Classroom Student Posture Recognition Based on an Improved High-Resolution Network

Yiwen Zhang

University of South China

Tao Zhu (✉ tzhu@usc.edu.cn)

University of South China

Huansheng Ning

University of Science and Technology Beijing

Zhenyu Liu

University of South China

Research

Keywords: Pose Estimation, Support Vector Machine, High-Resolution Networks, Squeeze-and-Excitation Networks, Object Detection

Posted Date: January 11th, 2021

DOI: <https://doi.org/10.21203/rs.3.rs-72287/v2>

License: © ⓘ This work is licensed under a Creative Commons Attribution 4.0 International License.

[Read Full License](#)

Abstract

Due to the large number of students in a university classroom and crowded seating, most features of student posture are typically obscured, making it difficult to balance accuracy in identifying students' postures with computational speed. To solve these problems, a classroom student posture recognition method is proposed in this article. First, because we must recognize the poses of multiple people in the classroom, we use the you-only-look-once (YOLOv3) algorithm for object detection and retrain it to detect human objects that are hunching on a table, creating the pose estimation network. Then, to improve the accuracy of the pose estimation network, we use the squeeze-and-excitation network (SENet) structure that is embedded in the residual structure of high-resolution networks (HRNet). Finally, using the improved HRNet algorithm output of key human body points, we design a pose classification algorithm based on a support vector machine (SVM), which is used to classify human poses in the classroom. Experiments show that the improved HRNet multi-person pose estimation algorithm yields the best mean average precision (mAP) performance of 73.76 points in the common objects in context (COCO) validation dataset. We tested the proposed posture classification algorithm using a custom dataset collected in a classroom and were able to achieve a high recognition rate (90.1%) and robustness. Thus, the proposed method can effectively measure student postures.

1. Introduction

In recent years due to the growth of surveillance systems for both public and personal usage, pose estimation and detection methods have been developed to meet the emerging needs of various industries. There are many problem behaviors that university students have in class, such as sleeping, playing on mobile phones and chatting. These inappropriate behaviors affect students' classroom learning. Classroom learning efficiency is an important factor that affects academic performance. Therefore, in a university classroom, student pose estimation and detection can be measured using computer vision technology and have important research significance and application value.

There are two candidates for this task. One is based on object detection algorithms. [1][2] used the improved Faster R-CNN [3] model for object recognition to detect student postures in classrooms. [4] only detected students sleeping based on improved R-FCN [5]. These methods can detect poses in small, low-quality pictures of a classroom where students are concentrated and the collected pictures have low resolution. However, using object detection methods to estimate the poses of each individual on this occasion is hindered by two primary obstacles. First, each method cannot be scaled. If we must detect a new pose, the model must retrain the entire network. Second, each method only identifies poses that significantly differ from others, such as sitting and standing; other less distinct poses, such as reading, chatting and raising your hand, are typically not recognized.

The other method is based on the pose estimation network. [6] uses OpenPose [7] to estimate the location of key points on the human body and then uses a classifier to classify the collected key points. The advantage of this method is that it is easy to use and has a fast calculation speed; however, its

disadvantage is low accuracy in its results. [8] uses pose estimation maps (heatmaps), the byproduct of pose estimation, to recognize human action. This method is effective when used with behaviors that have large movements; however, it cannot learn behaviors with small movements.

Due to the large number of students in classrooms and bodies being covered by objects such as tables or other bodies, there are four primary obstacles in recognizing students' postures in a classroom.

1. The estimation of human-body key points has a high error rate and low accuracy rate.
2. Some human joints are invisible to cameras due to occlusions, relying on only a few unreliable features to estimate human body key points. Therefore, it is difficult to recognize the hunched posture.
3. Most top-down pose estimation methods have low calculation speeds; thus, the results cannot be shown in real time.
4. If only object detection is used to detect poses in a classroom, problems will occur, including only being able to detect a single gesture or poor scalability.

Faced with these difficulties, we propose a posture recognition method for use in a classroom that combines the pose estimation algorithm and the object detection algorithm. The contributions of this paper are threefold. First, we capitalize on the you-only-look-once (YOLOv3) model [9] for object detection to detect human objects and students hunching on tables. Second, we propose an improved HRNet model as the pose estimation algorithm to reduce the error rate of estimating human body key points. We term the proposed improved HRNet as SE-HRNet, which is constructed by embedding the SENet [10] structure into the HRNet model [11]. Finally, we design a multiposture classification network based on a support vector machine (SVM) [12]. Experimental results show that the mean average precision (mAP) that uses YOLOv3 to detect the hunching pose is 91.6 points, the mAP that uses the SE-HRNet model to detect key points of the human body is 73.7 points, the accuracy of the pose classification is 88.6%, and the computation speed of the proposed classroom student posture recognition method is 7 images per second.

