Real-time Human Pose Estimation with Convolutional Neural Networks

Marko Linna[1], Juho Kannala[2] and Esa Rahtu[3]
[1]*University of Oulu, Finland*
[2]*Aalto University, Finland*
[3]*Tampere University of Technology, Finland*

Keywords: Human Pose Estimation, Person Detection, Convolutional Neural Networks.

Abstract: In this paper, we present a method for real-time multi-person human pose estimation from video by utilizing convolutional neural networks. Our method is aimed for use case specific applications, where good accuracy is essential and variation of the background and poses is limited. This enables us to use a generic network architecture, which is both accurate and fast. We divide the problem into two phases: (1) pre-training and (2) finetuning. In pre-training, the network is learned with highly diverse input data from publicly available datasets, while in finetuning we train with application specific data, which we record with Kinect. Our method differs from most of the state-of-the-art methods in that we consider the whole system, including person detector, pose estimator and an automatic way to record application specific training material for finetuning. Our method is considerably faster than many of the state-of-the-art methods. Our method can be thought of as a replacement for Kinect in restricted environments. It can be used for tasks, such as gesture control, games, person tracking, action recognition and action tracking. We achieved accuracy of 96.8% (PCK@0.2) with application specific data.

1 INTRODUCTION

Human pose estimation in unconstrained environment is a problem where humans yet perform better than computers. In recent years, the research has moved from traditional methods (Felzenszwalb et al., 2008; Andriluka et al., 2009; Yang and Ramanan, 2011; Sapp and Taskar, 2013) towards convolutional neural networks (ConvNets) (Jain et al., 2013; Toshev and Szegedy, 2014; Pfister et al., 2014; Jain et al., 2014; Carreira et al., 2015; Pishchulin et al., 2015; Pfister et al., 2015; Tompson et al., 2015; Lifshitz et al., 2016; Wei et al., 2016; Newell et al., 2016; Charles et al., 2016). Due to this, significant improvements in accuracy have been accomplished. ConvNets became popular, when AlexNet (Krizhevsky et al., 2012) was introduced. AlexNet could classify images on different categories. Since then, several more efficient network architectures have been proposed, for both classification and human pose estimation.

Many state-of-the-art ConvNet human pose estimation methods uses more complex network architectures and they perform considerably well in unconstrained environments (Lifshitz et al., 2016), (Newell et al., 2016), (Insafutdinov et al., 2016), where large variations in pose, clothing, view angle and background exists. While these methods have high accuracy, they are usually slow considering real time pose estimation. Recent research (Toshev and Szegedy, 2014), (Pfister et al., 2014) shows that by using a generic ConvNet architecture, a competitive accuracy can be achieved, while still maintaining a fast forward pass time. This is the main motivation of our research. With our method, we don't aim for overall human pose estimation in diverse input data, but rather target to specific use cases where high accuracy and speed are required. In such cases, the problem is different, because the environment is usually constrained, persons are in close proximity of the camera and poses are restricted. Possible application for our method are, for instance, gesture control systems and games.

Our method is a multi-person human pose estimation system, targeted for use case specific applications. In order to support multiple people, we use a person detector, which gives locations and scales of the persons in the target image. This brings our method towards the practice, since person location and scale are not expected to be known, which is the case with many state-of-the-art methods (Lifshitz et al., 2016), (Wei et al., 2016), (Newell et al., 2016). We use a generic ConvNet architecture, with eight layers. The key idea of our method is to pre-train the network with highly diverse input data and then finetune

Linna, M., Kannala, J. and Rahtu, E.
Real-time Human Pose Estimation with Convolutional Neural Networks.
DOI: 10.5220/0006624403350342
In *Proceedings of the 13th International Joint Conference on Computer Vision, Imaging and Computer Graphics Theory and Applications (VISIGRAPP 2018) - Volume 5: VISAPP*, pages 335-342
ISBN: 978-989-758-290-5

it with use case specific data. We show that competitive accuracy can be achieved in application specific pose estimation, while operating in real-time. Our method can be used for higher level tasks, for example, gesture control, gaming, action recognition and action tracking.

