
A COMPUTER VISION SYSTEM TO HELP PREVENT THE TRANSMISSION OF COVID-19

Fevziye Irem Eyiokur

Karlsruhe Institute of Technology, Karlsruhe, Germany
fevziye.yaman@kit.edu

Hazım Kemal Ekenel

Istanbul Technical University, Istanbul, Turkey
ekenel@itu.edu.tr

Alexander Waibel

Karlsruhe Institute of Technology, Karlsruhe, Germany
Carnegie Mellon University, Pittsburgh PA, USA
alexander.waibel@kit.edu

May 5, 2021

ABSTRACT

The COVID-19 pandemic affects every area of daily life globally. To avoid the spread of coronavirus and retrieve the daily normal worldwide, health organizations advise social distancing, wearing face mask, and avoiding touching face. Based on these recommended protective measures, we developed a deep learning-based computer vision system to help prevent the transmission of COVID-19. Specifically, the developed system performs face mask detection, face-hand interaction detection, and measures social distance. For these purposes, we collected and annotated images that represent face mask usage and face-hand interaction in the real world. We introduce two different face datasets, namely, Interactive Systems Labs Unconstrained Face Mask Dataset (ISL-UFMD) and Interactive Systems Labs Unconstrained Face Hand Interaction Dataset (ISL-UFHD). We trained the proposed models on our own datasets and evaluated them on both our datasets and already existing datasets in the literature without performing any adaptation on these target datasets. Besides, we proposed a distance measurement module to track social distance between people. Experimental results indicate that ISL-UFMD and ISL-UFHD represent the real-world's diversity well. The proposed system achieved very high performance and generalization capacity in a real-world scenario for unseen data from outside the training data to detect face mask usage, face-hand interaction detection, and measuring social distance. The ISL-UFMD and ISL-UFHD datasets will be available at <https://github.com/iremeyiokur/COVID-19-Preventions-Control-System>.

Keywords COVID-19 · Face Mask Detection · Face-hand Interaction Detection · Social Distance Measurement · CNN

1 Introduction

The COVID-19 pandemic, which is caused by severe acute respiratory syndrome coronavirus (SARS - CoV-2) [1], has affected the whole world since the beginning of 2020 and it still continues to affect worldwide daily life. In order to decrease the transmission of the COVID-19 disease, many health institutions, particularly the World Health Organization (WHO), have recommended serious constraints and preventions [2]. The fundamental precautions that individuals can carry out, are to keep the distance from others (practicing social distance) [3], wear a face mask properly (covering mouth and nose), pay attention to personal hygiene, especially hand hygiene, and avoid touching faces with hands without cleanliness [2].

Convolutional Neural Networks (CNNs), which are introduced in late 80's [4, 5], have gained popularity during the last decade with the help of the deep architectures, powerful hardware, and big data. Based on the increasing popularity and success of deep learning in computer vision, novel research topics that emerged as a consequence of the COVID-19

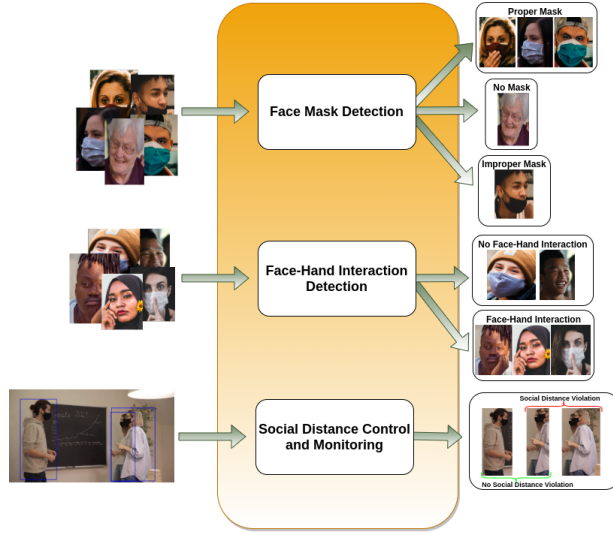


Figure 1: Overview of the executed tasks to develop the proposed system to avoid transmission of COVID-19.

pandemic are handled in this context by researchers. These studies focus on diagnosing COVID-19 disease [6, 7, 8, 9], adjusting the already existing surveillance systems to COVID-19 conditions [10, 11, 12, 13, 14, 15], and building systems to control the preventions [16, 17, 18, 11, 19, 20, 21, 22, 23, 24, 25, 26, 27, 28]. While some of the studies employ CT scans [6, 7] to diagnose COVID-19 disease by detecting typical features that are common in almost all the COVID-19 patients, the others benefit from chest X-ray images [8, 9]. Additionally, face detection and recognition systems' performance deteriorates when subjects wear face masks. Thus, novel face recognition and detection studies [10, 12, 13] try to improve the performance under the condition in which subjects use face masks to cover most parts of their faces. Besides, the age prediction [14] is investigated when face mask is used. Moreover, in order to track the execution of preventions against the spread of COVID-19, several works investigate the detection of wearing a mask suitably [16, 17, 18, 11, 19, 20, 21, 22, 23, 24] and keeping the distance from other people [20, 22, 25, 26, 27, 28]. In addition to these studies, there is a study [29] that focuses on face-hand interaction with the concept of the detection of COVID-19 prevention.

In this work, we focus on building a computer vision system to help prevent the spread of COVID-19. In this context, we present a deep learning-based system that controls preventions based on the advice of the health institutions. Three crucial preventions that can be covered using a computer vision application are to detect whether people wear a face mask, keep away from touching their faces, and to monitor whether the social distance is preserved between people or not. To investigate the detection of face mask and face-hand interaction subtasks and improve the performance of the system for these subtasks, we present two novel face datasets, namely Interactive Systems Labs Unconstrained Face Mask Dataset (ISL-UFMD) and Interactive Systems Labs Unconstrained Face Hand Interaction Dataset (ISL-UFHD), that are collected from the web to provide a significant amount of variation in terms of pose, illumination, resolution, and ethnicity. We implemented the system as the combination of three submodules, corresponding to face mask detection, face-hand interaction detection, and social distance measurement tasks, respectively. First, we trained two separate deep CNN models to classify face images for the face mask detection and face-hand interaction detection tasks. While the first model classifies the face image as wearing a mask properly, wearing a mask improperly, or not wearing a mask, the second model classifies face images as touching the face or not touching the face. To investigate the performance of different models, we employed several well-known deep CNN architectures and analyzed their performance for the different cases. Besides, we provided class activation maps (CAM) [30] to investigate the trained models' activations for both tasks. In the end, we evaluated our best models on existing face mask datasets in the literature without training on them. We also proposed an approach to measure the social distance which is based on a person detector. Overview of the proposed system is shown in Fig. 1. Our contributions can be summarized as follows:

- We provide a vision-based system to help people to follow the recommended protective measures –wearing a face mask properly, not touching faces, and having social distance between each other– to avoid spread of COVID-19. Face mask and face-hand interaction detection modules consist of a face detector and trained CNN models to predict related detection. Social distance measurement module is based on a deep learning-based person detector's outputs –bounding box and shoulder points of the same person.

- We present two novel datasets, namely, ISL Unconstrained Face Mask Dataset (ISL-UFMD) and ISL Unconstrained Face Hand Interaction Dataset (ISL-UFHD) for face mask and face-hand interaction detection tasks to evaluate proposed prevention tracking and control system for COVID-19. ISL-UFMD is one of the largest face mask dataset that includes images from unconstrained real-world scenes. In addition, the ISL-UFHD is the first dataset that contains images from unconstrained real-world scenes, and it has large number of subjects and images from various conditions.
- We provide a comprehensive analysis of the experimental results. We extensively investigate several deep CNN models trained on our proposed datasets and also tested them on publicly available masked face datasets without training on them to demonstrate the generalization capacity of our models. We achieved very high classification accuracies on these experiments, which indicates the collected datasets' capability to represent real-world cases and trained models' ability to generalize. Moreover, in order to evaluate the overall system, we utilized six different short real-world video recordings.

The rest of the paper is organized as follows. In Section 2, we review some recent papers, which investigate vision-based problems on COVID-19 research area. In Section 3, we present the details of the proposed datasets and explain how we collected and annotated them. In Section 4, we introduce our whole system, used CNN methods, and training procedures. Then, in Section 5, we provide information about the used publicly available datasets for the tests, experimental setups and results, and overall system performance on test videos. Finally, Section 6 concludes the paper.

