



# Estimating pose from pressure data for smart beds with deep image-based pose estimators

Vandad Davoodnia<sup>1</sup> · Saeed Ghorbani<sup>2</sup> · Ali Etemad<sup>1</sup>

Accepted: 4 April 2021 / Published online: 4 June 2021

© The Author(s), under exclusive licence to Springer Science+Business Media, LLC, part of Springer Nature 2021

## Abstract

In-bed pose estimation has shown value in fields such as hospital patient monitoring, sleep studies, and smart homes. In this paper, we explore different strategies for detecting body pose from highly ambiguous pressure data, with the aid of pre-existing pose estimators. We examine the performance of pre-trained pose estimators by using them either directly or by re-training them on two pressure datasets. We also explore other strategies utilizing a learnable pre-processing domain adaptation step, which transforms the vague pressure maps to a representation closer to the expected input space of common purpose pose estimation modules. Accordingly, we used a fully convolutional network with multiple scales to provide the pose-specific characteristics of the pressure maps to the pre-trained pose estimation module. Our complete analysis of different approaches shows that the combination of learnable pre-processing module along with re-training pre-existing image-based pose estimators on the pressure data is able to overcome issues such as highly vague pressure points to achieve very high pose estimation accuracy.

**Keywords** Smart beds · Pressure-sensing mattress · Pose estimation · Machine learning

## 1 Introduction

Sleep studies have recently attracted considerable attention due to the availability and popularization of sensing and processing tools for monitoring users with smart bed technologies. Such technologies play a critical role towards pervasive and unobtrusive sensing and analysis of people in smart homes as well as clinical settings, which in turn can have implications for health, quality of life, and even security. For example, it has been previously demonstrated that different sleeping poses can impact certain conditions

or disorders such as sleep apnea [26], pressure ulcers [43], and even carpal tunnel syndrome [33, 34]. As another example, in specialized units, the movements of hospitalized patients are monitored to detect critical events and to analyze parameters such as lateralization, movement range, or the occurrence of pathological patterns [11]. Moreover, patients are usually required to maintain specific poses after certain surgeries or procedures to obtain better recovery results. Therefore, long-term in-bed monitoring and automated detection of poses is of critical interest in health-care applications [28, 30].

Currently, most in-bed examinations are performed with manual visual inspections by caretakers or reports from patients themselves, which are prone to subjective prognosis and user errors. To address the underlying problems in subjective and manual inspections, automated in-bed pose estimation systems are needed in clinical and smart home settings. A number of different learning-based approaches have recently been developed to minimize manual involvement and provide more consistent and accurate results [1].

Automatic in-bed pose monitoring can be achieved by Deep Neural Networks (DNN), which provide rich information using convolutional operations for feature extraction on different modalities, such as pressure mapping

---

✉ Vandad Davoodnia  
vandad.davoodnia@queensu.ca

Saeed Ghorbani  
saeed@eecs.yorku.ca

Ali Etemad  
ali.etemad@queensu.ca

<sup>1</sup> Department of Electrical and Computer Engineering, Queen's University, Kingston, Canada

<sup>2</sup> Department of Electrical Engineering and Computer Science, York University, Toronto, Canada

sensors [13, 23] or camera-based systems [31]. Camera-based systems, suffer from a range of implementation issues when trying to address challenging situations including blanket occlusions [1], lighting variations [32], and concerns regarding privacy for in-home and clinical use. Additionally, accurate visual monitoring may require advanced sensors such as infrared [31, 44], time-of-flight [37], and depth cameras [18]. The disadvantages of pressure-based systems, on the other hand, are the high cost and need for calibration. Nonetheless, they are not subject to occlusion or point-of-view problems, complications caused by lighting variations, and privacy issues. Moreover, textile-based pressure sensors can be seamlessly embedded into mattresses to construct unobtrusive smart beds [27].

Recent studies on pressure mapping systems have generally been limited to coarse posture identification (i.e. left, right and supine) [12, 16, 35]. Moreover, the notion of body pose estimation using pressure arrays has rarely been explored [4, 14, 29]. The limited scope of existing works on pressure data is mainly due to the lack of extensive datasets that span the pose and body distributions required to learn generalized models for pose estimation.


With recent advances in machine and deep learning, a large number of advanced data-driven pose estimation methods have been developed [3, 6–8, 21, 22, 24, 39, 45] to be used in a variety of different applications such as animation, clinical monitoring, human-computer interaction, and robotics [20, 25, 40]. While these models have the potential to be used for pressure-based in-bed pose estimation, they have all been designed and trained for *images* with naturally appearing human figures, and mostly in upright postures. Moreover, *weak pressure areas* resulting from supported body parts are perceived as body occlusions in the pressure maps, which can cause low fidelity with existing models.

