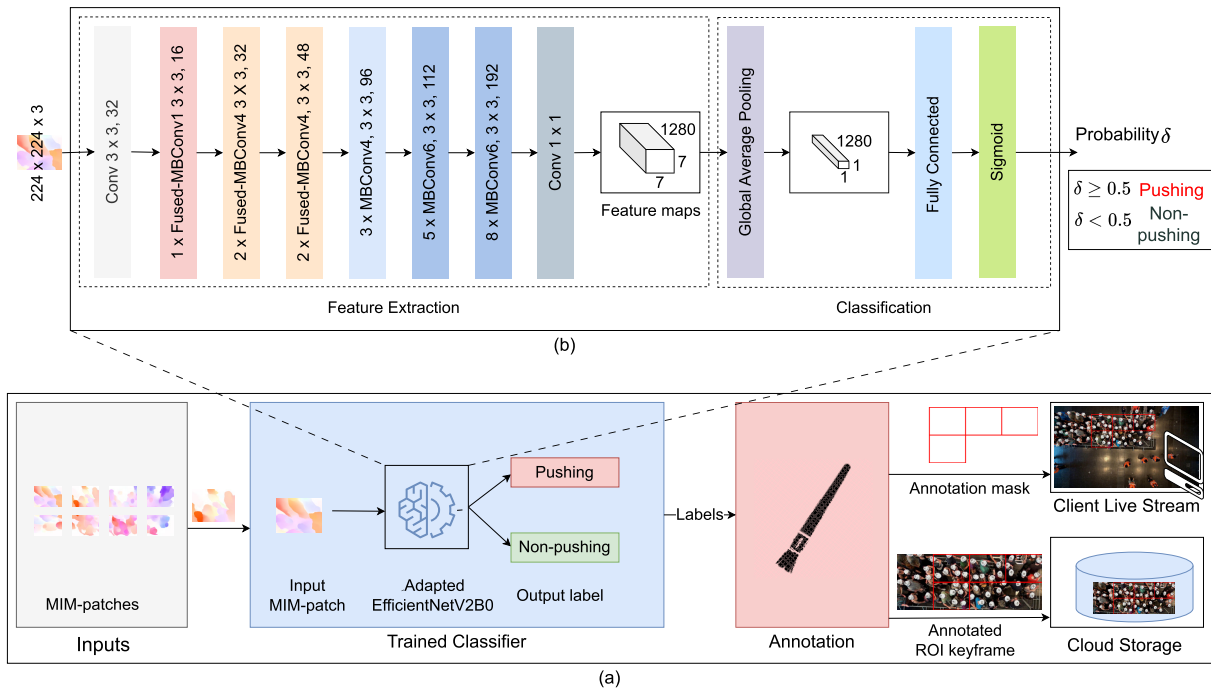


FIGURE 4. (a) The pipeline of pushing detection and annotation component. (b) Adapted EfficientNetV2B0 Architecture.

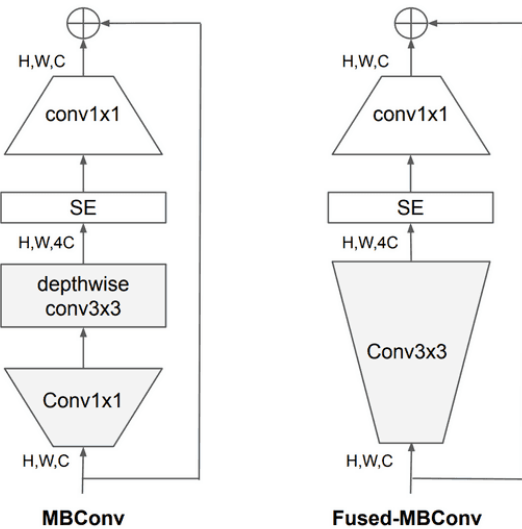


FIGURE 5. Structure of MBConv and Fused-MBConv [22].

adapting and training the EfficientNetV2B0 CNN architecture [22] from scratch, which is then utilized to label the MIM-patches.

1) ADAPTED EfficientNetV2B0 ARCHITECTURE

EfficientNetV2B0 is a convolutional neural network belonging to the EfficientNetV2 family, designed by the Google Brain team [22]. Such a family outperforms state-of-the-art accuracy in different classification tasks with a far smaller model and faster converging speed. EfficientNetV2B0 is the smallest model in this family and achieves high accuracy with minimal computational cost.

TABLE 1. The hyperparameter values used in the training process.

Parameter	Value
Optimizer	Adam
Loss function	Binary cross-entropy
Learning rate	0.001
Batch size	32
Epoch	100

Fig. 4b depicts the overall architecture of the modified EfficientNetV2B0, which firstly performs a 3×3 convolution operation on the input image, which has dimensions of $224 \times 224 \times 3$. Then it utilizes a combination of 5 Fused-MBConv (Fused Mobile Inverted Residual Bottleneck Convolution) [36] and 16 MBConv [37] modules for extracting the feature maps ($7 \times 7 \times 1280$) from the input image. The model then employs a global average pooling layer and a fully connected layer with a Sigmoid activation function for binary classification. The global average pooling layer transforms the dimensions of the stacked feature maps to $1 \times 1 \times 1280$ and assigns them to the fully connected layer. Finally, the fully connected layer with a Sigmoid activation function finds the probability δ of the label of the input MIM-patch. Then, the classifier uses the threshold to determine the class of the MIM-patch as Eq. (5):

$$Class(MIM - patch) = \begin{cases} \text{pushing} & \text{if } \delta \geq 0.5 \\ \text{non-pushing} & \text{if } \delta < 0.5 \end{cases} \quad (5)$$

It's important to note that the classification part of this model differs from the original EfficientNetV2B0, which was designed to classify images into 1,000 categories. However,

pushing detection requires labeling the input image into one of two possible classes.

