

This WACV 2021 paper is the Open Access version, provided by the Computer Vision Foundation. Except for this watermark, it is identical to the accepted version; the final published version of the proceedings is available on IEEE Xplore.

SuPEr-SAM: Using the Supervision Signal from a Pose Estimator to Train a Spatial Attention Module for Personal Protective Equipment Recognition

Adrian Şandru^{1,2}, Georgian-Emilian Duță^{1,2}, Mariana-Iuliana Georgescu^{1,2}, Radu Tudor Ionescu^{1,2}, ¹SecurifAI, ²University of Bucharest

Abstract

We propose a deep learning method to automatically detect personal protective equipment (PPE), such as helmets, surgical masks, reflective vests, boots and so on, in images of people. Typical approaches for PPE detection based on deep learning are (i) to train an object detector for items such as those listed above or (ii) to train a person detector and a classifier that takes the bounding boxes predicted by the detector and discriminates between people wearing and people not wearing the corresponding PPE items. We propose a novel and accurate approach that uses three components: a person detector, a body pose estimator and a classifier. Our novelty consists in using the pose estimator only at training time, to improve the prediction performance of the classifier. We modify the neural architecture of the classifier by adding a spatial attention mechanism, which is trained using supervision signal from the pose estimator. In this way, the classifier learns to focus on PPE items, using knowledge from the pose estimator with almost no computational overhead during inference.

1. Introduction

In many workplace environments, such as construction sites or oil and gas factories, there are potentially hazardous materials, dangerous chemicals or dangerous equipment (large moving parts, operating cranes), that present health or safety risks for the workers. In order to protect workers against accidents caused in hazardous or dangerous areas, the workers may be required to wear personal protective equipment (PPE). The personal protective equipment may include, but is not limited to, helmets, masks, glasses, uniforms, safety vests, boots, gloves, etc. In many countries, there are governmental regulations that require certain employees to be trained in the proper and effective use of PPE in hazardous or dangerous working environments. Nevertheless, in many cases, the workers choose not to wear PPE items because these items are uncomfortable, potentially causing sweat or skin irritations, or because these items prevent the workers from executing their tasks with precision, for example screwing or unscrewing a small screw with gloves on. Managing such risks is of utter importance to employers that aim at preventing unnecessary health or safety accidents. In this context, a system, such as the one proposed in this work, to automatically monitor the workers and verify if they are permanently wearing the mandatory personal protective equipment is extremely useful. We also recognize a broader application of our system in the context of the COVID-19 pandemic, which forced governments to impose new regulations requiring people to wear surgical masks in public places in order to reduce the spread of the SARS-CoV-2 virus. Our system could also be used to check for compliance with such regulations, as shown in our experiments.

Typical deep learning approaches for PPE detection are (*i*) to train an object detector [9, 12, 28, 30, 31, 32], resulting in a single-stage pipeline, or (ii) to train a person detector and a classifier that takes the bounding boxes predicted by the detector and discriminates between people wearing and people not wearing the corresponding PPE items [3, 6, 21], resulting in a two-stage pipeline. A recent work [21] demonstrated that the latter methodology produces more accurate results, mainly due to the poor ability of object detectors of detecting small objects. Another advantage of the two-stage approach is that it eliminates the necessity to annotate PPE items with bounding boxes. Given an accurate pre-trained person detector, we only need binary labels for each detected person with respect to each PPE item. This significantly reduces the annotation labor, allowing us to easily construct large data sets typically required for training deep learning models.

We propose a novel and accurate approach that uses three components: a person detector, an articulated body pose estimator and a classifier. Our novelty consists in employing the pose estimator only at training time, providing pseudoground-truth attention maps to improve the performance of the classifier. We modify the neural architecture of the classifier by adding a spatial attention mechanism (SAM), which is trained using **su**pervision signal from the **p**ose **e**stimator (SuPEr), as shown in Figure 1. In this way, the classifier learns to focus on PPE items, using knowledge



Figure 1. A neural network classifier based on SuPEr-SAM, trained with classification loss and attention loss to discriminate people wearing helmets from people not wearing helmets. The supervision signal for the attention loss is provided by an articulated body pose estimator.

from the pose estimator. At inference time, the pose estimator is removed from our processing pipeline, since it is only required during training to provide supervision to the attention module. Hence, the only computational overhead during inference is due to the spatial attention module. We term our approach SuPEr-SAM.

We conduct experiments on three data sets, comparing our approach with two baselines, namely the standard twostage approach and a two-stage approach that integrates a standard spatial attention module. Each data set includes people wearing a certain PPE item type, namely helmets, masks or boots. We show that our approach significantly outperforms both baselines, irrespective of the PPE type.

The rest of this paper is organized as follows. We present the related works on PPE detection in Section 2. Our approach based on spatial attention trained with supervision signal from a pose estimator is described in Section 3. The comparative experiments are presented in Section 4. Finally, we draw our conclusion in Section 5.