The main contributions of our method are: (1) utilization of person detector to crop person centered images in both training and testing, thus enabling multi-person pose estimation in real world images, (2) ability to learn from heterogeneous training data, where the set of joints is not the same in all the training samples, thus enabling to use more varied datasets in training, (3) utilization of Kinect for automatic training data generation, thus making it easy to generate large amount of annotated training data, (4) somewhat slower and less accurate depth sensor free alternative for Kinect (Shotton et al., 2013) in restricted environments. The frame rate of our method is about 13 Hz, when with Kinect it is 15 or 30 Hz, depending on the lighting conditions. Our method works with RGB cameras while Kinect needs also a depth sensor.

2 RELATED WORK

Jain et al. (Jain et al., 2013) demonstrated that ConvNet based human pose estimation can meet the performance, and in many cases outperform, traditional methods, particularly deformable part models (Felzenszwalb et al., 2008) and multimodal decomposable models (Sapp and Taskar, 2013). Their network architecture consisted of three convolutional layers, followed by three fully connected layers. They trained the network for each body part (e.g. wrist, shoulder, head) separately. Each network was applied as sliding windows to overlapping regions of the input image. A window of pixels was mapped to a single binary output: the presence or absence of that body part. This made possible to use much smaller network, at the expense of having to maintain a separate set of parameters for each body part.

Another application to human pose estimation was presented by Toshev and Szegedy (Toshev and Szegedy, 2014). Their network architecture was similar to AlexNet (Krizhevsky et al., 2012), but the last layer was replaced by a regression layer, which output joint coordinates. In addition to this, they trained a cascade of pose regression networks. The cascade started off by estimating an initial pose. Then at subsequent stages, additional regression networks were trained to predict a transition of the joint locations from previous stage to the true location. Thus, each subsequent stage refined the currently predicted pose. Similar idea is applied in more recent work by Carreira et al. (Carreira et al., 2015).

A video based human pose estimation method was introduced by Pfister et al. (Pfister et al., 2014). Their method utilized the temporal information available in constrained gesture videos. This was achieved by training the network with multiple frames so that the frames were inserted into the separate color channels of the input. The network architecture was similar to AlexNet, having five convolutional layers, followed by three fully connected layers, from which the last one was a regression layer. However, there were some differences compared to the previous architectures. Some of the convolutional layers were much deeper and pooling was non-overlapping, when in most of the previous architectures it was overlapping. The network produced significantly better pose predictions on constrained gesture videos than the previous work. For this reason, we base our method to this network architecture.

3 METHOD

Our method is targeted for video inputs. The rough steps for a single video frame in testing are: (1) detect persons, (2) crop person centered images, (3) feedforward person images to the pose estimation network. We use an object detector (Ren et al., 2015) to solve person bounding boxes from the input frame. The pose estimation is done for each person individually. As a result of the pose estimation, our network outputs locations of body keypoints.

We pre-train our network by using data from multiple publicly available datasets, thus offering good initialization values for finetuning. We evaluate pretraining and finetuning separately. For the evaluation of the finetuning, we use data recorded with Kinect. As for ConvNet framework, we use Caffe (Jia et al., 2014) with small modifications.

3.1 Person Detection

Our method utilize Faster R-CNN (F-RCNN) (Ren et al., 2015) to detect persons from training and testing images. The forward pass time of the F-RCNN is 60ms or 200ms, depending on the used network. We use the slower and more accurate model.

We noticed that sometimes F-RCNN gives false positives. This is not a problem in training, since we use both the ground truth and the F-RCNN together to crop the training image. But in testing, the pose estimation is also performed for false positives. However,

Table 1: Overview of used datasets in pre-training. Only the training set of the MPII Human Pose is used, because the annotations are not available for the test set. In the BBC Pose, the training set is annotated semi-automatically (Buehler et al., 2011), while the test set is manually annotated. We use only manually annotated data from the BBC pose. We use data augmentation to expand the number of training images.

