

# Learning Human Poses from Actions

Aditya Arun<sup>1</sup>  
aditya.arun@research.iiit.ac.in  
C.V. Jawahar<sup>1</sup>  
jawahar@iiit.ac.in  
M. Pawan Kumar<sup>2</sup>  
pawan@robots.ox.ac.uk

<sup>1</sup> IIIT Hyderabad  
<sup>2</sup> University of Oxford &  
The Alan Turing Institute

---

## Abstract

We consider the task of learning to estimate human pose in still images. In order to avoid the high cost of full supervision, we propose to use a diverse data set, which consists of two types of annotations: (i) a small number of images are labeled using the expensive ground-truth pose; and (ii) other images are labeled using the inexpensive action label. As action information helps narrow down the pose of a human, we argue that this approach can help reduce the cost of training without significantly affecting the accuracy. To demonstrate this we design a probabilistic framework that employs two distributions: (i) a conditional distribution to model the uncertainty over the human pose given the image and the action; and (ii) a prediction distribution, which provides the pose of an image without using any action information. We jointly estimate the parameters of the two aforementioned distributions by minimizing their dissimilarity coefficient, as measured by a task-specific loss function. During both training and testing, we only require an efficient sampling strategy for both the aforementioned distributions. This allows us to use deep probabilistic networks that are capable of providing accurate pose estimates for previously unseen images. Using the MPII data set, we show that our approach outperforms baseline methods that either do not use the diverse annotations or rely on pointwise estimates of the pose.

## 1 Introduction

Current methods for learning human pose estimation from still images require the collection of a fully annotated data set, where each training sample consists of an image of a person, together with its ground-truth joint locations. The collection of such detailed annotations is onerous and expensive, which makes this approach unscalable. We propose to alleviate the deficiency of fully supervised learning by using a diverse data set. Part of the images of the data set are labeled with expensive pose annotations, while the remaining images are labeled with inexpensive action annotations.

Throughout the paper, we assume that the distribution of the images labeled with different types of annotations is the same (which is a necessary assumption for learning) and that the annotations themselves are noise-free. Under these assumptions, we argue that action information can be used to learn pose estimation. Note that earlier works have exploited the relationship between action and pose for action recognition. However, our problem is

significantly more challenging due to the high uncertainty in pose given the action. In order to model this uncertainty, we propose to use a probabilistic learning formulation. A typical probabilistic formulation would learn a joint distribution of the pose and the action given an image. In order to make a prediction on a test sample, where action information is not known, it would marginalize over all possible actions. In other words, it would use one set of parameters for two distinct tasks: (i) model the uncertainty in the pose for every action; and (ii) predict the pose given an image.

As our goal is to make an accurate pose prediction, we argue that such an approach would waste the modeling capability of a distribution in representing pose uncertainty in the presence of action information. In other words, the parameters of the distribution will be tuned to perform well in the presence of action information, which will not be available during testing. Instead, we use two different distributions for the two different tasks: (i) a *conditional distribution* of the pose given the image and the action; and (ii) a *prediction distribution* of the pose given the image.

We jointly estimate the parameters of the two distributions by minimizing their dissimilarity coefficient [29], which measures the distance between two distributions using a task-specific loss function. By transferring the information from the conditional distribution to the prediction distribution, we learn to estimate the pose of a human using a diverse data set. Figure 1 shows the necessity of using a probabilistic model. Specifically, the figure shows the average entropy of each joint as predicted by our model on test images. We observe that the most articulate joints like wrists and ankles have highest entropy, which a non probabilistic network does not explicitly model.

While our approach can be used in conjunction with any parametric family of distributions, in this work we focus on the state of the art deep probabilistic networks. Specifically, we model both the conditional and the prediction distributions using a DISCO Net [6], which allows us to efficiently sample from the two distributions. As will be seen later, the ability to sample efficiently is sufficient to make both training and testing computationally feasible.

We demonstrate the efficacy of our approach using the publicly available MPII Human Pose data set [2]. We discard the pose information of a portion of the training samples but retain the action information for all the samples in order to generate a diverse data set. We provide a thorough comparison of our probabilistic approach with two natural baselines. First, a fully supervised approach, which discards the weakly supervised samples that have been labeled using only the action information. Second, a pointwise model that uses a self-paced learning [24] strategy by first learning from easy samples and then gradually increase the difficulty of the training samples. We show that, by explicitly modeling the uncertainty on the pose of diverse supervised samples, our approach significantly outperforms both the baselines under various experimental settings.