The remainder of this paper is organized as follows. In Section 2, we introduce related work. Then, we propose the classroom student postures recognition method in Section 3. In Section 4, we introduce the proposed dataset and then discuss the proposed experimental results in detail. Conclusions are summarized in Section 5.

2. Related Works

2.1 Pose estimation methods

Currently, multi-person human pose estimation methods can be divided into two categories:

Top-down approaches: First, object detection is performed on human bodies in an image, and each human body is cropped into single images. Then, single-person pose estimation is used for each cropped

human body. Therefore, for each detection, a single-person pose estimator is run, and the more people there are, the greater the computational cost. However, the accuracy of the top-down method is typically higher; the common models include CPN [13], hourglass [14], CPM [15], alpha pose [16], etc.

Bottom-up approaches: First, the model detects all key points of the human body in the picture and then matches these points to different individuals; thus, this method is faster in calculation but yields marginally lower accuracy than that of the top-down method. The most common bottom-up model is OpenPose [7]

The high-resolution network (HRNet) [11] is a human body pose estimation method that is an example of a top-down method. High-resolution network pose estimation can maintain high-resolution representations and begins from a high-resolution subnetwork in its first stage, gradually adding high-to-low resolution subnetworks one by one to form more stages and connecting the multiresolution subnetworks in parallel [11]. The network then performs multiple multiscale fusions by repeatedly exchanging information across parallel multiresolution subnetworks and estimates the key points of the human body via the high-resolution representations of the network output. The architecture of HRNet is shown in **Figure. 1**.

HRNet has two benefits compared to the common pose estimation networks [13-16]. First, this approach connects high- to low-resolution subnetworks in parallel, rather than serial, as most existing networks do. Therefore, HRNet can maintain high resolution rather than restore resolution using a low- to high-resolution process. Thus, the predicted heatmap is spatially more precise. Second, most existing fusion schemes combine low- and high-level representations [11]. Conversely, this method uses the low-resolution representation of the same depth and a similar level to perform multiple multiscale fusions to improve the high-resolution representation, and vice versa, giving the high-resolution representation detailed pose estimation data. Thus, this method yields more accurate heatmaps.

HRNet can maintain high-resolution features without the need to recover the high resolution. HRNet also fuses parallel multiresolution representations repeatedly, enhancing the reliability of high-resolution representations and yielding accurate and spatially precise point heatmaps. However, because HRNet is a top-down method, its image processing speed is typically slower than that of a bottom-up method. Additionally, to achieve HRNet multi-person pose estimation, the object detection algorithm must process the image first. Therefore, the detection speed of the object detection algorithm has a strong impact on the pose estimation speed.

2.2 Squeeze-and-Excitation Networks

Squeeze-and-Excitation Networks (SENet) [10] introduced a new architectural unit called squeeze-and-excitation (SE) blocks with the goal of improving the quality of representations produced by a network by explicitly modeling the interdependencies between the channels of conventional features [10]. In this structure, squeeze and excitation are two critical operations. A new "feature recalibration" strategy is used,

through which networks can learn to use global information to selectively emphasize informative features and suppress less useful features.

The structure of the SE building block is shown in **Figure. 2**, where SE represents the SENet blocks. The first block passes through a squeeze operation, which first performs global average pooling on the input feature map to obtain a feature map of size $C \times 1 \times 1$, where C is the number of feature map channels, allowing information from the global receptive field of the network to be used by all its layers. Aggregation is followed by an excitation operation, through which the parameter is used to generate weights for each feature channel, where the parameter is learned from the correlation between the feature channels. After two fully connected layers (first dimensionality reduction and then dimensionality increase), the method uses the sigmoid activation function to obtain a weight of $C \times 1 \times 1$, followed by a re-weighting operation. We regard the output weight as the importance of each feature channel after feature selection and then weight the previous features one by one via multiplication to complete the feature recalibration. The output of the SE blocks can be fed directly into subsequent layers of the network. Both BasicBlock and Bottleneck are the classic residual modules used in ResNet. SE-BasicBlock embeds the SE structure into the regular BasicBlock unit, and SE-Bottleneck embeds the SE structure into the regular Bottleneck unit.

The structure of the SE block is simple and can be directly embedded into existing network architectures, which markedly improves results, is computationally lightweight, and imposes only a marginal increase in model complexity and computational burden.