		Person boxes we use from the dataset			Person boxes we use for pre-training and validation	
Dataset	Annotated points	Train	Test	Total	Train (aug.)	Validation
MPII Human Pose (Andriluka et al., 2014)	1-16	28821	0	28821	71018	1160
Fashion Pose (Dantone et al., 2013)	13	6530	765	7295	14538	694
Leeds Sports Pose (Johnson and Everingham, 2010)	14	1000	1000	2000	5074	146
FLIC (Sapp and Taskar, 2013)	11	3987	1016	5003	14780	0
BBC Pose (Charles et al., 2013)	7	0	2000	2000	6764	0
		40338	4781	45119	112174	2000

most likely these false positives could be filtered, especially with use case specific images, by adjusting the parameters of the F-RCNN. In the evaluation, we use also the ground truth to decide if the frame has a person or not, so it is guaranteed that all the evaluation frames contain a person. Apart from this, we ran the F-RCNN for the original finetuning evaluation frames, where the ground truth was not yet used for the frame selection. This resulted in false positive rate of 2.86% and false negative rate of 0.65%. In all of the original evaluation frames, there is one fully visible person making gestures in constrained environment. Person detection was considered false if the resulted bounding box did not contain a person, or if it had partially visible person on the edges of the bounding box. In other words, if the intersection-over-union (IoU) ratio between the detection and the ground truth was 0.5 or less.

3.2 Data Augmentation

The F-RCNN person detector is applied for each training image. For each detected person, the IoU between the detected person bounding box and the expanded ground truth bounding box is calculated. The expanded ground truth person box is the tightest bounding box, including all the joints, expanded by a factor of 1.2. The person box having the biggest IoU is selected as the best choice. Based on the best IoU, the training image is augmented by using either of the person bounding boxes, or both (see Table 2).

Table 2: The relation between the person box overlapping ratio and the data augmentation.

	Person box type used in augmentation	
Overlapping ratio	F-RCNN	Ground truth
IoU > 0.7	X	
IoU < 0.5		X
$0.5 \geq$ IoU ≤ 0.7	X	X

In practice, this means that if the detected person box is near to the ground truth expanded person box, only the former is used to crop the person image. And if the detected person box is far from the ground truth expanded person box, only the latter is used to crop the person. And in between of these, both person boxes are used to crop the person, resulting in two training images, where both have small differences in translation and scale. The shortest side of a person box is expanded to equal the longest side, resulting a square crop area, defining the person image used in training. Zero padding is added where needed. A single cropped person image is rescaled to size 224×224 before feeding it to the network.

In addition to aforementioned, a training image is augmented by doing a horizontal flip. All in all, a single person image from a source dataset can result in either two or four augmented person centered training images.

3.3 Pre-training

We pre-train the model from scratch by using several publicly available datasets (see Table 1). The number of annotated joints varies between the datasets. The MPII Human Pose (Andriluka et al., 2014), Fashion Pose (Dantone et al., 2013) and Leeds Sports Pose (Johnson and Everingham, 2010) have full body annotations, while the FLIC (Sapp and Taskar, 2013) and BBC Pose (Charles et al., 2013) have only upper body annotated. Since we use a single point for the head, and because the MPII Human Pose and Leeds Sports Pose have annotations for the neck and head top, we take the center point of these and use it as a head point.

As we aim to study that whether additional partially annotated training data brings improvement over

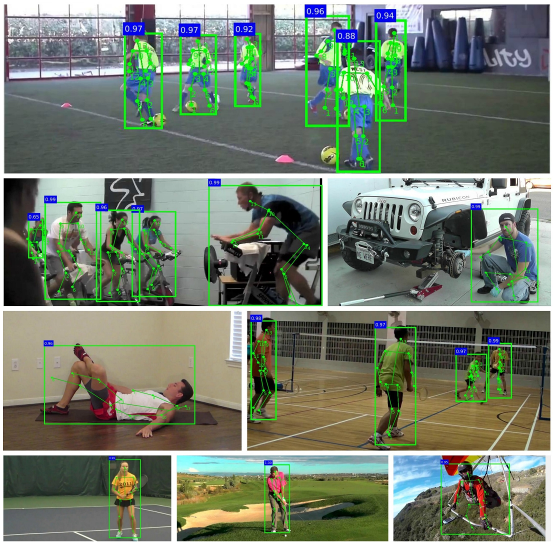

Figure 1: Example pose estimations with the pre-trained network. Samples are taken randomly from the testing set of the MPII Human Pose dataset. The green bounding boxes are the results of person detection and the number on the top-left corner is the probability of a box containing a person.