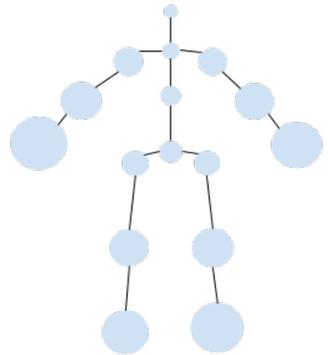


Figure 1: *Average entropy of joints in test images over a stick figure. The radius of circle around a joint is proportional to the joint's entropy.*

## 2 Related Work

With the introduction of “DeepPose” by Toshev *et al.* [35], research on human pose estimation began to shift from classic approaches based on pictorial structures [1, 2, 11, 12, 16, 25, 28, 30, 40] to deep networks. Subsequent methods include [64], which simultaneously captures features at a variety of scales using heatmaps, and [38], which employs a hierarchical model to capture the relationships between joints. A popular approach by Newell *et al.* [19] uses conv-deconv architecture and residual model to efficiently generate the heatmap without the need for any hierarchical processing. This approach has been further extended by using visual attention [8] and feature pyramid [40]. However, these methods rely on the network capacity to capture the highly articulated human pose and to handle occlusion, without modeling the uncertainty in pose explicitly.

Modeling the uncertainty over the human pose becomes crucial in a diverse data setting, where some of the training samples only provide action information. While pose has often been used to predict action [17, 52, 56, 57], the use of action for pose estimation has largely been explored for either 3D human pose [42], or for videos where there is temporal information available [11, 27, 39, 44]. To the best of our knowledge, our work is the first to exploit action information for 2D pose estimation in still images.

While the specific problem of pose estimation using action information has not been the subject of much attention, the general problem of diverse data learning has a rich history in machine learning and computer vision. Most of the traditional approaches relied on the use of simple parametric structured models such as conditional random fields, or structured support vector machines [6, 15, 18, 23, 31, 43]. These methods framed the task of predicting the missing information as estimating latent variables, and employed either the maximum likelihood or the max-margin formulation to efficiently estimate the parameters of the corresponding models. However, as the traditional structured prediction models have now been replaced by deep learning, the aforementioned formulations would need to be adapted for parameter estimation of neural networks. Indeed, our work can be viewed as a natural generalization of [15] for deep probabilistic models that admit efficient sampling mechanisms.

The deep learning community also realizes the importance of using diverse data sets to scale-up the data hungry neural network based approaches. This has led to recent research in deep multiple instance learning [9, 21, 24], as well as expectation-maximization based methods [20, 22]. However, most of the deep diverse data learning approaches have been designed to work for a specific task, such as semantic segmentation [13, 33]. It is not clear how the proposed methods can be adapted to learn human poses from action labels. In contrast, our general formulation (presented in the next section) can be easily adapted to any task by simply specifying a task-specific loss function. While we are primarily interested in pose estimation, our formulation may be of interest to the broader audience working on diverse data deep learning.

## 3 Problem Formulation

Our approach uses the recently proposed deep probabilistic network, DISCO Net [6]. The DISCO Net framework allows us to adapt a pointwise network (that is, a network that provides a single pointwise prediction) to a probabilistic one by introducing a noise filter in the pointwise network.

As a concrete example, consider the modified stacked hourglass network in Figure 2,

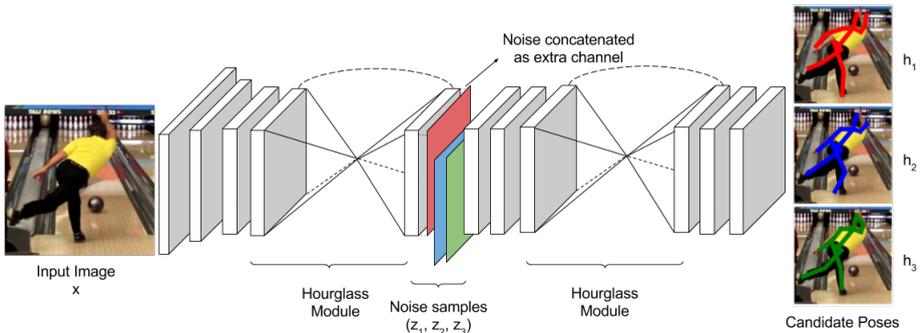


Figure 2: For a single input image  $\mathbf{x}$  and three different noise samples  $\{\mathbf{z}_1, \mathbf{z}_2, \mathbf{z}_3\}$  (represented as red, green, blue matrix respectively), DISCO Nets produces three different candidate poses  $\{\mathbf{h}_1, \mathbf{h}_2, \mathbf{h}_3\}$ . Here each block is a residual layer and two hourglass shaped blocks represent the hourglass module proposed by Newell et al. [19]. Best viewed in color.

which can be used for human pose estimation. The colored filters in the middle of the network represent the noise that is sampled from a uniform distribution. Each value of the noise filter results in a different pose estimate for the same image, thereby enabling us to generate samples from the underlying distribution encoded by the network parameters. Note that obtaining a single sample is as efficient as a forward pass through the network. By placing the filters sufficiently far away from the output layer of the network, we can learn a highly non-linear mapping from the uniform distribution (used to generate the noise filter) to the output distribution (used to generate the pose estimates).

In [9], the parameters of a DISCO Net were learned by minimizing the dissimilarity of the network distribution and the true distribution (as specified by fully supervised training samples). However, we show how the DISCO Net framework can be extended to enable diverse data learning.