2 Related Work

During the 2020 pandemic, numerous studies have been published in the computer vision research field to prevent, control, and diagnose COVID-19 infection using various imagery. Most of the studies mainly focused on making a diagnosis through CT and X-Ray scans and tracking COVID-19 preventions [31]. In [6], authors constructed a deep learning-based model using around 46000 computed tomography (CT) images of 106 approved patients to diagnose COVID-19 disease. In the experiments, it is stated that the performance for diagnosing COVID-19 per person reached 95.24% accuracy. In [7], the authors implemented a ResNet50 [32] backbone deep CNN model, namely COVID-19 Detection Neural Network (COVNet), and they used 4356 chest CT images and obtained over 90% accuracy for diagnosing COVID-19. In [8], 16756 chest radiography images that belong to 13645 patients are utilized to fine-tune a deep CNN based on ResNet50 [32] using a pre-trained model.

On the other hand, some of the studies on COVID-19 concentrated to prevent and control the spread of infection in the social areas and they utilized images, video footage, and sensor data. The studies herein can be categorized as face mask detection, social distance tracking, or generic COVID-19 warning systems.

In [11], a novel masked face recognition dataset is published for improving the face recognition performance in the case of occlusion due to face masks. This dataset contains three different sub-datasets which are Masked Face Detection Dataset (MFDD), Real-world Masked Face Recognition Dataset (RMFRD), and Simulated Masked Face Recognition Dataset (SMFRD). Each part contains 24771 masked face images collected from other publicly available datasets and the internet, 5000 masked face and 90000 non-masked face images that belong to 525 celebrities collected with web crawling, and 500000 masked face images of 10000 subjects constructed with a mask generation tool for including mask on the mouth of the subjects of popular face recognition datasets, respectively. All sub-datasets except MFDD, which is more comparable to our dataset, are publicly available. In [16], a large scale masked face dataset, named as MaskedFace-Net, which includes Correctly Masked Face Dataset (CMFD) and Incorrectly Masked Face Dataset (IMFD), is presented and there are 137016 images in total. Besides, the authors presented a mask-to-face deformable model to create this masked face dataset from the Flickr-Faces-HQ3 (FFHQ) face dataset [33]. Unlike the other publicly available datasets on masked face detection, MaskedFace-Net contains sample images to represent improper usage of a mask and these images can be used for the classification of proper/improper mask usage. Joshi et. al [17] proposed a framework to detect whether people are wearing a mask or not in public areas. They utilized MTCNN [34] and MobileNetV2 [35] to detect faces and classify them on their own video dataset. In [10], a one-stage detector is proposed to detect faces and classify them whether they contain masks or not. The proposed RetinaFaceMask model is designed based on RetinaFace detector [36] that is ResNet-backboned [32] and MobileNet-backboned [37]. In [18], the authors proposed a real-time face mask detector framework named SSDMNV2, which is composed of Single Shot Multibox Detector [38] as a face detector and MobileNetV2 [35] as a mask classifier. The system is designed to be used suitably in embedded devices like NVIDIA Jetson Nano by taking advantage of light-weight MobileNetV2 architecture. In the experiments, a novel face mask dataset that contains 5521 masked and unmasked face images is presented and 92.6% accuracy is achieved.

A recent study [29] investigated the face-hand touching behavior. In this study, the authors presented face-hand touching interaction annotations on 64 video recordings which are collected for investigating social interactions on a small group



Figure 2: Example images from ISL Unconstrained Face Mask Dataset (ISL-UFMD). This figure shows the sample images that belong to three different classes; no mask, face mask, improper face mask.



Figure 3: Example images from ISL Unconstrained Face Hand Interaction Dataset (ISL-UFHD). This dataset contains images that represent face-hand interaction and no interaction. Besides, there are occluded face images by face mask.

of people (four-people meeting). In addition to these annotations of face-hand touching interactions, they evaluated the annotated 2M and 74K frames with rule-based, hand-crafted feature-based, and CNN learned feature-based models. As a result of evaluations, CNN based model obtained the best results with 83.76% F1-score.

Different from these works, we collected real-world datasets with a high amount of variety in terms of subject diversity, head pose, illumination, and ethnicity. In addition to face mask detection task, we investigated the face-hand interaction detection. We presented the first work that collected unconstrained face-hand interaction dataset under the real-world conditions, and integrated face-hand detection task in a comprehensive prevention system in the context of COVID-19. Moreover, we address the problem of measuring social distance.

3 The Proposed Datasets

To train our system, we collected both face masked images and face-hand interaction images. Recently published datasets on the tracking of COVID-19 preventions, which are presented in Table 1, fundamentally focused on collecting face mask images to develop a system that examines whether there is a mask on the face or not. Most of them contain a limited amount of images or include synthetic images generated with putting a mask on the face using landmark points around the mouth and nose. Besides, the variety of subjects' ethnicity, image conditions such as environment, resolution, and particularly different head pose variations are limited in the existing datasets. Thus, these limitations led us to collect a dataset to overcome all these drawbacks. In addition to face mask, there is only one dataset [29] that is recently annotated to investigate face-hand interaction in the literature. However, these face-hand interaction annotations are also limited based on the number of subjects and the dataset is collected in an indoor environment under the controlled conditions. Therefore, we present the first work that collected images from unconstrained real world




Figure 4: Example images from existing face mask datasets. Different from ISL-UFMD dataset which is shown in Fig. 2, most of the examples have Asian subjects and head poses of subjects are close to the frontal.

Table 1: Comparison of the face mask datasets. (*) Although it is stated that RMFD dataset [26] contains 5000 face images with mask, there are only 2203 face images with mask in the publicly available version.

Dataset name	No mask	Mask	Improper Mask	Data Type	Ethnicities	Head Pose
UFMD	10698	10618	500	Real	Various	Various
RMFD [26]*	90468	2203	-	Real	Asian	Frontal to Profile
RWMFD [26]	858	4075	238	Real	Mostly Asian	Frontal to Profile
Face mask [39]	718	3239	123	Real	Mostly Asian	Various
MaskedFace-Net [16]	-	67049	66734	Artificial	Various	Mostly Frontal

scenes to present a face-hand interaction dataset to track whether the people are touching their faces. We collected and then annotated face images for both tasks to train our proposed system’s components. Moreover, we gathered some videos to evaluate and further analyze the integrated system.

3.1 Data Collection

We collected a large amount of face images to represent the masked face and face-hand interaction from several different resources such as publicly available face datasets, – FFHQ [33], CelebA [40], LFW [41] –, YouTube videos, and web crawling from websites that contain free licensed images. These various sources enable us to collect a significant variety of human face images in terms of ethnicity, age, and gender. In addition to the subject diversity, we obtained images from indoor and outdoor environments, under different light conditions and resolutions to cover unconstrained conditions. We also considered ensuring large head pose variations to represent real-world scenarios and make the system more robust against these conditions. Moreover, another important key point that we take into account is to leverage the performance of our COVID-19 prevention system for the common scenario, e.g., determining mask usage in the case of touching faces or detecting face-hand interaction in the case of wearing a mask. Besides, our images include different sorts of occlusion that make the dataset more challenging. In the end, we collected 21316 face images for the face-mask detection scenario, 10618 face images with masks and 10698 images without a mask. In addition to these two base classes, we gathered additional 500 images that represent improper mask usage. The improper face mask class has a relatively small number of images compared to no mask and mask classes. One of the main reasons for this is the difficulty of finding images for improper mask usage. We named our face mask dataset as Interactive Systems Labs Unconstrained Face Mask Dataset (ISL-UFMD).

The other dataset that we proposed in this work is the Interactive Systems Labs Unconstrained Face Hand Interaction Dataset (ISL-UFHD). This dataset is composed of face images that represent the interaction between the face and hand of the subjects. We collected 22289 negative samples (no face-hand interaction) and 10004 positive samples (face-hand interaction). Please note that, even if the hand is around the face without touching it, we annotated it as a no interaction. Therefore, the model should be able to distinguish whether the hand in the image is touching the face (or very close to the face) or not.

