



A View Independent Classification Framework for Yoga Postures

Mustafa Chasmai¹ · Nirjhar Das² · Aman Bhardwaj³ · Rahul Garg¹

Received: 14 April 2022 / Accepted: 13 August 2022 / Published online: 13 September 2022
© The Author(s), under exclusive licence to Springer Nature Singapore Pte Ltd 2022

Abstract

Yoga is a globally acclaimed and widely recommended practice for a healthy living. Maintaining correct posture while performing a *Yogasana* is of utmost importance. In this work, we employ transfer learning from human pose estimation models for extracting 136 key-points spread all over the body to train a random forest classifier which is used for estimation of the *Yogasanas*. The results are evaluated on an in-house collected extensive yoga video database of 51 subjects recorded from four different camera angles. We use a three step scheme for evaluating the generalizability of a Yoga classifier by testing it on (1) unseen frames, (2) unseen subjects, and (3) unseen camera angles. We argue that for most of the applications, validation accuracies on unseen subjects and unseen camera angles would be most important. We empirically analyze over three public datasets, the advantage of transfer learning and the possibilities of target leakage. We further demonstrate that the classification accuracies critically depend on the cross validation method employed and can often be misleading. To promote further research, we have made key-points dataset and code publicly available.

Keywords Machine learning · Pose estimation · *Yogasana* · Transfer learning

Introduction

Human activity recognition (HAR) is one of the important problems in computer vision.¹ It includes accurate identification of the activity being performed by a person or a group with the help of images, videos or raw data collected from Sensors. Wide variety of applications of HAR include but are not limited to active and assisted living (AAL), health-care monitoring, security and surveillance, tele-immersion, and smart home automation [44].

These advances in HAR have been backed by large scale open datasets which have been broadly categorized as action level, behavior level, interaction level and group activities level. A comprehensive category wise list of the datasets along with their details can be referred to in survey [6].

Yoga is a group of physical and mental practices that originated in ancient India. It aims to bring body, mind, and soul in unison. Over the past few decades Yoga has gained a tremendous popularity across the globe as an art and science of healthy living. United Nations has declared 21st June as ‘*International day of Yoga*’. Several researchers have studied and evaluated the medical benefits of practicing Yoga [8, 29, 30, 45, 53, 54]. Including the recent COVID-19 pandemic times study [46] that shows, people who practiced Yoga during the lockdowns experienced lower levels of stress, anxiety, and depression (Fig. 1).

‘*Asana*’ is the literal Sanskrit translation of posture. To get the maximum benefit from *Yogasanas* and to prevent any adverse effects and injuries during practice, one need to perform them correctly. This makes the problem of recognition and correction of Yoga pose, an important one. This requires access to a Yoga trainer, which may not be widely available. Therefore, use of ML based Yoga trainer is proposed.

✉ Mustafa Chasmai
cs1190341@cse.iitd.ac.in

Nirjhar Das
ee3190585@iitd.ac.in

Aman Bhardwaj
aman.bhardwaj@cse.iitd.ac.in

Rahul Garg
rahulgarg@cse.iitd.ac.in

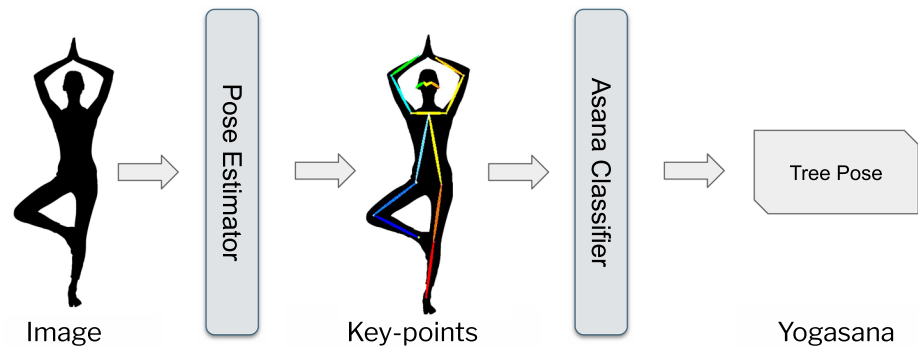
¹ Computer Science and Engineering, Indian Institute of Technology Delhi, Delhi, India

² Electrical Engineering, Indian Institute of Technology Delhi, Delhi, India

³ School of Information Technology, Indian Institute of Technology Delhi, Delhi, India

¹ The dataset and code can be found here: <https://github.com/mustafa1728/Yogasana-Classifier>.

Fig. 1 A simplified view of the 2-stage architecture of the Yogasana Classifier



However, Yogasanas involve complex body postures which are not seen in general activities. Hence, the existing training datasets available for HAR are likely to fail in accurate estimation of Yoga poses.

Several studies have created customized dataset for yoga pose estimation [9, 10, 31, 35, 41, 43, 51, 52, 55]. YogaNet [41] is based on 3D joint angular displacement maps and collects 3D information using complex 3D motion capture systems, whereas the rest employ 2D datasets in the form of either videos [31, 35, 51, 55] or images [9, 10, 43, 52].

Only a few of the listed datasets have been made publicly available [35, 43, 52, 55]. Yoga-82 [52] is a large scale dataset created from 28.4K images of people doing Yogasanas in the wild, collected from various internet sources with the images. These asanas have been classified into 82 poses. Yoga-82 is the most challenging database available till date for Yogasana classification. The other datasets which have been made open are smaller in size, therefore cannot be utilized for Deep Learning model training. For such cases **Transfer Learning** (TL) proves to be an effective technique and is likely to give better performance. The models using TL are expected to generalize well as compared to the ones that are using limited datasets.

Human pose estimation is a key attraction in computer vision and has been an active area of research for past couple of decades. Recently, several novel methods have been proposed that can be widely categorized into two types - (i) top-down approach and (ii) bottom-up approach. In the first approach, the architectures first predict a bounding-box around the humans and these bounding boxes are separately processed to predict the joint keypoint coordinates on the human. Prominent works under this approach are AlphaPose [11], Simple Baseline [12], HRNet [13], DARK [19] and DC-Pose [20].

In the second approach, the keypoints are first predicted without being assigned to a particular human (among other humans) in the image. Instead, after predicting the keypoints, they are grouped together and assigned to the different humans in the image. Prominent works under

this method include OpenPose [22], MultiPoseNet [21], HigherHRNet [23], PifPaf [24], SIMPLE [25]. Since yoga data is limited in size, transfer learning from these models trained on large scale public pose estimation datasets can be employed to improve yoga pose estimation performance.

In this paper, we employ a powerful and efficient pose estimation model, AlphaPose [11], for pose estimation of the Yoga performer, followed by a simple random forest for classification of the performed yoga. We also experiment with the more recent methods DCPose [20] and KAPAO [18] on the Yoga-82 [52] dataset to compare the effect of different pose estimation methods on our overall pipeline.

Choice of the pose estimation model affects the first stage of the pipeline and good performance here propagates to overall good performance of the pipeline. Our pipeline design allows the choice of pose estimation models and can leverage the developments in this domain.

We also collect an extensive in-house dataset, a subset of which has been used to evaluate the performance of our model using a three-stage evaluation strategy consisting of evaluation on (1) unseen frames of Yogasanas, (2) unseen subjects, and (3) unseen angles. In summary, the major contributions of this work are three fold:-

- An extensive dataset explicitly capturing yoga poses from four camera angles for each of the 51 subjects performing 20 asanas. To promote future research, we make the AlphaPose [11] inferred body key points dataset and code openly available.
- A view independent classification framework, orthogonal to the choice of pose estimation or classification algorithms used, employing a three-stage evaluation strategy to provide a more accurate estimate of generalization capacity of a model.
- A simple but effective pipeline involving less computational complexity and real-time inference along with competitive performance exhaustively evaluated on existing datasets.

In the rest of the paper, we first discuss related work in the subsequent section. In “Methodology”, we discuss about

our data collection method, the two stage architecture for yogasana estimation, possibilities of target leakage and our novel evaluation techniques designed to eliminate target leakage. Experimental results of our and three other² openly available datasets can be found in “Experimental Results” followed by discussions on these results in the next section, ending with concluding remarks and future scope in the subsequent section.