2. Related Work

Some of the first systems developed for PPE detection [1] are based on radio-frequency identification (RFID). Workers have to carry RFID devices or similar sensing devices which are programmed to identify the equipment based on proximity. For example, Barro-Torres *et al.* [1] presented an architecture formed of a wireless local area network and a body area network. Workers carry a microcontroller-based device that detects the presence of the PPE, sending reports to a central unit where alerts are generated. We note that such systems fail if the workers decide not to carry the RFID sensors, just as they would do with the personal protective equipment. Another problem with systems based on sensing devices is that they are costly to implement, since every piece of protective equipment needs to be tagged with a transmitting device. Unlike such systems, our approach employs a computer vision system that detects and recognizes PPE items in images captured by surveillance cameras. Since one surveillance camera can cover an area with multiple workers, a system based on image analysis is more cost effective than a system based on RFID sensors. Furthermore, the workers cannot trick the system by not wearing the sensing devices. Other than being applied for the same task, our method has no connection to works such as [1].

Before deep learning [13] was widely-adopted outside the computer vision community, researchers developed systems for PPE detection based on conventional machine learning or statistical methods [5, 7, 14, 16, 17, 18, 19, 20, 27, 29]. It seems that the majority of systems, even from the recent literature [14, 16, 17, 18, 29], still rely on handcrafted methods. In order to detect helmets, Shrestha et al. [27] employed a standard face detector. The detected faces are passed to a contour extraction method, which checks that the contour forms a semi-circle to confirm the presence of a helmet. The method is not reliable under low-resolution or blurry images or when the faces or helmets are partially occluded. Dahiya et al. [5] compared three methods based on handcrafted feature extraction: Histogram of Oriented Gradients (HOG), Scale-Invariant Feature Transform (SIFT) and Local Binary Patterns (LBP). Their method is designed to detect bike or motorbike riders without helmet. Wu et al. [29] and Mneymneh et al. [16] also employed a set of handcrafted image descriptors, which are combined together instead of being used independently. Nevertheless, these methods are tested on very few images, making it hard to draw generic conclusions. Moohialdin et al. [18] employed decision trees for PPE and posture detection. Decision trees are part of the standard machine learning methods, which cannot model the complexity of data affected by large variations. Li et al. [14] used a background subtraction algorithm to detect people. Their system is not reliable when objects other than people enter the analyzed scene, because these objects belong to the foreground. Another system based on background subtraction and handcrafted features is proposed in [17]. The system is trained and tested on a few hundred images collected in ideal conditions, which raises questions about the robustness of the system with respect to illumination changes, pose variation and occlusion. Mosberger et al. [19] handcrafted a method to detect reflective safety garments. The method is specifically designed to detect the reflective patterns using thresholding and contour extraction. This limitation makes the method inapplicable to other PPE items that are not reflective, e.g. glasses or masks. Moreover, their system requires an expensive camera mount, thus being less cost effective than our approach. Mosberger et al. [20] enhanced their previous method by employing multi-band Hough Forests. While being more accurate than their former system [19], their latter system [20] suffers from the same limitations indicated above.

We note that, in general, the systems based on conventional machine learning methods, e.g. Support Vector Machines [14] or decision trees [18], and handcrafted feature extraction, e.g. HOG [7, 16, 29], SIFT [5] or LBP [5, 16, 29], exhibit poor modeling capacity. Such systems do not cope well under illumination changes, aspect variations or occlusion, among other factors of variation in images. To this end, researchers and engineers turned their attention to deep learning methods, which provide superior results due to their capacity to learn a hierarchy of features from large amounts of data [13]. Our system is based on deep learning, thus being able to achieve much better accuracy rates than systems based on conventional machine learning or handcrafted features. Our method is fundamentally different from [5, 7, 14, 16, 17, 18, 19, 20, 27, 29], since, in our case, the features are automatically learned instead of being designed by hand.

The PPE detection approaches based on deep learning are divided into two categories: (i) training an object detector for PPE item detection or (ii) training a person detector and a classifier that takes the bounding boxes predicted by the person detector as input. The majority of deep learning systems [9, 12, 28, 30, 31, 32] perform PPE detection in a single step, by training an object detector for the corresponding object classes. For example, Fang et al. [9] trained a Faster R-CNN object detector [24] on more than 100,000 images to detect people not wearing helmets. Another approach for helmet detection based on Faster R-CNN is presented in [32]. In other works [12, 31], the authors approached the task of helmet detection using the YOLO object detector [22]. Another approach based on a fast deep object detector is that of Wang et al. [28]. For a fast processing time, their detector relies on a lightweight convolutional neural network as backbone. Wu et al. [30] proposed a deep object detector based on multi-level features to detect people wearing helmets. To progressively encode the multi-level features, the authors employed reverse progressive attention. Their attention mechanism is trained in a conventional way. Unlike Wu *et al.* [30], we train the attention mechanism in a novel and different way. While they use attention in the detector, we employ attention in the classifier. Our novelty consists in training the attention maps by minimizing the mean absolute error with respect to a pseudo-ground-truth attention mask provided by a pose estimator. By using supervision from a pose estimator, our method is able to obtain better accuracy rates compared to attention mechanisms based on conventional training (without additional components added to the loss).