3.2 Data Annotation

For labelling the ISL-UFMD and ISL-UFHD, we designed a web-based image annotation tool. We utilized crowd-sourcing to annotate each image and after examining these annotations, we decided each image’s final label. Since we formulated our tasks as classification problems, we annotated our images in that manner. While we have three classes –mask, no mask, improper mask– for the mask detection task, we have two for the face-hand interaction detection task.

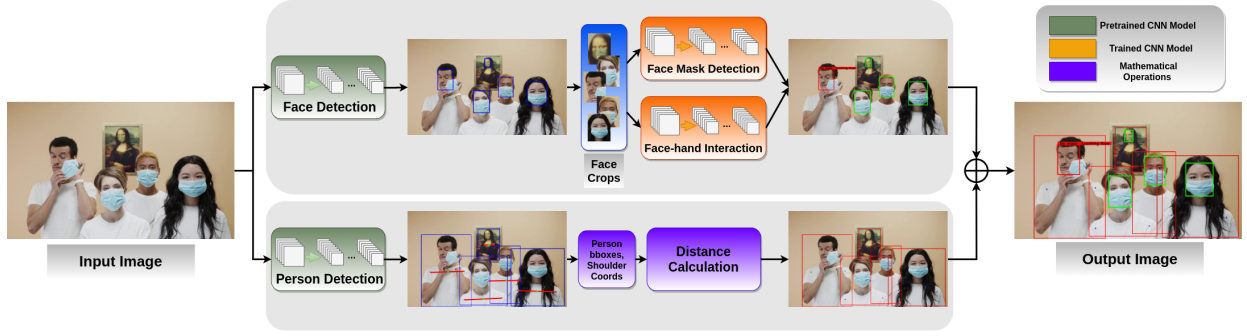


Figure 5: Visualization of the proposed system for face mask detection, face-hand interaction detection, and social distance controlling tasks. Firstly, the input data is sent to the face detection and person detection models separately yet simultaneously. After face detection model detects faces and these detections are used to obtain face crops with adequate margins, they are sent to the face mask model and face-hand model separately. At the same time, the detected people are sent to the pose estimation model to get shoulder points. Later, the Euclidean distance between each person is calculated based on the center point of the detected shoulder points. In the end, the system decides based on the distance and adaptively calculated threshold for each pair of persons individually. In the overall system, all outcomes are shown on the input data.

During annotation, we considered the advice of health institutions for both face mask and face-hand interaction tasks. For instance, if a person has a mask under the chin, we annotated the image with the no mask label instead of the improper mask label, since it is clear that the person’s intent is not to wear a mask. The images that include the face without a fully covered nose and mouth by the mask are annotated with the improper mask label. In the face-hand annotation, we considered the direct contact or too close to contact as the existence of face-hand interaction. Many examples of annotated face images for face mask and face-hand interaction tasks are shown in Fig. 2 and Fig. 3. It can be clearly seen from the Fig. 2 and Fig. 3 that our proposed datasets have large amount of variations especially for gender, ethnicity, and head pose. Also, the examples have diversity in terms of position of hand upon face and usage of face mask. In Fig. 4, we showed some examples of publicly available face mask datasets. When the sample images are examined, it can be seen that Asian people are in the majority in the first three datasets, RMFD [26], RWTFD [26], and Face-Mask (Kaggle) [39]. Although MaskedFace-Net dataset [16] includes variation in terms of ethnicity, it consists entirely of images with artificial face masks. While all face mask datasets have limited head poses mostly from frontal view to profile view in yaw axis, our proposed datasets contain face images with more head pose variations in terms of covering all the yaw, pitch, and roll axis.

4 Methodology

We proposed a deep learning based system to protect people from COVID-19 transmission. In the following subsections, we first give an overview of the developed system and then present its components in detail.

4.1 The Proposed System

In this paper, we proposed a comprehensive COVID-19 prevention control system which is illustrated in Fig. 5. The proposed system consists of three sub-modules and each module utilizes deep CNN models to obtain predictions. The system performs person detection and calculates distances between detected subjects on input image/video frame. Meanwhile, the same input is also used to detect and crop faces of subjects to perform the face mask and face-hand interaction detections. While the face mask model decides whether a person wears a mask (properly) or not, the face-hand interaction model identifies whether a hand touches the subject’s face. We decided to conduct person detection and face detection separately on the input image/video frame to eliminate the effect of missing modality. For instance, although a person’s body is occluded and social distancing cannot be measured with this person, system can still detect the face of the corresponding subject to perform face mask and face-hand interaction tasks. Similarly, if the subject’s face is occluded or not turned to the camera, system can be able to capture the person’s body to perform the social distance task.

4.1.1 Face mask detection and face-hand interaction detection

To execute face mask and face-hand interaction tasks, firstly, we performed face detection using the pretrained ResNet50 [32] backbone RetinaFace model [36] that was trained on the large-scale Wider-Face dataset [42]. When choosing a proper face detector, we executed MTCNN [34] and Dlib [43] face detectors to obtain face crops as alternative to the RetinaFace detector. However, according to the experiments, we found that MTCNN and Dlib models have two main drawbacks: First, they are not robust against tiny faces. Second, detection performance of models is significantly decreasing when subjects wear face masks on different head poses. On the contrary, the RetinaFace model [36] is robust against tiny faces and faces with a mask. For that reason, we decided to use RetinaFace model as a face detection model. After detection, we cropped detected faces with a 20% margin for each side since the face detector’s outputs are quite tight. To perform face mask and face-hand interaction detections, we employed several different deep CNN architectures that are ResNet50 [32], Inception-v3 [44], MobileNetV2 [35], and EfficientNet [45]. We decided to use EfficientNet since it is the state-of-the-art model and we preferred to use MobileNetV2 since it is a light-weight deep CNN model. Finally, we chose ResNet and Inception-v3 models based on their accurate performances in the literature. ResNet50 includes 50 convolutional layers with residual connections followed by a global average pooling layer. The necessity of fully-connected layers is eliminated in ResNet architecture and it is used only for producing output in the end. The other architecture that we used is Inception-v3, which is a more accurate and computationally efficient version of the original Inception (GoogLeNet) architecture [46] with the help of the factorization of the convolutional layers. In GoogLeNet architecture, inception modules are presented and these modules apply different size convolution filters on the same level (wider instead of deeper). The following architecture is MobileNetV2 that is one of the most light-weight architectures. It reduces the complexity of the model with depthwise factorized convolutional layers. It also has novel bottleneck modules which are represented as inverted residual blocks, and applies depthwise convolutions to the high dimensional representation that is obtained from low dimensional representation. The last architecture, EfficientNet, is based on a light-weight backbone model with the compound scaling approach to obtain the best performance by scaling different dimensions of depth, width, and resolution. With the help of this approach, EfficientNet achieved state-of-the-art results on ImageNet [47] and several other datasets for the image classification task. In the training, we benefited from transfer learning and initialized our networks with the weights of the pretrained models that were trained on ImageNet dataset [47]. We employed softmax loss at the end of each network. In EfficientNet and MobileNetV2, we utilized dropout with a 0.2 probability rate to avoid overfitting. The input sizes of the networks are 224×224 , 256×256 , 299×299 for MobileNetV2, ResNet50, and Inception-v3, respectively. For the EfficientNet, we employed four different versions (EfficientNet-b0, b1, b2, b3) with 224×224 , 240×240 , 260×260 , and 300×300 input sizes.

We addressed the mask classification task as a multi-class classification –*improper mask, proper mask, no mask*– and a binary classification task –*mask, no mask*– separately. While our goal is to identify the existence and proper usage of the mask by focusing around the mouth and nose of the face with our model, the model should discriminate non-mask occlusion from the mask and improper usage from the proper usage. Further, we handled the face-hand interaction detection task as a two class classification –*interaction, no interaction*. We aim to identify whether the hand touches the face using 2D images without using predefined or estimated depth information. Fig. 5 represents the overall pipeline of our proposed system. The upper part of the figure shows the face mask and face-hand interaction tasks. At first, the input data passes through the face detector, which is the RetinaFace model [36], to detect bounding box coordinates of the faces. Then, these predicted bounding box coordinates are used to obtain face crops with suitable margins. Afterward, the face mask detection and face-hand interaction detection models are used to predict on acquired face crops. While the face mask model classifies the face images based on whether they contain masks (proper or improper) or not, the face-hand model performs similar action based on whether there is an interaction between face and hand.