As mentioned above, the main fundamental blocks in EfficientNetV2B0 for feature extraction are MBConv and Fused MBConv [22]. As shown in Fig. 5, MBConv firstly uses a 1×1 convolution operation to expand the input activation maps to increase the depth of the feature maps. Next, 3×3 depthwise convolutions are applied to reduce the computational complexity and the number of parameters. Then, a Squeeze-and-Excitation (SE) block enhances the representation power of the architecture. Finally, another 1×1 convolution is employed to reduce the dimensionality of the output feature maps, producing the final output of this block. Moreover, A residual connection is added to enhance the performance further. Despite depthwise convolutions having fewer parameters, they can not often fully utilize modern accelerators. In contrast, the Fused-MBConv tries to solve this problem by replacing the depthwise and expansion conv 1×1 in MBConv conv 3×3 with a single regular conv 3×3 , resulting in a faster training process (see Fig. 5). It is worth mentioning that using only Fused-MBConv in the architecture increases parameters while slowing down the training. Therefore, EfficientNetV2B0 applied a combination of MBConv and Fused-MBConv to improve training speed with a small overhead on parameters and enhance the feature extraction process [22].

The following subsection will discuss the training process for the adapted EfficientNetV2B0 model to classify MIM patches into pushing and non-pushing.

2) ADAPTED EfficientNetV2B0 TRAINING

To classify the MIM-patches into pushing and non-pushing categories, we trained the adapted EfficientNetV2B0 model (Fig. 4b) using new training and validation sets comprising both types of MIM-patches (details about the dataset can be found in Section IV-A). The model parameters used during the training process are listed in Table 1, and were chosen based on experimentation to obtain optimal performance with the given dataset. To prevent overfitting, we halted the training if the validation accuracy did not improve after 20 epochs.

Fig. 4a shows the pipeline of the pushing detection and annotation component. Firstly, the trained classifier labels MIM-patches $p_{i,k}$ received from the previous MIM component. Then, the current component displays an annotation mask of the pushing patches in the live stream on the web client. Simultaneously, it blurs and annotates the corresponding ROI keyframe \tilde{f}_i before saving it in the cloud storage. Notably, web clients can access this storage via an internet connection.

IV. EVALUATION AND RESULTS

This section introduces the dataset, implementation details, and performance metrics utilized in evaluating the proposed framework. The results of various experiments conducted to

assess the performance of our classifier and the proposed framework are also discussed.

A. DATASET PREPARATION

Here, we explain how we prepared the labeled dataset (training, validation, and test sets) for training and evaluating the adapted EfficientNetV2B0 as well as all models used in the evaluation. The dataset contains two classes of MIM-patches, which are pushing and non-pushing.

1) DATA COLLECTION

In this section, we discuss the data sources used to obtain our dataset. The sources are mainly based on video experiments of crowded event entrances, trajectory data, and ground truth data for pushing behavior. Five video experiments with their trajectory data are chosen from the data archive hosted by Forschungszentrum Jülich under CC Attribution 4.0 International license [33], [38]. Static top-view cameras were used to record the videos with a frame rate of 25 frames per second. It is worth mentioning that the selected experiments contain varied characteristics, which help to improve the generality of the dataset, as seen in Table 2. The ground truths for the last data source were manually created by social psychologists, who established the definition of pushing behavior in forward motion among crowds [2]. These ground truths indicate whether the behavior of each pedestrian in every frame is classified as either pushing or non-pushing.

2) DATASET GENERATION

The methodology of the labeled dataset generation, as seen in Fig. 6, includes three steps: (1) MIM-patches generation, (2) MIM-patches labeling and (3) Labeled dataset generation.

In the MIM-patches generation step, the motion descriptor component was employed (Fig. 2) on the video experiments and their $n \times m$ patches (Table 2) to produce MIM-patches. To increase the number of patches, the component is applied four times for each video with a different commencement; half a second is the delay duration of each time compared to the previous time. According to [7], half a second delay helps to generate diverse MIM-patches, while less than this period may result in redundant samples. Based on the trajectory and ground truth data, the second step labels the patches as pushing and non-pushing. Patches are classified as pushing if it contains at least one pushing behavior, and non-pushing if no pedestrians engage in pushing behavior. On the other hand, the patches that only show a portion of one pedestrian pushing are discarded; because they do not offer complete information about pushing or non-pushing behavior. According to the labels of the patches, the last step stores the patches in pushing and non-pushing directories to create the labeled dataset. At the end, the generated dataset consists of 2257 pushing and 1684 non-pushing samples. To generate the holdout data, the produced dataset is randomly divided into three sets: 70% for training, 15% for validation, and 15% for testing. This split ratio is one of the most commonly used splitting methods in

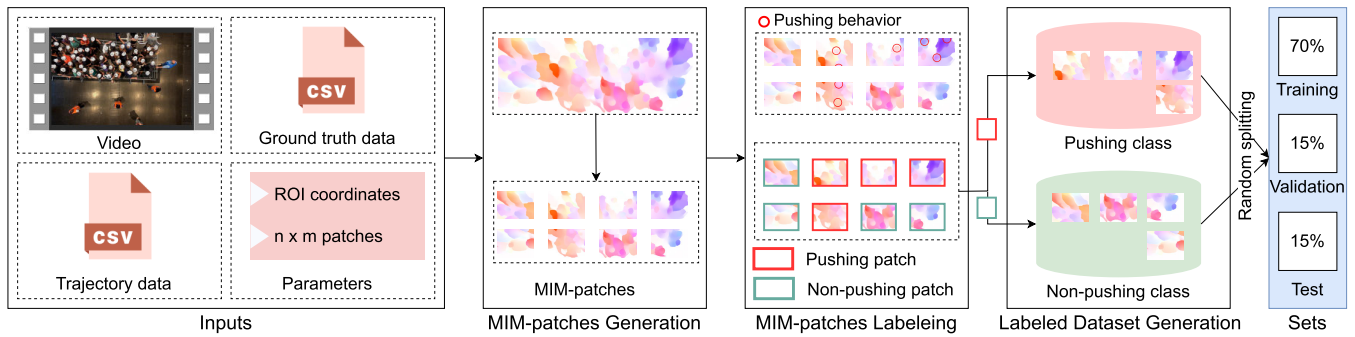


FIGURE 6. The methodology of the labeled dataset generation.

TABLE 2. Characteristics of the selected video experiments.