Related Work

Human pose estimation, being one of the most popular problems in computer vision, has many large scale and openly available benchmark datasets like COCO [40], Halpe [11, 39], MPII [1], CrowdPose [2] and HiEve [3]. Microsoft COCO [40] dataset is one of the most popular datasets in human pose estimation. It consists of 200,000 images with pose annotation of 17 joint keypoints. MPII [1] dataset contains about 29,000 images of humans performing various activities captured from different angles. This dataset is annotated with 15 body joint keypoints along with their visibility flags. In CrowdPose [2], the dataset contains around 20,000 images with about a total of 80,000 humans spanning across all these images. The images of this dataset are sampled from three existing datasets based on a metric called Crowd Index. HiEve [3] dataset is the largest human pose dataset with a total of >1 M poses spanning across 31 videos. This dataset specially focuses on complex and crowded events like entrance/exits at subways, collision, fighting etc. In the absence of a Yoga-pose benchmark, the performance on these datasets can be taken as a proxy for the performance of various pose estimation methods on yoga poses.

Under the top-down approach as described in Section Introduction, the authors of Simple Baseline [12] propose a simple but effective pipeline so as to provide a strong baseline for pose estimation methods. Their architecture is based on ResNet [26] followed by a few deconvolution layers. With this simple architecture, they achieve 73.7 mAP on COCO dataset and 74.6 mAP and 57.8 MOTA³ score on PoseTrack dataset [4]. In HRNet [13], the authors propose a deep architecture with parallel high-to-low resolution subnetworks with repeated information exchange across multi-resolution subnetworks. They achieve 75.5 mAP on COCO dataset and 74.9 mAP and 57.9 MOTA on PoseTrack dataset. Somewhat orthogonal to these methods, DARK [19] investigates the keypoint coordinate representation in human pose estimation architectures. They propose an efficient

Taylor-expansion based decoding process of predicted joint heatmap to the keypoint coordinate in the original image space and an unbiased sub-pixel centred coordinate encoding scheme. With an HRNet-W48 [13] backbone, DARK [19] achieves 76.2 mAP on COCO dataset. In a different direction, Lite-HRNet [14] proposes a lightweight High Resolution network that focuses on reducing computational overhead while not degrading the performance significantly. They apply the shuffle block from ShuffleNet [5] to HRNet [13] to show performance gain. Thereafter, they replace computationally intensive pointwise (1×1) convolutions in shuffle blocks with conditional channel weighting, wherein weights are learnt across all the channels over multiple resolutions. They achieve 69.7 mAP on the COCO dataset and 87.0 PCKh⁴ on MPII dataset.

There are several methods that rely on the bottom up approach as it reduces the computational overhead in multi-person pose estimation as the model does not need to process each human in the image separately. OpenPose [22] proposes *part affinity fields* (PAFs) which is a method representing the unstructured pairwise relationship between the body parts of the different humans in the image. They achieve 61.8 mAP on COCO dataset using the vanilla method while the foot+body model achieves 65.3 mAP on the COCO dataset. In MultiPoseNet [21], the authors design a deep architecture that consists of a shared backbone of feature extractor which is then fed into two parallel subnets—one is a person detection/segmentation subnet and another is the keypoint detection subnet. The outputs of these two subnets are then fed into a network called Pose Residual Network that assigns the keypoints to the detected persons. This architecture achieves 69.6 mAP on the COCO dataset. HigherHRNet [23] tackles the scale variation of humans in the images. It generates high resolution feature pyramid with multi-resolution supervision during training and multi-resolution heatmap aggregation during inference. The pipeline particularly aims at small humans in images and crowded scenes. It achieves 70.5 mAP on COCO dataset. PifPaf [24] relies on *part intensity fields* (pif) and *part association fields* (paf) to localize body parts and to associate the body parts among each other to fully form the human poses. They also use Laplace loss for regression to encode the uncertainty. This method achieves 66.7 mAP on the COCO dataset. In SIMPLE [25], the authors aim to close the gap in performance in terms of accuracy between the top-down approaches and the bottom-up approaches. The pipeline employs mimicking the estimated heatmaps of a high performance top-down approach to transfer the knowledge of high level features from the top-down model to the bottom-up model. The human detection

² Images in [43] dataset are very similar to those in Yoga-82 [52], but [43] is relatively smaller in size. Thus, we experiment on Yoga-82 [52] and do not explore [43].

³ mean Average Precision and Multiple Object Tracking Accuracy.

⁴ Percentage of Correct Keypoints based on the head.

and pose estimation modules share the same backbone and are unified by treating the problems as point learning problems so that both tasks can benefit each other. It achieves 71.1 mAP on the COCO dataset and a 69.5 mAP and a 55.7 MOTA on the PoseTrack dataset. DEKR [15] is a simple yet effective pipeline that uses adaptive convolutions and a multi-branch structure for attending to the different pixel regions relevant for different keypoints. The authors argue that to learn keypoints by regression, the model needs to focus on the keypoint regions. This is achieved by the pixel-wise extension of spatial transformer network which activate the pixels near a keypoint allowing the model to learn rich representations from these activated pixels. Further, the multi-branch structure helps the model to focus on the pixels relevant to each keypoint separately, thus learning a disentangled representation. This method achieves 71.0 mAP on the COCO dataset.

Another paradigm in pose estimation is that of regression-based methods, where the keypoint coordinates are directly treated as targets and the model is made to learn a regression-based mapping to the pixel coordinates. Although these methods are computationally less expensive compared to heatmap-based methods, their performances are lower. This is because these models fail to incorporate contextual information which present around the keypoint and also fail to capture the inherent uncertainty in the keypoint annotations, especially in the cases of occlusion and motion blur. A notable method in this area is the Residual Log-likelihood Estimation [16] that actually achieves performance higher than the SOTA on the COCO dataset. In [16], the authors propose a novel and effective regression paradigm with the re-parameterization design and Residual Log-likelihood Estimation (RLE) which, instead of learning the actual distribution of keypoint coordinates, attempts to learn the change of the distribution from the presumed distribution. With an HRNet-w48 backbone and RLE, the authors achieve 75.7 mAP on the COCO dataset. In a similar vein, authors of [18] propose a heatmap-free method which they name as KAPAO (Keypoints And Poses As Objects), where individual keypoints and sets of related keypoints (poses) are modelled as objects within a dense single-stage anchor-based detection framework. KAPAO solves the problem of single-stage multi-person human pose estimation by simultaneously detecting human pose and keypoint objects and fusing the detections to exploit the strengths of both object representations. KAPAO achieves 70.3 mAP on COCO dataset with the pipeline being 1–2 orders of magnitude faster. We use KAPAO as a pose estimation method in testing on Yoga-82 [52] dataset to observe the effect of various pose estimation models on our pipeline.

AlphaPose [11] framework consists of three modules—(i) Spatial Symmetric Transformer Network (SSTN) (ii) Parametric Pose Non-Max Suppression (NMP) and (iii) Pose

Guided Proposals Generator (PGPG). The SSTN is used to extract high-quality single person area in the image from an inaccurate bounding box which can come from a sub-optimal object detector. The Parametric Pose NMS eliminates redundant poses using a novel pose distance metric. The PGPG is used to augment the training data to generate (sub-optimal) bounding boxes based on the given pose which are used to train the SSTN. It achieves strong results on the MPII [1] benchmark (76.7 mAP) and the COCO dataset (72.3 mAP), and is able to provide a frame rate of 23 fps when fed with video data. Hence we choose AlphaPose [11] as our pose estimation method.

DCPose [20] aims to solve the problem of multi-person pose estimation in video data. The framework encodes the spatial-temporal keypoint context into localized search scopes, computes pose residuals, and subsequently refines the keypoint heatmap estimations. In particular, the pipeline consists of three task-specific modules—(i) Pose Temporal Merger (PTM) network, which performs keypoint aggregation over three consecutive frames with group convolution, thereby localizing the search range for the keypoint (ii) Pose Residual Fusion (PRF) network, which efficiently obtains the pose residuals between the current frame and adjacent frames and (iii) Pose Correction Network (PCN) comprising five parallel convolution layers with different dilation rates for resampling keypoint heatmaps in the localized search range. This method achieves 79.2 mAP on the PoseTrack dataset. We also experiment with DCPose on Yoga-82 [52] to observe the effect of a different pose estimation method on our pipeline.