using only fully annotated samples, our validation samples should be fully annotated. Thus, we put all the fully annotated (13 joints) person images to a single pool and sample 2000 images randomly for validation. The validation images are then removed from the pool. Next, we put all the partially annotated images to the same pool so that it eventually contains person images with heterogeneous set of annotated joints. Then we use the pool in training. The purpose of the pre-trained model is to offer a good weight initialization values for finetuning. Pre-training takes 23 hours on three NVIDIA Tesla K80 GPUs. Fig. 1 contains example pose estimations with the pre-trained network.

3.4 Finetuning

The purpose of the finetuning is to adapt the pre-trained model for the particular use case. For instance a gesture control system or a game. The pre-trained model alone is not a good enough pose estimator for our use cases, because the shallow network we use lacks the capacity to perform well with highly diverse training data. More complicated network architectures, such as (Newell et al., 2016), (Lifshitz et al., 2016) would certainly give better results, but then the speed gain achieved with shallow network architecture would most likely be lost.

In finetuning, the pre-trained model is used for weight initialization. When the network is finetuned

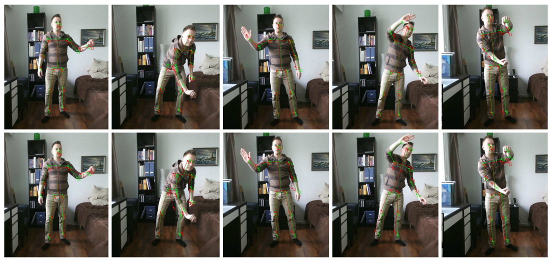

Figure 2: Example pose estimations with the finetuned network. Predictions are in red and Kinect ground truth in green. On the columns are five different frames from the evaluation data. The first row shows results of the full finetuning (experiment 3) and the second row shows results of the phase 1 (experiment 1). Experiments are explained later in Section 4. Full videos are available at https://youtu.be/qjD9NBEHapY and https://youtu.be/e-P5SYL-Aqw.

with use case specific data, for example to estimate poses in gesture control system, the training data is most likely consistent. This is a good thing when thinking of accuracy. Even a shallow network can produce very good estimations, if the training data is limited to particular use case. Using more complicated, and potentially slower, network architectures in these situations is therefore not necessary. We use Kinect in our experiments to produce annotations for the finetuning data, but alternative methods can be considered as well. Fig. 2 contains example pose estimations with the finetuning evaluation data.

3.5 Network Architecture

Our method utilizes generic ConvNet architecture, having five convolutional layers followed by three fully connected layers, from which the last layer is regression layer (see Fig. 3). The regression layer produces (x, y) position estimates for human body joints. More closely, one estimation for head, six for arms and six for legs, a total of 13 position estimations. The network input size is $224 \times 224 \times 3$. We use generic ConvNet architecture, because it has shown to perform well in human pose regression tasks (Toshev and Szegedy, 2014), (Pfister et al., 2014). The forward pass time of the network is 16ms on Nvidia GTX Titan GPU, which makes it highly capable for real-time tasks.

3.6 Training Details

In model optimization, the network weights are updated using batched stochastic gradient descent (SGD) with momentum set to 0.9. In pre-training, where the network is trained from scratch, the learning rate is set

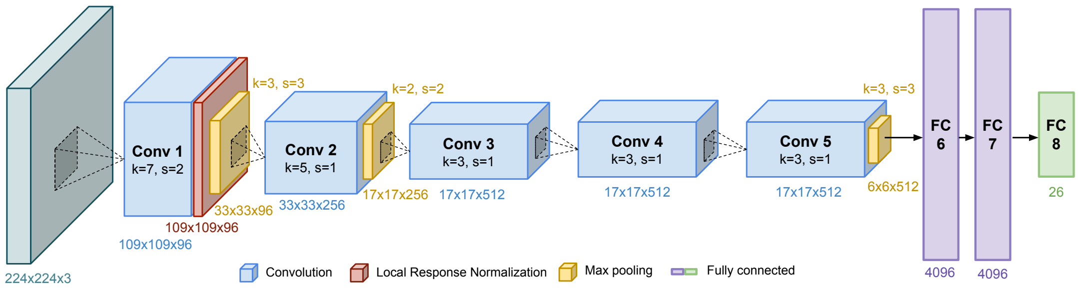

Figure 3: The architecture of the pose estimation network. Letters k and s denotes kernel size and stride.