### 3.1 Model

Due to the uncertainty inherent in the task of pose estimation (occlusion of joints, articulation of human body) as well as the uncertainty introduced by the use of a diverse data set during training, we advocate the use of a probabilistic formulation. To this end, we define two distributions. The first is the prediction distribution that models the probability of a pose  $\mathbf{h}$  given an image  $\mathbf{x}$ . As the name suggests, this distribution is used to make a prediction during test time. In this work, we model the prediction distribution  $\text{Pr}_{\mathbf{w}}(\mathbf{h}|\mathbf{x})$  as a DISCO Net, where  $\mathbf{w}$  are the parameters of the network.

In addition to the prediction distribution, we also model a conditional distribution of the pose given the image and the action label. As the conditional distribution contains additional information, it can be expected to provide better pose estimates. We will use this property during training to learn an accurate prediction distribution using the conditional distribution. As will be seen shortly, the conditional distribution will not be used during testing. Similar to the prediction distribution, the conditional distribution  $\text{Pr}_{\boldsymbol{\theta}}(\mathbf{h}|\mathbf{x}, \mathbf{a})$  is modeled using a DISCO Net, with parameters  $\boldsymbol{\theta}$ . Note that, while we do not have access to the partition function of the two aforementioned distributions, the use of a DISCO Net ensures that we can efficiently sample from them. This property will be exploited to make both the testing and the training computationally feasible.

## 3.2 Prediction

Throughout the rest of the paper, we will assume a task-specific loss function  $\Delta(\cdot, \cdot)$  that measures the difference between two putative poses of an image. Given an image  $\mathbf{x}$  containing a human, we would like to estimate the pose  $\mathbf{h}$  of the human such that it minimizes the risk of prediction (as measured by the loss function  $\Delta$ ). Since the ground-truth pose is unknown, we use the principle of maximum expected utility (MEU) [26]. The MEU criterion minimizes the expected loss using a set of samples  $\mathcal{H} = \{\mathbf{h}^k, k = 1, \dots, K\}$  obtained from the distribution  $\Pr_{\mathbf{w}}(\mathbf{h}|\mathbf{x})$ .

Formally, given an image  $\mathbf{x}$ , we provide a pointwise prediction of the pose in two steps. First, we estimate  $K$  pose samples using  $K$  different noise filters, each of which is sampled from a uniform distribution. Second, we use the MEU criterion to obtain the prediction as,

$$\mathbf{h}_{\Delta}^*(\mathbf{x}; \mathbf{w}) = \arg \min_{k \in [1, K]} \sum_{k'=1}^K \Delta(\mathbf{h}^k, \mathbf{h}^{k'}). \quad (1)$$

As can be seen, the above criterion can be easily applied for any loss function. For human pose estimation, we adopt the commonly used loss function that measures the mean squared error between the belief maps of two poses over all the joints [19, 62, 68]. The belief map  $b_{\mathbf{h}}(j)$  of a joint  $j$  is created by defining a 2D Gaussian whose mean is at the estimated location of the joint, and whose standard deviation is a fixed constant.

## 3.3 Diverse Data Set

In order to learn the parameters  $\mathbf{w}$  of the prediction distribution, we require a training data set. Current methods rely on a fully supervised setting, where each training sample is labeled with its ground-truth pose. In order to avoid the cost of such detailed annotations, we advocate the collection of a diverse data set, with a small number of fully supervised samples and a large number of weakly supervised samples. The presence of fully supervised samples helps disambiguate the problem of pose estimation from the problem of action classification.

Formally, we denote our training data set as  $\mathcal{D} = \{\mathcal{W}, \mathcal{S}\}$ , where  $\mathcal{W} = \{(\mathbf{x}_i, \mathbf{a}_i), i = 1 \dots n\}$  is the weakly annotated data set, and  $\mathcal{S} = \{(\mathbf{x}_j, \mathbf{a}_j, \mathbf{h}_j), j = 1 \dots m\}$  is the strongly annotated data set and  $m < n$ . Here  $\mathbf{x}_i$  refers to the  $i$ -th training image and  $\mathbf{a}_i$  denotes its action. We denote the underlying pose of the image  $\mathbf{x}_i$  as the latent variable  $\mathbf{h}_i$ . Note that we do not assume a single underlying pose. Instead, we model the distribution over all putative poses given the image and the action.

## 3.4 Learning Objective

Given the diverse data set  $\mathcal{D}$ , our goal is to learn the parameters  $\mathbf{w}$  such that it provides an accurate pose estimate  $\mathbf{h}_{\Delta}^*(\mathbf{x}; \mathbf{w})$  (specified in equation (1)) for a test image  $\mathbf{x}$ . A typical learning objective for this purpose would estimate the joint distribution  $\Pr_{\mathbf{w}}(\mathbf{h}, \mathbf{a}|\mathbf{x})$  using expectation-maximization or its variants [4]. Given an image  $\mathbf{x}$ , the pose would then be obtained by marginalizing over all actions  $\mathbf{a}$ . However, we argue that this approach needlessly places the burden of accurately representing the uncertainty of the pose and the action of an image on a single distribution. Since the action information would not be provided during testing, such an approach may fail to fully utilize the modeling capacity of the distribution parameters to obtain the best pose.