In this paper, we explore the use of pre-existing *image-based pose estimators*, namely OpenPose [3] and Cascaded Pyramid Network (CPN) [7], for in-bed pressure-based pose estimation, with the goal of detecting keypoint locations reliably in challenging conditions such as weak pressure areas. See Fig. 1 (top row) for a few examples of the

pressure data used in this study. To this end, we exploit four general approaches: *i)* First, as a baseline, we use the pose estimators pre-trained on their respective image datasets without any modification or training with the pressure data; *ii)* Next, we manually label body joint locations for the pressure dataset and re-train the pose estimators from the previous pipeline on pressure datasets. *iii)* We propose a fully-convolutional network called PolishNetU, to learn a pre-processing step such that the polished outputs are consistent with the data on which the pose estimators have been originally trained. PolishNetU is then followed by the frozen pose estimators (not re-trained or modified for pressure data). *iv)* Finally, we re-train the entire pipeline *iii)* (consisting of PolishNetU and a pose estimator) on the labeled pressure data. Our analysis shows that method *(ii)* and *(iv)* perform superior compared to other approaches, followed by solution *(iii)*, which achieves significantly better results compared to the mere use of a pose recognition model directly on the inputs (pipeline *(i)*). In addition, we show that while most pose estimation models require large datasets for training, PolishNetU is capable of generalizing when trained on a dataset with a limited number of subjects. In summary, our contributions are as follows:

- We explore the notion of using existing pose recognition networks for in-bed pose estimation using pressure data and improve upon the previous studies.
- We propose PolishNetU, a deep UNet-style neural network for taking pressure images and transforming them into an embedding close to real human figures captured in images.
- We compare and provide comprehensive insights into several strategies for using pre-existing pose estimators namely OpenPose and CPN for pressure-based pose estimation.
- Finally, we conclude that fine-tuning pre-existing pose estimators along with using PolishNetU as a pre-processing domain adaptation step can perform pose recognition very effectively with a detection rate of over 99%.

**Fig. 1** Some examples of pressure-based in-bed pose estimation are presented. The top row shows the input pressure maps recorded using a mattress with embedded sensors, while the bottom row presents the estimated poses using an end-to-end method consisting of our proposed PolishNetU network with OpenPose



## 2 Related work

Generally, in-bed pose estimation methods can be divided into two main categories based on the input modality: camera-based and pressure-based. The former are a group of techniques that make use of different types of cameras such as infrared, range, or normal digital cameras, while the latter use a matrix of pressure sensors. While our work focuses on pressure data, we review both approaches in this section for completeness and providing a complete picture of solutions used for in-bed pose estimation.

### 2.1 Camera-based pose estimation

These methods suffer from occlusion and lighting variation conditions. As a result, one of the main focuses of these approaches is to address such problems. For example, [32] proposed a novel recording method, called Infrared Selective (IRS) image acquisition, to address the problem of lighting variations caused by the daylight cycle. Then an  $n$ -end Histogram of Oriented Gradients (HOG) feature extraction followed by a Support Vector Machine (SVM) classifier was used to align the orientation such that the rectified images were consistent with common pose detection methods. Finally, a few layers of a convolutional pose machine were fine-tuned on the in-bed dataset. Their pose rectification and estimation blocks assume no occlusion, for example by a blanket, and use an on-demand trigger to reduce the high computational cost of the pipeline. In another camera-based work, [41] developed a video-based monitoring approach to estimate human pose in conditions with occluded body parts. The proposed method comprised two main blocks: a weak human model and a modified pose matching algorithm. First, in order to reduce the search space of poses, a weak human model was used to quickly generate soft estimates of obscured upper body parts. The obscured parts were then detected using edge information in multiple stages. Next, an enhanced human pose matching algorithm was introduced to address the problem of weak image features and obstruction noise. This was used as a subsequent fine-tuned block to be optimized in the constrained space.

End-to-end deep learning methods have also been explored for in-bed camera-based pose estimation. Achilles et al. [1] trained a deep model to infer body pose from RGB-D data, while the ground truth was provided by a synchronized optical motion capture system. The model was constructed by a convolutional neural network (CNN) followed by a recurrent neural network (RNN) to capture the temporal consistency. Since it was impossible to track the markers while occluded by a blanket, the RGB-D data were

augmented with a virtual blanket to simulate the conditions where body parts were occluded.

In [5], a semi-automatic approach was proposed for upper-body pose estimation using RGB video data. The video data were normalized in a pre-processing step using contrast-limited adaptive histogram equalization, making the processed data invariant to lighting variations. Then, a CNN model was trained on the subsequent data for each subject outputting 7 heatmaps for 7 upper-body joint locations. Finally, a Kalman filter was applied as a post-processing step to refine the predicted joint trajectories and achieve a more temporally consistent estimation.

### 2.2 Pressure-based pose estimation

Pressure-based approaches have recently attracted attention as they avoid some of the problems that camera-based systems suffer from, for example, occlusion, lighting variations, and subject privacy. In [35], subject classification was performed with pressure data in three standard postures, namely supine, right side, and left side. Eighteen statistical features were extracted from the pressure distributions in each frame of each posture and fed to a dense network. Hidden layers were pre-trained by incorporating restricted Boltzmann machines into the deep belief network to find the proper initial weights.

In [17], a generative inference approach was proposed similar to [38]. However, pressure data were used as the input modality, and the body was simulated using a less sophisticated human body model. The pipeline included two main blocks. First, the patient orientation was detected and then the coarse body posture was classified using a  $k$ -nearest-neighbor classifier by comparing the query pressure distribution to the labeled training data. In the second step, a cylindrical 3D human body model was used in a generative inference approach to synthesize pressure distributions. The body model parameters (shape and pose) were iteratively optimized using Powell's method, minimizing the sum of squared distances between the synthetic pressure distribution and observed distribution.

In a more recent study, [10] proposed PressureNet, a pressure-based 3D pose and shape reconstruction network, which was trained on synthetic data and tested on real pressure images. Their method consisted of two modules, first, to encode shape, pose, and global transformation from the gender and pressure data, and second, for reconstructing the 3D model and consequently estimating the pressure images from first module's input pose information. By incorporating pose information loss for the first module and heatmap loss for the second, they were able to achieve a 3D pose recognition error of less than 75mm.