Video	Entrance type	Gates	Width (m)	Ped.	Dur.	Resolution	ROI coordinates (pixel)	n × m patches *
110	Straight	1	1.2	63	53	1920 × 1440	(374, 548), (1382, 864)	1 × 3
150	Straight	1	5.6	57	57	1920 × 1440	(364, 200), (1378, 1250)	3 × 3
270	Straight	1	3.4	67	59	1920 × 1440	(374, 330), (1390, 1070)	2 × 3
280	Straight	1	3.4	67	67	1920 × 1440	(374, 330), (1390, 1070)	2 × 3
Entrance_2	90° Corner	2	2	123	125	1920 × 1080	(213, 110), (1337, 540)	2 × 4

The video experiments' names are the same as reported in [33, 38]. "Dur." means duration. "Ped." is an abbreviation for the number of pedestrians. ROI coordinates: left-top and bottom-right coordinates of ROI in the pixel unit. n × m: number of rows and columns that are used to divide ROI into n × m regions, which are required for dividing MIM_i into n × m MIM_i-patches. * These values ensure that the dimensions of each region on the ground are greater than one meter, which is enough to accommodate a group of pedestrian [7].

the deep learning field [39]. Table 3 shows the number of pushing and non-pushing samples in the training, validation, and test sets.

TABLE 3. A number of samples in training, validation, and test sets in the generated dataset.

	Video	110	150	270	280	E_2	Total
Training	P	122	182	215	258	808	1585
	NP	72	206	197	182	525	1182
	Total	194	388	412	440	1333	2767
Validation	P	26	38	45	55	172	336
	NP	15	44	42	38	112	251
	Total	41	82	87	93	284	587
Test	P	26	38	45	55	172	336
	NP	15	44	42	38	112	251
	Total	41	82	87	93	284	587
All	Total	276	552	586	626	1901	3941

"All" refers to all sets. P means pushing. NP is non-pushing. E_2 stands for Entrance_2.

B. IMPLEMENTATION DETAILS AND EVALUATION METRICS

In this article, all the experiments and implementations were conducted on Google Colaboratory Pro (with a GPU NVIDIA of 15 GB and system RAM of 12.7 GB), utilizing JavaScript and Python 3 programming languages along with Keras, TensorFlow 2.0, and OpenCV libraries. Furthermore, all models in the experiments were trained using the same hyperparameter values utilized in the training of our adapted version of EfficientNetV2B0 (Table 1).

In order to evaluate the performance of our framework, we utilized a combination of metrics, including accuracy, macro F1-score, and area under the receiver operating charac-

teristic curve (AUC) over the test set. This set of metrics was necessary due to the imbalanced nature of our dataset [40]. In addition to these metrics, computational time was also measured as an essential performance metric. The following provides a detailed explanation of these metrics.

Accuracy: the ratio of successfully classified MIM-patches to the total number of samples in the test set, and mathematically can be defined as

$$accuracy = \frac{TP + TN}{TP + FP + TN + FN}, \tag{6}$$

where TP and TN denote correctly classified pushing and non-pushing patches, respectively. FP and FN represent incorrectly predicted pushing (P) and non-pushing (NP) samples. Accuracy is not enough to evaluate the classifier's performance over an imbalanced dataset, such as our used dataset. Therefore, we used the macro F1-score and AUC metrics, which are valuable for evaluating imbalanced classification problems.

Macro F1-score: the mean of class-wise F1-scores as described in the formula below:

$$Macro\ F1 - score = \frac{F1 - score(P) + F1 - score(NP)}{2}, \tag{7}$$

where F1-score is the harmonic average of precision and recall as described in:

$$F1 - score = \frac{2 \times precision \times recall}{precision + recall}, \tag{8}$$

where recall of pushing class is the ratio of correctly classified pushing MIM-patches to all pushing samples, while precision of pushing class is the ratio of correctly classified pushing patches out of all the samples labeled as pushing by the classifier. Recall and precision are defined in Eq. (9) and Eq. (10), respectively.

$$\text{recall} = \frac{TP}{TP+FN}, \quad (9)$$

$$\text{precision} = \frac{TP}{TP+FP}. \quad (10)$$

AUC is the area under the Receiver Operating Characteristics (ROC) curve. ROC is a graph showing the performance of a classification model at all thresholds. The ROC curve plots the false positive rate on the horizontal axis and the true positive rate on the vertical axis. The AUC value ranges from 0 to 1, while a model with an AUC of 1 is considered perfect, while a value of 0.5 indicates that the model performs no better than random guessing.

Computational time: this metric was employed to calculate how long the proposed framework takes to read, analyze and annotate every input, which is two seconds of stream. In other words, computational time determines whether our framework can detect pushing patches within a reasonable time or not.

C. EVALUATION OF OUR CLASSIFIER PERFORMANCE

We conducted three main comparative empirical experiments to evaluate the effect of our modified EfficientNetV2B0 classifier on the performance of the proposed framework. The first experiment compared the proposed classifier against eleven of the most popular CNN architectures. In the second experiment, we compared it to two custom CNN architectures designed for detecting abnormal behavior. Lastly, it was compared to CNN architecture used for pushing detection. Our classifier and all other models were implemented, trained, and assessed utilizing the same MIM-patches dataset, environment, and settings. Moreover, we utilized accuracy, F1-score, and AUC metrics to measure each model's performance.

1) A COMPARISON WITH ELEVEN POPULAR CNNs

Table 4 depicts the popular CNN architectures used in the first experiment, as well as the comparison results. It is clear that the adapted version of the EfficientNetV2B0 classifier outperformed the rest of the exploited classifiers. In particular, the proposed classifier achieved 87% accuracy and 86% F1-score, whereas the second top model in this comparison, DenseNet169, produced an 83% level of both accuracy and F1-score. This finding is primarily attributable to EfficientNetV2B0's superior efficiency for feature extraction compared to earlier CNN architectures. The main reason for this efficiency is the combination of MBCConv and Fused-MBCConv blocks used in EfficientNetV2B0.

Furthermore, as shown in Fig. 7, the proposed classifier obtained the highest AUC score (93%) among all the models tested, while the next best model achieving 85%.

TABLE 4. Comparison results to the well-known CNN-based classifiers.

CNN	Acc. %	Pre. %	Rec. %	F1. %
Xception [41]	81	81	81	81
VGG16 [42]	57	36	29	50
VGG19 [42]	61	61	62	62
ResNet50 [43]	80	79	81	79
ResNet50V2 [44]	77	76	77	75
ResNet101 [43]	72	70	72	70
ResNet101V2 [44]	72	72	72	71
ResNet152V2 [44]	74	73	73	73
DenseNet121 [45]	79	79	79	79
DenseNet169 [45]	83	83	83	83
NASNetMobile [46]	57	56	56	56
Our classifier	87	87	86	86

"Acc." refers to Accuracy. "Pre." stands for Precision. "Rec." means Recall. "F1." means F1-score.