Inspired by the work of Kumar *et al.* [15], we design a joint learning objective that minimizes the dissimilarity coefficient between the prediction distribution and the conditional distribution. Briefly, the dissimilarity coefficient between two distributions  $\Pr_1(\cdot)$  and  $\Pr_2(\cdot)$  is determined by measuring their diversities. The diversity coefficient of a distribution  $\Pr_1(\cdot)$  and a distribution  $\Pr_2(\cdot)$  is defined as the expected difference between their samples, where the difference is measured by any task-specific loss function  $\Delta'(\cdot, \cdot)$ . Formally, we define the diversity coefficient as,

$$\text{DIV}_{\Delta'}(\Pr_1, \Pr_2) = \sum_{\mathbf{y}_1, \mathbf{y}_2 \in \mathcal{Y}} \Delta'(\mathbf{y}_1, \mathbf{y}_2) \Pr_1(\mathbf{y}_1) \Pr_2(\mathbf{y}_2). \quad (2)$$

where  $\mathcal{Y}$  is the space over which the distributions are defined. Using the definition of diversity, the dissimilarity coefficient of  $\Pr_1$  and  $\Pr_2$  is given by,

$$\text{DISC}_{\Delta'}(\Pr_1, \Pr_2) = \text{DIV}_{\Delta'}(\Pr_1, \Pr_2) - \gamma \text{DIV}_{\Delta'}(\Pr_1, \Pr_1) - (1 - \gamma) \text{DIV}_{\Delta'}(\Pr_2, \Pr_2). \quad (3)$$

In other words, the dissimilarity between  $\Pr_1$  and  $\Pr_2$  is the difference between the diversity of  $\Pr_1$  and  $\Pr_2$  and an affine combination of their self-diversities. In our experiments, we use  $\gamma = 0.5$ , which results in a symmetric dissimilarity coefficient between two distributions.

Given the above definition, we can now specify our learning objective as,

$$\arg \min_{\mathbf{w}, \boldsymbol{\theta}} \sum_{i=1}^n \text{DISC}_{\Delta}(\Pr_{\mathbf{w}}(\cdot | \mathbf{x}_i), \Pr_{\boldsymbol{\theta}}(\cdot | \mathbf{x}_i, \mathbf{a}_i)). \quad (4)$$

In other words, our learning objective encourages the prediction distribution and the conditional distribution to agree with each other (that is, have a small dissimilarity coefficient) for all training samples. Intuitively, the conditional distribution  $\Pr_{\boldsymbol{\theta}}(\cdot | \mathbf{x}, \mathbf{a})$  would be able to significantly narrow down the set of probable poses for a given image using the action information. By minimizing the dissimilarity between the prediction distribution and the conditional distribution, our learning objective will encourage the prediction to assign a high probability to the set of poses that are compatible with the given action. During testing, only the prediction distribution will be used to obtain the pose of a given image.

Computationally, the main challenge of employing the learning objective (4) is that its value can only be determined by estimating the loss function over all possible pairs of poses. However, the key observation that enables its use in practice is that we can obtain an unbiased estimate of its value, as well as its gradient, by sampling from the distributions  $\Pr_{\mathbf{w}}$  and  $\Pr_{\boldsymbol{\theta}}$ . In other words, given samples  $\{\mathbf{h}_k, k = 1, \dots, K\}$  from the prediction distribution, and samples  $\{\mathbf{h}'_k, k = 1, \dots, K\}$  from the conditional distribution, the unbiased estimated value of the learning objective (4) can be computed as,

$$\frac{1}{K^2} \left( \sum_{k, k'} \Delta(\mathbf{h}_k, \mathbf{h}'_{k'}) - \gamma \sum_{k, k'} \Delta(\mathbf{h}_k, \mathbf{h}_{k'}) - (1 - \gamma) \sum_{k, k'} \Delta(\mathbf{h}'_k, \mathbf{h}'_{k'}) \right). \quad (5)$$

### 3.5 Optimization

As a DISCO Net provides an efficient sampling mechanism, it is ideally suited to stochastic gradient descent. In order to make the most use of the diverse nature of the data set, as well as the learning objective, we estimate the parameters of the two networks in three stages. First, we use supervised training for the two networks using the small amount of the ground truth