Methods that detect PPE items in a single stage [9, 12, 28, 30, 31, 32] are typically less accurate than two-stage methods performing people detection and PPE recognition [3, 6, 21]. The main issue with methods based solely on object detectors is the inability of object detectors to detect small objects, e.g. helmets or boots, from a far range. Instead of trying to detect these small objects relative to the entire scene, two-stage methods detect people, which are proportionally larger. For this reason, the latter approaches have superior accuracy rates. This statement is confirmed by the experimental results presented in [21]. Our method is based on a two-stage approach, hence it is substantially different from methods performing detection in a single stage [9, 12, 28, 30, 31, 32]. Some recent works [3, 6] focused on helmet detection of bike or motor riders. These systems use the YOLO object detector to detect people. Then, the authors trained convolutional neural networks to classify the detected people in two classes: violation (without helmet) or non-violation (with helmet). Different from such methods, our approach relies on an additional system, namely an articulated body pose estimator, that provides supervision to a spatial attention module inserted in the neural network classifier. This leads to accuracy improvements with an imperceptible computational time increase. To our knowledge, we are the first to train a spatial attention mechanism with supervision signal from a pose estimator.

3. Method

3.1. Overview of Processing Pipelines

The training and inference pipelines of the proposed method are slightly different, since the pose estimator is not employed during inference. We next present the steps involved in each of the two pipelines of our method.

Training pipeline. During training, the input images are processed by three neural models: a person detector, a pose estimator, one or more PPE classifiers. The number of classifiers is determined by the number of distinct types of PPE that need to be detected. We assume that the person detector and the body pose estimator are pre-trained models.

We still need a set of labeled images to train our classifiers. First, we provide the training images to the person detector that outputs bounding boxes enclosing the detected people. Second, we apply the pose estimator on each image region (bounding box), providing as output a skeleton of the body indicating the location of body parts. From an entire bounding box, we automatically crop out a sub-region that can potentially contain a specific PPE item, e.g. a helmet, and provide it to the corresponding classifier, e.g. the helmet classifier. From the corresponding skeleton, we construct a pseudo-ground-truth attention mask indicating the location of the body part of interest inside the crop. For example, the body part of interest for the helmet classifier is the head. In summary, the input for a classifier is formed of an image, for example representing the upper body, and an attention mask, typically indicating the location of the body part of interest, for example the head. We process the input image through a series of convolutional blocks, as illustrated in Figure 1. We pass the activation maps from the last convolutional layer to a spatial attention module. We train each classifier by jointly minimizing an objective composed of a classification loss and an attention loss.

Inference pipeline. During inference, we apply the person detector on test images, obtaining a set of bounding boxes. From each bounding box, we crop out sub-regions that can potentially contain certain PPE items and provide them to the corresponding classifiers. We process each image crop through the corresponding convolutional neural network (CNN) classifier and its spatial attention module. The classifier provides a binary prediction for each input image, indicating whether the (partially-truncated) person in the image wears the corresponding PPE item or not. Additionally, the predicted attention map can be used to explain the model's decision for a given input image.

3.2. Neural Models

Person detector. To detect people in images, we employ the YOLOv3 [23] object detector, which is pre-trained on the MS COCO data set [15]. YOLOv3 [23] is a single-stage object detector that offers superior detection performance and high speed (more than 30 frames per second on a single GPU). For each input frame, the object detector returns the bounding boxes of the objects together with the corresponding class labels. From these, we are only interested in the bounding boxes from the *person* category.

Pose estimator. To obtain the body pose of each detected person, we apply a pre-trained pose estimator based on Part Affinity Fields [2]. The chosen pose estimator relies on global context, jointly learning part locations and their associations by using two branches in the same sequential estimation process. For each input image, the pose estimator processes the image to recover the pose of the articulated bodies depicted in the image, representing the pose as a

skeleton. The skeleton is a tree-structured chain with nodes representing body joints and edges representing rigid body segments. Each edge is associated with a specific body part, enabling us to build pseudo-ground-truth attention maps for the body parts of interest to the PPE classifiers. We note that the computational time of the pose estimator is not of utter importance, as it is used only at training time.

PPE classifier. Our PPE classifier is a lightweight CNN based on the MobileNetV2 architecture [26], which is equivalent with AlexNet [11] in terms of accuracy, yet offers faster processing times, thus being suitable even for mobile devices. In the standard MobileNetV2 model, which is pre-trained on ImageNet [25], we integrate SuPEr-SAM, a spatial attention module that receives supervision signal from the pose estimator. As illustrated in Figure 1, the spatial attention module is introduced after the last convolutional layer. Since the activation maps from the last convolutional layer are passed to the spatial attention module, the neural network is forced to focus on the body part of interest with respect to the classification task. The proposed spatial attention module is composed of two depth-wise reduction layers. The goal of each reduction layer is to reduce the depth of the input tensor (a set of activation maps) to one, while preserving the spatial dimensions. Each reduction layer applies a different rule for the depth-wise reduction operation, *i.e.* one selects the maximum value and the other computes the average. Each reduction layer produces a 2D activation map as output. The activation maps from the two reduction layers are stacked in the third dimension to form a 3D tensor. The tensor is then provided as input to a convolutional layer with a single 7×7 filter. Being formed of a single filter, the output of the convolutional layer is a single 2D activation map. The resulting activation map is passed through a sigmoid gate, scaling the values inside the map to the interval [0, 1]. The final activation map is interpreted as an attention mask denoted by \hat{A} . The attention mask is multiplied element-wise with each activation map that was provided as input to the attention module. The result is a set of activation maps with high filter responses only in the region indicated by the attention mask. To focus on the correct region, the attention mask \hat{A} is trained in order to minimize an attention loss with respect to a pseudo-ground-truth attention mask A, which is automatically computed with the help of the pose estimator. Since each classifier's goal is to detect a specific kind of PPE covering a certain body part, the ground-truth activation mask is generated such that it contains a region of interest in the location where the pose estimator found the body part of interest. The region of interest is indicated as a rectangle filled with the value 1 in the ground-truth attention mask, the rest of the attention mask being filled with the value 0. Since the size of the ground-truth attention mask does not match the size of the predicted attention mask coming out of the spatial attention module, the ground-truth attention mask is rescaled such that the masks have the same size.