Extensive research has been done in the application of pose estimation and classification in Yoga as well (Table 1). Islam et al. [34] used keypoints to get selected joint angles. They used the deviation in these from a set of reference angles as the asana's accuracy. Although they were not involved in classifying an asana, their experiments and results demonstrated that key-points detected from pose estimation are indeed relevant features for asanas.

YogaNet [41] extended this work using JADMs⁵ instead of selecting special angles. Using angles instead of keypoints allowed them to improve their method's position and size invariance to some extent. However, both this and the previous work relied heavily on key-points detected by Microsoft Kinect [56], which we found to be under-performing compared to more recent pose estimation frameworks like AlphaPose [11] and OpenPose [22]. This deep learning approach for extracting key points is a relatively inexpensive alternative which only requires RGB images as compared to the depth and infra-red based approach used by Kinect.

Yadav et al. [55] used OpenPose [22] for key-point extraction and LSTMs [32] for exploiting temporal information,

⁵ Joint Angular Displacement Maps.

Table 1 Summary of related work in chronological order

Work	No. of asanas	Subj	Cam	Key points	Code open	Dataset open	Modality ^a	Target leakage	TL
Islam et al. [34]	–	5	1	20	No	No	Frames	Yes	No
Maddala et al. [41]	42	10	1	25	No	No	Videos	No	No
Yadav et al. [55]	6	12	1	18	Yes	Yes	Frames	Yes	Yes
Verma et al. [52]	82	–	–	–	No	Yes	Images	No	No
Jain et al. [35]	10	27	1	–	No	Yes	Videos	No	No
Gupta et al. [31]	1	20	1	–	No	No	Frames	Yes	No
Ours	20	51	4	136	Yes	Yes	Frames	No	Yes

^aModality can be interpreted as (1) frames - extract input frames from a video, (2) images - still images not taken from any video, (3) video - use complete video segment as input. Note that frames modality with random splits will lead to target leakage

building a complete end-to-end pipeline for classification from yoga videos. Instead of using the keypoint co-ordinates, they used intermediate features learned by OpenPose and processed them with CNNs [38] before putting them into LSTMs [32]. We improve over this work using a better performing pose estimator, larger dataset and more concrete evaluations.

Yoga-82[52] is one of the first large scale Yoga classification dataset openly available. It comprises of 28.4K images of people performing one of the 82 specified yogasanas collected from web search engines. A three-level hierarchical structure for labels, with 20 and 6 super-classes on subsequent levels has been provided. However, they formulated their problem as pose *classification* instead of *estimation*, and thus, do not contain key-point annotations. Also, a significant portion of their data consists of clip art images and diagrams, instead of actual human images, which could lead to a poor performance when using TL from methods trained on real-world examples.

Jain et al. [35] used a complete end to end architecture for this problem. To utilise the spatio-temporal relationship effectively, they proposed to use 3D CNNs [36]. However, they have used a small dataset with only 261 videos.

YogaHelp [31] took a slightly different approach and used various motion sensors to provide feedback and instructions for improvement of a person performing yoga. They extensively explored and tested a single asana for different subjects of varying expertise. They designed different parameters to define the correctness of the asana, and used them for evaluation. They found that using the feedback system, beginners in yoga showed considerable improvement over a small period of four weeks. Their findings demonstrate the overall benefits of this work, and show that a yoga trainer would indeed be very helpful, especially for beginners in the field.

Yadav et al. [55] used pre-trained OpenPose [22] model for extracting their key-points. This was the first work which used Transfer Learning for better performance on yoga classification. Using the pre-trained model allowed them to work

with a deep learning framework even with a relatively small sized dataset. They used a custom yoga dataset having only 6 asanas, 15 subjects and a uniform camera angle. However, for their frame-by-frame evaluation, they used randomly split data for validation, likely leading to target leakage, allowing them to obtain 100% accuracy in three out of the six asanas considered.


One of the major concerns we found in some existing works was that of target leakage which has been formally defined in section [Target Leakage](#). In general, target leakage leads to higher accuracy measures during testing although the underlying model generalization capability may not be that high.

In our work, we try to address some of these limitations we found in existing works. We collect Yoga data in a more systematic way. Larger number of subjects and variations across different camera angles allow us to better generalize our models, and obtain more realistic evaluations. Using a three level evaluation strategy, we address target leakage that was prevailing in most existing works. Lastly, we explored a two-stage architecture instead of an end-to-end approach. This allowed us to use Transfer Learning and leverage existing large scale data and extensive research in the form of pre-trained models. In this domain, where large scale datasets are not available, using TL allowed us to improve our performance and paved the way for future works.

Methodology

Our methodology can broadly be divided into 3 parts. Firstly, we collect an extensive yoga pose dataset while keeping in mind the limitations and shortcomings of existing open datasets. Secondly, we use a 2-stage architecture with human pose estimation for feature extraction followed by decision tree based classification to get asana prediction for each of the frames given as image input. Finally, we use a tri-level evaluation strategy to evaluate the performance of our model.

Fig. 2 An overview of the **a** data collection setup and the **b** 136 key-points detected by AlphaPose trained on Halpe full body dataset



Data Collection

Using Microsoft Kinect [56], we collect data as videos of subjects performing yogasanas. For a better generalization of our models across individuals, 51 volunteers were contacted as subjects for our data. They were each asked to perform 20 yoga poses and 1 still pose (to signify no asana being performed) in a controlled environment. Each video lasts about 2–5 min. Some asanas were bi-lateral, and the subjects performed these twice, one with each side. The two sides were labelled differently.⁶ For such asanas, the subjects were asked to perform the asana in both directions one after the other, and video timestamps for the start and end time of each direction were recorded as well.

Some asanas may be harder to classify from the front facing direction, while they may be easily classified from some other angle. Also, while operating with users in a real-world scenario, the classifier could be fed images from many different orientations. To make our models robust across different view angles, every asana performed by each subject was recorded from 4 different cameras situated in the corners of the room, as shown in Fig 2a.

After all the yoga videos were collected, the pose estimator AlphaPose [11] was first run on all of them together, storing the key-points for each frame in each video. We used model weights pre-trained on the Halpe Full Body [39] dataset directly for this first stage. Once all the videos were inferred, we started processing it for the second stage of classification.

For this work, we consider frame-by-frame inference, and so, need frame-wise data for the second classification stage. While collecting the data, we also recorded timestamps corresponding to start and end times of the asana. From each asana, we uniformly sample some frames between the start and end times, such that the frames are equally spaced. We observed that subjects took some time to get into the final

pose of the asana. So, we used the frames from the beginning of the video till the start of an asana as ‘Still’ frames, adding an extra class in our classifier. Rest of the frames were labelled with the respective asana. This extra class helped improve the overall performance, since these frames would otherwise be wrongly classified as any one of the asanas. However, we estimate approximately 2–3% of mis-labeled data points where the subject has lost balance while performing the yogasana. The key-points detected by AlphaPose in all the frames and the corresponding asana labels formed the training dataset for our classifier.

In summary, videos of 51 subjects were systematically recorded from 4 different camera angles. From these videos, frames in which actual asana was being performed were uniformly sampled, with a maximum of 200 frames per video segment. Katichakrasana was performed by maximum number of subjects, i.e. 47, and different asanas were recorded for different number of subjects. To the best of our knowledge, this is the first yoga dataset explicitly considering asanas from different camera angles, and has the largest number of subjects being recorded. A more detailed description of the number of recorded subjects and extracted frames for each yoga pose and camera, as used in the subsequent sections, can be found in the supplementary material Appendix A

Pose Estimation

Pose estimation is the first and most challenging stage in our model pipeline. A human wanting to distinguish different yogasanas would primarily use the pose and posture of the subject for comparison. Thus, it is only natural to use pose estimation for extracting rich features characterising the pose of the subject. There has already been extensive research in this field and we leverage existing architectures and datasets in our work. Transferring knowledge from openly available large scale datasets allows us to overcome the limitations of scarcity of annotated data to some extent.

AlphaPose [11] is a popular method for Pose Estimation. It follows a two-step framework where it first detects human

⁶ Bi-lateral asanas were labelled as *asana_left* and *asana_right* separately. Unilateral asanas were labelled normally as *asana*.