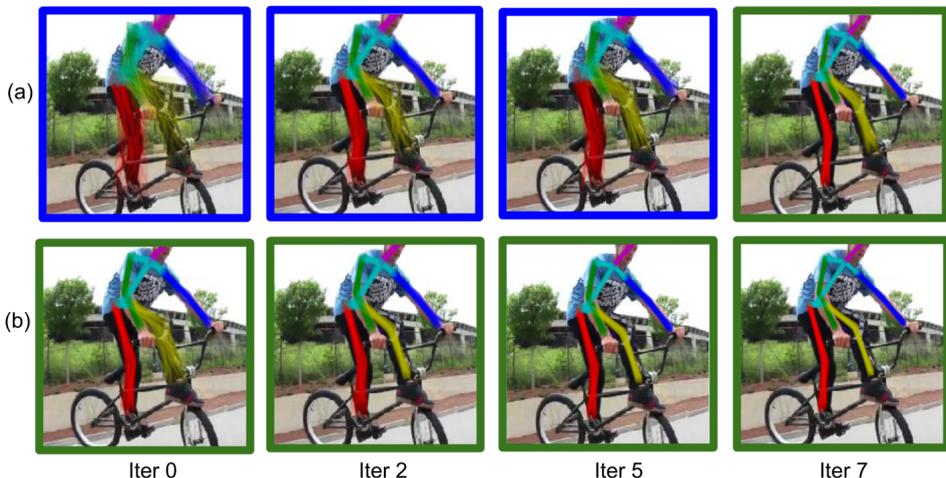


Figure 3: Example of superimposed pose predictions by DISCO Nets illustrating the uncertainty in the pose across training iterations. The blue box around the images represent a high diversity coefficient value, and the green box around them represents low diversity coefficient value. Row (a) represents outputs from the prediction network and row (b) represents outputs from the conditional network. The first column shows the initial prediction of the networks; columns 2 through 4 shows prediction of the networks at second, fifth and final iteration respectively. The images show a common action of riding a bike where the conditional network performs well from the beginning of the optimization procedure and transfers its knowledge to the prediction network. Best viewed in color.

pose data. Second, we perform iterative training of the two networks, that is, we update one network while keeping the other fixed. Third, we jointly optimize both the networks together. At each stage, we use stochastic gradient descent in a similar manner to [9]. Joint training of the two networks is expensive in terms of memory and time. However, by first training the two networks using strong supervision and then using iterative optimization strategy, we significantly reduce the number of iterations required in the third stage of the optimization. We provide further details in the supplementary.

The prediction of the two networks during the iterative training stage is visualized in figure 3. For a commonly occurring action of riding a bike, we depict the hundred different pose estimates from the prediction and the conditional network by superimposing them. Hence, if all the pose estimates agree with each other, the lines depicting the samples will be thin and opaque. In order to represent the low uncertainty in the pose estimates of this image, we will draw a green bounding box around the image. In contrast, if the pose estimates vary significantly from each other, then the lines depicting the samples will be spread out and less opaque. In order to represent the high uncertainty in the pose estimates of this image, we will draw a blue bounding box around the image.

Here, we observe that initially  $\text{Pr}_w$  has high uncertainty for the predicted pose, but  $\text{Pr}_\theta$  is confident about the predictions. However, after several iterations of the optimization algorithm, the information present in the conditional network is successfully transferred over to the prediction network. This is shown in the last column, where both the networks start to agree with each other (that is, have a low self diversity coefficient). For difficult images where both prediction and conditional distribution are highly uncertain at the beginning,

the networks learns from other easy examples that may be present in the data set. Further visualizations of the learning process are provided in the supplementary material.

## 4 Experiments

**Data set.** We use the MPII Human pose data set [14], which consists of 17.4k images with publicly available action and ground-truth pose annotations. We split the images into  $\{70, 15, 15\}\%$  training, validation and test sets, which corresponds to 12,156 images in the training set and 2605 images each in the testing and the validation set. In order to obtain a diverse data set, we discard the pose information for a random subset of training examples, while retaining action labels for all samples. This results in (i) a fully annotated training set, which contains both the ground truth pose annotations and the action labels; and (ii) a weakly annotated training set, which only contains action labels.

To obtain the tasks of varying levels of difficulty, we choose three different data splits,  $\{25 - 75, 50 - 50, 75 - 25\}\%$ , where we randomly discard 75%, 50%, and 25% of the pose annotations from the training images respectively. We note here that for each split, we augment our training set by rotating the images with an angle ( $+/- 30^\circ$ ) and by horizontal flipping the original image.

**Implementation Details.** In order to implement our probabilistic DISCO network, shown in figure 2, we adopt the popular stacked hourglass network [19] for human pose estimation, which stacks 8 hourglass modules. For the prediction network, a noise filter of size  $64 \times 64$  is added to the output of the penultimate hourglass module, which itself consists of 256  $64 \times 64$  filters. The 257 channels are convolved with a  $1 \times 1$  filter to bring the number of channels back to 256. This is followed by a final hourglass module as shown in figure 2 (closely following the approach of [19]). As noise is treated as input, all parameters of the network remain differentiable and hence can be trained via backpropagation. Our conditional network is modeled exactly as the prediction network, except that there are  $\mathbf{a}$  different output branches, one for each possible action class, stacked on top of penultimate hourglass module. Each output branch has its own noise filter followed by the final hourglass module as described before. We present additional implementation details in the supplementary.