## 3 Methodology

### 3.1 Problem setup

Our goal is to explore possible approaches for in-bed pose estimation using pre-existing pose recognition networks. Our problem can be formulated as desiring a set of 14 keypoints, indicating different limb positions, by taking an input pressure image and passing it through a deep neural network. In this paper, we analyze a set of different solutions where pose estimation is achieved by utilizing off-the-shelf pose estimators. To this end, the explored networks may include a learnable pre-processing network (PolishNetU), which aims to edit the pressure images to prepare them for pose identification. In the PolishNetU model, the latent features are learned through the multi-scale architecture of the U-Net style network. On the encoder part of PolishNetU, the latent features of each step are pooled with the previous blocks to be used to generate a polished and pre-processed image in the decoder part. The polished image is then passed to a pose estimation block, which locates the position of different joints. As mentioned in Section 1 (Introduction), we consider OpenPose and CPN for pose estimation. These methods are from a family of convolutional pose machines [42], where they use a VGG and a ResNet backbone as feature extractors respectively, and their goal is to find the keypoints by estimating the probability of the existence of the limb in the image. As a result, both CPN and OpenPose generate 14 heatmaps, each corresponding to one limb position. Furthermore, OpenPose incorporates additional 28 output channels called Part Affinity Fields (PAF) corresponding to the connection of the adjacent limbs and their difference in position. See Fig. 2 for an overview of the explored approaches.

### 3.2 Solutions

We consider 4 possible solutions: (i) first, the off-the-shelf pose estimators can be used without any further training or addition of additional modules; (ii) the pose estimation networks can be re-trained on the pressure data; (iii) a domain adaptation network can be designed and trained to pre-process the vague input pressure maps to then be used with the frozen pose estimators; and (iv) the entire pipeline consisting of the domain adaptation network and the pose estimator can be re-trained end-to-end with the pressure data. Following we describe the details of each of these possible solutions.

**Frozen pose estimators** utilized directly on the pressure images. Let  $I \in \mathbb{R}^{W \times H \times 3}$  be the unpolished input pressure data. Our objective would then be to estimate a set of keypoints  $\hat{K} = Q(I; \theta_Q)$ , where  $Q$  is the approximated function by the pre-existing pose recognition networks

and  $\theta_Q$  is its pre-trained parameters. In this scenario, the pre-trained parameters  $\theta_Q$  are optimized on the original image-pose dataset that the pose estimator has been trained on.

**Re-training Pose Estimators** is the second possible solution which is built upon the previous one by adding the objective of optimizing  $\theta_Q$  on the pressure data such that  $Q$  closely estimates the keypoints  $K$ . We define the objective as:

$$\max_{\theta_P} J(\hat{K}, K), \quad (1)$$

where  $J$  is defined as a function of similarity between the predicted and ground truth keypoints. Accordingly, We define two loss functions with a heatmap term and a PAF term. The heatmap term,  $E_{heatmap}$  is defined as:

$$E_{heatmap} = \frac{1}{KHW} \sum_{k=1}^K \sum_{i=1}^H \sum_{j=1}^W V_k \|C_k - C'_k\|_2^2, \quad (2)$$