3.3. Loss Function

Let $X \in \mathbb{R}^{h \times w}$ be an input image of $h \times w$ pixels, *e.g.* representing the upper body of a person, and $y \in \{0, 1\}$ a ground-truth label associated to the input image. The classifier acts as a function $f : \mathbb{R}^{h \times w} \to [0, 1]$ with parameters (weights) θ , providing as output a class probability \hat{y} indicating whether the input X belongs to the class y. Then, the parameters θ are optimized such that the classifier f is minimizing a classification loss function. In our case, the classification loss is the binary cross-entropy:

$$\mathcal{L}_{class} = -\sum_{i=1}^{n} y_i \cdot \log(\hat{y}_i) + (1 - y_i) \cdot \log(1 - \hat{y}_i), \quad (1)$$

where $\hat{y}_i = f(\theta, X_i)$, $S = \{(X_i, y_i) | X_i \in \mathbb{R}^{h \times w}, y_i \in \{0, 1\}, \forall i = \overline{1, n}\}$ is a training set of images and n is the size of the training set. In addition to the classification loss function, our method optimizes the classifier towards minimizing an attention loss function. Since our attention mask A contains binary values, we opt to train the attention module using binary cross-entropy:

$$\mathcal{L}_{attention} = -\sum_{i=1}^{n} \sum_{j=1}^{h'} \sum_{k=1}^{w'} A_{i,j,k} \cdot \log(\hat{A}_{i,j,k}) + (1 - A_{i,j,k}) \cdot \log(1 - \hat{A}_{i,j,k}),$$
(2)

where *n* is the size of the training set, h' and w' are the height and the width of the predicted attention mask, A_i is the ground-truth attention mask associated to the example X_i , \hat{A}_i is the predicted attention mask for the example X_i , and *j* and *k* iterate through the components of the matrices A_i and \hat{A}_i .

We combine the classification loss and the attention loss into a joint loss function expressed as follows:

$$\mathcal{L}_{joint} = \lambda \cdot \mathcal{L}_{class} + (1 - \lambda) \cdot \mathcal{L}_{attention}, \qquad (3)$$

where λ is a parameter between 0 and 1 that controls the balance between the classification loss and the attention loss. Each PPE classifier is trained using the Adam optimizer [10], minimizing the joint loss function defined in Equation (3). It is important to note that, during inference, the ground-truth attention masks are no longer needed, since these are used only at training time to compute the attention loss defined in Equation (2). Therefore, the articulated body pose estimator is not required at inference time, significantly reducing the processing time of the proposed method. Since only the person detector and the classifiers are employed during inference, our entire architecture can process an input video stream in real-time at about 30 frames per second (32.6 milliseconds per frame, with an average of one person per frame) using one GPU (see Table 3).

4. Experiments

4.1. Data Sets

We evaluate our model on three types of PPE, namely helmets, boots and surgical masks.

Helmets data set. The helmets data set is composed of people wearing and people not wearing helmets. We collected a total of 13,481 images of people wearing helmets and 4,170 images of people not wearing helmets. We split the data set into 12,286 examples for training, 2,628 examples for validation and 2,737 examples for testing, using stratified sampling. Hence, the ratio of people wearing helmets to people not wearing helmets is roughly 3:1 in all three subsets.

Boots data set. The boots data set contains images with people wearing and people not wearing protective boots. We gathered a total of 4,784 images of people wearing protective boots and 988 images of people wearing other types of shoes. We divide the data set in 3,988 examples for training, 866 examples for validation and 918 examples for testing. Since we used stratified sampling for splitting the data, the ratio of people wearing protective boots to people not wearing them is nearly 4.5:1 in all three subsets.

Surgical Masks data set. The surgical masks data set is formed of images of people with and without masks. We collected a total of 1,535 images of people wearing surgical masks and 5,473 images of people that do not wear masks. Using stratified sampling, we split the data set in 3,421 examples for training, 1,710 examples for validation and 1,877 examples for testing. The ratio of people wearing to people not wearing masks in each subset it is 1:3.5.