Notice that when drawing  $K$  samples from this modified stacked hourglass architecture for the same input image, we can reuse the output of the penultimate layer of the 8-stacked hourglass net. We only need to recompute the final hourglass module  $K$  times to generate  $K$  samples which greatly reduces our runtime complexity. In practice, a single forward pass to draw  $K = 100$  samples from our probabilistic net takes 114 ms. compared to 68 ms. for the vanilla stacked hourglass network on NVIDIA GTX 1080Ti GPU.

Initialization of the prediction network is done by training it using a small number of fully annotated training data, while for the conditional networks, we initialize them by fine tuning the prediction network weights with the small number of action specific fully annotated training data. We then optimize our two set of networks by first, iterative optimization procedure, and then through joint optimization, as described in the previous section.

**Methods.** We compare our proposed probabilistic method, learned with diverse data, with two baselines: (i) a fully supervised human pose estimation network, the stacked hourglass network [19], which we refer to as FS Net; and (ii) a non-probabilistic pointwise network trained with diverse data, which uses the same architecture as shown in figure 2 but provides

Method	Split	Head	Sho.	Elb.	Wri.	Hip	Knee	Ank.	Total
Supervised Subset	100%	98.16	96.22	91.23	87.08	90.11	87.39	83.55	90.92
FS	25%	59.17	46.98	30.00	21.33	36.32	20.05	23.93	37.54
	50%	90.18	80.60	64.29	52.43	67.44	55.41	51.30	67.88
	75%	94.61	90.56	81.28	74.15	81.86	73.20	67.19	80.88
PW	25-75	73.77	55.69	37.21	25.32	43.24	28.01	30.82	45.16
	50-50	92.97	83.56	71.08	59.18	72.56	60.49	57.27	73.11
	75-25	95.46	93.50	86.47	81.05	85.58	80.98	76.81	85.89
Pr <sub>w</sub> (iterative)	25-75	78.21	60.98	42.01	28.75	42.37	29.07	33.54	48.12
	50-50	93.42	86.91	75.03	66.56	77.22	67.38	60.96	76.43
	75-25	96.28	94.53	88.36	83.31	87.54	82.45	79.48	88.16
Pr <sub>w</sub> (joint)	25-75	79.54	62.87	43.38	29.38	43.38	30.91	34.86	<b>49.41</b>
	50-50	94.07	88.32	75.93	67.53	78.20	67.80	61.49	<b>78.01</b>
	75-25	97.45	95.87	90.21	86.09	89.42	86.26	82.92	<b>90.21</b>

Table 1: Results on MPII Human Pose (PCKh@0.5), where FS is trained on varying percentages of fully annotated data and PW and Pr<sub>w</sub> are trained on varying splits of fully annotated and weakly annotated training data. Here FS and PW are the fully supervised and the pointwise networks respectively, and Pr<sub>w</sub> (iterative) and Pr<sub>w</sub> (joint) is our proposed probabilistic network trained with iterative optimization and joint optimization respectively. The supervised subset is the fully supervised stacked hourglass net [14] trained with all the available labels and defines the upper bound on the total accuracy that can be achieved through this architecture.

a single prediction. We refer this pointwise network as PW Net. The first baseline helps us to compare the performance of a fully supervised network with a network trained on the diverse collection of data, and the second baseline demonstrates the benefit of our probabilistic network when compared to a non probabilistic pointwise network.

We train FS net on the fully annotated data set using stochastic gradient descent, as discussed in [19]. The PW net is trained using diverse data, making use of the action annotations. We provide the detailed training setup of the FS and the PW net in the supplementary.

**Results.** We evaluate the three trained networks, FS, PW and Pr<sub>w</sub>, by computing their accuracy on the held out test set. We use the normalized “Probability of Correct Keypoint” (PCKh) metric [60] to report our results. Table 1 shows the performance of the three networks when trained on varying splits of the training set.

Here, we observe that, for all the data splits, our proposed probabilistic network Prob<sub>w</sub> outperforms the other baseline networks FS and PW. This superior performance is seen consistently across predictions of all joints as well as on the overall pose prediction.

Performance of the three networks, FS, PW and Pr<sub>w</sub>, increases with the increase in level of supervision. In the more challenging 25 – 75 split, there are far fewer fully supervised examples present for each action category which results in PW and Pr<sub>w</sub> to learn a poor initial estimate of action specific pose from diverse data. This leads to overall poor performance when compared to 50 – 50 or 75 – 25 split case, where we have more supervised data.

Moreover, both the methods trained using diverse data, PW and Pr<sub>w</sub>, show a significant gain in accuracies when compared to the fully supervised network, FS. This empirically shows us that the action information present in the weakly annotated set is helpful for predicting pose.

As our proposed probabilistic network Pr<sub>w</sub> performs better than the pointwise network PW, we see the significance of modeling uncertainty over pose. Though the proposed proba-

bilistic network only marginally improves the prediction for joints with low uncertainty, like the head, shoulder and hips, the difference in the accuracies of the two networks is due to better performance of the probabilistic network  $\text{Pr}_w$  on difficult joints like wrists, elbows, knees and ankles. We see that the  $\text{Pr}_w$  network provides a significant improvement of up to 5% improvement in accuracies over the PW Net on joints with high uncertainty (wrists, elbows, ankles and knees).