to 10^{-2}, weights are initialized randomly using Xavier algorithm (Glorot and Bengio, 2010) and biases are set to zero. In finetuning, the learning rate is set to 10^{-3}. The loss function we use in optimization, penalizes the distance between predictions and ground truth. We use weighted Euclidean (L2) loss

$$E = \frac{1}{2N}\sum_{i=1}^{N} w_i \left\| x_i^{gt} - x_i^{pred} \right\|_2^2 \tag{1}$$

where vectors w, x^{gt} and x^{pred} holds joint coordinates and weights in form of $(x_1, y_1, x_2, y_2, \ldots, x_{13}, y_{13})$. Weight w_i is set to zero if the ground truth of the joint coordinate x_i^{gt} is not available. Otherwise it is set to one. This way only the annotated joints contribute to the loss. This enables training the network using datasets having only the upper body annotations, along with datasets having full body annotations. Ability to utilize heterogeneous training data, where the set of joints is not the same in all training samples, potentially leads to better performance as more training data can be used.

As for comparison, we train the pre-trained model also without using the weighted Euclidean loss. In this case, we use only images with fully annotated joint positions (13 joints), so that the training data is homogenous regarding to joint annotations. Doing this reduces the size of the training data from 112174 to 66598 images. The average joint prediction error with heterogeneous and homogenous data are 15.7 and 16.6 pixels on 224×224 images. With heterogeneous data, there is about 5% improvement on prediction error.

In batched SGD, we use batch size of 256. Each iteration selects images for the batch randomly from the full training set. A training image contains roughly centered person of which joints are annotated. The training images are resized to 224×224 before feeding to the network. Mean pixel value of 127 is reduced from every pixel component and the pixel components are normalized to range [-1, 1]. Joint annotations are normalized to range [0, 1], according to the cropped person centered image.

3.7 Testing Details

The person detector is applied for an image from which poses are to be estimated. Person images are cropped based on detections as described earlier. In addition, for each person image, a horizontally flipped double is created. Both the original and the doubled person images are fed to the network. The final joint prediction vector is average of the estimations of these two (the predictions of the doubled image are flipped so that they correspond predictions of the original image). By doing this, a small gain in accuracy is achieved.

4 EVALUATION

We evaluate pre-training and finetuning with the percentage of correct keypoints (PCK) metric (Sapp and Taskar, 2013), where the joint location estimate is considered correct, if its L2 distance to the ground truth is at most 20% of the torso length. The torso length is the L2 distance between the right shoulder and the left hip.

We use 2000 randomly taken samples for the evaluation of the pre-training. For finetuning, we record data with Kinect for Windows v2 (see Table 3). We use the joint estimates produced by Kinect as a ground truth. We made sure that the data was recorded in a such way, that the error in the joint estimations is minimal. Practically this means good lightning conditions, no extremely rapid movements and no major body part occlusions. The gestures performed in the data tries to mimic different gesture control events, where the hands are used for tasks like object selection, moving, rotating and zooming, in addition to hand drawing and wheel steering.

For the evaluation of the finetuning, we record additional 4000 frames with identical clothing . We do three finetuning experiments, using different set of training frames in each case (See Table 4). The experiments 1 and 2 together uses the same training frames as the experiment 3. Basically, the experiment 3 is the same as the experiments 1 and 2 performed consecutively. The purpose of this divide is to see the effect of using the same/different clothing between the training and testing data. The experiment 1 express more of the ability of generalization (for all people) while the experiments 2 and 3 of specificity (for certain people).