2.3 Object detection method

Existing object detection algorithms are primarily divided into two types: the two-stage method (i.e., the region proposal method) and the one-stage method (i.e., the regression method). Two-stage object detection algorithms include RCNN [18], Fast RCNN [19] and Faster-RCNN [20]. Faster R-CNN tends to be a slower but more accurate model [21] and consists of two stages. In the first stage, called the region proposal network (RPN), images are processed by RPN to predict class-agnostic box proposals. In the second stage, these proposal boxes are used to crop features from the same intermediate feature maps, which are then entered into the feature extractor to predict a class for each proposal box and to optimize each proposal box.

In a one-stage object detection algorithm (e.g., SSD [22] and YOLO [23]), object classification and bounding-box regression are conducted concurrently without a region proposal stage [24]. YOLO converts object detection into regression work. Based on a single end-to-end network, calculations are completed from the original image to the output of the object position and category. These one-stage methods typically exhibit a high detection speed and high efficiency but low accuracy. YOLOv3 [9] can detect multiple objects with a single inference; thus, its detection speed is high. Additionally, using a multistage detection method, YOLOv3 improves upon the low accuracies of YOLO and YOLOv2 [25]. However, YOLOv3 yields lower detection accuracy than Faster-RCNN with small targets. However, the detection

speed of YOLOv3 is marked higher than Faster-RCNN [20]. Therefore, YOLOv3 is suitable for many engineering applications.

Considering the detection speed and accuracy of the algorithm, this paper uses YOLOv3 for object detection to detect the human body and hunched postures in the classroom, and provides the foundation for the proposed real-time classroom human posture recognition method based on SE-HRNet.

3. Methodology

3.1 Overview of the framework

An overview of the classroom student postures recognition method proposed in this paper is shown in **Figure. 3**. First, we use pretrained YOLOv3[9] to detect images collected from classrooms. Results fall into two categories: one is the human body object provided for SE-HRNet pose estimation, and the other is the hunched posture object. Then, the results of the hunched posture object are directly output, and the results of the human body object are cropped from the image. The cropped images are input into SE-HRNet for pose estimation, which detects the locations of 17 key points of the human body. The next step is to preprocess the output key points. We thus design an SVM classifier to classify the preprocessed key points of the human body, and then output the classification results. Finally, the proposed method is used for real-time posture recognition of student images in a classroom.

3.1 YOLOv3 application

Due to the limitations of classroom usage scenarios, the results of student posture recognition must be shown in real time and must be as accurate as possible. Therefore, we must address the slow estimation speed of HRNet, which is a top-down pose estimation method. The original object detection network used by HRNet is Faster R-CNN [20]. Based on the discussion in section 2.3 of this paper, we propose replacing Faster R-CNN with YOLOv3 [9] for object detection in the proposed method.

Among the three poses we proposed to recognize, the hunched posture is the most difficult to recognize and classify using a pose estimation network. Because the hunched posture means that the person is hunched over the table, usually only the top of his head is shown in the photo, key points of human body will be seriously lost if we try to use the pose estimation network to estimate the hunched posture. This fact makes pose estimation of hunched posture impossible.

To solve this problem, we use an object detection network to detect the hunched posture. We also must use an object detection network to detect human objects in the classroom for the pose estimation network. Therefore, we use the datasets we collected to retrain YOLOv3 to detect the hunched posture and improve the accuracy of human object detection in the classroom.

The proposed method uses YOLOv3 to detect the hunched posture and to improve the existing OpenPose method, which cannot estimate the hunched posture and has difficult in recognizing heavily occluded human body key point postures.

3.2 Designing the improved HRNet

When human bodies overlap, many human body features are occluded; this is particularly true in crowded and complex places, such as classrooms. Conventional pose estimation networks output feature maps that have high confidence in the key points of the overlapping parts. The network mistakenly believes that the overlaps or missing key points are also part of the human body. This unbalanced confidence distribution causes many misidentifications [17]. To enable the network to learn more global features, enhancing the receptive field can be used to balance the confidence of the heat map in different positions. Therefore, we propose embedding the SENet structure into HRNet to increase the global information of HRNet.

The squeeze operation in the SENet structure converts a feature map into a number, which has a global receptive field, and two fully connected layers serve to reduce the number of parameters. HRNet for feature extraction is the key to accurately estimating the key points, which the residual layer fuses into multiple layer features. Therefore, we propose embedding the SENet structure into BasicBlocks and Bottleneck of the HRNet to obtain the SE-BasicBlock and SE-Bottleneck substructures (see **Figure. 2**), thereby expanding the receptive range of the feature map to global information.