Fig. 3 Sample images visualising keypoints detected by AlphaPose. **a** Has the same asana and same subject from 4 different camera angles, **b** has the same subject performing different asanas while **c** has different subjects in the ‘still’ pose



bounding boxes and then estimates the pose within each box independently. AlphaPose [11] is comparatively fast, which makes it ideal for real-time tasks. DCPose [20] is one of the latest methods that achieves very high performance on video datasets though the model has greater computational overhead as it processes three consecutive frames of a video at a time. KAPAO [18] is another recently published method that involves a single stage approach and is able to obtain commendable performance (Fig. 3).

AlphaPose [11] pre-trained on the Halpe Full Body [39] dataset, detects 136 keypoints over a person. As can be seen in Fig. 2b, it detects 68 key-points on the person’s face, 42 key-points on both hands together and 26 key-points scattered across the rest of the body. Initially, we used the coordinates of all these 136 key-points as our feature vectors. However, since a large number of keypoints in the face and hands are not very relevant for classifying yoga poses, we observed that replacing these sets⁷ with the mean, minimum and maximum values of their 10 highest confidence key-points allowed us to obtain better performance. In addition to the keypoints, AlphaPose [11] also detects bounding boxes around the subject. We found that including the aspect ratio of these bounding boxes as a feature also increased the performance. We also normalise the key-point coordinates with respect to the bounding box, for better size invariance. Thus, we obtain a 71 dimensional (35 key-points x 2 coordinates + aspect ratio) feature vector from our first stage and pass this to the second stage of classification.

We further consider latest pose estimation methods DCPose [20] and KAPAO [18] as the first stage instead of AlphaPose [11] for Yoga-82 [52] classification. We use DCPose pre-trained on the PoseTrack dataset and KAPAO pre-trained on the COCO dataset, with 17 keypoints each. Unlike AlphaPose, we do not aggregate or reduce the already low number of keypoints here.

⁷ 21 left hand key-points, 21 right hand key-points and 68 face key-points are the three key-point sets considered.

Classification

This second stage of the pipeline is a general classification task. The keypoints detected by the pose estimator are used as features by the classifier. To test the validity of our approach, we train a Random Forest [7] classifier on the inferred keypoints data. We further explore other boosting methods like ADABOOST [48], Gradient Boosting [37] and Bagging [49] classifiers and also an ensemble of the best performing methods. After preparing the dataset of frames and corresponding classes, as explained in Sect. Data Collection, we first trained a Random Forest classifier for the second stage of classification.

Evaluation

Target Leakage

When splitting a dataset randomly into train and test folds, if the data points vary continuously,⁸ then the train and test data points will have very little variation between them. Formally, if we split the data in ratio $m : n$, an expected m data points out of any set of $(m + n)$ data points will be in train set, while the rest will be in the test set. Now, if this subset of $(m + n)$ data points are very similar, then the test offers no challenge to the model as it has already memorized the m training data points similar to the n data points in test set. Hence, the model is observed to perform very well, sometimes even with 100% accuracy. This is particularly common in time series and video data, and is formally termed as *Target Leakage*. Since target leakage will effectively take the testing accuracy closer to training accuracy, it can lead to a false measure of the model’s generalizability.

In the case of Yoga pose classification, target leakage would be possible if frames from a single video are randomly

⁸ Here, we refer to the mathematical continuity of the function describing frame dynamics.

split into train and test. In most previous works that use video data, the frame-wise train and test splits are random, and thus target leakage is inevitable. We use video data too, and so, target leakage is a major concern for us. To eliminate the possibility of target leakage and thus, get a more accurate estimate of generalisation, we must ensure that train and test data come from different videos. We experiment with two such methods in addition to frame-wise evaluation as described below.

Three Stage Evaluation

First, we consider the normal evaluation strategy of frame-wise testing, where the data is randomly split into train and test sets.

Second, we consider dividing the data subject-wise. Since each video in our data records only a single subject, this strategy would ensure that the train and test samples are never contiguous and thus, eliminate the possibility of target leakage. In addition to this, it will also allow us to test the generalizability of our models across unseen subjects.

Lastly, we consider dividing the data camera-wise. Again, as each video only has data from one camera, target leakage would be absent if the data is split camera-wise. Further, this strategy would allow us to test the generalizability of our model across unseen camera orientations. The position of the key points would be very different for data captured from different angles, and a good performance across cameras would be a more concrete measure of the model's actual performance for samples in the wild.

Implementation Details

We use a single RTX 5000 GPU with 16 GB memory for most of our experiments. We use weights trained on Halpe full body [39], PoseTrack[4] and COCO [40] for Alpha-pose, DCPose and KAPAO respectively, with the models detecting 136,⁹ 17 and 17 keypoints respectively. We used hyperparameters recommended in the corresponding paper themselves and further details regarding their methodology can be found in their work. For our second stage classifier, we use off the self implementations provided by the library sklearn. For experiments on the Yoga-82 dataset, we use an ensemble of Histogram Gradient Boosting, LightGBM and Random forest, while for the remaining experiments we use a standard Random Forest. We use the Gini criterion with total 500 trees and allow trees to grow until all leaves are pure. Further details on setting up and running the experiments can be found in our code repository.

⁹ 136 keypoints subsequently reduced to 35, as described in Methodology.

Table 2 Frame-wise results in tenfold cross-validation. This is equivalent to training and testing on all 4 cameras in Table 4

ID	Yoga pose	Precision	Recall	F1 score
0	Garudasana left	98.19%	99.22%	98.70%
1	Garudasana right	99.57%	99.92%	99.74%
2	Gorakshasana	99.88%	99.98%	99.93%
3	Katichakrasana	100.00%	99.87%	99.93%
4	Natavarasana left	99.93%	99.88%	99.91%
5	Natavarasana right	99.75%	99.92%	99.83%
6	Pranamasana left	99.87%	99.83%	99.85%
7	Pranamasana right	99.83%	99.98%	99.91%
8	Tadasana	99.82%	99.88%	99.85%
9	Vrikshasana left	100.00%	99.92%	99.96%
10	Vrikshasana right	99.83%	98.32%	99.07%
11	Still	99.72%	99.65%	99.68%
Average		99.70%	99.70%	99.70%

Experimental Results

Although we recorded data for 20 asanas, the data corresponding to some asanas was of smaller size, and so we only use a subset of these for our evaluation to make the dataset more balanced. We use 72k extracted frames spanning 11 asanas and a still class, giving a total of 12 classes. These 12 classes include left and right asanas for bilateral ones. Each asana has a total of 6000 samples, leading to a highly balanced dataset. For still class, we take all the frames in a video before the start of the actual asana, with a buffer of 1 s. We realise that this buffer may not be sufficient, and as a result, a part of our data may be mislabelled. In all tables, bold numbers indicate the highest score obtained in corresponding column when comparing with different methods and indicate our combined scores when only reporting our results. More details about the size of our dataset and the camera and subject wise distributions can be found in the supplementary material Appendix A.

Frame-Wise Evaluation

Frame-wise evaluation is the standard evaluation strategy used by most of the existing works. We evaluate our model using vanilla 10-fold cross validation. The class-wise results can be seen in Table 2. The random forest classifier obtained mean precision, recall and F1 scores of 99.70% each. Such a high performance due to the explained problem of target leakage (Section [Target Leakage](#)) can be misleading and can also be witnessed in other studies reporting near perfect accuracy for their classifier.

Table 3 Subject-wise results in 10 fold cross-validation

ID	Yoga pose	Precision	Recall	F1 score
0	Garudasana left	88.32%	99.02%	93.36 %
1	Garudasana right	96.32%	99.58%	97.93%
2	Gorakshasana	99.85%	98.08%	98.96%
3	Katichakrasana	99.97%	99.53%	99.75%
4	Natavarasana left	99.97%	97.47%	98.70%
5	Natavarasana right	98.41%	99.22%	98.81%
6	Pranamasana left	99.61%	94.38%	96.93%
7	Pranamasana right	98.93%	98.40 %	98.66%
8	Tadasana	98.16%	98.78%	98.47%
9	Vrikshasana left	100.0%	99.87%	99.93%
10	Vrikshasana right	99.00%	93.70%	96.28%
11	Still	98.69%	97.72 %	98.20%
Average		98.10%	97.97%	97.99%

Subject-Wise Evaluation

Our dataset has 51 subjects performing different yoga poses. The subjects are of varying build, ages and genders. Irrespective of the characteristics of the subject, the pose they take during an asana should be similar. Thus, the classifier should be able to generalize across subjects. To test this, we create folds in which the dataset is split subject-wise into a 9:1 ratio. After training our classifier on the first set of subjects, we test on the remaining. These test subjects would be unseen by the model, and thus, this strategy prevents the model to simply memorise the train subjects, and eliminates the possibility of target leakage (Section [Target Leakage](#)).