The results are displayed in Fig. 4 and Table 5. In full finetuning (experiment 3), with the use case specific data, the accuracy of 96.8% is achieved. In finetuning phase 1 (experiment 1), where no same clothing occurs between the training and testing data, the accuracy is 90.6%. However, if we look at the accuracy of wrist (pre-train: 24.5%, phase 1: 67.4%, full: 89.2%), which is the most challenging body joint to estimate, but perhaps also the most important one considering a gesture control system, we can see that additional case specific training data can significantly improve the accuracy and make the system usable in practice. This originates partially from the finetuning data, where the wrist location variation is biggest. We believe, that if more training data would be used, and perhaps a better data augmentation, a better wrist accuracy could be achieved with the current network architecture. After all, the wrist accuracy is still decent, making our method useful for many use cases.

The results indicate that a trade-off between generalization and specificity exists between pre-training and finetuning. This can be seen by comparing accuracies between the pre-trained and finetuned networks, first with the pre-train validation samples and then with the finetuning validation samples. The pre-train validation samples express the case of generalization as they contain a large variation of persons and poses in unconstrained environment. On the contrary, the finetuning validation samples reflects the case of specificity as they have restricted poses in constrained environment. After the full finetuning, the accuracy on the pre-train validation set drops from 63.1% to 44.2% (light red and dark red curves in Fig. 4), while in the same time, the use case specific accuracy increases from 69.6% to 96.8% (blue and magenta curves). In certain cases, the loss in generalization is acceptable, if at the same time, gain in specificity is achieved. One example of a such case is a gesture control system set up in a factory, where all the persons wear identical clothing. Most importantly, while generic person detection in highly varying poses and contexts is an important and challenging problem, our results show that in some use cases the state-of-the art for the generic problem may produce inferior results compared to a simpler approach which has been specifically trained for the problem at hand.

Table 3: Kinect recorded finetuning data for training. All the frames have similar background, person and gestures, but clothing differs. For the evaluation, we additionally record 4000 frames, which have identical clothing (clothing number 1).

Clothing	Frames
1	27222
2	18760
3	20244
4	20726
5	10560
6	11666
7	10136
	119314

5 CONCLUSION

We introduced a real-time ConvNet based system for human pose estimation and achieved accuracy of 96.8% (PCK@0.2) by finetuning the network for specific use case. Our method can be thought of as a replacement for Kinect in restricted environments. It can be used in various tasks, like gesture control, gaming, person tracking, action recognition and action tracking. Our method supports heterogeneous training data, where the set of joints is not the same in all the training samples, thus enabling utilization of different datasets in training. The use of a separate person detector brings our method towards the practice, where the person locations in the input images are not expected to be known. In addition, we demonstrated an automatic and easy way to create large amounts of annotated training data by using Kinect. The network forward time of our method is 16ms, without the person detector and with the person detector, either 60+16=76ms or 200+16=216ms.

Table 4: Finetuning experiments. The training data have (1) different clothing from the testing data in every frame, (2) the same clothing as the testing data in every frame, (3) the same clothing as the testing data in some of the frames. In phase 2, the finetuning is done over already finetuned network of the phase 1. Otherwise it is done over the pre-trained network.

#	Name	Initialization network	Clothing in training frames
1	phase 1	pre-train	2,3,4,5,6,7
2	phase 2	phase 1	1
3	full	pre-train	1,2,3,4,5,6,7

Table 5: The results of pose estimation (PCK@0.2). The first three cases uses pre-train validation samples (2000 images) in testing, while other models use finetuning validation samples (4000 frames).