Joint training of the two set of networks improves our prediction by around 1.5%. We also note that, while the supervised subset, which is the fully supervised stacked hourglass network [19] trained using all available labels in the training set, achieves 90.9% [19], our probabilistic network provides comparable results when trained only on 75% pose annotations and 25% action annotations. Note that the supervised subset defines the upper bound on accuracy that can be achieved through this architecture.

We argue that the relative position of joints like head, shoulder and hip remains largely in similar spatial location with respect to each other across various actions and therefore have low entropy, whereas, joints like wrists, elbows, knees and ankles not only show huge variations in their relative spatial location across various action categories but also within same action category, resulting in large entropy. Therefore, even though pointwise network PW does a good job of estimating pose locations for joints with low uncertainty, it fails to capture the high inter-class and intra-class variability of joints with high uncertainty. On the other hand,  $\text{Pr}_w$  explicitly models uncertainty over joint locations as can be seen in figure 1.

Our method was implemented using PyTorch library<sup>1</sup>. Further details of the experimental setup, full PcKh curves, and results for additional experiments using a different architecture [9], demonstrating the generality of our method, are included in the supplementary material.

## 5 Discussion

We presented a novel framework to learn human pose using diverse data set. Our framework uses two separate distributions: (i) a conditional distribution for modeling uncertainty over pose given the image and the action during training time; and (ii) a prediction distribution to provide pose estimates for a given image. We model the two aforementioned distributions using a deep probabilistic network. We learn these separate yet complimentary distributions by minimizing a dissimilarity coefficient based learning objective. Empirically, we show that: (i) action serves as an important cue for predicting human pose; and (ii) modeling uncertainty over pose is essential for its accurate prediction.

Our approach can be easily adapted to other diverse learning tasks by specifying an appropriate loss function for the evaluation of the diversity coefficient. This may be of interest to a wider machine learning and computer vision audience. We would also like to investigate the the use of active learning, so that our network benefits the most in terms of accuracy from the fully supervised annotations. The diversity of the pose samples, which can be computed efficiently in our framework, can provide a useful cue to enable active learning.

## 6 Acknowledgements

This work is partially funded by the EPSRC grants EP/P020658/1 and TU/B/000048 and a CEFIPRA grant. Aditya is supported by Visvesvaraya Ph.D. Fellowship program.

<sup>1</sup>The code and the pre-trained model is available at <http://bit.ly/poses-from-actions>

## References

- [1] Mykhaylo Andriluka, Stefan Roth, and Bernt Schiele. Pictorial structures revisited: People detection and articulated pose estimation. In *CVPR*, 2009.
- [2] Mykhaylo Andriluka, Leonid Pishchulin, Peter Gehler, and Bernt Schiele. 2d human pose estimation: New benchmark and state of the art analysis. In *CVPR*, 2014.
- [3] Vasileios Belagiannis, Christian Rupprecht, Gustavo Carneiro, and Nassir Navab. Robust optimization for deep regression. In *ICCV*, 2015.
- [4] C Bishop. Pattern recognition and machine learning (information science and statistics), 1st edn. 2006. corr. 2nd printing edn, 2007.
- [5] Diane Bouchacourt, Sebastian Nowozin, and M Pawan Kumar. Entropy-based latent structured output prediction. In *ICCV*, 2015.
- [6] Diane Bouchacourt, M Pawan Kumar, and Sebastian Nowozin. Disco nets: Dissimilarity coefficient networks. In *NIPS*, 2016.
- [7] Lubomir Bourdev and Jitendra Malik. Poselets: Body part detectors trained using 3d human pose annotations. In *ICCV*, 2009.
- [8] Xiao Chu, Wei Yang, Wanli Ouyang, Cheng Ma, Alan L. Yuille, and Xiaogang Wang. Multi-context attention for human pose estimation. In *CVPR*, 2017.
- [9] Thibaut Durand, Nicolas Thome, and Matthieu Cord. Weldon: Weakly supervised learning of deep convolutional neural networks. In *CVPR*, 2016.
- [10] Vittorio Ferrari, Manuel Marin-Jimenez, and Andrew Zisserman. Progressive search space reduction for human pose estimation. In *CVPR*, 2008.
- [11] Umar Iqbal, Martin Garbade, and Juergen Gall. Pose for action - action for pose. *CoRR*, abs/1603.04037, 2016. URL <http://arxiv.org/abs/1603.04037>.
- [12] Sam Johnson and Mark Everingham. Learning effective human pose estimation from inaccurate annotation. In *CVPR*, 2011.
- [13] Alexander Kolesnikov and Christoph H Lampert. Seed, expand and constrain: Three principles for weakly-supervised image segmentation. In *ECCV*, 2016.
- [14] M. P. Kumar, B. Packer, and D. Koller. Self-paced learning for latent variable models. In *NIPS*, 2010.
- [15] M Pawan Kumar, Ben Packer, and Daphne Koller. Modeling latent variable uncertainty for loss-based learning. In *ICML*, 2012.
- [16] Lubor Ladicky, Philip HS Torr, and Andrew Zisserman. Human pose estimation using a joint pixel-wise and part-wise formulation. In *CVPR*, 2013.
- [17] Ivan Lillo, Juan Carlos Niebles, and Alvaro Soto. A hierarchical pose-based approach to complex action understanding using dictionaries of actionlets and motion poselets. In *CVPR*, 2016.