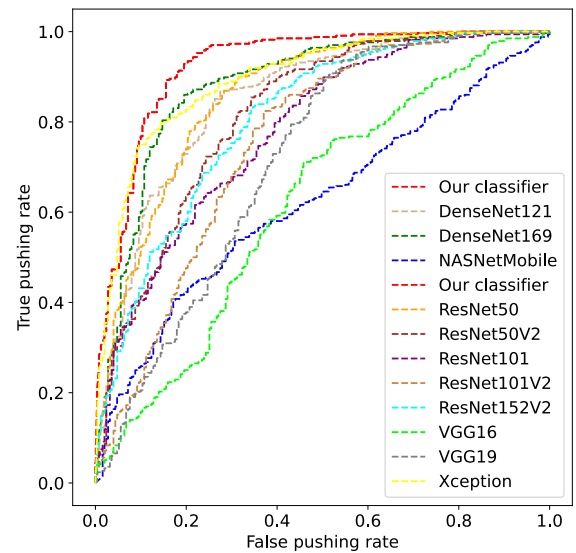


FIGURE 7. ROC curves of our classifier and eleven popular CNN models.

2) A COMPARISON WITH CUSTOMIZED CNNs IN ABNORMAL BEHAVIOR DETECTION

Here, we have two objectives, 1) Evaluating the performance of some existing CNN models developed to detect abnormal human behavior for pushing detection purposes. 2) Further evaluation of our classifier. The customized architectures are CNN-1 [10] and CNN-2 [24]. The first architecture, CNN-1, employed product 75×75 pixels as an input image. Furthermore, three convolutional layers, batch normalization, and max pooling operations were used for feature extraction. The developers of this model utilized a fully connected layer with a softmax activation function for classification. The second architecture, CNN-2, downsized the input images to product 32×32 pixels before employing three convolutional layers with three max-pooling layers. For classification, it used two fully connected layers, with the first layer based on a ReLU activation function and the second layer employing a softmax activation function.

The results in Fig. 8 and Fig. 9 show that our classifier surpassed the two classifiers in terms of accuracy, F1-score and AUC. Furthermore, as pushing detection in crowded

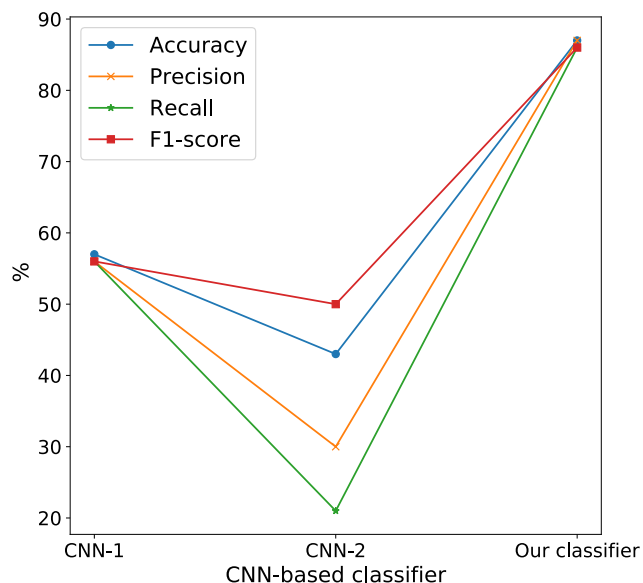


FIGURE 8. Comparison results of our classifier and two customized CNNs.

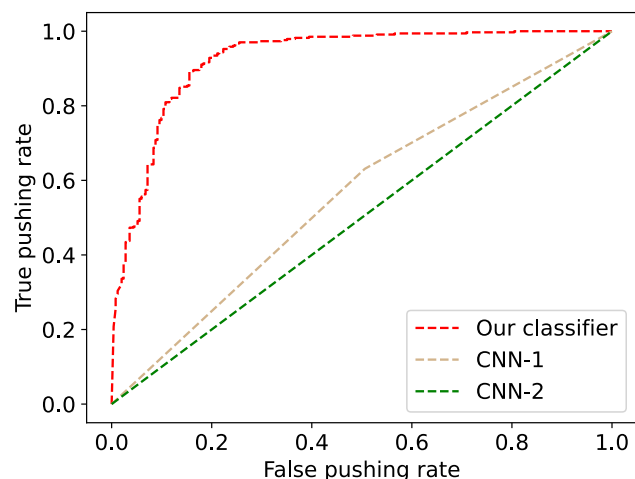


FIGURE 9. ROC curves of our classifier and the two customized CNNs.

scenarios is highly complex, CNN-1 and CNN-2’s simple architectures failed to identify pushing MIM-patches. In particular, CNN-1 outperformed CNN-2, but it still produced unsatisfactory outcomes with accuracy, F1-score, and AUC values of 57%, 56%, and 56%, respectively.

3) A COMPARISON WITH RELATED WORK IN PUSHING DETECTION

Here, we compare the proposed classifier with the CNN architecture (EfficientNetV1B0) employed in [7], which is the only published work for detecting pushing behavior for forward motion. Notably, this work does not meet the early identification requirements. As demonstrated in Table 5 and Fig. 10, our combination of adapted EfficientNetV2B0 and MIMs achieved better performance than integrating EfficientNetV1B0 with MIMs by a margin of at least 3% in accuracy and F1-score. While EfficientNetV1B0 achieved 91% AUC,

TABLE 5. Comparison results to the state-of-art pushing detection approach.

CNN	Acc. %	Pre. %	Rec. %	F1. %
State-of-art approach [7]	83	83	84	83
Our hybrid approach	87	86	87	86

“Acc.” refers to Accuracy. “Pre.” stands for Precision. “Rec.” means Recall. “F1.” means F1-score.




FIGURE 10. ROC curves of our classifier and EfficientNetV1B0.

our classifier achieved 93%. This comparison highlights that our hybrid approach surpassed the state-of-art method in pushing detection regarding the accuracy, F1-score, and AUC metrics. In Section IV-D2, we will analyze the computational time of both approaches.

To summarize the three comparisons, the new hybrid approach based on adapted EfficientNetV2B0 and MIMs outperformed all other tested combinations of CNN models and MIMs in the experiments. This superiority is due to the power of MBConv and Fused-MBConv blocks used in EfficientNetV2B0 for learning the features. Based on the experiments, it can be concluded that our classifier enhanced the performance of the proposed framework.