Network	Head	Wrist	Elbow	Shoulder	Hip	Knee	Ankle	All
Mean pose	31.1	18.9	8.5	11.8	10.0	40.8	33.5	21.4
Pre-train	84.2	41.6	60.5	76.9	72.8	62.6	53.7	63.1
Finetune (full)	77.5	22.2	42.9	49.8	52.5	42.6	38.6	44.2
Pre-train	86.1	24.5	64.1	86.8	88.0	82.5	64.0	69.6
Finetune (phase 1)	95.3	67.4	87.3	98.4	96.3	96.4	95.5	90.6
Finetune (phase 2)	**99.6**	88.1	95.6	**99.9**	97.3	**98.5**	98.5	96.6
Finetune (full)	99.3	**89.2**	**95.9**	99.7	**97.6**	**98.5**	**98.6**	**96.8**

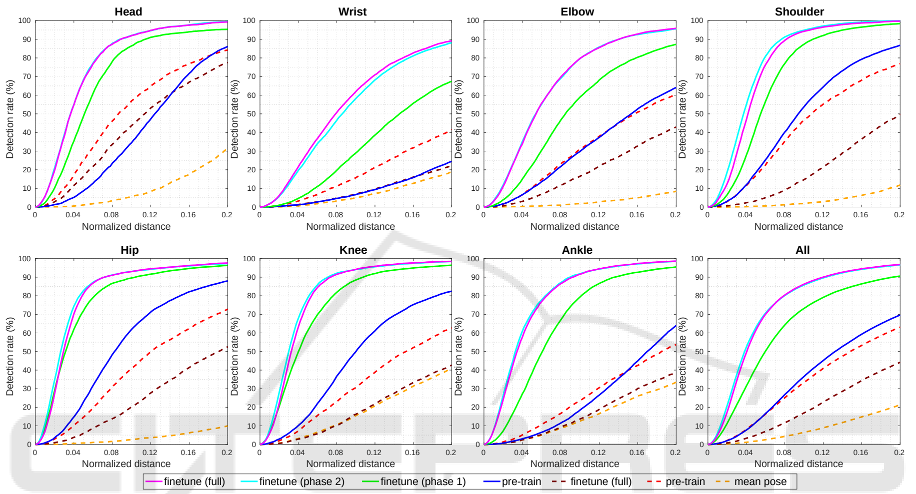

Figure 4: The results of pose estimation (PCK@0.2). The dashed lines uses pre-train validation samples (2000 images) in testing, while the solid lines use finetuning validation samples (4000 frames). To put it other way, the dashed lines represent the accuracy of generalization, while the solid lines represent the use case specific accuracy. The label indicates which network is used in testing.

As for future work, there are several things that could be considered in order to get better accuracy. One option would be to use current network as a coarse estimator and use another network for refining the pose estimation. In addition, as our method is targeted for video inputs, the utilization of the spatiotemporal data would most likely give accuracy boost. The network forward time of the person detector is relatively slow compared to the pose estimation network (16ms vs. 60ms/200ms). While the person detector works well with diverse input data, perhaps, with most pose estimation use cases, that is not necessary. By using more restricted and possibly faster person detector, a good enough performance in more constrained environments could be most likely achieved. Also, with ConvNets, generally, holds that if more data used in training, the better performance gained. Hence, the use of more advanced data augmentation methods, such as (Pishchulin et al., 2012), especially in the finetuning, would most probably lead to better accuracy. Advanced data augmentation could, for example, change colors of the clothes, adjust limb poses and change backgrounds.

REFERENCES

Andriluka, M., Pishchulin, L., Gehler, P., and Schiele, B. (2014). 2d human pose estimation: New benchmark and state of the art analysis. In *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*.

Andriluka, M., Roth, S., and Schiele, B. (2009). Pictorial structures revisited: People detection and articulated pose estimation. In *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 1014–1021.

Buehler, P., Everingham, M., Huttenlocher, D. P., and Zisserman, A. (2011). Upper body detection and tracking in extended signing sequences. *International Journal of Computer Vision (IJCV)*, 95(2):180–197.

Carreira, J., Agrawal, P., Fragkiadaki, K., and Malik, J. (2015). Human pose estimation with iterative error feedback. *arXiv preprint arXiv:1507.06550*.

Charles, J., Pfister, T., Everingham, M., and Zisserman, A. (2013). Automatic and efficient human pose estimation for sign language videos. *International Journal of Computer Vision (IJCV)*.

Charles, J., Pfister, T., Magee, D., Hogg, D., and Zisserman, A. (2016). Personalizing video pose estimation. In *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*.