4.1.2 Social distance controlling

As the application of face mask detection and face-hand interaction detection preventions, keeping the social distance from others is also a crucial measurement to avoid spreading of COVID-19 disease. To address this issue, we try to measure the distance between people based on the Euclidean distance metric instead of utilizing a learning-based system. For this, we benefited from 2D images, $I \in \mathbb{R}^{W \times H \times 3}$. Firstly, we detect each person on the image using a pretrained person detection model, DeepHRNet [48]. Thus, we obtain bounding boxes around the people and estimated pose information of each person, $p = (b_1, b_2, b_3, b_4, s_1, s_2)$, where (b_1, b_2, b_3, b_4) represent bounding box coordinates of a detected person and (s_1, s_2) represent estimated shoulder points of the same person. Principally, we focus on the shoulders’ coordinates to measure the approximate body width of a person on the projected image. In many studies, measurements are calculated based on the bounding box around the person. However, when the angle of the body joints and pose of the person are considered, changes on the bounding boxes may reduce the precision of the measurements. To prevent this, we propose to use shoulders’ coordinates to measure the width of the body and identify the middle point of shoulders line as center of the body. This makes the representation of the person’s width more robust according to our empirical evaluation as well. After performing detection and pose estimation, we generated pairs based on the

combination of the detected persons, e.g., $P(p_i, p_j)$. Then, we calculated the Euclidean distance between the centers of shoulder points of each pair of persons. The overall formula is shown in Equation 1,

$$D_{dist(p_i, p_j)} = \left\| \frac{(p_{i_{s_1}} + p_{i_{s_2}})}{2} - \frac{(p_{j_{s_1}} + p_{j_{s_2}})}{2} \right\|_2 \quad (1)$$

where p_i represents each person in a pair and s_i represents shoulder points. In order to decide whether these corresponding persons keep social distance between each other, we adaptively calculate a threshold for each pair individually based on the average of their body width. Since the represented measurement of the real world that is expressed by pixels in the image domain constantly changes as depth increases, we represent the mapping between real-world measurement and pixel domain measurement by calculating the average of the body widths of two people in order to express this effect. Since the average distance between shoulder points of an adult is around 40-50 cm in the real-world and the average social distance between two persons is 1.5-2.0 meters, we empirically decide to select λ coefficient as 3 when calculating threshold for social distance in the pixel domain as in Equation 2.

$$T_{p_i, p_j} = \frac{\lambda \times (\|p_{i_{s_1}} - p_{i_{s_2}}\|_2 + \|p_{j_{s_1}} - p_{j_{s_2}}\|_2)}{2} \quad (2)$$

Finally, if the Euclidean distance between two persons is lower than the calculated threshold for this pair, we decide that these people do not keep sufficient social distance as in Equation 3.

$$M_{p_i, p_j} = \begin{cases} 1, & D_{dist_{p_i, p_j}} < T_{p_i, p_j} \\ 0, & D_{dist_{p_i, p_j}} \geq T_{p_i, p_j} \end{cases} \quad (3)$$

where M_{p_i, p_j} represents the decision for person p_i and person p_j , $D_{dist_{p_i, p_j}}$ states the calculated distance between corresponding two persons, and T_{p_i, p_j} expresses the threshold between them. Fig. 5 is the visualization of the overall pipeline of the proposed model. The below part of the figure shows the social distance module. First of all, the input data is given to the person detection and pose estimation model. After that, pairs are created using a combination of the detected people bounding boxes. The distance between people is calculated using Euclidean distance and the adaptive threshold is calculated for each pair individually. In the end, the system decides whether the corresponding two persons keep social distance between them according to the threshold value and calculated distance.

5 Experimental Results

In this section, we briefly explained the experimental setups and test datasets for evaluating our model. Besides we explored the performance of each model and overall system separately.

5.1 Datasets

We used publicly available datasets to evaluate the generalization capacity of our system and also compared our mask detection models with the previous works.

RMFD [11] This dataset is presented to investigate face recognition performance when subjects wear face masks. Although the paper indicates that there are 5000 face mask images belonging to 525 subjects, the publicly available version includes around 2203 masked face images and 90000 face images without a mask.

RWMFD [11] Another dataset is presented ¹ by Wang et al. We executed RetineFace detector to detect faces from presented 4343 images and we obtained 5171 face images in the end. Then, we annotated face crops to use in test phase.

Face-mask dataset (Kaggle) [39] This dataset contains 853 images and we used provided annotations to crop face images and obtain labels. In the end, we acquired 4080 face images. We included margin around face bounding boxes when we cropped annotated faces as we added margin to the face crops on our proposed datasets.

MaskedFace-Net dataset [16] This dataset is created by using face images from FFHQ dataset [33]. It contains 130000 images and is generated by including artificial face masks on the FFHQ face images. While the half of the dataset (CMFD) has correctly worn face masks, the remaining half (IMFD) has incorrectly worn face masks.

¹<https://github.com/X-zhangyang/Real-World-Masked-Face-Dataset>

Table 2: Face mask detection results on proposed ISL-UFMD dataset for three classes case.

Model	Accuracy	Precision			Recall		
		No Mask	Mask	Improper Mask	No Mask	Mask	Improper Mask
Inception-v3	98.20%	0.985	0.986	0.833	0.988	0.984	0.800
ResNet50	95.63%	0.965	0.954	0.636	0.973	0.973	0.389
MobileNetV2	97.91%	0.988	0.975	0.842	0.983	0.992	0.640
EfficientNet-b0	97.82%	0.973	0.984	0.929	0.992	0.986	0.520
EfficientNet-b1	97.91%	0.979	0.986	0.800	0.990	0.984	0.711
EfficientNet-b2	97.91%	0.990	0.977	0.792	0.977	0.992	0.760
EfficientNet-b3	98.19%	0.988	0.990	0.733	0.986	0.982	0.880

Table 3: Face mask detection results on proposed UFMD dataset for two classes scenario.

Model	Accuracy	Precision	Recall	Improper mask images
Inception-v3	99.22%	0.9942	0.9903	×
ResNet-50	99.12%	0.9883	0.9941	×
MobileNetV2	99.41%	0.9981	0.9903	×
EfficientNet-b0	98.53%	0.9786	0.9922	×
EfficientNet-b1	99.22%	0.9942	0.9903	×
EfficientNet-b2	99.51%	0.9942	0.9961	×
EfficientNet-b3	99.31%	0.9981	0.9884	×
Inception-v3	99.31%	0.9903	0.9961	✓
ResNet-50	99.02%	0.9961	0.9846	✓
MobileNetV2	99.12%	0.9903	0.9922	✓
EfficientNet-b0	98.53%	0.9786	0.9921	✓
EfficientNet-b1	99.12%	0.9903	0.9922	✓
EfficientNet-b2	99.12%	0.9942	0.9884	✓
EfficientNet-b3	97.75%	0.9823	0.9728	✓

5.2 Experimental setups

We split our proposed face mask dataset into training, validation, and test sets. While 90% of the data is reserved for training, the remaining data is split equally for validation and testing. We followed the same strategy for face-hand interaction dataset. Additionally, before creating train-val-test splits, we put aside around 5000 images from no face-hand interaction class to obtain balanced dataset to execute face-hand interaction detection. On the contrary, we used all images from the existing face mask datasets which are mentioned in Section 5.1 while evaluating our face mask detection model.

In the face mask detection task, we performed experiments based on two different setups. While the first setup contains improper mask images as an additional class, we handled face mask detection task as a binary classification by eliminating improper mask label in the second setup. One of the main reasons for this approach is the lack of images for improper mask class. To discard improper mask label, we followed two different approaches. In the first one, we did not include the images of this class in training data and test data. In the second approach, we included these images in both training data and test data as no mask class. With the help of these scenarios, we tried to investigate how the model behaves for the improper mask class without seeing these images or with learning improper usage as a no mask label.

5.3 Results

In this section, we conveyed the evaluation results of our models for each three task. We discussed the model’s predictions based on class activation maps of several positive and negative examples. Further, we presented evaluation results on collected videos and discussed the overall system performance.