The structure of SE-HRNet is shown in **Figure. 4**. SE-HRNet consists of four stages with four parallel subnetworks: the resolution is gradually reduced to half, and the width (i.e., number of channels) is correspondingly increased twice. This paper embeds the SENet structure into the first stage (Stage 1), which contains 4 SE bottleneck units and is composed of an SE bottleneck with a width of 64. The first stage is followed by one 3×3 convolution feature map to reduce the width to C (i.e., the number of channels). The second, third, and fourth stages contain 1, 4, and 3 exchange blocks, respectively. One exchange block contains 4 SE-BasicBlocks embedded in the SENet structure at each resolution, where each contains two 3×3 convolutions and exchanges units across resolutions. Thus, there are a total of 8 exchange units (i.e., a total of 8 multiscale fusions are conducted).

The SE structures introduce primitive information into deep layers, inhibit information degradation, expand the receptive field by pooling, and then integrate shallow information with deep information from multiple dimensions so that the combined output contains multiple levels of information, enhancing the feature map's expression ability.

The results of the SE-HRNet tested on the proposed dataset are compared to OpenPose, as shown in **Figure. 6**. OpenPose mistakenly identified wall patterns as human bodies, and the human pose estimation accuracy was poor. SE-HRNet yields significant improvements over OpenPose, reducing the human-body-object-detected error rate and increasing the accuracy of estimated key points.

The experimental results show that the detection accuracy of HRNet is significantly enhanced by introducing the SE structure, which reduces high detection error rates compared to existing methods that use OpenPose.

3.3 Designing the classification method

Data Preprocessing.

To reduce the number of calculations, speed up convergence, and improve accuracy, the human body key point data output from SE-HRNet must be preprocessed. First, because the coordinate origin of the key point data output by SE-HRNet is in the upper left corner of the image, and each image contains multiple human bodies, it is necessary to shift the coordinate origin to the nose position of the 17 points in each human body. Then, we normalize the data and scale the coordinate data to between 0 and 1 based on the image resolution.

Designing the SVM classifier.

The classifier structure is a simple four-layer fully connected network. Each layer has 125 neurons and uses a rectified linear unit (ReLU) as the activation function. We use the Adam optimizer, the base learning rate is set as 1e-3, and the training process is terminated after 150 epochs. We only use a classifier to classify two types of actions in the classroom: reading and looking. The loss function used is the hinge loss form of the Support Vector Machine (SVM; see Eq. (4)) [12]. The simplest way to extend SVMs for multiclass problems is using the so-called one-vs-rest approach [26]:

$$\min_w \frac{1}{2} w^T w + C \sum_{n=1}^N \max(1 - w^T x_n t_n, 0)^2 \quad (1)$$

The classroom student posture recognition method combines object detection, pose estimation and key point classification. Therefore, if we want to recognize a new pose in the classroom, we only retrain the key point classification network instead of retraining all networks in the method. Therefore, the proposed method improves scalability compared to existing methods by combining three different models.

4. Experimental

4.1 Dataset and setup

Dataset.

In this experiment, the COCO2017 dataset is used to train and validate the improved HRNet [27]. This dataset includes 149,808 pictures and over 250,000 person instances labeled with 17 key points. We evaluate the improved HRNet with the val2017 data set, which contains 6384 pictures.

To describe the situation in the real classroom environment as much as possible, we collected a dataset that included pictures of students in several classrooms during class. This dataset contains many images with different degrees of occlusion and light changes (see **Figure. 6**). This dataset was created using Dahua network dome surveillance cameras. Two dome cameras were installed in a large

classroom of 120 people, and one was installed in a small classroom. A total of 10 cameras were installed in 4 large rooms and 2 small rooms in 6 classrooms of the university. If we take a picture of the entire classroom directly, details are lost; however, we also cannot take a picture of everyone in the classroom. Therefore, we used a sampling method. First, we set a fixed number of cruising points (30 points for a camera in a large classroom and 20 points for a small classroom), and the photos collected from these points covered every seat in the classroom. During a class, the camera moves to a cruising point every 15 s to collect two photos at a resolution of 2592×1520. The advantage of this method is that it can cover every seat in the class, and the collected high-resolution data can be kept intact. The disadvantage of this method is that each collected picture inevitably contains incomplete human body data. As shown in **Figure. 7**, there is always a certain number of incomplete human body data around the upper, lower, left and right sides of the picture. We collected more than 40,000 images; removed many poor-quality images and images with no information about the human body; and ensured that the training and test sets received the same proportion of categories as the dataset. The data collected in this paper contained 1,951 training samples and 943 test samples. Each picture has an average of 5 people, and we labeled a total of 14,470 student body postures. There are three types of poses annotated in this dataset, including reading, hunching and

Experimental environment.