D. THE OVERALL FRAMEWORK EVALUATION

To evaluate the quality of the proposed framework, we not only evaluated its accuracy and F1-score, but also measured the computational time required for each framework component.

1) PERFORMANCE IN TERMS OF ACCURACY AND F1-SCORE

The evaluation methodology used comprises several steps as follows: 1) To simulate acquiring the actual inputs, we created a live video stream of crowded event entrances using video recordings of entrances (Table 2) and a virtual camera on a web client. In this context, we changed the camera’s input to the video recordings. Moreover, we down-scaled the dimensions of each video to half their original resolution to reduce the computational time of the framework. 2) We executed the cloud-based framework to display the live camera stream on the web client, detect pushing patches and record the

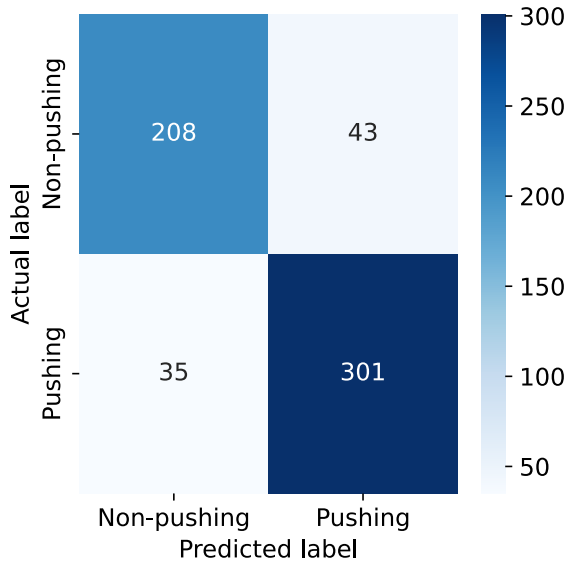


FIGURE 11. Confusion matrix for the proposed framework on all videos.

predicted labels for the test patches in a file. 3) We counted the number of true pushing, false pushing, true non-pushing, and false non-pushing for all videos by comparing the ground truth data with the predicted labels for the test patches, Fig. 11 exhibits the confusion matrix that presents them. 4) Finally, we computed the accuracy and F1-score metrics. After computing the accuracy and F1-score metrics from the values in the confusion matrix (as shown in Fig. 11), our proposed framework achieved an accuracy of 87%, precision of 87%, recall of 86%, and F1-score of 86%. These results are consistent with the corresponding quantitative outcomes in our adapted EfficientNetV2B0 classifier over the test set.

2) COMPUTATIONAL TIME ANALYSIS

In order to evaluate the overall computational time of the proposed framework, we computed the required time for each component in the framework. Then, we compared the results against the corresponding parts in the baseline framework [7]. After running both the proposed and baseline frameworks in the same environment using twenty inputs, where each input is a two-second video stream, we calculated the average run-time of all runs. As mentioned in the previous paragraph, our framework read the videos using a live camera, while the baseline framework read the same videos directly. Fig. 12 depicts the time of each component in the proposed framework over every experiment. As indicated by the produced results, the preprocessing component took more than 50% of the overall time to collect and process keyframes from the client camera stream. The resolution of the keyframes primarily determines the time required for this component. For example, experiments 110, 150, 270, and 280 took roughly the same time because they have the same resolution, while experiment entrance_2 needed less time because its resolution is lower. In contrast, the motion descriptor component took the least time compared to others, where the ROI res-

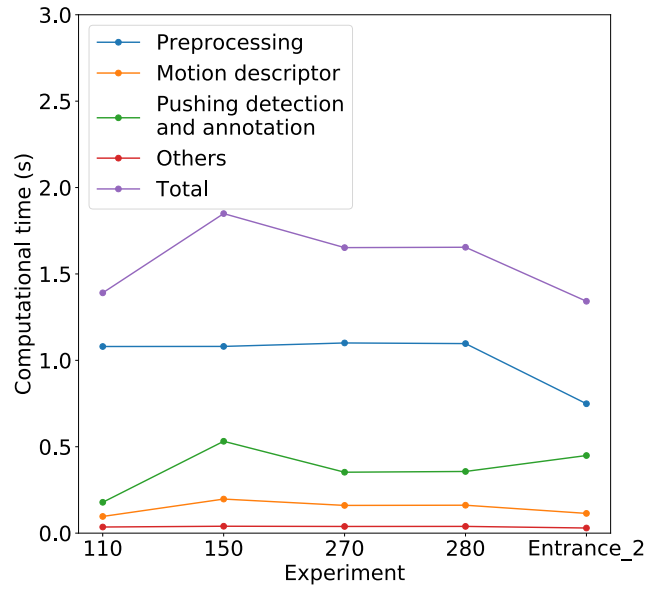


FIGURE 12. Computational time of each component of the proposed framework for annotating two seconds of stream.

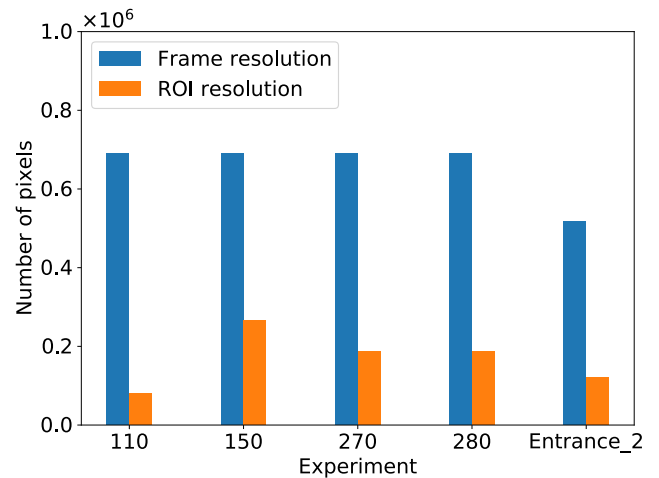


FIGURE 13. Comparison between frame and ROI resolutions for each experiment.

olution plays the most critical role in this component speed. While in the pushing detection and annotation component, the number of patches affects the computational time of this component because each patch requires one classification process. Fig. 13 and Table 2 display the ROI resolution and the number of patches in each experiment, respectively.