4.2. Evaluation Metrics

In the evaluation, we employ several metrics, such as the accuracy, the macro-averaged precision, the macroaveraged recall and the macro-averaged F_1 score. As the accuracy is equivalent to the micro-averaged F_1 score, we do not report micro-averaged metrics. A macro-averaged metric computes the score for each class, then averages the results. For example, in the binary case, the macro-averaged precision is computed by calculating the precision for class 0 and class 1, then averaging the resulting scores. Since the macro-averaged metrics do not take into account the imbalanced class distributions in our data sets, we also report the weighted precision, the weighted recall and the weighted F_1 score. For instance, the weighted precision is computed by calculating the precision for class 0 and class 1, but instead of averaging the resulting scores, it weights the scores by the class probabilities.

In the last experiment, we compare the two-stage approaches, including our own, with an approach that performs detection and recognition in a single step, considering the mean Average Precision (mAP). The mAP score is the mean of the AP scores corresponding to all object classes. The AP score for one class is represented by the

Data Set	Method	Accuracy	Macro			Weighted		
			Precision	Recall	F1-score	Precision	Recall	F1-score
Helmets	MobileNet	$96.90{\pm}0.9$	96.96 ± 1.4	95.01 ± 1.1	95.93 ± 1.2	$96.91 {\pm} 1.0$	$96.90{\pm}0.9$	96.87 ± 0.9
	MobileNet+SAM	95.76 ± 1.0	$96.53 {\pm} 0.5$	92.56 ± 2.2	94.28 ± 1.5	$95.88 {\pm} 0.9$	95.76 ± 1.0	95.65 ± 1.1
	MobileNet+SuPEr-SAM	98.92 ±0.4	98.89 ±0.2	98.33 ±0.9	98.60 ±0.5	98.93 ±0.4	98.92 ±0.4	98.92 ± 0.4
Boots	MobileNet	$95.69{\pm}0.5$	89.89 ± 2.2	$91.27 {\pm} 1.7$	90.47 ± 0.8	$95.85 {\pm} 0.4$	$95.69{\pm}0.5$	95.74 ± 0.4
	MobileNet+SAM	95.48 ± 0.9	93.20±2.1	85.69 ± 3.7	88.83 ± 2.6	$95.37 {\pm} 0.9$	95.48 ± 0.9	95.26 ± 1.0
	MobileNet+SuPEr-SAM	96.54 ± 0.2	91.48 ± 1.1	92.93 ±0.9	92.17 ± 0.5	96.63 ±0.2	96.54 ± 0.2	96.57 ± 0.2
Masks	MobileNet	$93.90 {\pm} 1.1$	88.59 ± 1.6	$95.69 {\pm} 0.6$	91.42 ± 1.3	$95.16 {\pm} 0.6$	93.90 ± 1.1	94.17 ± 1.0
	MobileNet+SAM	94.74 ± 1.1	89.88 ± 1.8	96.23 ± 0.7	92.50 ± 1.5	$95.70 {\pm} 0.7$	94.74 ± 1.1	94.94 ± 1.1
	MobileNet+SuPEr-SAM	95.26 ± 1.0	90.92 ±1.9	96.08 ± 0.5	93 .14±1.4	95.92 ± 0.7	95.26 ± 1.0	95.41 ± 1.0

Table 1. Results of the MobileNet and the MobileNet+SAM baselines versus our MobileNet+SuPEr-SAM on the Helmets, the Boots and Surgical Masks data sets. Reported values are averaged over 5 runs. The best score on each data set and each metric is highlighted in bold.

area under the precision-recall (PR) curve. The PR curve is obtained by mapping each detected bounding box to the most-overlapping ground-truth bounding box, considering a minimum Intersection over Union (IoU) score of 0.5 [8].

4.3. Baselines and Models

MobileNet. As our first baseline, we consider a two-stage approach that detects people using a pre-trained YOLOv3 model [23], then classifies the cropped detections using a fine-tuned MobileNetV2 architecture [26].

MobileNet+SAM. As our second baseline, we take into consideration a two-stage approach based on YOLOv3 and MobileNetV2, which introduces a spatial attention module (SAM) in the MobileNetV2 architecture. The module is placed immediately after the last convolutional layer.

MobileNet+SuPEr-SAM (ours). Our model is most similar to the second baseline, the only difference being that the spatial attention module (SAM) is trained using **su**pervision signal from a **p**ose **e**stimator (SuPEr). Our approach is therefore called SuPEr-SAM.

YOLO on faces. Another relevant baseline is to fine-tune the YOLOv3 model to detect and recognize PPE items directly. Since this baseline requires manual bounding box labeling, we test it on a single data set, namely the Surgical Masks data set. Hence, the model detects and recognizes two object categories: faces with masks and faces without masks. Both categories are required, since we also need to know when people do not wear the mask.

4.4. Parameter and Implementation Choices

We trained the MobileNetV2 [26] classification models in Keras [4]. We employed the Adam optimizer [10] with the learning rate set to 10^{-4} , while keeping the default values for the other parameters of Adam. Each model was trained for a maximum of 150 epochs using early stopping with respect to the validation set. In each experiment, we set the mini-batch size to 32 samples. We tuned the parameter λ in Equation (3) on each validation set, considering values between 0.1 and 1, using a step of 0.1. We obtained



Figure 2. Validation accuracy rates on the Helmets data set, while tuning λ in Equation (3) between 0.1 and 1.

optimal results using $\lambda = 0.5$. In Figure 2, we illustrate the parameter tuning process on the Helmets data set.