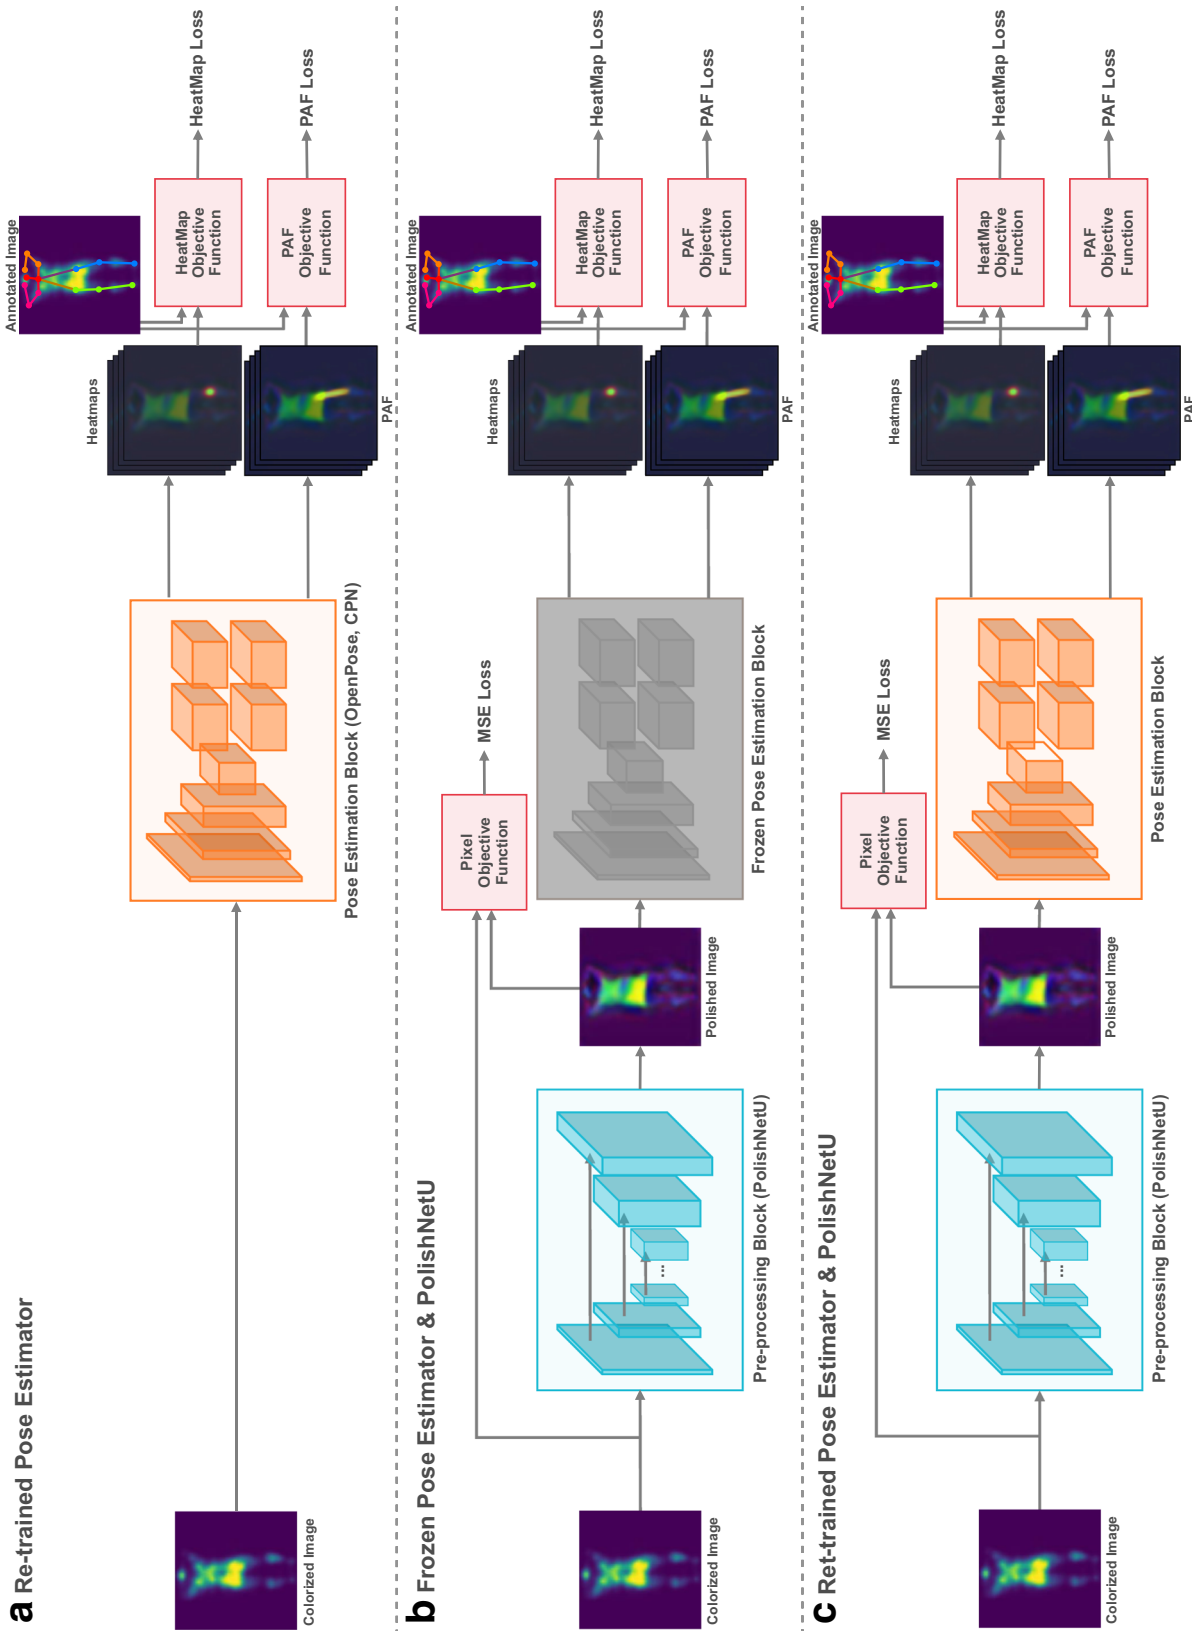
where  $C_k$  and  $C'_k$  are the corresponding ground-truth and predicted heatmaps for keypoint  $k$ ,  $K = 14$  is the number of visible keypoints, and  $V_k$  is 1 if the  $K_{th}$  limb is visible and 0 if it's not. Next, the PAF term,  $E_{PAF}$ , is defined as:

$$E_{PAF} = \frac{1}{LHW} \sum_{l=1}^L \sum_{i=1}^H \sum_{j=1}^W V_l \|F_l - F'_l\|_2^2, \quad (3)$$

where  $F_l$  and  $F'_l$  are the corresponding ground-truth and predicted PAFs for limb  $l$ ,  $L = 28$  is the number of connections between the limbs in  $y$  or  $x$  axis, and  $V_l$  is 1 if both limbs producing the  $l_{th}$  connection are visible and 0 otherwise. The final objective would then be optimized by minimizing the sum of the two loss functions.

**Image space representation learning** is the next possible solution. Here, our goal is to implement a learnable pre-processing step that receives the pressure data  $I$  as inputs and synthesizes colored images close to a pre-trained pose estimation network's learnt data. Therefore, the output data from the learner should lie on the data manifold by which the pose estimation module was trained. This learnable pre-processing step, which we call PolishNetU, converts the pressure data to polished images that better resemble human figures as expected by common pose estimation models.