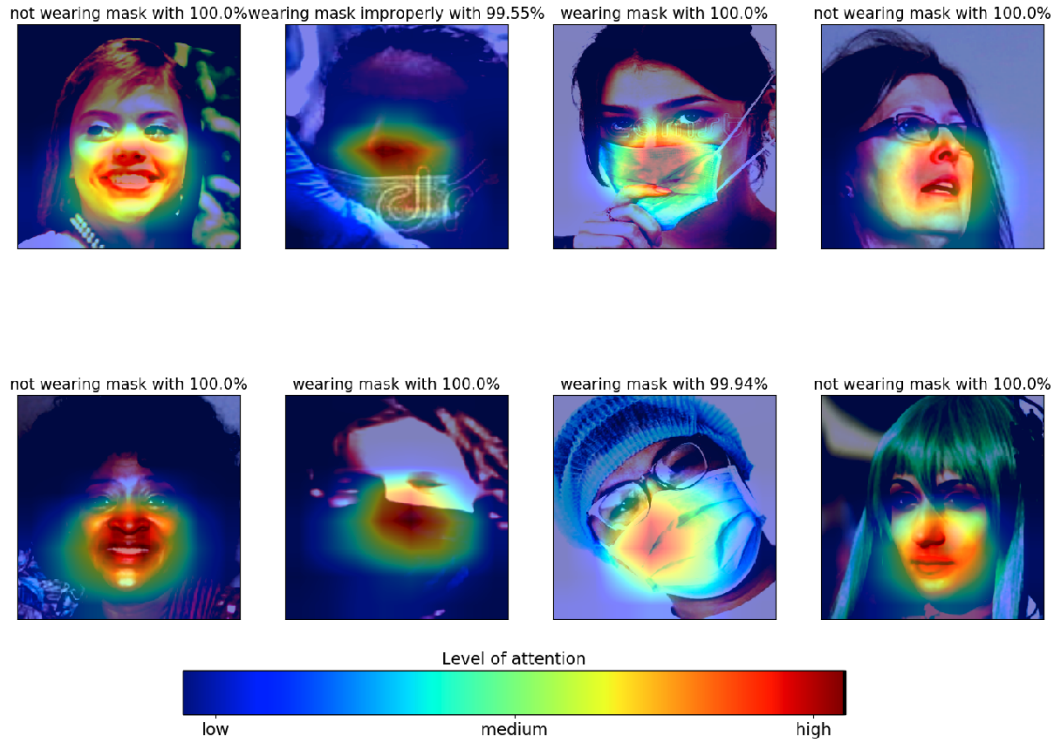


Figure 6: Class activation map (CAM) for the face mask detection task (3 classes case). CAM samples indicate that our model can focus on the mouth and nose of the subject and effectively decide whether corresponding subject wears a mask or not.

5.3.1 Face mask detection

In Table 2 and 3, we presented various evaluation results using different metrics, namely classification accuracy, precision, and recall. We showed face mask detection experiments for three classes in Table 2, while we presented experiments for two classes in Table 3. In the tables, while the first column indicates the employed deep CNN models, the following columns represent evaluation results for face mask detection with these models. According to the experimental results in Table 2, although all employed models achieved significantly high performance, the best one is Inception-v3 model with 98.20% classification accuracy. On the other hand, we achieved very similar classification performance for binary classification task and we obtained the best performance as 99.51% with EfficientNet-b2 model for without improper face mask images and 99.31% classification accuracy with Inception-v3 for with improper face mask images. In Table 3, the performance of all models except Inception-v3 is decreased when improper face mask images are employed in the no mask class. Intuitively, the possible reason of this outcome is that when improper face mask images are used in the no mask class, the model is deceived due to the similarity between these images and face mask images. Besides, the small amount of improper face mask images in the no mask class does not provide well enough feature representation to the model in order to distinguish these images from the mask class. In the end, we decided to employ the model which is trained for three class classification scenario in our system.

In addition to the classification accuracy, we also present precision and recall measurements for each class separately to demonstrate the performance of the models individually. In Table 2, although the precision and recall values are significantly accurate for no mask and mask classes, these results for improper mask class are slightly lower than these two classes. Even though improper face mask can be confusing in terms of discrimination from mask class (proper), the more probable reason behind this outcome is the lack of images for improper mask usage.

Discussion of results In Fig. 6, we present class activation maps for the face mask detection task to investigate the model's activation on the test image. It is clearly seen that the model focuses on the bottom part of the faces, particularly on the nose and mouth. In the second image of the first row, the model identified improper mask usage since the nose of the subject is not covered by the face mask even though the mouth is covered. In the first row of Fig. 8, we present some misclassified images for the face mask detection task. Although the model classified the images incorrectly, the

Table 4: Results for cross-dataset experiments. All models are trained on corresponding training sets that are shown in the second column and tested on evaluation sets that are presented in the third column. Please note that all experiments are conducted on the 3-class classification setup to perform fair comparison.

Architecture	Training Set	Test Set	# Images		Accuracy
			Train	Test	
MobileNetV2	ISL-UFMD	RMFD [11]	20764	92671	91.4%
MobileNetV2	ISL-UFMD	RWMFD [11]	20764	5171	94.7%
MobileNetV2	ISL-UFMD	MaskedFace-Net [16]	20764	133782	88.11%
MobileNetV2	ISL-UFMD	Face-mask [39]	20764	4080	95.71%
Inception-v3	ISL-UFMD	RMFD [11]	20764	92671	95.91%
Inception-v3	ISL-UFMD	RWMFD [11]	20764	5171	95.9%
Inception-v3	ISL-UFMD	MaskedFace-Net [16]	20764	133782	91.42%
Inception-v3	ISL-UFMD	Face-mask [39]	20764	4080	94.7%
MobileNetV2	RMFD + RWMFD	ISL-UFMD	97842	21816	86.59%
MobileNetV2	RMFD + RWMFD	Face-mask[39]	97842	4080	91.07%
MobileNetV2	MaskedFace-Net + FFHQ	ISL-UFMD	211936	21816	51.49%
MobileNetV2	MaskedFace-Net + FFHQ	Face-mask[39]	211936	4080	20.4%
Inception-v3	RMFD + RWMFD	ISL-UFMD	97842	21816	88.92%
Inception-v3	RMFD + RWMFD	Face-mask[39]	97842	4080	88.4%
Inception-v3	MaskedFace-Net + FFHQ	ISL-UFMD	211936	21816	51.39%
Inception-v3	MaskedFace-Net + FFHQ	Face-mask[39]	211936	4080	19.2%

prediction probabilities of model are not as high as in correct predictions. This outcome indicates that the model did not confidently misclassify images. In the second and third images of the first row, the model classified the images incorrectly due to the difficulty in the head pose of the subject, while illumination is other challenging factor for second image. Although the correct label is improper mask in the third image since the nose is not covered by the mask, the model could not identify this since it focused on the mouth area and decided as proper mask. In the last image of the first row, the model interpreted the hair of the front subject as a mask since it covered the mouth and the nose of the subject.

Cross-dataset experiments In Table 4 we presented cross-dataset experiments on used face mask datasets to investigate the effect of the datasets on the generalization capacity of the proposed model. First, we evaluated our MobileNetV2 and Inception-v3 models on four different public face mask datasets. Additionally, we finetuned the MobileNetV2 and Inception-v3 models with two different training setups to compare our approach. The first setup contains 97842 images from the combination of RMFD and RWMFD datasets that are presented by the same authors [11]. We used them together since RMFD dataset has no improper mask class. The second setup includes 211936 images from the MaskedFace-Net dataset [16] with FFQH dataset [33]. We used FFHQ dataset as a no mask data due to the absence of no mask class on MaskedFace-Net dataset. We conducted all experiments as three class classification task. While we selected RMFD, RWMFD, MaskedFace-Net, and Face-mask (Kaggle) [39] datasets as target for our model, we used the proposed ISL-UFMD dataset and Face-mask (Kaggle) dataset as target datasets for other models. The models that are trained on ISL-UFMD achieved more than 90% with all models except MobileNetV2 model which are evaluated on MaskedFace-Net dataset. These results indicate that our ISL-UFMD dataset is significantly representative to provide a well generalized models for face mask detection task. We employed two different architectures to endorse this outcome. Otherwise, the combination of RMFD and RWMFD provide accurate results although they are not as precise as our results. On the contrary, the models that are trained on MaskedFace-Net dataset show the worst performance. The possible reason of this outcome is that the artificial dataset is not as useful as the real data in terms of providing robust models.