Dantone, M., Gall, J., Leistner, C., and van Gool, L. (2013). Human pose estimation using body parts dependent joint regressors. In *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 3041–3048, Portland, OR, USA.

Felzenszwalb, P., McAllester, D., and Ramanan, D. (2008). A discriminatively trained, multiscale, deformable part model. In *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 1–8.

Glorot, X. and Bengio, Y. (2010). Understanding the difficulty of training deep feedforward neural networks. In *International Conference on Artificial Intelligence and Statistics (AISTATS'10)*.

Insafutdinov, E., Pishchulin, L., Andres, B., Andriluka, M., and Schiele, B. (2016). Deepercut: A deeper, stronger, and faster multi-person pose estimation model. *arXiv preprint arXiv:1605.03170*.

Jain, A., Tompson, J., Andriluka, M., Taylor, G. W., and Bregler, C. (2013). Learning human pose estimation features with convolutional networks. *arXiv preprint arXiv:1312.7302*.

Jain, A., Tompson, J., LeCun, Y., and Bregler, C. (2014). Modeep: A deep learning framework using motion features for human pose estimation. In *Asian Conference on Computer Vision (ACCV)*, pages 302–315.

Jia, Y., Shelhamer, E., Donahue, J., Karayev, S., Long, J., Girshick, R., Guadarrama, S., and Darrell, T. (2014). Caffe: Convolutional architecture for fast feature embedding. *arXiv preprint arXiv:1408.5093*.

Johnson, S. and Everingham, M. (2010). Clustered pose and nonlinear appearance models for human pose estimation. In *British Machine Vision Conference (BMVC)*. doi:10.5244/C.24.12.

Krizhevsky, A., Sutskever, I., and Hinton, G. E. (2012). Imagenet classification with deep convolutional neural networks. In *Advances in Neural Information Processing Systems (NIPS)*, pages 1097–1105.

Lifshitz, I., Fetaya, E., and Ullman, S. (2016). Human pose estimation using deep consensus voting. *arXiv preprint arXiv:1603.08212*.

Newell, A., Yang, K., and Deng, J. (2016). Stacked hourglass networks for human pose estimation. *arXiv preprint arXiv:1603.06937*.

Pfister, T., Charles, J., and Zisserman, A. (2015). Flowing convnets for human pose estimation in videos. In *International Conference on Computer Vision (ICCV)*.

Pfister, T., Simonyan, K., Charles, J., and Zisserman, A. (2014). Deep convolutional neural networks for efficient pose estimation in gesture videos. In *Asian Conference on Computer Vision (ACCV)*.

Pishchulin, L., Insafutdinov, E., Tang, S., Andres, B., Andriluka, M., Gehler, P., and Schiele, B. (2015). Deepcut: Joint subset partition and labeling for multi person pose estimation. *arXiv preprint arXiv:1511.06645*.

Pishchulin, L., Jain, A., Andriluka, M., Thormählen, T., and Schiele, B. (2012). Articulated people detection and pose estimation: Reshaping the future. In *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 3178–3185.

Ren, S., He, K., Girshick, R., and Sun, J. (2015). Faster r-cnn: Towards real-time object detection with region proposal networks. *arXiv preprint arXiv:1506.01497*.

Sapp, B. and Taskar, B. (2013). Modec: Multimodal decomposable models for human pose estimation. In *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 3674–3681.

Shotton, J., Sharp, T., Kipman, A., Fitzgibbon, A., Finocchio, M., Blake, A., Cook, M., and Moore, R. (2013). Real-time human pose recognition in parts from single depth images. *Communications of the ACM*, 56(1):116–124.

Tompson, J., Goroshin, R., Jain, A., LeCun, Y., and Bregler, C. (2015). Efficient object localization using convolutional networks. In *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 648–656.

Toshev, A. and Szegedy, C. (2014). Deeppose: Human pose estimation via deep neural networks. In *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*.

Wei, S.-E., Ramakrishna, V., Kanade, T., and Sheikh, Y. (2016). Convolutional pose machines. *arXiv preprint arXiv:1602.00134*.

Yang, Y. and Ramanan, D. (2011). Articulated pose estimation with flexible mixtures-of-parts. In *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 1385–1392.