The software environment used in this study included Ubuntu 18.04 based on PyTorch 1.4 with CUDA 10.1. The hardware environment included an Intel Core i7 7820X CPU, 64 GB RAM, and an Nvidia TITAN X (Pascal) 12G graphics card.

Evaluation Metrics.

The standard evaluation metric of the pose estimation experiment is based on object key point similarity (OKS) in Eq. (5), where d_i is the Euclidean distance between the detected key points and the corresponding ground truth; v_i is the visibility flag of the ground truth; s_i is the object scale; and k_i is a per-key-point constant that controls falloff. We report standard average precision and recall scores [11], where average precision (AP) stands for the mean of AP scores at 10 positions, OKS = 0.50, 0.55...0.9, 0.95, and the same applies to the average recall (AR).

The evaluation measurements of the object detection algorithm are the average precision (AP), which was proposed in [28]. A common judgment for the correctness is the intersection-over-union (IOU) between the detection result and ground truth. If the IOU is greater than a threshold percentage of the ground truth size, the result is considered correct [1]. To obtain a higher recall, we set the IOU threshold to 0.5:

$$OKS = \frac{\sum_i \exp(-d_i^2 / 2s_i^2 k_i^2) \delta(v_i > 0)}{\sum_i \delta(v_i > 0)} \quad (2)$$

5. Results

5.1 Training YOLOv3 on the proposed dataset

YOLOv3 uses the Darknet-53 backbone and pretrains the backbone on the COCO dataset. We set the network input resolution to 416×416 and use multiple scale training. The dataset we use to retrain and test YOLOv3 is the dataset containing real classroom images that we collected. We retrain YOLOv3 to detect the hunched posture and student bodies, and the training process is finished within 150 epochs. After retraining YOLOv3, the best result on the test set for the hunched posture is AP=91.6 points.

Table 1. Comparison of different methods using the COCO validation set

Method	Input size	#Params	GFLOPs	AP	AR
OpenPose [7]	368x368	—	—	61.8	66.5
Baseline ResNet-50[29]	256×192	34.0 M	8.90	70.4	76.3
HRNet-W32 (paper) [11]	256×192	28.5 M	7.10	73.4	78.9
HRNet-W32 (our implement)	256×192	28.54 M	7.20	73.1	78.7
SE-HRNet-W32 (our)	256×192	28.75 M	7.21	73.8	79.2

5.2 Comparing different pose estimation methods

To verify the effectiveness of the improved HRNet, several pose estimation frameworks are investigated, including OpenPose [7], original HRNet [11], and ResNet [29], for comparison.

Both the original HRNet [11] and SE-HRNet were trained on COCO2017 [27] with an input size of 256×192 . The learning and dropout rates remain unchanged based on the settings in [11], and the training is set for a total of 210 epochs.

Table 1 shows the results of the proposed improved HRNet compared to other multi-person pose estimation methods on the COCO verification set. The improved HRNet with an embedded SENet structure achieves an AP score of 73.7, outperforming other methods with the same input size (256×192) except OpenPose [7]. The OpenPose uses an input size of 368×368 and is a bottom-up approach. The proposed approach yields much more accurate results than the bottom-up approach: the proposed improved network improves AP by 11.9 points compared to OpenPose. Compared to SimpleBaseline-ResNet-50 [29], the proposed model yields marked improvements: a gain of 3.3 points with a smaller model size and fewer GFLOPs.

Based on these results, the proposed HRNet [11] training results exhibited a marginal decrease (0.3 points) compared to the training results provided in [11]. HRNet's GFLOPs and number of parameters were also marginally higher. The results of the SE-HRNet compared to those of the original HRNet yielded

0.7 points of improvement, and its model size (#Params) and GLOPs did not increase significantly (approximately 1%).

5.3 Comparing different methods

To verify the effectiveness of the proposed method, we try to combine different pose estimation and object detection algorithms with pose classification algorithms. The results are shown in **Table 2**. First, we tried to use OpenPose+SVM as the classroom student posture recognition method. However, OpenPose’s pose estimation is not sufficiently accurate and yielded many classification errors, preventing the hunched posture from being recognized. Because the human body features of the hunched posture are frequently occluded, OpenPose cannot output any useful human body key points. Second, the method using Faster RCNN+HRNet+SVM exhibited certain improvements compared to the method using OpenPose and could recognize the hunched posture because Faster RCNN was used to detect the hunched posture. Additionally, the accuracy of Faster RCNN in detecting the hunched posture was high because Faster RCNN is a two-stage object detection algorithm. Finally, the YOLOv3+SE-HRNet+SVM proposed in this paper yielded significant improvements (8.3%) compared to other methods, reaching 90.1% accuracy. Although YOLOv3 is a one-stage object detection algorithm, the accuracy of detecting the hunched posture is similar to Faster RCNN.