Lets define  $I' \in \mathbb{R}^{W \times H \times 3}$  as the output of our PolishNetU  $P$ , in other words  $I' = P(I; \theta_P)$  and  $\hat{K} = Q(I'; \theta_Q)$ , where  $P$  is the pre-processing function, and  $\theta_P$  is its set of trainable parameters. Using the  $E_{heatmap}$  and  $E_{PAF}$ , a pre-trained pose estimation module can force PolishNetU to synthesize entirely new images. To prevent pose deviations, we also added a third term, pixel loss, to the objective function. This term acts as a regularizer, which penalizes the distance between the input pressure maps and



**Fig. 2** The explored methods for in-bed pose estimations and respective loss functions are presented: **a** The off-the-shelf pose estimation network is used either as is or trained on the pressure data and annotated keypoints. **b** PolishNetU is utilized to pre-process the input pressure data and generate outputs close to realistic images to feed into the pre-existing pose estimation networks. **c** Similar to previous approach, but the pose estimation network is re-trained, either alone or together with PolishNetU

the synthesized polished images, which we defined as:

$$E_{pixel} = \frac{1}{HW} \sum_{i=1}^H \sum_{j=1}^W \|I - I'\|_2^2. \quad (4)$$

Finally, we optimize  $\theta_P$  using the objective function:

$$E(\theta_P) = \lambda_{heatmap} E_{heatmap} + \lambda_{PAF} E_{PAF} + \lambda_{pixel} E_{pixel}, \quad (5)$$

where we chose  $\lambda_{heatmap} = \lambda_{PAF}$  and adjusted  $\lambda_{pixel}$  to achieve the optimum representations. It is important to note that the pixel loss weight  $\lambda_{pixel}$  has a large impact on the representations, where large values can result in PolishNetU producing outputs overly similar to the unpolished inputs. In contrast, smaller values caused the heatmap and PAF losses to overlook this regularizer, resulting in significant deviations from the input images, generating non-human like images. Therefore in this setup, we first train the network by setting a high value for  $\lambda_{pixel}$ , and then slowly decrease its weight to enable our networks to focus on reconstructing weak pressure points after stabilizing.

**End-to-end re-training of PolishNetU and pose estimators** are the final possible solution. Since PolishNetU is acting as the pre-processing step, it is intuitive that the main pose estimation network would have a large impact on the performance. Therefore, after learning the real image representations  $I'$  by training the PolishNetU, we move on to optimizing  $\theta_Q$  by using the same objective function. Additionally, we can also choose to fine-tune PolishNetU alongside our pose estimation network, which we refer to as fine-tuned PolishNetU in the future sections.

## 4 Experiments and results

### 4.1 Data preparation

We use two different datasets to train and test our pressure-based pose estimation approach, PmatData [35] available in the PhysioNet repository [15] and HRL-ROS dataset [9]. These datasets are briefly described as follows:

**PmatData** Ostadabbas et al. [35] was recorded by a force sensitive pressure mapping mattresses. Each mattress contained 2048 sensors spread on a 32 by 64 grid with each sensor being 25.4 mm apart. The recording was performed with a frequency of 1 Hz for a pressure range of 0 – 100 mmHg. Data were recorded from 13 healthy subjects in 8 standard postures and 9 further sub-postures, for a total of 17 unique pose classes. Subjects were within a height range of 169 – 186 cm, a weight range of 63 – 100 Kg, and an age range of 19 – 34 years. We developed and utilized a tool in MATLAB for annotating the body