4.5. Quantitative Results

We present the quantitative results on the three data sets in Table 1. For each system, we report the average scores and the standard deviations for 5 runs, in order to account for the variation that commonly appears in the initialization and the stochastic training process of neural networks.

Results on Helmets. On the Helmets data set, we outperform the two baselines, namely MobileNet [26] and MobileNet+SAM, by large margins, regardless of the metric. For instance, we surpass the standard MobileNet by 2.02%in terms of the accuracy rate and by 2.67% in terms of the macro-averaged F_1 score, respectively. Interestingly, the MobileNet+SAM model attains lower performance levels than the standard MobileNet. By adding SAM, we observe performance drops higher than 1% for almost all metrics. This shows that the spatial attention module was not able to learn relevant spatial information without direct supervision. It is worth noting that our approach is more stable, having lower standard deviations over the 5 runs compared to the two baselines. For example, our standard deviation in terms of the macro-averaged precision is 0.2%, which is significantly lower than the standard deviation of the baseline MobileNet architecture, which is equal to 1.4%. In sum-

Method	No Mask	Mask	mAP
YOLO on faces	91.93%	83.90%	87.91%
MobileNet	90.15%	86.39%	88.27%
MobileNet+SAM	90.33%	86.18%	88.25%
MobileNet+SuPEr-SAM	90.31%	87.19%	88.75%

Table 2. AP and mAP scores of a single-stage pipeline (YOLO on faces) versus three two-stage pipelines based on YOLO on people and various MobileNet classifiers, with and without attention.

mary, the results clearly show the benefit of using the supervision signal from the pose estimator to train the spatial attention module.

Results on Boots. The results presented in Table 1 indicate that we generally outperform both baselines on the Boots data set. Our approach attains a macro-averaged precision of 91.48%, being surpassed by the MobileNet+SAM, which yields a macro-averaged precision of 93.20%. Even though SAM seems to perform better than SuPEr-SAM in terms of the macro-averaged precision, when we consider the macro-averaged recall, our method attains a score that is 7.24% higher. Consequently, the macro-averaged F_1 score of MobileNet+SuPEr-SAM is superior. Moreover, we note that MobileNet+SAM attains generally lower performance scores, even compared to the standard MobileNet architecture. Once again, this proves that the integration of the standard spatial attention module is rather harmful. On the contrary, SuPEr-SAM brings performance gains, regardless of the evaluation metric. Our model is also more stable, having a maximum standard deviation of 1.1%, compared to 3.7% obtained by MobileNet+SAM and 2.2% obtained by MobileNet, respectively.

Results on Surgical Masks. The macro-averaged recall of MobileNet+SAM is slightly higher (0.15%) than the macroaveraged recall obtained by MobileNet+SuPEr-SAM, but our macro-averaged precision is superior, with a difference of 1.04%. Consequently, in terms of the macro-averaged F_1 score (which combines precision and recall), we obtain an improvement of 0.64% compared to MobileNet+SAM. We also surpass the standard MobileNet model, regardless of the evaluation metric. Hence, we conclude that our approach outperforms the two baselines on the Surgical Masks data set. We also note that, on this data set, MobileNet+SAM outperforms the standard model. This is the only case in which we observe performance gains due to the integration of standard spatial attention.

One-stage versus two-stage pipelines. In Table 2, we compare the results of the one-stage pipeline (YOLO on faces) with the results of the two-stage pipelines on the Surgical Masks data set. We present the results of the one-stage pipeline only on one data set, because this approach requires human labor to annotate images at the bounding box level. This represents a disadvantage compared to the two-stage approaches, which require only class label anno-

Method	Detector Time	Classifier Time
YOLO on faces	23.3 ms	-
MobileNet	23.3 ms	8.7 ms
MobileNet+SAM	23.3 ms	9.3 ms
MobileNet+SuPEr-SAM	23.3 ms	9.3 ms

Table 3. Inference times in milliseconds (ms) per image of MobileNet+SuPEr-SAM versus three baselines. Reported times are measured on a machine with an Intel i9-9900K 3.6GHz CPU, 64GB of RAM and an NVIDIA GeForce RTX 2080Ti 11GB GPU.

tations. We want to underline, once again, that the twostage pipelines, denoted by MobileNet, MobileNet+SAM and MobileNet+SuPEr-SAM, are based on detecting people with YOLOv3 and on classifying cropped body parts with MobileNetV2. The Average Precision (AP) of YOLO on detecting faces with masks is 83.90%. When we employ MobileNet+SuPEr-SAM on the detected people, we obtain a much higher AP score (87.19%) for the same class (people with mask). In general, the two-stage pipelines exhibit lower AP scores for people without masks, when we compare them to the single-stage pipeline. However, the mAP scores of the two-stage pipelines are always higher than the mAP score of the single-stage pipeline. Our results confirm the already known fact [21] that two-stage approaches outperform single-stage models.