Table 2. Comparisons of different methods on the proposed dataset.

Method	Reading	Looking	Hunching	Accuracy
OpenPose +SVM	67.3%	61.4%	—	64.2%
Faster RCNN +HRNet-W32 +SVM	83.4%	84.5%	92.4%	81.8%
YOLOV3 +SE-HRNet-W32 +SVM (our)	88.6%	89.2%	91.6%	90.1%

5.4 Computational costs

To evaluate computational costs, we tested different approaches on a PC with the same configuration as described above. These results are shown in **Table 3**, where no method includes the last step of pose classification.

The average running time of each image in the proposed method is 0.142 s. Thus, the proposed method is marginally slower than OpenPose because OpenPose is a bottom-up method. However, the proposed method yields a similar time to the HRNet+ YOLOV3 method and a significantly faster time than the HRNet-W32+Faster RCNN method by 404%. These results also show that the proposed improved HRNet does not add much computational cost based on its addition of the SENet structure, indicating that the proposed approach works well in real classroom environments.

Table 3. Comparison of the processing time of each image of these two methods on the proposed dataset

Methods	Time (s)	Frames Per Second (FPS)
OpenPose	0.11	10
HRNet-W32+ Faster RCNN	0.321	3
HRNet + YOLOV3	0.136	7
SE-HRNet + YOLOV3(our)	0.142	7

6. Discussion

This paper introduces a new approach to multi-student posture recognition in a classroom environment based on an improved high-resolution network. Specifically, this method combines the YOLOV3 object detection algorithm and the HRNet pose estimation network and further enhances HRNet with the SENet structure, leading to SVM-based pose classification algorithms. The proposed approach and methods have been tested and evaluated using the COCO validation dataset and a custom dataset, and both yield impressive results. Most existing classroom posture recognition methods are object detection or pose estimation methods; however, the proposed method combines object detection, pose estimation and neural network classification algorithms for multi-student posture recognition. The method proposed in this paper is also compared to the methods used in [1][2][4], which use object detection technology for multi-student posture recognition in a classroom environment. The proposed method can recognize more postures and exhibits better scalability and robustness. Compared to [6], which uses the multi-person pose estimation method, the proposed method improves HRNet using the SENet structure and yields high-accuracy pose estimation in complex cluttered classroom environments.

Figure. 7 shows the representative recognized results of the proposed method from sparse to dense situations. Each pose is labeled immediately next to body key points, and the hunched posture is identified by a bounding box in the images. The proposed method can manage sparse and concentrated distributions of students, and can clearly locate students and recognize their poses in difficult situations, even with occlusions. In another difficult situation, where certain hunch postures can easily be confused with the looking pose, the proposed method still frequently predicts the correct label. In other cases, such as with occlusions or background noise, the proposed method also exhibited more robust results than other methods.

Certain limitations in this study do exist. Due to limited manpower, we collected a small amount of data in the custom dataset used in this study. The amount of data used to test the proposed methodology is not large and thus cannot represent all classroom environments.

7. Conclusions

In this paper, we propose a classroom student posture recognition method. The proposed method combines pose estimation, object detection and posture classification to complete classroom posture

recognition. This paper compares the speed and accuracy of different methods, and then chooses YOLOv3 as the proposed object detection network to detect hunch postures. Then, due to the high error rate of other common pose estimation methods, we propose embedding the SENet structures into HRNet. Experiments tested the proposed improved HRNet on the COCO dataset, and results show that the AP of the proposed method reaches 73.8 points, which is only 0.7 points higher than the original HRNet. Additionally, the GFLOPs and #Params of the proposed method increased marginally. Finally, we design a posture classification algorithm based on SVM. The accuracy of the proposed method can reach 90.1%, outperforming other traditional methods. This method combines three different modules to recognize student postures, yielding strong robustness and scalability. In future work, we will redesign the posture classification algorithm to recognize more student postures. Labeling more data from different students in different environments and at different times will also improve the adaptability of the proposed approach.

Abbreviations

HRNet: high-resolution network

SENet: squeeze-and-excitation networks

YOLO: you-only-look-once

SVM: Support vector machine

COCO: common objects in context

Declarations

Availability of data and materials

The datasets used and/or analyzed during the current study are available from the corresponding author on reasonable request.