In general, the computation time increases as the number of patches, frames resolution, and ROI size increase. Fig. 12 shows that our framework needed less than two seconds to collect, process, detect and annotate each input from the live stream camera. This means that our framework can annotate the live camera stream within 4 seconds; two seconds for the input duration and lower than two seconds for identifying the pushing patches.

TABLE 6. The computational time of motion descriptor and detection components in our and the baseline frameworks.

Experiment	Baseline framework		Our framework	
	Motion information extraction (s)	Detection (s)	Motion descriptor (s)	Detection and annotation (s)
110	19.42	0.15	0.10	0.19
150	19.86	0.47	0.20	0.53
270	19.38	0.32	0.16	0.35
280	19.30	0.30	0.16	0.36
Entrance_2	13.51	0.41	0.11	0.45

The results in Table 6 show the comparisons between the motion descriptor, and pushing detection and annotation components in our framework with the corresponding parts in the baseline framework. The motion information extraction part in the baseline framework is similar to the motion descriptor component in our framework, whereas the motion information extraction is slow; it needs more than 13.5 seconds to generate MIM-patches from two seconds of the video stream. The main reason for this slowness is that it employed CPU-based RAFT to estimate the optical flow vectors for all pixels in the frame. To address this problem, the motion descriptor in our framework implemented RAFT on GPU to calculate the optical flow vectors for each pixel in ROIs instead of all pixels in the frame. As shown in Fig. 13, the number of pixels in ROIs is lesser than 40% of the total pixels in the corresponding frames. As a result, the new component took 0.2 seconds or less to produce MIM-patches from the two seconds of the live stream. On the other hand, the baseline framework's detection part is slightly faster than the detection and annotation component in the proposed framework. For example, the previous and new components required 0.47 and 0.53 seconds to work with one input from experiment 150, respectively. It is important to highlight that the detection part in the baseline framework only finds the labels of the patches, whereas the component in our framework labels, annotates, blurs, and stores the inputs.

In summary, the proposed cloud-based framework can annotate the pushing patches in the live camera stream within four seconds and an accuracy rate of 87%.

V. CONCLUSION

This paper proposed a novel automatic framework for the early detection of pushing patches in crowded event entrances. The proposed framework is based on live camera streaming technology, cloud environment, visualization method, and deep learning algorithms. The framework first displays the live camera stream of the entrances on the web client in real-time. Then, it relies on the color wheel method and pre-trained RAFT model to extract the visual motion information from the live stream. After that, the EfficientNetV2B0-based classifier is adapted and trained to identify pushing patches from the extracted information. Finally, the framework annotates the pushing patches in the live stream on the web client. Additionally, it stores the annotated data in the cloud storage, where the stored data is blurred to protect people's privacy. In order to train and evaluate the classifier, a new dataset was generated using

five real-world video experiments and their associated ground truth data. The experimental results show that the framework identified pushing patches from the live camera stream with 87% accuracy rate within a reasonable time delay.

One of the current limitations of the proposed framework is that it is only compatible with a fixed and top-view camera.

In future, the plan is to develop a new pushing data representation method for machine learning. This method aims to generate dynamic patches based on temporal, spatial, and size dimensions, focusing on one pedestrian for labeling each patch. This could potentially help to generate a large dataset with a more efficient sample representation.

ACKNOWLEDGMENT

The authors would like to thank Anna Sieben, Helena Lügering, and Ezel Üsten for the valuable discussions, manual annotation of the pushing behavior in the video experiments, and for revising the proposed framework's output.

ETHICAL APPROVAL

The experiments used in the dataset were conducted according to the guidelines of the Declaration of Helsinki and approved by the ethics board at the University of Wuppertal, Germany. Informed consent was obtained from all subjects involved in the experiments.

DATA AND CODE AVAILABILITY

All videos and trajectory data used in generating the patch-based dataset were obtained from the data archive hosted by the Forschungszentrum Jülich under CC Attribution 4.0 International license [33], [38]. The undistorted video experiments, implementation of the proposed framework, as well as codes used for building and training the models, are publicly available at: <https://github.com/PedestrianDynamics/CloudFast-DL4PuDe> (accessed on 15 Jan 2023). The generated patch-based dataset is available from the corresponding authors upon request.

REFERENCES

- [1] J. Adrian, A. Seyfried, and A. Sieben, "Crowds in front of bottlenecks at entrances from the perspective of physics and social psychology," *J. Roy. Soc. Interface*, vol. 17, no. 165, Apr. 2020, Art. no. 20190871.
- [2] E. Üsten, H. Lügering, and A. Sieben, "Pushing and non-pushing forward motion in crowds: A systematic psychological observation method for rating individual behavior in pedestrian dynamics," *Collective Dyn.*, vol. 7, pp. 1–16, Aug. 2022.
- [3] M. Haghani, M. Sarvi, and Z. Shahhoseini, "When 'push' does not come to 'shove': Revisiting 'faster is slower' in collective egress of human crowds," *Transp. Res. A, Policy Pract.*, vol. 122, pp. 51–69, Apr. 2019.