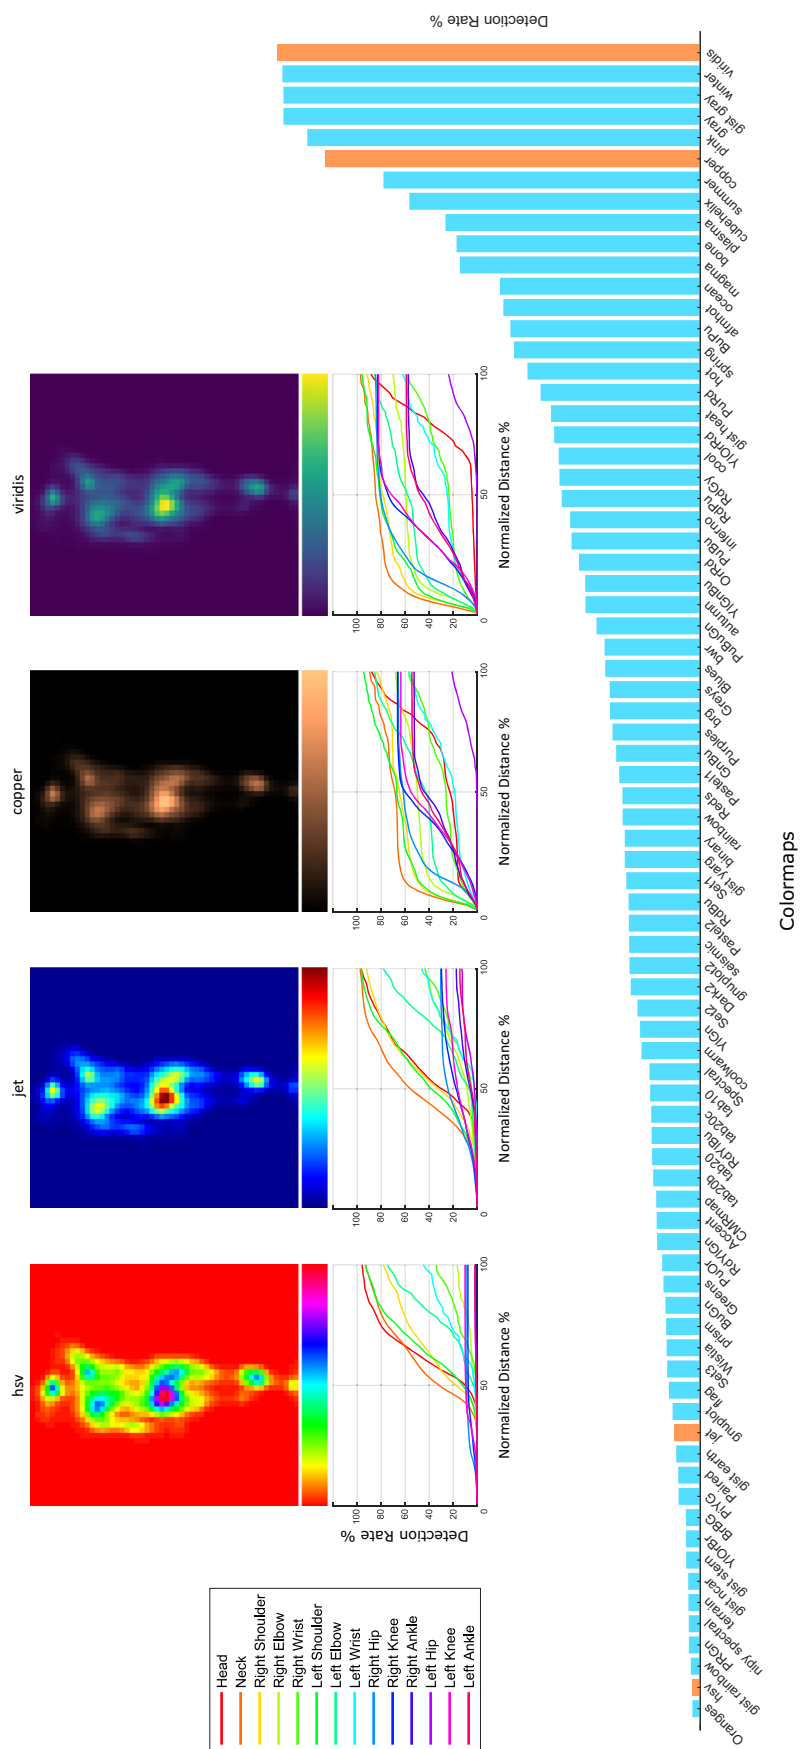
part keypoints in 18256 data samples. The annotating procedure was carried out by two researchers and then cross-checked to ensure consistency. To perform very rigorous evaluation experiments, we employ a leave-some-subjects-out validation strategy, training the network on 9 subjects and testing it on the remaining 4.

**HRL-ROS** Clever et al. [9] was collected for kinematic-based 3D pose recognition using a configurable bed embedded with pressure sensors and motion capture cameras. The bed was equipped with an array of  $27 \times 64$  sensors distributed 28.6 mm apart. A total of 17 subjects were asked to lie or sit in different postures and move a body limb in a specific path, while their limb position was being tracked using motion capture cameras. Subjects were within a height range of 160 – 185 cm, a weight range of 45.8 – 94.3 Kg, and an age range of 19 – 32 years. For our purposes we use the data from all of the subjects in all 13 lying postures, resulting in a total of 39095 pressure maps. Similar to PmatData, we leave the last 4 subjects for testing while keeping the rest for training.

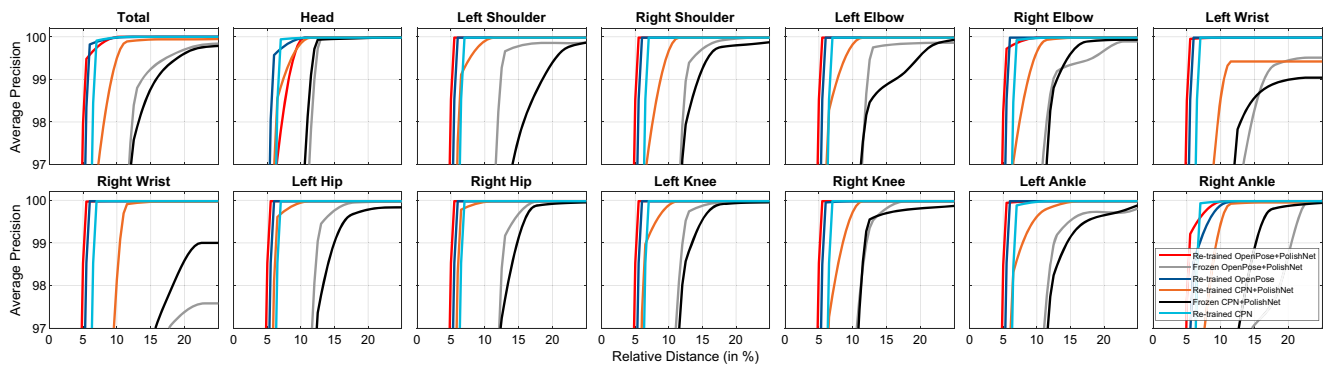
### 4.2 Pre-processing

First, we remove the noise caused by occasional malfunctioning pressure sensors. These artifacts usually occur when certain individual sensors become subject to pressure values outside the calibrated voltage range. To clean up the pressure values, we use a  $3 \times 3 \times 3$  spatio-temporal median filter. We tune the filter size by evaluating the pose estimation performance of a frozen OpenPose on the PmatData dataset, showing that larger filter sizes do not improve the performance. Following previous studies [12–14], we also remove the first 3 frames of each sequence, which in some cases are transition frames where pressure maps are not clear. We identified the frames and cross-checked them visually to remove the outliers using the histogram of the dataset based on the average pressure of each image.