Competing interests

The authors declare that there is no conflict of interest.

Funding

This work is supported by the National Natural Science Foundation of China (No. 61872038), Natural Science Foundation of Hunan Province (No. 2019JJ50499)

Authors' contributions

Yiwen Zhang performed the experiments and was a major contributor in designing the new method and writing the manuscript, Tao Zhu contributed to the proposed methods and the experiments. Huansheng Ning and Zhenyu Liu funded for and conceived the idea of the work. All authors read and approved the final manuscript.

Acknowledgements

Not applicable

References

1. Tang L., Gao C., Chen X., Pose detection in complex classroom environment based on improved Faster R-CNN. *IET Image Processing*. 13(3),451-457 (2019).
2. Bin T, Shu-Han Y., Research on the Algorithm of Students' Classroom Behavior Detection Based on Faster R-CNN. *Modern Computer*. (2018).
3. Ren S, He K, Girshick R., Faster R-CNN: Towards Real-Time Object Detection with Region Proposal Networks. *IEEE Transactions on Pattern Analysis & Machine Intelligence*. 39(6),1137-1149 (2017).
4. Li W, Jiang F, Shen R., in *International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, 2019 IEEE Conference. Sleep Gesture Detection in Classroom Monitor System (Brighton, 2019), pp. 7640-7644.
5. Jifeng Dai, Yi Li, Kaiming He, Jian Sun, R-FCN: Object detection via region-based fully convolutional networks. *Advances in Neural Information Processing Systems*, 379–387 (2016).
6. J. Zaletelj, in *Proceedings of the 10th International Symposium on Image and Signal Processing and Analysis. Estimation of students' attention in the classroom from kinect features (Ljubljana, 2017)*, pp. 220-224.
7. Zhe Cao, Tomas Simon, Shih-En Wei, Yaser Sheikh, in *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2017 IEEE Conference. Realtime Multi-Person 2D Pose Estimation Using Part Affinity Fields (Hawaii, 2017), pp. 7291-7299.
8. Mengyuan Liu, Junsong Yuan, In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2018 IEEE Conference. Recognizing Human Actions as the Evolution of Pose Estimation Maps (Salt Lake, 2018), pp. 1159-1168.
9. Joseph Redmon, Ali Farhadi, YOLOv3: An Incremental Improvement. (arXiv, 2018)<https://arxiv.org/abs/1804.02767>. Accessed 8 April 2018.
10. Jie Hu, Li Shen, Gang Sun, in *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2018 IEEE Conference. Squeeze-and-Excitation Networks (Salt Lake, 2018), pp. 7132-7141.
11. Ke Sun, Bin Xiao, Dong Liu, Jingdong Wang, in *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2019 IEEE Conference. Deep High-Resolution Representation Learning for Human Pose Estimation (Long Beach, 2019), 5693-5703.

12. Boser, B.E., A training algorithm for optimal margin classifiers. Paper presented at ACM Fifth Workshop on Computational Learning Theory, Pittsburgh, 1992.
13. Yilun Chen, Zhicheng Wang, Yuxiang Peng, in Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR), 2018 IEEE Conference. Cascaded Pyramid Network for Multi-person Pose Estimation (Salt Lake, 2018), 7103-7112.
14. Newell A., Yang K., Deng J., Stacked Hourglass Networks for Human Pose Estimation. by Leibe B., Matas J., Sebe N., Welling M. (eds) Computer Vision – ECCV 2016. ECCV 2016. Lecture Notes in Computer Science, vol 9912 (Springer, Cham ,2016), pp. 483-499.
15. Wei S. E., Ramakrishna V., Kanade T, et al, in Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR), 2016 IEEE Conference. Convolutional Pose Machines (Las Vegas, 2016), pp. 4724-4732.
16. Fang H., Xie S., Tai Y., Lu C., in Proceedings of the IEEE International Conference on Computer Vision (ICCV), 2017 IEEE Conference. RMPE: Regional Multi-person Pose Estimation (Venice, Italy, 2017), pp. 2334-2343.
17. Xueping Liu, Li Yuqian, Liu Li, Improved YOLOV3 Object Recognition Algorithm with Embedded SENet Structure. Computer Engineering. 45(11), 243-248 (2019).
18. Uijlings J R R, Sande K E A V D, Gevers T, Selective Search for Object Recognition. International Journal of Computer Vision. 104(2), 154-171 (2013).
19. Girshick R., in Proceedings of the IEEE International Conference on Computer Vision (ICCV), 2015 IEEE Conference. Fast R-CNN (Santiago, Chile, 2015), pp. 1440-1448.
20. Ren S, He K, Girshick R, et al., Faster R-CNN: Towards Real-Time Object Detection with Region Proposal Networks. IEEE Transactions on Pattern Analysis & Machine Intelligence. 39(6), 1137-1149(2017).
21. Jonathan Huang, Vivek Rathod, Chen Sun, in Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR), 2017 IEEE Conference. Speed/Accuracy Trade-Offs for Modern Convolutional Object Detectors (Hawaii, 2017), pp. 7310-7311.
22. Liu W. et al: SSD: Single Shot MultiBox Detector. by Leibe B., Matas J., Sebe N., Welling M. (eds) Computer Vision – ECCV 2016. ECCV 2016. Lecture Notes in Computer Science, vol 9905 (Springer, Cham ,2016).
23. Joseph Redmon, Santosh Divvala, Ross Girshick, in Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR), 2016 IEEE Conference. You Only Look Once: Unified, Real-Time Object Detection (Las Vegas, 2016), pp. 779-788.
24. Jiwoong Choi, Dayoung Chun, Hyun Kim, in Proceedings of the IEEE/CVF International Conference on Computer Vision (ICCV), 2019 IEEE Conference. Gaussian YOLOv3: An Accurate and Fast Object Detector Using Localization Uncertainty for Autonomous Driving. (Seoul, Korea, 2019), pp. 502-511.
25. Joseph Redmon and Ali Farhadi, in Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR), 2017 IEEE Conference. Yolo9000: better, faster, stronger (Hawaii, 2017), pp. 7263–7271.