5.3.2 Face-hand interaction detection

In Table 5, we show similar evaluation results that we did for the face mask detection task. While the first column of table represents the utilized deep CNN model, the further columns indicate the evaluation metrics. As we mentioned in the Section 4, we handled this task as a binary classification task *–touch, not touch*. As in the face mask detection, all of the employed models have considerably high performance to discriminate whether there is an interaction between face and hand. The best classification accuracy is obtained as 93.35% using EfficientNet-b2. The best recall and precision

Table 5: Face-hand interaction detection results on proposed ISL-UFHD dataset.

Model	Accuracy	Precision	Recall
Inception-v3	93.20%	0.932	0.932
ResNet50	91.76%	0.918	0.918
MobileNetV2	92.37%	0.924	0.924
EfficientNet-b0	92.37%	0.926	0.924
EfficientNet-b1	92.90%	0.929	0.929
EfficientNet-b2	93.35%	0.933	0.934
EfficientNet-b3	92.44%	0.925	0.924

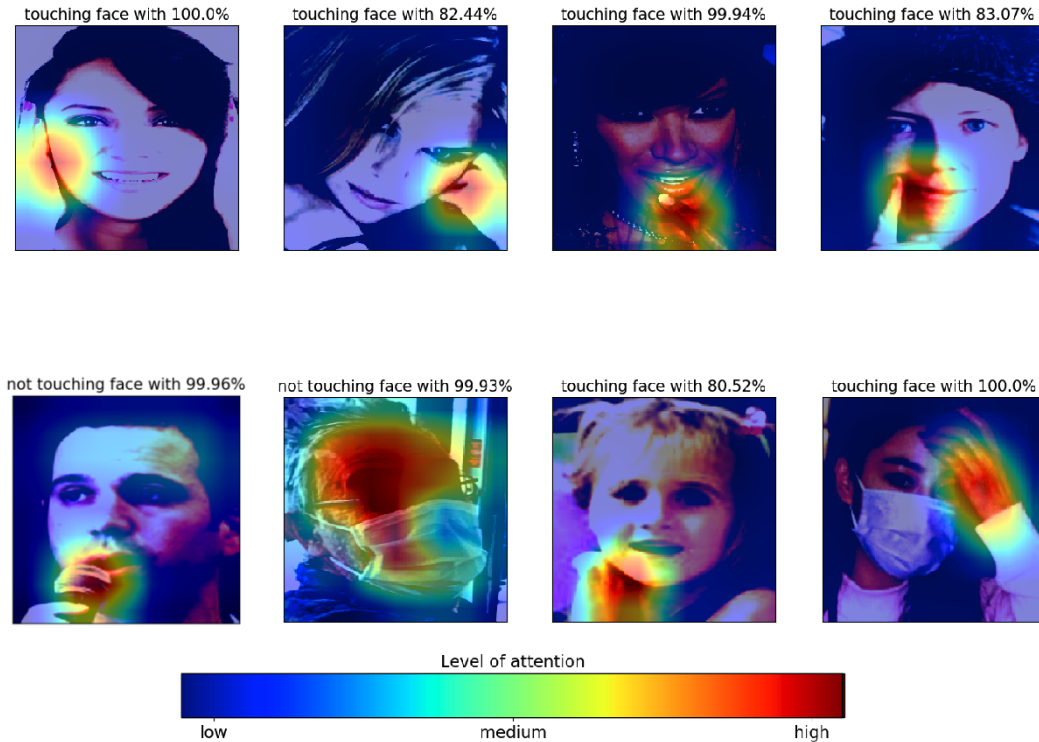


Figure 7: Class activation map (CAM) for the face-hand interaction detection task. CAM samples indicate that our model can focus on the region that face and hand are in an interaction and effectively identify this interaction. If we compare the third image of the first row and the first image of the second row, we can obviously see that our face-hand model can distinguish the difference between the hand and another objects, e.g., microphone for this sample.

results are achieved by EfficientNet-b2. However, almost all results in the table is considerably similar to each other. Precision and recall metrics are balanced and compatible with the accuracies.

Discussion of results In Fig. 7, we provide class activation maps for the face-hand interaction detection task to further analyze the model’s activation on the evaluation data. It is explicitly seen that the model focuses on the hand and around the hand to decide whether there is an interaction between the hand and the face of the subject. If the image does not contain any hand, then the model focuses all of the images uniformly as in the second image of the second row of Fig. 7. In the second row of Fig. 8, we present some misclassified images for the face-hand interaction detection task. In the first image, the model could not identify the interaction since the face is not completely visible. In the second and fourth images, there are interactions between faces and other subjects’ hands and the angles of the heads and hands are challenging. In the third image, although the model can detect hand and face, it cannot identify the depth between the face and the hand due to the position of the hand.

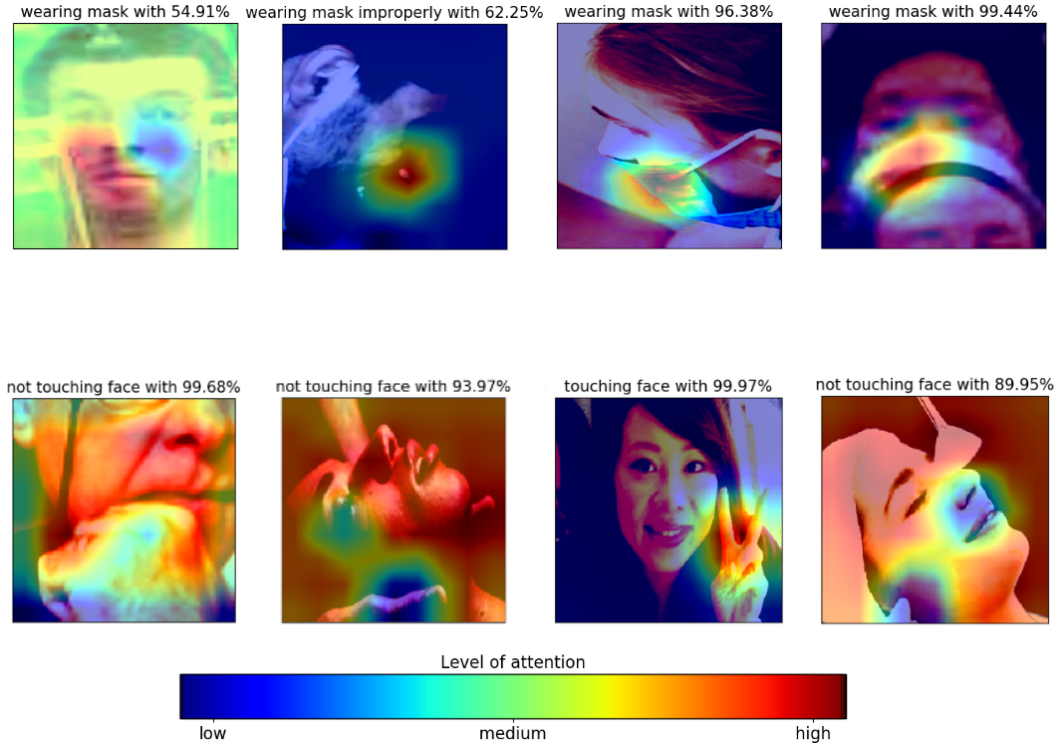


Figure 8: Class activation map (CAM) for the face mask detection and face-hand interaction detection tasks. All these samples represent misclassified images. While first row shows misclassified images for face mask task, the second row contains misclassified images for face-hand task.

5.4 Social distance controlling

We utilized six different videos that we collected from the web in order to evaluate proposed social distancing control module. These videos have different number of frames and they were recorded in various environments with different camera angles. The test results on the video set are presented in Table 6. The last column in Table 6 represents the performance of the social distance controlling algorithm. During the calculation of the accuracy of the social distance algorithm, we utilized the annotations that we decided based on subject pairs and existing distance between each other. Person detector could not detect some of the subjects in the scene if they are not visible in the camera due to the occlusion by other people or objects. For that reason, we ignored the missing detections when we annotated the videos' frames and calculated the accuracies. According to the results in the Table 6, we achieved very high accuracies to detect whether people keep a social distance. However, the fundamental problem, especially occurred in the last video, is caused by a lack of depth information. In order to adapt our algorithm to different camera views and scenes while calculating the distance between people, we project real-world lengths to the image pixel with a rule-based approach without using camera angle, camera intrinsic parameters, and any reference points. Because of this lack of information, depth perception can be problematic for a specific type of angle that is shown in Fig. 9.