4.6. Running Time

In Table 3, we compare our model to the baselines in terms of running time. Since MobileNet+SAM and MobileNet+SuPEr-SAM use the same spatial attention module, there is no difference in terms of running time between the two models. MobileNet requires 8.7 milliseconds (ms) to infer the class for one image, while MobileNet+SAM and MobileNet+SuPEr-SAM require 9.3 ms for the same task. Hence, the spatial attention module increases the running time by less than 1 ms (about 6.5%). The YOLOv3 detector requires 23.3 ms per image, being more computationally intensive than MobileNetV2. With an average of one person per image, our complete two-stage pipeline requires 23.3 + 9.3 = 32.6 ms.

4.7. Qualitative Results

Figure 3 illustrates examples of cropped images that are provided to MobileNet+SAM and MobileNet+SuPEr-SAM during inference. For the selected examples, we computed the pseudo-ground-truth attention masks, using the pose estimator to determine the location of body parts. This enables us to compare the pseudo-ground-truth attention masks with the masks predicted by MobileNet+SAM and by MobileNet+SuPEr-SAM, respectively. In general, the attention masks predicted by SuPEr-SAM are significantly more correlated to the pseudo-ground-truth attention masks than the attention masks predicted by SAM. For SAM, we observe visual correlations for the sixth image on the first



Figure 3. Examples of cropped input images and associated pseudo-ground-truth attention masks versus the attention masks predicted by SAM and SuPEr-SAM, respectively. The examples are chosen from the Boots, the Helmets and the Surgical Masks test sets. The labels predicted by MobileNet+SAM or by MobileNet+SuPEr-SAM are placed immediately below the corresponding attention masks. Correct labels are colored in green and wrong labels are colored in red. Best viewed in color.

row (counting from left to right) and for the last two images on the second row. Although, in general, the attention masks produced by SAM seem uncorrelated with the locations of the PPE items, the predicted class labels are still correct in most cases. Nevertheless, we conclude that the illustrated examples clearly show that SuPERr-SAM helps the MobileNet classifier to focus on the right regions.

5. Conclusion

In this paper, we presented a two-stage approach for PPE recognition. Our contribution consists in introducing a pose

estimator at training time, which provides a supervision signal to the spatial attention module incorporated in the classifier. Since the pose estimator is not required during inference, our pipeline adds a negligible margin (less than one millisecond) to the overall computational cost. The empirical results reported on three PPE data sets show that our contribution provides important accuracy gains. We conclude that our novel contribution is both effective and efficient. In future work, we aim to extend our work beyond the task of PPE recognition, integrating SuPEr-SAM in other neural architectures.

References

- Santiago Barro-Torres, Tiago M. Fernández-Caramés, Héctor J. Pérez-Iglesias, and Carlos J. Escudero. Real-time personal protective equipment monitoring system. *Computer Communications*, 36(1):42–50, 2012.
- [2] Zhe Cao, Tomas Simon, Shih-En Wei, and Yaser Sheikh. Realtime Multi-Person 2D Pose Estimation using Part Affinity Fields. In *Proceedings of CVPR*, pages 7291–7299, 2017.
- [3] Aphinya Chairat, Matthew Dailey, Somphop Limsoonthrakul, Mongkol Ekpanyapong, and Dharma Raj KC. Low Cost, High Performance Automatic Motorcycle Helmet Violation Detection. In *Proceedings WACV*, pages 3560–3568, 2020.
- [4] François Chollet et al. Keras. https://keras.io, 2015.
- [5] Kunal Dahiya, Dinesh Singh, and C. Krishna Mohan. Automatic detection of bike-riders without helmet using surveillance videos in real-time. In *Proceedings of IJCNN*, pages 3046–3051, 2016.
- [6] Madhuchhanda Dasgupta, Oishila Bandyopadhyay, and Sanjay Chatterji. Automated Helmet Detection for Multiple Motorcycle Riders using CNN. In *Proceedings CICT*, pages 1– 4, 2019.
- [7] Romuere Rodrigues Veloso e Silva, Kelson Rômulo Teixeira Aires, and Rodrigo de Melo Souza Veras. Helmet Detection on Motorcyclists Using Image Descriptors and Classifiers. In *Proceedings of SIBGRAPI*, pages 141–148, 2014.
- [8] Mark Everingham, S. M. Eslami, Luc Gool, Christopher K. Williams, John Winn, and Andrew Zisserman. The Pascal Visual Object Classes Challenge: A Retrospective. *International Journal of Computer Vision*, 111(1):98–136, 2015.
- [9] Qi Fang, Heng Li, Xiaochun Luo, Lieyun Ding, Hanbin Luo, Timothy M. Rose, and Wangpeng An. Detecting nonhardhat-use by a deep learning method from far-field surveillance videos. *Automation in Construction*, 85:1–9, 2018.
- [10] Diederik P. Kingma and Jimmy Ba. Adam: A method for stochastic optimization. In *Proceedings of ICLR*, 2015.
- [11] Alex Krizhevsky, Ilya Sutskever, and Geoffrey E. Hinton. ImageNet Classification with Deep Convolutional Neural Networks. In *Proceedings of NIPS*, pages 1106–1114, 2012.
- [12] Shubham Kumar, Nisha Neware, Atul Jain, Debabrata Swain, and Puran Singh. Automatic Helmet Detection in Real-Time and Surveillance Video. In *Proceedings of ICM-LIP*, pages 51–60, 2020.
- [13] Yann LeCun, Yoshua Bengio, and Geoffrey Hinton. Deep learning. *Nature*, 521(7553):436–444, 2015.
- [14] Jie Li, Huanming Liu, Tianzheng Wang, Min Jiang, Shuai Wang, Kang Li, and Xiaoguang Zhao. Safety Helmet Wearing Detection Based on Image Processing and Machine Learning. In *Proceedings of ICACI*, pages 201–205, 2017.
- [15] Tsung-Yi Lin, Michael Maire, Serge Belongie, James Hays, Pietro Perona, Deva Ramanan, Piotr Dollár, and C Lawrence Zitnick. Microsoft COCO: Common Objects in Context. In *Proceedings of ECCV*, pages 740–755, 2014.
- [16] Bahaa Eddine Mneymneh, Mohamad Abbas, and Hiam Khoury. Automated hardhat detection for construction safety applications. In *Proceedings of CCC*, volume 196, pages 895–902, 2017.