26. Yichuan Tang, Deep Learning using Linear Support Vector Machines. (arXiv, 2013), <https://arxiv.org/abs/1306.0239>. Accessed 2 June 2013.
27. Lin, M. Maire, S. J. Belongie, Microsoft COCO: common objects in context (2014), <https://cocodataset.org>. Accessed 2014.
28. Padilla, S. L. Netto and E. A. B. da Silva, in the 27th International Conference on Systems, Signals and Image Processing (IWSSIP). A Survey on Performance Metrics for Object-Detection Algorithms (Online, 2020), pp. 237-242.
29. XIAO B, WU H, WEI Y., in 15th European Conference on Computer Vision (ECCV). Simple Baselines for Human Pose Estimation and Tracking (Munich, Germany, 2018), pp. 472-487.
30. Machado E., Carrillo I., Collado M., Chen L., Visual Attention-Based Object Detection in Cluttered Environments, Proc. of 2019 IEEE Smart World Congress, pp.133-139, 2019.

Figures

Figure 1

The architecture of the HRNet. The architecture of the HRNet. It consists of parallel high-resolution and low-resolution subnetworks with repeated information exchange between multi-resolution subnetworks (multi-scale fusion).

Figure 2

The architecture of SENet. BasicBlock and Bottleneck are original Residual modules. SE-BasicBlock and SE-Bottleneck are obtained by embedding the SE block into the residual modules.

Figure 3

The overview of our method First, we use pre-trained YOLOv3[9] to detect the images we collect from classrooms. The output results fall into two categories, one is the human body object provided for SE-

HRNet pose estimation, and the other is the Groveling pose object. Then the results of the Groveling pose object are directly output, and the results of the human body object are cropped from the image. The cropped images are input into SE-HRNet for pose estimation. SE-HRNet detects the locations of 17 keypoints of the human body. The next step is to preprocess the data of output keypoints. We design an SVM classifier to classify the preprocessed keypoints of the human body. Then output the classification results. Finally, the proposed method applied to online de-tECTION of real surveillance images of the classroom.

Figure 4

The architecture of the SE-HRNet SE-HRNet consists of four stages with four parallel subnetworks, the resolution is gradually reduced to half and the width (number of channels) is correspondingly in-creased twice.

Figure 5

The type of poses annotated in our dataset To identify students' learning status during class. We labeled a total of three class-room postures of the students, Reading(lift) Groveling(middle) Looking(right).

Figure 6

Comparison of the pose estimation results The pose estimation results of our method (right two images) and OpenPose method (left two images). OpenPose mistakenly identified the patterns of the wall as human bodies and the human pose estimation accuracy is not good. SE-HRNet compares to OpenPose obtains significant improvements: reduced the human body object detectd error rate and increased the accuracy of estimate keypoints.

Figure 7

The results of our proposed method. It shows the representative recognized results of our proposed method of sparse to dense situations. In which each pose is labeled right next to body keypoints, and groveling pose is labeled the bounding-box in the images.