5.5 Overall system performance

In order to evaluate the overall system, we utilized six different videos to test all three subtasks. The corresponding experiments are presented in Table 6. While the second column shows the number of frames in each video, the fourth column represents the number of subject in each frame. Last three columns list the classification accuracies for three classes face mask detection, face-hand interaction detection, and social distance measurement. We evaluated the each social distance prediction as follows: If a subject keeps the distance from others and the algorithm decides as the corresponding person keeps the distance based on the calculated distances, we assume that the prediction is correct, and vice versa. When we examined the face-hand interaction and face mask detection performance of our system, the results on videos that contains various people and cases indicate that system can reach the very high performance similar to the

Table 6: Evaluation of the overall system on the test videos.

Video	# frames	FPS	# subject	Mask acc.	Face-hand acc.	Distance acc
Video 1	179	0.75	2	100%	99.16%	98.32%
Video 2	307	0.71	2	99.51%	96.25%	100%
Video 3	303	0.78	3	96.91%	89.43%	96.69%
Video 4	192	0.69	3	100%	86.97%	97.22%
Video 5	207	0.85	5	99.03%	95.45%	100%
Video 6	105	0.67	7	87.07%	99.86%	74.55%
Total	1293	0.74	22	97.95%	93.84%	96.51%




Figure 9: Sample frame from test videos to visualize social distancing between people. While red boxes represent the subjects that violate social distance with others, green boxes represent the subjects that keep social distance.

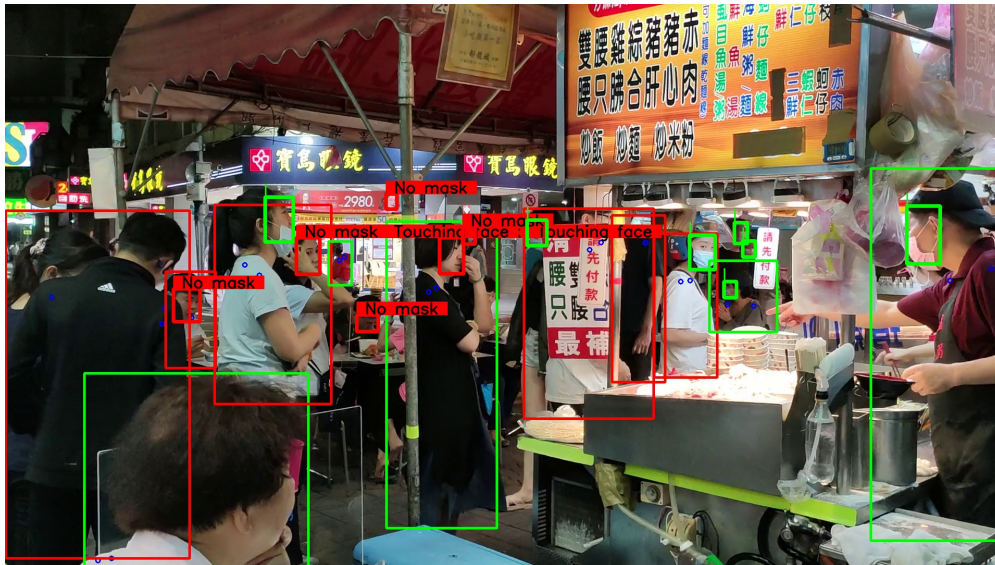


Figure 10: Sample frame from one of the test video to visualize the detection and prediction results for all three tasks of proposed system.

ones that are obtained by the models on individual test sets. Fig. 10 visualizes all tasks on a sample video frame that contains people in a crowded street food area.

6 Conclusion

In this paper, we presented two datasets, ISL-UFMD and ISL-UFHD, with high diversity to examine essential COVID-19 preventions and we proposed a system to track these preventions: proper face mask usage, avoiding face-hand interaction, and keeping social distance in a public area. While we employed several different deep CNN-based models to perform face mask detection and face-hand interaction detection tasks, we benefited from a rule-based method to track the social distance between people. Besides, we presented an end-to-end prevention control system to perform all these three tasks. To evaluate our trained models, we performed several tests on both our proposed datasets and on existing publicly available datasets in the literature. Experimental results showed our proposed models' significantly high performance on all datasets with the help of our proposed datasets, since they contain a large amount of variation and they represent various cases in a real-world scenario. The cross-dataset experiments indicate the generalization capacity of our proposed models on unseen data. The proposed system can be effectively utilized to track all preventions against the transmission of COVID-19.

7 Acknowledgement

The project on which this report is based was funded by the Federal Ministry of Education and Research (BMBF) of Germany under the number 01IS18040A. The authors are responsible for the content of this publication.

References

- [1] Catharine I Paules, Hilary D Marston, and Anthony S Fauci. Coronavirus infections—more than just the common cold. *Jama*, 323(8):707–708, 2020.
- [2] Coronavirus disease (covid-19) advice for the public. <https://www.who.int/emergencies/diseases/novel-coronavirus-2019/advice-for-public>. Accessed: 2020-10-25.
- [3] Covid-19: physical distancing. <https://www.who.int/westernpacific/emergencies/covid-19/information/physical-distancing>. Accessed: 2020-10-25.
- [4] Alex Waibel, Toshiyuki Hanazawa, Geoffrey Hinton, Kiyohiro Shikano, and Kevin J Lang. Phoneme recognition using time-delay neural networks. *IEEE transactions on acoustics, speech, and signal processing*, 37(3):328–339, 1989.
- [5] Yann Le Cun, Bernhard Boser, John S Denker, Donnie Henderson, Richard E Howard, Wayne Hubbard, and Lawrence D Jackel. Handwritten digit recognition with a back-propagation network. In *Proceedings of the International Conference on Neural Information Processing Systems (NIPS)*, pages 396–404, 1989.
- [6] Jun Chen, Lianlian Wu, Jun Zhang, Liang Zhang, Dexin Gong, Yilin Zhao, Qiuxiang Chen, Shulan Huang, Ming Yang, Xiao Yang, et al. Deep learning-based model for detecting 2019 novel coronavirus pneumonia on high-resolution computed tomography. *Scientific reports*, 10(1):1–11, 2020.
- [7] Lin Li, Lixin Qin, Zeguo Xu, Youbing Yin, Xin Wang, Bin Kong, Junjie Bai, Yi Lu, Zhenghan Fang, Qi Song, et al. Using artificial intelligence to detect covid-19 and community-acquired pneumonia based on pulmonary ct: evaluation of the diagnostic accuracy. *Radiology*, 296(2):E65–E71, 2020.
- [8] Muhammad Farooq and Abdul Hafeez. Covid-resnet: A deep learning framework for screening of covid19 from radiographs. *arXiv preprint arXiv:2003.14395*, 2020.
- [9] Ali Narin, Ceren Kaya, and Ziyet Pamuk. Automatic detection of coronavirus disease (covid-19) using x-ray images and deep convolutional neural networks. *arXiv preprint arXiv:2003.10849*, 2020.
- [10] Mingjie Jiang and Xinqi Fan. Retinamask: a face mask detector. *arXiv preprint arXiv:2005.03950*, 2020.
- [11] Zhongyuan Wang, Guangcheng Wang, Baojin Huang, Zhangyang Xiong, Qi Hong, Hao Wu, Peng Yi, Kui Jiang, Nanxi Wang, Yingjiao Pei, et al. Masked face recognition dataset and application. *arXiv preprint arXiv:2003.09093*, 2020.
- [12] Aqeel Anwar and Arijit Raychowdhury. Masked face recognition for secure authentication. *arXiv preprint arXiv:2008.11104*, 2020.