We generate 10 such folds, with randomly selected subjects, in rotation, such that each subject gets tested exactly once. The class wise results of these experiments can be found in Table 3. As can be seen, the random forest performs

well even across subjects, with an average precision and recall of 98.10% and 97.97%, leading to an F1 score of 97.99%. Compared to the frame wise results, these results are consistently around 1–2% lower for all three metrics we use for evaluation.

Camera-Wise Evaluation

We expect camera-wise evaluation to be more challenging than the previous two methods. Firstly, many asanas can be very difficult to identify from side or back views. Secondly, the same asana, viewed from different cameras would have very different key-point co-ordinates, and generalizing to such wide range of variations is challenging. Lastly, some angles would have much higher occlusion than others, leading to poorer pose estimation performance as well. The poor performance in the first stage is propagated to the second stage classifier as well.

The results can be seen in Table 4. The results are consistently lower than those obtained from earlier methods, with average precision and recalls of only 81.94% and 81.37% respectively in the case training on 3 camera angles. We also observe that changing the number of angles being used in training directly impacts the performance. A clear decreasing trend can be seen in the performances of training on 3, 2 and 1 cameras in Table 4. This pattern suggests that including a wider variety of camera angles while training tends to give better results. Including all four camera angles would be equivalent to the earlier two evaluation strategies, both of which gave significantly better results.

We perform the same tri-level testing with key-points detected by Microsoft Kinect [56] as well. Using these key-points, the mean F1 scores for frame-wise, subject-wise and camera-wise were 69.49%, 57.91% and 35.21% respectively.

Table 4 Camera-wise results with varying number of Views trained upon

ID	Yoga pose	Train on 3 views			Train on 2 views			Train on 1 view		
		Prec	Rec	F1	Prec	Rec	F1	Prec	Rec	F1
0	Garudasana left	78.4	81.1	79.7	75.8	78.6	77.2	59.2	76.6	64.3
1	Garudasana right	83.4	87.7	85.5	71.0	82.8	76.2	67.1	79.7	72.3
2	Gorakshasana	77.9	68.8	72.6	73.8	77.0	75.3	74.5	57.6	64.9
3	Katichakrasana	80.2	71.1	75.3	61.1	77.5	67.1	73.1	64.4	64.4
4	Natavarasana left	99.9	99.4	99.7	99.6	99.4	99.5	96.8	99.4	98.0
5	Natavarasana right	68.8	47.0	55.1	64.4	46.0	53.2	54.6	37.1	43.7
6	Pranamasana left	60.1	82.2	69.3	52.3	63.2	57.1	37.3	50.9	42.9
7	Pranamasana right	77.7	89.1	83.0	74.9	73.5	74.0	55.0	55.7	55.2
8	Tadasana	98.0	96.5	97.2	96.6	90.3	93.3	78.3	85.7	79.8
9	Vrikshasana left	95.0	93.0	93.9	96.1	90.2	93.0	94.2	78.4	84.6
10	Vrikshasana right	76.6	74.4	75.4	69.9	59.9	64.1	68.8	44.7	51.1
11	Still	86.7	85.6	85.9	92.3	70.5	78.9	76.0	65.6	69.5
Average		81.94	81.37	81.11	77.37	75.79	75.79	69.63	66.37	65.93

Table 5 Frame-wise results for Yadav et al. [55] dataset

ID	Yoga pose	Precision	Recall	F1 score
0	Bhujangasana	98.85%	99.44%	99.14%
1	Padamasana	98.38%	99.79%	99.08%
2	Shavasana	99.59%	98.29%	98.93%
3	Tadasana	99.87%	99.13%	99.50%
4	Trikonasana	99.81%	99.58%	99.69%
5	Vrikshasana	99.37%	99.40%	99.38%
Average (Ours)		99.31%	99.27%	99.29%
Average (Yadav et al. [55])		98.97%	99.11%	99.03%

Table 6 Subject-wise results for Yadav et al. [55] dataset

ID	Yoga pose	Precision	Recall	F1 score
0	Bhujangasana	94.41%	95.83%	95.12%
1	Padamasana	92.33%	94.96%	93.63%
2	Shavasana	98.73%	95.51%	97.09%
3	Tadasana	92.24%	94.30%	93.26%
4	Trikonasana	98.73%	98.80%	98.77%
5	Vrikshasana	94.76%	91.07%	92.88%
Average		95.20%	95.07%	95.12%

These are considerably lower than the corresponding scores we obtained using AlphaPose [11] key-points, with the same second stage classifier. Thus, a good pose estimator is essential for good performance in our method.

Experiments on Yadav et al. Dataset

Yadav et al. [55] collected an in-house dataset consisting of 88 videos, with 15 subjects being recorded across 6 asanas. We suspect that their frame-wise evaluation may have target leakage, and evaluate our model on their data both frame-wise and subject-wise.

Using all the frames extracted from the videos together, the authors were able to achieve 99.04% accuracy on their test set. As can be seen in Table 5, we obtain similar results while training on only 200 uniformly extracted frames from each video. However, going through the extracted frames, we observed that some frames were having the subject in *transition*, and thus would most likely be mislabelled. We estimate that around 3.5% of our extracted frames were mislabelled. With this in mind, an accuracy greater than 96.5% clearly indicates that the model is over-fitting, and that there may be target leakage.

The subject-wise results were, however, more acceptable. The exact results can be seen in Table 6. These results further demonstrate the robustness of our evaluation strategy. Some examples where the model performed badly can be seen in Fig 4. Note that in majority of the wrongly classified cases, the key-points detected by AlphaPose [11] are at fault.

Table 7 Frame-wise results for Jain et al. [35] dataset

ID	Yoga pose	Precision	Recall	F1 score
0	Garland pose	99.27%	98.27%	98.76%
1	Happy baby pose	99.39%	98.65%	99.02%
2	Head to knee pose	98.78%	98.59%	98.69%
3	Lunge pose	96.70%	96.83%	96.76%
4	Mountain pose	97.78%	99.10%	98.44%
5	Plank pose	97.88%	96.41%	97.14%
6	Raised arms pose	94.02%	97.21%	95.59%
7	Seated forward bend	98.63%	98.76%	98.69%
8	Staff pose	98.85%	98.47%	98.66%
9	Standing forward bend	94.57%	93.83%	94.20%
Average (ours)		97.59%	97.61%	97.59%
Average (Jain et al. [35])		91%	91%	91%

In some, there is a second human in sight, while the others are mislabelled. This comparatively poor performance of AlphaPose [11] on yogasana data points to a need for new key-point annotated datasets containing some of the difficult poses common here.

Experiments on Jain et al. Dataset

The Jain et al. [35] dataset comprises of 27 subjects performing 10 different poses. The authors propose the use of 3D CNNs on video segments of 16 frames each. Nonetheless, we analyse the performance of our image based classification on their dataset. Again, we sample 100 frames from each video and use these frames for further analyses. Although this dataset consists of different subjects being recorded, we could not determine which subjects were involved in which video, and thus, will not use our subject-wise evaluation strategy here. These videos were recorded from a single uniform camera angle, and thus, we cannot perform our camera-wise evaluation either.

For frame-wise evaluation, we extracted a total of 23090 high resolution frames and re-scaled them to lower resolution because of resource constraints. The frame-wise results of our method on this dataset can be found in Table 7. Even with image-based prediction instead of video segments, our method outperformed their 3D CNN-based method that gave an accuracy of 91%. This is a clear demonstration of the advantage of using transfer learning over end to end training in this data scarce domain.