- [4] V. Filingeri, K. Eason, P. Waterson, and R. Haslam, "Factors influencing experience in crowds—The participant perspective," *Appl. Ergonom.*, vol. 59, pp. 431–441, Mar. 2017.
- [5] N. R. Johnson, "Panic at 'the who concert stampede': An empirical assessment," *Social Problems*, vol. 34, no. 4, pp. 362–373, Oct. 1987.
- [6] B. Tyagi, S. Nigam, and R. Singh, "A review of deep learning techniques for crowd behavior analysis," *Arch. Comput. Methods Eng.*, vol. 29, no. 7, pp. 1–29, 2022.
- [7] A. Alia, M. Maree, and M. Chraibi, "A hybrid deep learning and visualization framework for pushing behavior detection in pedestrian dynamics," *Sensors*, vol. 22, no. 11, p. 4040, May 2022.
- [8] A. Mehmood, "Efficient anomaly detection in crowd videos using pre-trained 2D convolutional neural networks," *IEEE Access*, vol. 9, pp. 138283–138295, 2021.
- [9] A. Al-Dhamari, R. Sudirman, and N. H. Mahmood, "Transfer deep learning along with binary support vector machine for abnormal behavior detection," *IEEE Access*, vol. 8, pp. 61085–61095, 2020.
- [10] C. Direkoglu, "Abnormal crowd behavior detection using motion information images and convolutional neural networks," *IEEE Access*, vol. 8, pp. 80408–80416, 2020.
- [11] A. Alia and A. Taweel, "Enhanced binary cuckoo search with frequent values and rough set theory for feature selection," *IEEE Access*, vol. 9, pp. 119430–119453, 2021.
- [12] A. F. Alia and A. Taweel, "Feature selection based on hybrid binary cuckoo search and rough set theory in classification for nominal datasets," *Algorithms*, vol. 14, no. 21, p. 65, 2017.
- [13] H. Gan, C. Xu, W. Hou, J. Guo, K. Liu, and Y. Xue, "Spatiotemporal graph convolutional network for automated detection and analysis of social behaviours among pre-weaning piglets," *Biosyst. Eng.*, vol. 217, pp. 102–114, May 2022.
- [14] H. Gan, M. Ou, E. Huang, C. Xu, S. Li, J. Li, K. Liu, and Y. Xue, "Automated detection and analysis of social behaviors among preweaning piglets using key point-based spatial and temporal features," *Comput. Electron. Agricult.*, vol. 188, Sep. 2021, Art. no. 106357.
- [15] L. Li, S. Zhang, and B. Wang, "Apple leaf disease identification with a small and imbalanced dataset based on lightweight convolutional networks," *Sensors*, vol. 22, no. 1, p. 173, Dec. 2021.
- [16] P. Wang, E. Fan, and P. Wang, "Comparative analysis of image classification algorithms based on traditional machine learning and deep learning," *Pattern Recognit. Lett.*, vol. 141, pp. 61–67, Jan. 2021.
- [17] E. Duman and O. A. Erdem, "Anomaly detection in videos using optical flow and convolutional autoencoder," *IEEE Access*, vol. 7, pp. 183914–183923, 2019.
- [18] T. Alafif, A. Hadi, M. Allahyani, B. Alzahrani, A. Alhothali, R. Alotaibi, and A. Barnawi, "Hybrid classifiers for spatio-temporal abnormal behavior detection, tracking, and recognition in massive Hajj crowds," *Electronics*, vol. 12, no. 5, p. 1165, Feb. 2023.
- [19] Z. Teed and J. Deng, "RAFT: Recurrent all-pairs field transforms for optical flow," in *Proc. Eur. Conf. Comput. Vis.* Glasgow, U.K.: Springer, 2020, pp. 402–419.
- [20] D. F. T. Runia, "Optical flow visualization," Tech. Rep., 2020. [Online]. Available: https://github.com/tomrunia/OpticalFlow_Visualization
- [21] S. Baker, D. Scharstein, J. P. Lewis, S. Roth, M. J. Black, and R. Szeliski, "A database and evaluation methodology for optical flow," *Int. J. Comput. Vis.*, vol. 92, no. 1, pp. 1–31, Nov. 2010.
- [22] M. Tan and Q. Le, "EfficientNetV2: Smaller models and faster training," in *Proc. Int. Conf. Mach. Learn.*, 2021, pp. 10096–10106.
- [23] P. Kuppusamy and V. C. Bharathi, "Human abnormal behavior detection using CNNs in crowded and uncrowded surveillance—A survey," *Meas. Sensors*, vol. 24, Dec. 2022, Art. no. 100510.
- [24] N. C. Tay, T. Connie, T. S. Ong, K. O. M. Goh, and P. ShenTeh, "A robust abnormal behavior detection method using convolutional neural network," in *Computational Science and Technology*. Kota Kinabalu, Malaysia: Springer, 2019, pp. 37–47.
- [25] M. Sabokrou, M. Fayyaz, M. Fathy, Z. Moayed, and R. Klette, "Deep-anomaly: Fully convolutional neural network for fast anomaly detection in crowded scenes," *Comput. Vis. Image Understand.*, vol. 172, pp. 88–97, Jul. 2018.
- [26] M. Xu, X. Yu, D. Chen, C. Wu, and Y. Jiang, "An efficient anomaly detection system for crowded scenes using variational autoencoders," *Appl. Sci.*, vol. 9, no. 16, p. 3337, Aug. 2019.
- [27] S. Smeureanu, R. T. Ionescu, M. Popescu, and B. Alexe, "Deep appearance features for abnormal behavior detection in video," in *Proc. Int. Conf. Image Anal. Process.* Catania, Italy: Springer, 2017, pp. 779–789.
- [28] Z. Ilyas, Z. Aziz, T. Qasim, N. Bhatti, and M. F. Hayat, "A hybrid deep network based approach for crowd anomaly detection," *Multimedia Tools Appl.*, vol. 80, pp. 1–15, Mar. 2021.
- [29] S. S. Khan and M. G. Madden, "One-class classification: Taxonomy of study and review of techniques," *Knowl. Eng. Rev.*, vol. 29, no. 3, pp. 345–374, Jan. 2014.
- [30] Y. Hu, "Design and implementation of abnormal behavior detection based on deep intelligent analysis algorithms in massive video surveillance," *J. Grid Comput.*, vol. 18, pp. 227–237, Aug. 2020.
- [31] A. A. Almazroey and S. K. Jarraya, "Abnormal events and behavior detection in crowd scenes based on deep learning and neighborhood component analysis feature selection," in *Proc. Joint Eur.-U.S. Workshop Appl. Invariance Comput. Vis.* Cairo, Egypt: Springer, 2020, pp. 258–267.
- [32] S. Zhou, W. Shen, D. Zeng, M. Fang, Y. Wei, and Z. Zhang, "Spatial-temporal convolutional neural networks for anomaly detection and localization in crowded scenes," *Signal Process., Image Commun.*, vol. 47, pp. 358–368, Sep. 2016.
- [33] *Entrance 2, Entry With Guiding Barriers (Corridor Setup)*, IAS-7, Forschungszentrum Jülich, Jülich, Germany, 2013, doi: [10.34735/ped.2013.1](https://doi.org/10.34735/ped.2013.1).
- [34] G. Farneback, "Two-frame motion estimation based on polynomial expansion," in *Proc. Scand. Conf. Image Anal.* Halmstad, Sweden: Springer, 2003, pp. 363–370.
- [35] Z. Yin, T. Darrell, and F. Yu, "Hierarchical discrete distribution decomposition for match density estimation," in *Proc. IEEE/CVF Conf. Comput. Vis. Pattern Recognit. (CVPR)*, Jun. 2019, pp. 6044–6053.
- [36] S. Gupta and M. Tan, "EfficientNet-EdgeTPU: Creating accelerator-optimized neural networks with AutoML," *Google AI Blog*, vol. 2, p. 1, Aug. 2019. [Online]. Available: <https://ai.googleblog.com/2019/08/efficientnet-edgetpu-creating.html>
- [37] M. Sandler, A. Howard, M. Zhu, A. Zhmoginov, and L.-C. Chen, "MobileNetV2: Inverted residuals and linear bottlenecks," in *Proc. IEEE/CVF Conf. Comput. Vis. Pattern Recognit.*, Jun. 2018, pp. 4510–4520.
- [38] *Crowds in Front of Bottlenecks From the Perspective of Physics and Social Psychology*, IAS-7, Forschungszentrum Jülich, Jülich, Germany, 2018, doi: [10.34735/ped.2018.1](https://doi.org/10.34735/ped.2018.1).
- [39] B. Genç and H. Tunç, "Optimal training and test sets design for machine learning," *TURKISH J. Electr. Eng. Comput. Sci.*, vol. 27, no. 2, pp. 1534–1545, 2019.
- [40] Z. DeVries, E. Locke, M. Hoda, D. Moravek, K. Phan, A. Stratton, S. Kingwell, E. K. Wai, and P. Phan, "Using a national surgical database to predict complications following posterior lumbar surgery and comparing the area under the curve and F1-score for the assessment of prognostic capability," *Spine J.*, vol. 21, no. 7, pp. 1135–1142, Jul. 2021.
- [41] F. Chollet, "Xception: Deep learning with depthwise separable convolutions," in *Proc. IEEE Conf. Comput. Vis. Pattern Recognit. (CVPR)*, Jul. 2017, pp. 1251–1258.
- [42] K. Simonyan and A. Zisserman, "Very deep convolutional networks for large-scale image recognition," 2014, *arXiv:1409.1556*.
- [43] K. He, X. Zhang, S. Ren, and J. Sun, "Deep residual learning for image recognition," in *Proc. IEEE Conf. Comput. Vis. Pattern Recognit. (CVPR)*, Jun. 2016, pp. 770–778.
- [44] K. He, X. Zhang, S. Ren, and J. Sun, "Identity mappings in deep residual networks," in *Proc. Eur. Conf. Comput. Vis.* Springer, 2016, pp. 630–645.
- [45] G. Huang, Z. Liu, L. Van Der Maaten, and K. Q. Weinberger, "Densely connected convolutional networks," in *Proc. IEEE Conf. Comput. Vis. Pattern Recognit. (CVPR)*, Jul. 2017, pp. 4700–4708.
- [46] B. Zoph, V. Vasudevan, J. Shlens, and Q. V. Le, "Learning transferable architectures for scalable image recognition," in *Proc. IEEE/CVF Conf. Comput. Vis. Pattern Recognit.*, Jun. 2018, pp. 8697–8710.