- [13] Naser Damer, Jonas Henry Grebe, Cong Chen, Fadi Boutros, Florian Kirchbuchner, and Arjan Kuijper. The effect of wearing a mask on face recognition performance: an exploratory study. In *2020 International Conference of the Biometrics Special Interest Group (BIOSIG)*, pages 1–6. IEEE, 2020.
- [14] Rucha Golwalkar and Ninad Mehendale. Age detection with face mask using deep learning and facemasknet-9. Available at SSRN 3733784, 2020.
- [15] Senqiu Chen, Wenbo Liu, and Gong Zhang. Efficient transfer learning combined skip-connected structure for masked face poses classification. *IEEE Access*, 8:209688–209698, 2020.
- [16] Adnane Cabani, Karim Hammoudi, Halim Benhabiles, and Mahmoud Melkemi. Maskedface-net—a dataset of correctly/incorrectly masked face images in the context of covid-19. *Smart Health*, 19:100144, 2021.
- [17] Aniruddha Srinivas Joshi, Shreyas Srinivas Joshi, Goutham Kanahasabai, Rudraksh Kapil, and Savyasachi Gupta. Deep learning framework to detect face masks from video footage. In *2020 12th International Conference on Computational Intelligence and Communication Networks (CICN)*, pages 435–440. IEEE, 2020.
- [18] Preeti Nagrath, Rachna Jain, Agam Madan, Rohan Arora, Piyush Kataria, and Jude Hemanth. Ssdmnv2: A real time dnn-based face mask detection system using single shot multibox detector and mobilenetv2. *Sustainable cities and society*, 66:102692, 2021.
- [19] G Jignesh Chowdary, Narinder Singh Punn, Sanjay Kumar Sonbhadra, and Sonali Agarwal. Face mask detection using transfer learning of inceptionv3. In *International Conference on Big Data Analytics*, pages 81–90. Springer, 2020.
- [20] Moein Razavi, Hamed Alikhani, Vahid Janfaza, Benyamin Sadeghi, and Ehsan Alikhani. An automatic system to monitor the physical distance and face mask wearing of construction workers in covid-19 pandemic. *arXiv preprint arXiv:2101.01373*, 2021.
- [21] Zekun Wang, Pengwei Wang, Peter C Louis, Lee E Wheless, and Yuankai Huo. Wearmask: Fast in-browser face mask detection with serverless edge computing for covid-19. *arXiv preprint arXiv:2101.00784*, 2021.
- [22] Nenad Petrović and Đorđe Kocić. Iot-based system for covid-19 indoor safety monitoring. *preprint*, *IcETAN*, 2020:1–6, 2020.
- [23] Mohamed Loey, Gunasekaran Manogaran, Mohamed Hamed N Taha, and Nour Eldeen M Khalifa. Fighting against covid-19: A novel deep learning model based on yolo-v2 with resnet-50 for medical face mask detection. *Sustainable Cities and Society*, 65:102600, 2021.
- [24] Mohamed Loey, Gunasekaran Manogaran, Mohamed Hamed N Taha, and Nour Eldeen M Khalifa. A hybrid deep transfer learning model with machine learning methods for face mask detection in the era of the covid-19 pandemic. *Measurement*, 167:108288, 2021.
- [25] Adarsh Jagan Sathyamoorthy, Utsav Patel, Yash Ajay Savle, Moumita Paul, and Dinesh Manocha. Covid-robot: Monitoring social distancing constraints in crowded scenarios. *arXiv preprint arXiv:2008.06585*, 2020.
- [26] Dongfang Yang, Ekim Yurtsever, Vishnu Renganathan, Keith A Redmill, and Ümit Özgüner. A vision-based social distancing and critical density detection system for covid-19. *arXiv preprint arXiv:2007.03578*, pages 24–25, 2020.
- [27] Mahdi Rezaei and Mohsen Azarmi. Deepsocial: Social distancing monitoring and infection risk assessment in covid-19 pandemic. *Applied Sciences*, 10(21):7514, 2020.
- [28] Imran Ahmed, Misbah Ahmad, Joel JPC Rodrigues, Gwanggil Jeon, and Sadia Din. A deep learning-based social distance monitoring framework for covid-19. *Sustainable Cities and Society*, 65:102571, 2021.
- [29] Cigdem Beyan, Matteo Bustreo, Muhammad Shahid, Gian Luca Bailo, Nicolo Carissimi, and Alessio Del Bue. Analysis of face-touching behavior in large scale social interaction dataset. In *Proceedings of the 2020 International Conference on Multimodal Interaction*, pages 24–32, 2020.
- [30] Ramprasaath R Selvaraju, Michael Cogswell, Abhishek Das, Ramakrishna Vedantam, Devi Parikh, and Dhruv Batra. Grad-cam: Visual explanations from deep networks via gradient-based localization. In *Proceedings of the IEEE International Conference on Computer Vision (ICCV)*, pages 618–626, 2017.
- [31] Anwaar Ulhaq, Jannis Born, Asim Khan, Douglas Pinto Sampaio Gomes, Subrata Chakraborty, and Manoranjan Paul. Covid-19 control by computer vision approaches: A survey. *IEEE Access*, 8:179437–179456, 2020.
- [32] Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. Deep residual learning for image recognition. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 770–778, 2016.
- [33] Tero Karras, Samuli Laine, and Timo Aila. A style-based generator architecture for generative adversarial networks. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 4401–4410, 2019.

- [34] Kaipeng Zhang, Zhanpeng Zhang, Zhifeng Li, and Yu Qiao. Joint face detection and alignment using multitask cascaded convolutional networks. *IEEE Signal Processing Letters*, 23(10):1499–1503, 2016.
- [35] Mark Sandler, Andrew Howard, Menglong Zhu, Andrey Zhmoginov, and Liang-Chieh Chen. Mobilenetv2: Inverted residuals and linear bottlenecks. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 4510–4520, 2018.
- [36] Jiankang Deng, Jia Guo, Evangelos Ververas, Irene Kotsia, and Stefanos Zafeiriou. Retinaface: Single-shot multi-level face localisation in the wild. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, June 2020.
- [37] Andrew G Howard, Menglong Zhu, Bo Chen, Dmitry Kalenichenko, Weijun Wang, Tobias Weyand, Marco Andreetto, and Hartwig Adam. Mobilenets: Efficient convolutional neural networks for mobile vision applications. *arXiv preprint arXiv:1704.04861*, 2017.
- [38] Wei Liu, Dragomir Anguelov, Dumitru Erhan, Christian Szegedy, Scott Reed, Cheng-Yang Fu, and Alexander C Berg. Ssd: Single shot multibox detector. In *European conference on computer vision*, pages 21–37. Springer, 2016.
- [39] Face mask detection. <https://www.kaggle.com/andrewmvd/face-mask-detection>. Accessed: 2020-09-12.
- [40] Ziwei Liu, Ping Luo, Xiaogang Wang, and Xiaoou Tang. Deep learning face attributes in the wild. In *Proceedings of International Conference on Computer Vision (ICCV)*, December 2015.
- [41] Gary B Huang and Erik Learned-Miller. Labeled faces in the wild: Updates and new reporting procedures. *Dept. Comput. Sci., Univ. Massachusetts Amherst, Amherst, MA, USA, Tech. Rep*, 14(003), 2014.
- [42] Shuo Yang, Ping Luo, Chen-Change Loy, and Xiaoou Tang. Wider face: A face detection benchmark. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 5525–5533, 2016.
- [43] Davis E King. Dlib-ml: A machine learning toolkit. *The Journal of Machine Learning Research*, 10:1755–1758, 2009.
- [44] Christian Szegedy, Vincent Vanhoucke, Sergey Ioffe, Jon Shlens, and Zbigniew Wojna. Rethinking the inception architecture for computer vision. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 2818–2826, 2016.
- [45] Mingxing Tan and Quoc Le. Efficientnet: Rethinking model scaling for convolutional neural networks. In *International Conference on Machine Learning*, pages 6105–6114. PMLR, 2019.
- [46] Christian Szegedy, Wei Liu, Yangqing Jia, Pierre Sermanet, Scott Reed, Dragomir Anguelov, Dumitru Erhan, Vincent Vanhoucke, and Andrew Rabinovich. Going deeper with convolutions. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 1–9, 2015.
- [47] Jia Deng, Wei Dong, Richard Socher, Li-Jia Li, Kai Li, and Li Fei-Fei. Imagenet: A large-scale hierarchical image database. In *2009 IEEE conference on computer vision and pattern recognition*, pages 248–255. Ieee, 2009.
- [48] Jingdong Wang, Ke Sun, Tianheng Cheng, Borui Jiang, Chaorui Deng, Yang Zhao, Dong Liu, Yadong Mu, Mingkui Tan, Xinggang Wang, et al. Deep high-resolution representation learning for visual recognition. *IEEE transactions on pattern analysis and machine intelligence*, 2020.