Experiments on Yoga-82 Dataset

Yoga-82 [52] is one of the few openly available benchmarks for yoga classification. Although the authors have provided the links to download each image, some links were not reachable. We were able to extract 18k images out of the total 28k in the benchmark. We further curated the data to




Fig. 4 Some wrongly classified examples in Yoga-82 [52] benchmark and Yadav et al [55] dataset. **a** Some example images from Yoga-82 that were mis-classified. These include challenges like inversion, body part occlusion and low image quality and resolution. **b** Some

mislabelled transition frames. The subject’s pose does not match the labelled asana. **c** AlphaPose poor predictions. Part or complete human is missed or bounding box is incorrect

remove clipart-like images and finally had 12k images. We use this subset of 12k images for all our further analysis.

To better demonstrate the benefits of using transfer learning, we have kept the classifier stage as simple as possible. For all of the analysis above, we used a simple random forest classifier. However, random forests were performing quite poorly on the more challenging Yoga-82 [52] dataset. Using an *ensemble* of Random Forests [7], Gradient Boosting [28] and LightGBM [37], we obtained results comparable to those reported by the authors. A brief description of some wrongly classified samples can be seen in Fig. 4.

We also compare the method with the recently proposed methods DCPose [20] and KAPAO [18], but observe somewhat poorer results. We believe this may be because of lesser number of keypoints (17 over the body) and absence of temporal information required for DCPose. The key-points inferred by all these methods for Yoga-82 images will be made openly available for reproducibility. Our results along with those obtained by the Yoga-82 [52] authors and some baselines used by them can be found in Table 8.

In contrast to our dataset and Yadav et al [55] dataset, Yoga-82 consists of images in the wild. Along with a wider variety of subjects and poses, their images are also from many different camera angles. In addition to this, the images are not in a controlled environment, and background conditions are not uniform. All these factors make this a more challenging benchmark, and thus, our model has lower performance here compared to the other datasets. Other models in Table 8 are deep architectures with a high level of complexity. Leveraging transfer learning, we were able to outperform them by training relatively simpler classifiers. This is clear validation of the usefulness of transfer learning in this data-scarce field of yoga classification.

Discussion

While in our experiments we use AlphaPose [11], a similar approach can be applied using other, more recent pose estimation methods. We experiment with KAPAO [18] and

DCPose [20] on Yoga-82 to observe the effect of various pose estimation methods on the first stage of our pipeline. To demonstrate the benefit of transfer learning on this data-scarce domain, we show that even with a simple classifier, the keypoints learnt by AlphaPose lead to quite promising results. Consistent performance on the 4 considered datasets demonstrates the efficacy of this approach. While previous Yogasana classification methods employed complex algorithms with end-to-end training, our simpler method is able to achieve comparable results.

Out of the three evaluation strategies we explore, frame-wise evaluation is by far the most commonly used. Most existing work in this domain use this strategy itself. However, when using frames that have been extracted from videos, this strategy leads to the serious problem of target leakage. Yadav et al. [55] used this strategy and reported 100% accuracies in half of the asanas they experimented with. Our results in Table 2 are also extremely high, with 100% precision in two poses. A similar trend can be seen with Jain et al [35] dataset. Because of target leakage, we believe that frame-wise evaluation is not an ideal testing strategy.

Table 8 Accuracies for all levels of hierarchy in Yoga-82 [52]. The performance of baselines are taken directly from the Yoga-82 paper

Method	Conference	Top-1 accuracy		
		Level 1	Level 2	Level 3
ResNext-101 [27]	CVPR’17	–	–	65.24%
ResNet-101 [26]	CVPR’16	–	–	65.84%
MobileNet-V2 [47]	CVPR’18	–	–	71.11%
DenseNet-201 [33]	CVPR’17	–	–	74.91%
Yoga-82 Variant 1 [52]	CVPRW’20	83.84%	85.10%	79.35%
Yoga-82 Variant 2 [52]	”	89.81%	84.59%	79.08%
Yoga-82 Variant 3 [52]	”	87.20%	84.42%	78.88%
Ours (DCPose [20])*	CVPR’21	86.47%	85.05%	55.07%
Ours (KAPAO [18])*	ECCV’22	86.53%	83.75%	78.01%
Ours*		91.21%	87.91%	80.14%

*Our analysis is on a subset of original dataset. So, the performance is only representative, not for exact comparison

With subject wise evaluations, our model gets tested on subjects it has never seen during training. We obtain good results with this strategy as well, as can be seen in Table 3. This suggests that our model is robust over variations in the subject performing the asana. This is a direct influence of using Transfer Learning in our first stage pose estimator. Although the classifier has not been trained on a particular subject, the pose estimator has been trained on a much wider variety of humans, albeit from a different dataset. This allows it to be robust over variations in the subjects and since the classifier only requires key-points from this pose estimator, this good performance is directly propagated to the final results. Similarly other desirable properties of a pose estimator would be propagated to the final results and so, advances in the field of yogasana pose classification also encourages development and research in the field of pose estimation.

The third evaluation strategy we use is camera-wise evaluation. In Table 4, it can be seen that as the model is trained on lesser number of camera angles, it tends to perform worse. This suggests that inclusion of more camera angles during training has a positive impact on the model's performance and generalizability. Even for the same pose being considered, different camera angles would be producing very different key-point coordinates. We believe that including all these different coordinates should allow a model to learn better quality view independent models, thus improving its performance.

The results of using KAPAO [18] and DCPose [20] on Yoga-82 shows no gain in performance over our experiments with AlphaPose [11], though these two methods outperform AlphaPose significantly on standard datasets. We believe that this is because the model of AlphaPose that we use is pre-trained on the Halpe Full Body [39] dataset which has 136 keypoints annotated on the human body while the model of DCPose is pre-trained on the PoseTrack [4] dataset which contains 17 keypoints and KAPAO is pre-trained on the COCO [40] dataset which consists of 17 keypoints. Since Yoga poses involve the human body to take poses which are extremely non-linear, more keypoints on the body are required to successfully classify a Yogasana. Moreover, the performance of DCPose is lower possibly because DCPose is primarily a pose-tracking model suitable for video data, requiring the previous and the next frame of the current frame to improve the pose estimation of the current frame. Since Yoga-82 is not a video dataset, there is no temporal continuity from one image to the next which can possibly degrade the performance of DCPose on the dataset which then results in poor performance of the classifier.

Besides the model itself being robust to camera angle variation, a different direction could be direct learning of a view independent representation of pose. Ongoing research on 3D pose estimation using only video data provides a new direction where detected poses from very different camera

angles can have identical pose representations, thereby increasing generalizability over camera angles. Although particularly prominent in Yogasana postures, the need to generalize across different camera angles is also evident in many other applications of human pose estimation. A particularly challenging scenario where many pose estimation algorithm fail is occlusion, where part of a body part is hidden or not directly visible. It is noteworthy to see that a body part occluded when viewed from one camera angle may actually be clearly visible from another. If a model is truly view independent, it should be able to get the pose from non-occluded view and use it for the pose in the occluded view. Thus, the study of view independent methods of pose classification may help mitigate the problem of occlusion. A pose classification model that generalizes to unseen camera angles would be quite beneficial to the community.

Yoga-82 [52] has images in the wild, with significant variation in the camera angles, subjects, background lighting conditions and many other factors. The performance on this dataset for the level 3 classes (82 asanas) is very similar to our camera-wise evaluations, where variations other than that of the camera angle are minimal. This suggests that the camera angle variations have the most impact on the performance of a method, compared to frame wise and subject wise variations. Thus, a view independent classification framework that can effectively evaluate the generalizability of the model to different camera angles would provide a much better estimate of the performance of the model for a dataset in the wild. We hope this work can serve as a strong baseline for our and other researchers' subsequent works.

Conclusion and Future Work

In this paper, we propose a two-stage architecture for classification of yoga poses. We use AlphaPose [11], a well established pose estimation method as the feature extractor in the first stage and a random forest classifier in the second stage. In the absence of large scale datasets for yoga pose estimation, we use models pre-trained on existing large scale dataset Halpe Full Body [39]. Transfer learning from pose estimation models is expected to result in better performance for Yogasana classification. We create a new dataset for Yoga Pose classification that focuses on different views of a subject. To the best of our knowledge, ours is the first yoga dataset explicitly considering asanas from 4 different camera angles, and has the largest number of subjects being recorded systematically. The key-points data as well as model codes used will be made open source to further accelerate research in this area.