The input pressure maps are provided in the form of  $W \times H \times 1$  arrays. However, most existing pose estimation methods have been trained on color images. As a result in order to utilize existing frameworks for our goal of estimating poses from pressure data, the pressure maps need to be converted to color images. Consequently, we convert the pressure maps to color images using a colormap. To this end, we need to select a colormap with high compatibility with the pose estimation network, and in general, with natural images containing human figures. Our investigations show that the choice of colormap can play a considerable role in the performance. We investigate 38 different colormaps and evaluate the error rates when PolishNetU is excluded from the pipeline and only the frozen OpenPose model is used. Figure 3 shows the



**Fig. 3** We compared the effect of applying 38 different colormaps to the pressure data by benchmarking the performance of OpenPose on the colorized images. Subsequently, the Viridis colormap achieved the highest average detection rate among all the available schemes



**Fig. 4** AP over different thresholds is demonstrated for each body part from PmatData dataset

performance over different colormaps and illustrates the estimated pose accuracy for different body parts for 4 sample colormaps from the distribution, namely HSV, Jet, Copper, and Viridis. Eventually, as shown in Fig. 3, Viridis exhibits the best color mapping characteristics. The pose estimation evaluation method is presented in Section 4.4.

### 4.3 Implementation details

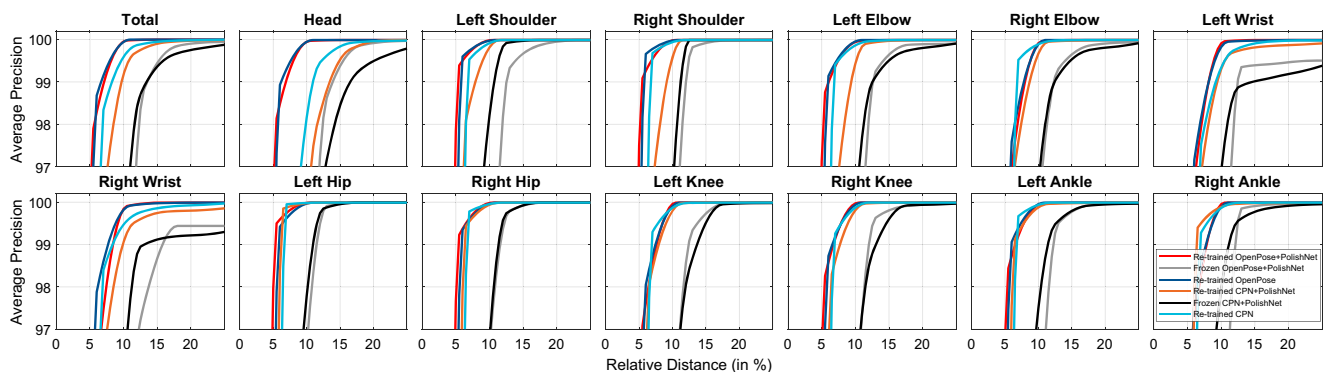
#### 4.3.1 PolishNetU

Polishing the colorized pressure maps and converting them to images consistent with the pose estimation module is done by a feed-forward network  $P$  called PolishNetU. Our proposed network is a variation of U-Net, consisting of a combination of encoders and decoders, including fully convolutional and deconvolutional layers, respectively. Given the colorized pressure data, a series of 8 encoder blocks of Conv-BatchNorm-LeakyReLU with a stride of 2 are exploited to encode the input to a data manifold, capturing pressure properties that incorporate pose information. The polished image space is then achieved by concatenating the residuals from the encoders with UpSample-DeConv-BatchNorm-LeakyReLU on the encoded latent space. Finally, using a  $\tanh$  activation

function, the last layer provides a polished image, which is compatible with the pose estimation module conditioned on the given pressure input. We use BatchNorm for faster and more reliable training, and upsampling layers instead of a deconvolution layer with a stride of 2 to avoid the deconvolution checkerboard artifact.