- [18] Kevin Miller, M Pawan Kumar, Benjamin Packer, Danny Goodman, Daphne Koller, et al. Max-margin min-entropy models. In *AISTATS*, 2012.
- [19] Alejandro Newell, Kaiyu Yang, and Jia Deng. Stacked hourglass networks for human pose estimation. In *ECCV*, 2016.
- [20] George Papandreou, Liang-Chieh Chen, Kevin P Murphy, and Alan L Yuille. Weakly and semi-supervised learning of a deep convolutional network for semantic image segmentation. In *ICCV*, 2015.
- [21] Deepak Pathak, Evan Shelhamer, Jonathan Long, and Trevor Darrell. Fully convolutional multi-class multiple instance learning. In *ICLR-W*, 2014.
- [22] Deepak Pathak, Philipp Krahenbuhl, and Trevor Darrell. Constrained convolutional neural networks for weakly supervised segmentation. In *ICCV*, 2015.
- [23] Wei Ping, Qiang Liu, and Alexander T Ihler. Marginal structured svm with hidden variables. In *ICML*, 2014.
- [24] Pedro O Pinheiro and Ronan Collobert. From image-level to pixel-level labeling with convolutional networks. In *CVPR*, 2015.
- [25] Leonid Pishchulin, Mykhaylo Andriluka, Peter Gehler, and Bernt Schiele. Strong appearance and expressive spatial models for human pose estimation. In *CVPR*, 2013.
- [26] Vittal Premachandran, Daniel Tarlow, and Dhruv Batra. Empirical minimum bayes risk prediction: How to extract an extra few% performance from vision models with just three more parameters. In *CVPR*, 2014.
- [27] Kumar Raja, Ivan Laptev, Patrick Pérez, and Lionel Oisel. Joint pose estimation and action recognition in image graphs. In *ICIP*, 2011.
- [28] Deva Ramanan. Learning to parse images of articulated bodies. In *NIPS*, 2006.
- [29] C Radhakrishna Rao. Diversity and dissimilarity coefficients: a unified approach. *Theoretical population biology*, 21(1):24–43, 1982.
- [30] Ben Sapp and Ben Taskar. Modec: Multimodal decomposable models for human pose estimation. In *CVPR*, 2013.
- [31] Alexander Schwing, Tamir Hazan, Marc Pollefeys, and Raquel Urtasun. Efficient structured prediction with latent variables for general graphical models. In *ICML*, 2012.
- [32] Christian Thureau and Václav Hlaváč. Pose primitive based human action recognition in videos or still images. In *CVPR*, 2008.
- [33] Pavel Tokmakov, Karteek Alahari, and Cordelia Schmid. Weakly-supervised semantic segmentation using motion cues. In *ECCV*, 2016.
- [34] Jonathan J Tompson, Arjun Jain, Yann LeCun, and Christoph Bregler. Joint training of a convolutional network and a graphical model for human pose estimation. In *NIPS*, 2014.

- [35] Alexander Toshev and Christian Szegedy. Deeppose: Human pose estimation via deep neural networks. In *CVPR*, 2014.
- [36] Raviteja Vemulapalli and Rama Chellapa. Rolling rotations for recognizing human actions from 3d skeletal data. In *CVPR*, 2016.
- [37] Raviteja Vemulapalli, Felipe Arrate, and Rama Chellappa. Human action recognition by representing 3d skeletons as points in a lie group. In *CVPR*, 2014.
- [38] Shih-En Wei, Varun Ramakrishna, Takeo Kanade, and Yaser Sheikh. Convolutional pose machines. In *CVPR*, 2016.
- [39] Bruce Xiaohan Nie, Caiming Xiong, and Song-Chun Zhu. Joint action recognition and pose estimation from video. In *CVPR*, 2015.
- [40] Wei Yang, Shuang Li, Wanli Ouyang, Hongsheng Li, and Xiaogang Wang. Learning feature pyramids for human pose estimation. In *ICCV*, 2017.
- [41] Yi Yang and Deva Ramanan. Articulated human detection with flexible mixtures of parts. *IEEE Transactions on PAMI*, 2013.
- [42] Angela Yao, Juergen Gall, and Luc Van Gool. Coupled action recognition and pose estimation from multiple views. *IJCV*, 2012.
- [43] Chun-Nam John Yu and Thorsten Joachims. Learning structural svms with latent variables. In *ICML*, 2009.
- [44] Tsz-Ho Yu, Tae-Kyun Kim, and Roberto Cipolla. Real-time action recognition by spatiotemporal semantic and structural forests. In *BMVC*, 2010.