We devise a 3 step evaluation scheme to evaluate a model's robustness to different sources of variation. In light of particular shortcomings with the individual evaluation

methods, we advocate the use of all three to effectively test generalizability of a model. We demonstrate that our model performs admirably on the first two parts and is competitive with previously published Yogasana classification methods on 3 publicly available datasets. Observing a noticeably lower performance when evaluated across camera angles, we discuss the particular challenges involved in this setting and argue that this camera wise evaluation can often be more reliable as a metric.

There are a few relevant limitations to our approach. Since we rely on pose estimation and accurate detection of keypoints in our first stage, most challenges present in pose estimation naturally reflect in our approach as well. Occlusion and inversion are two commonly studied failure cases. These cases are particularly challenging in the case of Yogasanas because of the complex poses involved here, where some body parts can often be difficult to distinguish or hidden from view because of other body parts. Most large scale pose estimation datasets include normal day to day postures that may not allow a model to perform well on such complex postures. Even in the cases where occlusion occurs, hidden body parts are usually not labelled. Our dataset also contains classification labels only and significantly higher effort and time would be required to collect accurate keypoint annotations for it. Besides scarcity of data, poor generalizability to camera angles is also a limitation for our approach. We believe future work can strive to alleviate many of these limitations. The limitations show significant scope for further research in view independent pose classification and a need for large scale pose estimation datasets with complex poses similar to Yogasanas. 3D pose estimation methods and use of depth information to label occluded or hidden body parts may also help mitigate some of these challenges.

In applications such as that of an automated Yoga trainer, it will be desirable to have models that are robust across different camera angles. To achieve this, either one needs to train models on larger data employing many different camera angles (our results indicate incorporating more cameras in training set improves the test performance), or one needs to incorporate computer vision methods to make these models independent of camera angle. Our results on unseen camera angles show that there is still significant scope for improvement in classification of Yogasanas from different camera angles. This is an exciting area for future study. We believe that subsequent research in camera view independent Yogasana classification will also find many other applications. It is also likely to advance the state-of-the-art in the area of human pose estimation.

Supplementary Information The online version contains supplementary material available at <https://doi.org/10.1007/s42979-022-01376-7>.

Acknowledgements We would like to thank Prof. Varsha Singh and Prof. Sanjeev Jain from IIT Delhi for conceptualization of the Yoga dataset collection and would like to thank Kunal Singh for his contribution in dataset collection. We would like to acknowledge D&L 1-2-3-4 projects (MI01722G D&L-2017) at IIT Delhi for motivating the analysis presented in this work. We would also like to acknowledge Siddhaant Priyam, Misha Mishra, and Ankita Mandal from IIT Delhi for their efforts in setting up multiple pose estimation models and generating AlphaPose inferences on the Yoga dataset.

Declarations

Conflict of interest The authors declare that they have no conflict of interest.

Funding The authors declare that they had no grant related to the funding of this work.

References

1. Andriluka Mykhaylo, Pishchulin Leonid, Gehler Peter, Schiele Bernt. 2d human pose estimation: New benchmark and state of the art analysis. In Proceedings of the IEEE Conference on computer Vision and Pattern Recognition. 2014;3686–3693.
2. Li Jiefenf, Wang Can, Zhu Hao, Mao Yihuan, Fang Hao-Shu, Lu Cewu. CrowdPose: Efficient Crowded Scenes Pose Estimation and a New Benchmark. In Proceedings of the IEEE Conference on computer Vision and Pattern Recognition 2019;10855–10864.
3. Lin Weiyao, Liu Huabin, Liu Shuzhan, Li Yuxi, Qian Rui, Wang Tao, Xu Ning, Xiong Hongkai, Qi Guo-Jun, Sebe Nicu. Human in Events: A Large-Scale Benchmark for Human-centric Video Analysis in Complex Events. arXiv 2020.
4. Andriluka Mykhaylo, Iqbal Umar, Insafutdinov Eldar, Pishchulin Leonid, Milan Anton, Gall Juergen, Schiele Bernt. PoseTrack: A Benchmark for Human Pose Estimation and Tracking. IEEE/CVF Conference on Computer Vision and Pattern Recognition. 2018;5167–5176
5. Ma Ningning, Zhang Xiangyu, Zheng Hai-Tao, Sun Jian. Shufflenet v2: Practical guidelines for efficient cnn architecture design. European conference on computer vision (ECCV). 2018;116–131
6. Beddiar Djamila Romaisa, Nini Brahim, Sabokrou Mohammad, Hadid Abdenour. Vision-based human activity recognition: a survey. Multimedia Tools and Applications 79, 41 (2020), 2020;30509–30555.
7. Breiman L. Random forests. Mach Learn. 2001;45(1):5–32.
8. Büssing Arndt, Michalsen Andreas, Khalsa Sat Bir S, Telles Shirley, Sherman Karen J. 2012. Effects of yoga on mental and physical health: a short summary of reviews. Evidence-based complementary and alternative medicine 2012 (2012).
9. Chen Hua-Tsung, He Yu-Zhen, Hsu Chun-Chieh. Computer-assisted yoga training system. Multimedia Tools and Applications 77, 18 (2018), 2018;23969–23991.
10. Chen Hua-Tsung, He Yu-Zhen, Hsu Chun-Chieh, Chou Chien-Li, Lee Suh-Yin, Lin Bao-Shuh P. Yoga posture recognition for self-training. In International Conference on Multimedia Modeling. Springer, 2014;496–505.
11. Fang Hao-Shu, Xie Shuqin, Tai Yu-Wing, Lu Cewu. RMPE: Regional Multi-person Pose Estimation. In Proceedings of the IEEE International Conference on Computer Vision. 2017;2334–2343.