- [17] Bahaa Eddine Mneymneh, Mohamad Abbas, and Hiam Khoury. Vision-Based Framework for Intelligent Monitoring of Hardhat Wearing on Construction Sites. *Journal of Computing in Civil Engineering*, 33(2):04018066, 2019.
- [18] Ammar Moohialdin, Fiona Lamari, Marc Miska, and Bambang Trigunarsyah. A Real-Time Computer Vision System for Workers' PPE and Posture Detection in Actual Construction Site Environment. In *Proceedings of EASEC*, 2019.
- [19] Rafael Mosberger, Henrik Andreasson, and Achim J. Lilienthal. A Customized Vision System for Tracking Humans Wearing Reflective Safety Clothing from Industrial Vehicles and Machinery. *Sensors*, 14(10):17952–17980, 2014.
- [20] Rafael Mosberger, Bastian Leibe, Henrik Andreasson, and Achim J. Lilienthal. Multi-band Hough Forests for Detecting Humans with Reflective Safety Clothing from Mobile Machinery. In *Proceedings of ICRA*, pages 697–703, 2015.
- [21] Nipun D Nath, Amir H Behzadan, and Stephanie G Paal. Deep learning for site safety: Real-time detection of personal protective equipment. *Automation in Construction*, 112:103085, 2020.
- [22] Joseph Redmon, Santosh Divvala, Ross Girshick, and Ali Farhadi. You Only Look Once: Unified, Real-Time Object Detection. In *Proceedings of CVPR*, pages 779–788, 2016.
- [23] Joseph Redmon and Ali Farhadi. YOLOv3: An incremental improvement. arXiv preprint arXiv:1804.02767, 2018.
- [24] Shaoqing Ren, Kaiming He, Ross Girshick, and Jian Sun. Faster R-CNN: Towards Real-Time Object Detection with Region Proposal Networks. In *Proceedings of NIPS*, pages 91–99, 2015.
- [25] O. Russakovsky, J. Deng, H. Su, J. Krause, S. Satheesh, S. Ma, Z. Huang, Karpathy A., A. Khosla, M. Bernstein, A. C. Berg, and L. Fei-Fei. ImageNet Large Scale Visual Recognition Challenge. *International Journal of Computer Vision*, 115(3):211–252, 2015.
- [26] Mark Sandler, Andrew Howard, Menglong Zhu, Andrey Zhmoginov, and Liang-Chieh Chen. MobileNetV2: Inverted Residuals and Linear Bottlenecks. In *Proceedings of CVPR*, pages 4510–4520, 2018.
- [27] Kishor Shrestha, Pramen P. Shrestha, Dinesh Bajracharya, and Evangelos A. Yfantis. Hard-hat detection for construction safety visualization. *Journal of Construction Engineering*, 2015(1), 2015.
- [28] Lu Wang, Liangbin Xie, Peiyu Yang, Qingxu Deng, Shuo Du, and Lisheng Xu. Hardhat-Wearing Detection Based on a Lightweight Convolutional Neural Network with Multi-Scale Features and a Top-Down Module. *Sensors*, 20(7):1868, 2020.
- [29] Hao Wu and Jinsong Zhao. An intelligent vision-based approach for helmet identification for work safety. *Computers in Industry*, 100:267–277, 2018.
- [30] Jixiu Wu, Nian Cai, Wenjie Chen, Huiheng Wang, and Guotian Wang. Automatic detection of hardhats worn by construction personnel: A deep learning approach and benchmark dataset. *Automation in Construction*, 106:102894, 2019.
- [31] Zaipeng Xie, Hanxiang Liu, Zewen Li, and Yuechao He. A convolutional neural network based approach towards real-

time hard hat detection. In *Proceedings of PIC*, pages 430–434, 2018.

[32] Faishal Zhafran, Endah Suryawati Ningrum, Mohamad Nasyir Tamara, and Eny Kusumawati. Computer Vision System Based for Personal Protective Equipment Detection, by Using Convolutional Neural Network. In *Proceedings of IES*, pages 516–521, 2019.