AHMED ALIA received the B.Sc. degree in computer science from An-Najah National University, Palestine, and the M.Sc. degree in computing from Birzeit University, Palestine. He is currently pursuing the Ph.D. degree in information technology with the Institute for Advanced Simulation, Forschungszentrum Jülich, in collaboration with the University of Wuppertal, Germany. From 2008 to 2020, he was a Teaching Assistant, a Web Developer, and a Trainer with An-Najah National University. His research interests include machine learning, deep learning, computer vision, data analysis, intelligent systems, and crowd behavior analysis.



MOHAMMED MAREE received the Ph.D. degree in information technology from Monash University. He began his career as a Research and Development Manager with gSoft Technology Solution Inc. Then, he was the Director of research and QA with Dimensions Consulting Company. Subsequently, he joined the Faculty of Engineering and Information Technology (EIT), Arab American University, Palestine (AAUP), as a full-time Lecturer. From September 2014 to August

2016, he was the Head of the Department of Multimedia Technology, and from September 2016 to August 2018, he was the Head of the Department of Information Technology. In addition to his work with AAUP, he was a Consultant for SocialDice and Dimensions Consulting companies. He is currently an Associate Professor of information technology and an Assistant to the Vice President for Academic Affairs with Arab American University. He has supervised a number of master's and Ph.D. students in the fields of knowledge engineering, data analysis, information retrieval, natural language processing, and hybrid intelligent systems. He has published papers in various high-impact journals and conferences, such as ICTAI, *Knowledge-Based Systems, Behaviour and Information Technology, Journal on Computing and Cultural Heritage, Information Development, and Journal of Information Science*. He is a committee member/reviewer of several conferences and journals.



MOHCINE CHRAIBI received the Diploma degree in computer sciences from the Technical University of Hamburg, Harburg, Germany, and the Ph.D. degree from the Institute for Theoretical Physics, University of Cologne, in 2012, under the supervision of Prof. A. Schadschneider and Prof. Armin Seyfried. He was a Research Fellow of JSPS with Tokyo University. Since March 2017, he is the Head of the Pedestrian Dynamics-Modelling Division, Research Centre Jich, Civil

Safety Research Institute. His current research interest includes the interdisciplinary study of the collective dynamics of self-driven particles and its jamming phenomena by means of experiments and simulations.



ANAS TOMA received the B.Sc. degree from An-Najah National University, Palestine, the M.Sc. degree from the Jordan University of Science and Technology (JUST), Jordan, and the Ph.D. degree in computer engineering from the Karlsruhe Institute of Technology (KIT), Germany. He was a Teaching and Research Assistant for one year and a Lecturer for three years with An-Najah National University. From 2016 to 2019, he was a Postdoctoral Researcher with TU-

Dortmund University, Germany. He is currently an Assistant Professor with the Department of Computer Engineering and also the Coordinator of the Artificial Intelligence Master's Program, An-Najah National University.



ARMIN SEYFRIED received the degree in theoretical physics from the University of Wuppertal. For his Diploma and Ph.D. thesis, he focused on many-particle systems, high-energy physics, and parallel computing. After the Ph.D. degree, he was responsible for the modeling and simulation of building evacuation in an engineering office. From 2004 to 2018, he was with the Jülich Supercomputing Centre, Forschungszentrum Jülich, developing models and simulations

for application in civil security and traffic planning, where he established a new research group for pedestrian and fire dynamics, which became the Institute for Advanced Simulation-7, in 2018. Since 2010, he is a Professor with the Department of Computer Simulations for Fire Protection and Pedestrian Traffic, University of Wuppertal.

...