12. Xiao Bin, Wu Haiping, Wei Yichen. Simple Baselines for Human Pose Estimation and Tracking. In Proceedings of the European Conference on Computer Vision (ECCV). 2018;466–481
13. Sun Ke, Xiao Bin, Liu Dong, Wang Jingdong. Deep high-resolution representation learning for human pose estimation. In Proceedings of the IEEE/CVF conference on computer vision and pattern recognition. 2019;5693–5703
14. Yu Changqian, Xiao Bin, Gao Changxin, Yuan Lu, Zhang Lei, Sang Nong, Wang Jingdong. Lite-HRNet: A Lightweight High-Resolution Network. In Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR). 2021;10435–10445
15. Geng Zigang, Sun Ke, Xiao Bin, Zhang Zhaoxiang, Wang Jingdong. Bottom-Up Human Pose Estimation via Disentangled Keypoint Regression. In Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR). 2021;14676–14686
16. Li Jiefeng, Bian Siyuan, Zeng Ailing, Wang Can, Pang Bo, Liu Wentao, Lu Cewu. Human pose regression with residual log-likelihood estimation. In Proceedings of the IEEE/CVF International Conference on Computer Vision. 2021;11025–11034
17. Li Shichao, Ke Lei, Pratama Kevin, Tai Yu-Wing, Tang Chi-Keung, Cheng Kwang-Ting. Cascaded Deep Monocular 3D Human Pose Estimation with Evolutionary Training Data. In Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. 2020;6173–6183
18. McNally William, Vats Kanav, Wong Alexander, McPhee John. Rethinking Keypoint Representations: Modeling Keypoints and Poses as Objects for Multi-Person Human Pose Estimation. In arXiv preprint [arXiv:2111.08557](https://arxiv.org/abs/2111.08557) (Accepted to ECCV 2022). 2021.
19. Zhang Feng, Zhu Xiattian, Dai Hanbin, Ye Mao, Zhu Ce. Distribution-Aware Coordinate Representation for Human Pose Estimation. In Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR). 2020;7093–7102
20. Liu Zhenguang, Chen Haoming, Feng Runyang, Wu Shuang, Ji Shouling, Yang Bailin, Wang Xun. Deep Dual Consecutive Network for Human Pose Estimation. In Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR). 2021;525–534
21. Kocabas Muhammed, Karagoz Salih, Akbas Emre. MultiPoseNet: Fast Multi-Person Pose Estimation Using Pose Residual Network. In European Conference on Computer Vision (ECCV). 2018;437–453
22. Cao Zhe, Hidalgo Gines, Simon Tomas, Wei Shih-En, Sheikh Yaser. OpenPose: realtime multi-person 2D pose estimation using Part Affinity Fields. *IEEE transactions on pattern analysis and machine intelligence* 43, 1 (2019), 2019;172–186.
23. Cheng Bowen, Xiao Bin, Wang Jingdong, Shi Honghui, Huang Thomas S., Zhang Lei. HigherHRNet: Scale-Aware Representation Learning for Bottom-Up Human Pose Estimation. In Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR). 2020;5386–5395.
24. Kreiss Sven, Bertoni Lorenzo, Alahi Alexandre. PifPaf: Composite Fields for Human Pose Estimation. In Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR). 2019;11969–11978.
25. Zhang Jiabin, Zhu Zheng, Lu Jiwen, Huang Junjie, Huang Guan, Zhou Jie. SIMPLE: SIngle-network with Mimicking and Point Learning for Bottom-up Human Pose Estimation. In Proceedings of the AAAI Conference on Artificial Intelligence. 2021;3342–3350.
26. He Kaiming, Zhang Xiangyu, Ren Shaoqing, Sun Jian. Deep Residual Learning for Image Recognition. In IEEE Conference on Computer Vision and Pattern Recognition (CVPR). 2016;770–778.
27. Xie Saining, Girshick Ross. Aggregated residual transformations for deep neural networks. In IEEE Conference on Computer Vision and Pattern Recognition (CVPR). 2017.
28. Friedman JH. Stochastic gradient boosting. *Computational statistics & data analysis*. 2002;38(2002):367–78.
29. Gothe Neha P, Khan Imadh, Hayes Jessica, Erlenbach Emily, Damoiseaux Jessica S. Yoga effects on brain health: a systematic review of the current literature. *Brain Plasticity* 5, 1 (2019), 2019;105–122.
30. Groessl Erik J, Chopra Deepak, Mills Paul J. An overview of yoga research for health and well-being. *Journal of Yoga & Physical Therapy* 5, 4 (2015), 2015;1.
31. Gupta Ashish, Gupta Hari Prabhat. YogaHelp: Leveraging Motion Sensors for Learning Correct Execution of Yoga with Feedback. *IEEE Transactions on Artificial Intelligence* (2021). 2021.
32. Hochreiter S, Schmidhuber J. Long short-term memory. *Neural Comput.* 1997;9(8):1735–80.
33. Huang Gao, Liu Zhuang, Maaten Laurens Van Der, Weinberger Kilian Q. Densely connected convolutional networks. In Proceedings of the IEEE conference on computer vision and pattern recognition. 2017;4700–4708.
34. Islam Muhammad Usama, Mahmud Hasan, Ashraf Faisal Bin, Hossain Iqbal, Hasan Md Kamrul. Yoga posture recognition by detecting human joint points in real time using microsoft kinect. In 2017 IEEE Region 10 humanitarian technology conference (R10-HTC). IEEE, 2017;668–673.
35. Jain Shrajal, Rustagi Aditya, Saurav Sumeet, Saini Ravi, Singh Sanjay. Three-dimensional CNN-inspired deep learning architecture for Yoga pose recognition in the real-world environment. *Neural Computing and Applications* 33, 12 (2021), 2021;6427–6441.
36. Ji Shuiwang, Xu Wei, Yang Ming, Yu Kai. 3D convolutional neural networks for human action recognition. *IEEE transactions on pattern analysis and machine intelligence* 35, 1 (2012), 2012;221–231.
37. Ke G, Meng Q, Finley T, Wang T, Chen W, Ma W, Ye Q, Liu T-Y. Lightgbm: A highly efficient gradient boosting decision tree. *Adv Neural Inf Process Syst*. 2017;30(2017):3146–54.
38. LeCun, Yann, Bottou, Léon, Bengio, Yoshua, Haffner, Patrick. Gradient-based learning applied to document recognition. *Proc. IEEE* 86, 11 (1998), 1998;2278–2324.
39. Li Yong-Lu, Xu Liang, Liu Xinpeng, Huang Xijie, Xu Yue, Wang Shiyi, Fang Hao-Shu, Ma Ze, Chen Mingyang, Lu Cewu. Pastanet: Toward human activity knowledge engine. In Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. 2020;382–391.
40. Lin Tsung-Yi, Maire Michael, Belongie Serge, Hays James, Perona Pietro, Ramanan Deva, Dollár Piotr, Zitnick C Lawrence. Microsoft coco: Common objects in context. In European conference on computer vision. Springer, 2014;740–755.
41. Maddala Teja Kiran Kumar, Kishore PVV, Eepuri Kiran Kumar, Dande Anil Kumar. YogaNet: 3-D Yoga Asana Recognition Using Joint Angular Displacement Maps With ConvNets. *IEEE Transactions on Multimedia* 21, 10 (2019), 2019;2492–2503.
42. Newell Alejandro, Yang Kaiyu, Deng Jia. Stacked hourglass networks for human pose estimation. In European conference on computer vision. Springer, 2016;483–499.
43. Pandit Niharika. New Yoga Pose Dataset. 2020. <https://www.kaggle.com/general/192938>
44. Ranasinghe Suneth, Al Machot Fadi, Mayr Heinrich C. A review on applications of activity recognition systems with regard to performance and evaluation. *International Journal of Distributed Sensor Networks* 12, 8 (2016), 2016;1550147716665520.
45. Ross Alyson, Thomas Sue. The health benefits of yoga and exercise: a review of comparison studies. *The journal of alternative and complementary medicine* 16, 1 (2010), 2010;3–12.

46. Sahni Pooja Swami, Singh Kamlesh, Sharma Nitesh, Garg Rahul. Yoga an effective strategy for self-management of stress-related problems and wellbeing during COVID19 lockdown: A cross-sectional study. *PLoS one* 16, 2 (2021), 2021:e0245214.
47. Sandler Mark, Howard Andrew, Zhu Menglong, Zhmoginov Andrey, Chen Liang-Chieh. Mobilenetv2: Inverted residuals and linear bottlenecks. In *Proceedings of the IEEE conference on computer vision and pattern recognition*. 2018;4510–4520.
48. Schapire Robert E. Explaining adaboost. In *Empirical inference*. Springer, 2013;37–52.
49. Skurichina M, Duin RPW. Bagging for linear classifiers. *Pattern Recogn.* 1998;31(7):909–30.
50. Toshev Alexander, Szegedy Christian. Deeppose: Human pose estimation via deep neural networks. In *Proceedings of the IEEE conference on computer vision and pattern recognition*. 2014;1653–1660.
51. Trejo Edwin W, Yuan Peijiang. Recognition of Yoga poses through an interactive system with Kinect device. In *2018 2nd International Conference on Robotics and Automation Sciences (ICRAS)*. IEEE, 2018;1–5.
52. Verma Manisha, Kumawat Sudhakar, Nakashima Yuta, Raman Shanmuganathan. Yoga-82: a new dataset for fine-grained classification of human poses. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition Workshops*. 2020;1038–1039.
53. Wang Donna. The use of yoga for physical and mental health among older adults: a review of the literature. *International Journal of Yoga Therapy* 19, 1 (2009), 2009;91–96.
54. Woodyard Catherine. Exploring the therapeutic effects of yoga and its ability to increase quality of life. *International journal of yoga* 4, 2 (2011), 2011;49.
55. Yadav Santosh Kumar, Singh Amitojdeep, Gupta Abhishek, Raheja Jagdish Lal. Real-time Yoga recognition using deep learning. *Neural Computing and Applications* 31, 2019;12(2019),9349–9361.
56. Zhang Zhengyou. Microsoft kinect sensor and its effect. *IEEE multimedia* 19, 2012;2 (2012), 4–10.
57. Zhu Rui, Galoogahi Hamed Kiani, Wang Chaoyang, Lucey Simon. Rethinking reprojection: Closing the loop for pose-aware shape reconstruction from a single image. In *Proceedings of the IEEE International Conference on Computer Vision*. 2017;57–65.

Publisher's Note Springer Nature remains neutral with regard to jurisdictional claims in published maps and institutional affiliations.

Springer Nature or its licensor holds exclusive rights to this article under a publishing agreement with the author(s) or other rightsholder(s); author self-archiving of the accepted manuscript version of this article is solely governed by the terms of such publishing agreement and applicable law.