#### 4.3.2 Pose estimation module

To train PolishNetU, we utilize OpenPose [3] or CPN [7] as our pose identification module. OpenPose is a well-known network developed for real-time pose estimation first, by which has been recently extended to support face landmark detection and hand gesture detection as well. On the other hand, CPN is a more recent and powerful pose recognition method which was able to achieve state-of-the-art in many of the pose recognition challenges. The authors of OpenPose defined heatmap and Part Affinity Fields (PAF) outputs which we utilize to define our pose estimation objectives. Each heatmap provides a 2D distribution of the belief that a keypoint is located on each pixel. Additionally, PAF is defined as a 2D vector field for each limb, where each 2D vector encodes both position and the orientation of the limb. Over the past few years, several versions of this network have been published, mostly with changes in



**Fig. 5** AP over different thresholds is demonstrated for each body part from HRL-ROS dataset



**Table 1** The AP5 values are presented for different methods on PmatData

|                                | H     | LS    | RS    | LE    | RE    | LW    | RW    | LH    | RH    | LK    | RK    | LA    | RA    |
|--------------------------------|-------|-------|-------|-------|-------|-------|-------|-------|-------|-------|-------|-------|-------|
| Re-trained OpenPose+PolishNetU | 99.4  | 100.0 | 100.0 | 100.0 | 100.0 | 100.0 | 100.0 | 100.0 | 100.0 | 100.0 | 100.0 | 100.0 | 99.8  |
| Re-trained OpenPose            | 91.3  | 100.0 | 100.0 | 100.0 | 100.0 | 100.0 | 100.0 | 100.0 | 100.0 | 100.0 | 100.0 | 100.0 | 100.0 |
| Frozen OpenPose+PolishNetU     | 92.2  | 90.1  | 45.7  | 80.8  | 88.5  | 52.5  | 43.2  | 48.1  | 36.9  | 80.2  | 98.6  | 100.0 | 96.0  |
| Frozen OpenPose                | 0.0   | 0.0   | 0.0   | 0.0   | 0.0   | 0.0   | 0.0   | 0.0   | 0.0   | 0.0   | 0.0   | 0.0   | 0.0   |
| Re-trained CPN+PolishNetU      | 99.6  | 99.5  | 99.5  | 99.5  | 99.6  | 99.4  | 99.4  | 99.6  | 99.5  | 99.6  | 99.6  | 99.5  | 99.5  |
| Re-trained CPN                 | 98.7  | 99.3  | 99.4  | 99.3  | 99.3  | 98.9  | 99.1  | 99.5  | 99.5  | 99.4  | 99.3  | 99.1  | 98.8  |
| Frozen CPN+PolishNetU          | 100.0 | 72.3  | 44.3  | 82.0  | 78.5  | 41.6  | 14.4  | 48.6  | 67.7  | 88.9  | 100.0 | 82.9  | 85.4  |
| Frozen CPN                     | 0.0   | 0.0   | 0.0   | 0.0   | 3.3   | 4.7   | 0.0   | 0.0   | 0.0   | 7.9   | 31.3  | 23.1  | 4.4   |

the final blocks for face and hand landmark detection and trade-offs between memory usage, performance, and speed. For our work, we choose the original version published in 2016 that includes 7 stages for refining the heatmaps and PAFs. We use the exact design parameters of the original OpenPose model. In our pipeline, neither OpenPose, nor CPN, do not use any non-linear activation functions, therefore vanishing gradients is not an issue. Furthermore, we neglect the loss for non-visible parts, eyes, and ears due to the nature of our pressure data. Consequently, we end up with 14 heatmaps for the head, neck, shoulders, elbows, wrists, ankles, knees, and the hip, as well as 28 PAFs connecting the body parts to be used in our cost functions.

### 4.3.3 Pipeline

The pipeline is implemented using TensorFlow on an NVIDIA Titan XP GPU. We use an Adam optimizer at the training stage with a learning rate of  $10^{-3}$ , which is decayed with a rate of 0.95 for every 1000 update iterations. The pipeline is trained for 40 epochs with a batch size of 16.  $\lambda_{PAF}$  and  $\lambda_{heatmap}$  are both set to 1 while  $\lambda_{pixel}$  is changed from 1 to 0.01 during training to first stabilize the model, and then allow it to interpolate the vivid body limbs without pixel loss penalty.


### 4.4 Performance evaluation

After obtaining the output heatmaps in the form of  $W_h \times H_h \times 14$ , we smooth them using a Gaussian kernel of  $3 \times 3$  across the spatial dimensions to reduce prediction noise in the output. Then, we perform a flip-test to reduce the model’s bias to left and right directions by obtaining the output from the original input and its flipped version and then averaging them. Finally, we take the location of the maximum of each channel as our predicted keypoints, obtaining a  $14 \times 3$  array containing the location and prediction scores of the body limbs.

To evaluate the performance of our pipeline on the annotated data, we use the Average Precision (AP), a measure of joint localization accuracy at 5% error margin, similar to the intersection-over-union (IoU) threshold in object detection research [19, 46, 47]. First, we sort the predictions by their scores. Then we measure the distances between the predicted and ground-truth keypoints. If this distance is below a threshold, we consider the prediction a true-positive. Finally, we calculate AP by measuring the area under the precision and recall curves. The threshold is defined as a fraction, here 5%, of the person’s size, where the size is defined as the distance between the person’s left shoulder and right hip [2]. In our implementation, the

**Table 2** The AP5 values are presented for different methods on HRL-ROS


|                                | H    | LS   | RS   | LE   | RE   | LW   | RW    | LH   | RH   | LK   | RK    | LA   | RA   |
|--------------------------------|------|------|------|------|------|------|-------|------|------|------|-------|------|------|
| Re-trained OpenPose+PolishNetU | 98.7 | 99.8 | 99.7 | 99.5 | 97.7 | 97.0 | 97.3  | 99.5 | 99.5 | 98.1 | 98.9  | 99.2 | 98.1 |
| Re-trained OpenPose            | 98.8 | 99.7 | 99.4 | 99.3 | 97.4 | 97.1 | 96.0  | 99.5 | 99.5 | 98.2 | 98.7  | 99.2 | 98.4 |
| Frozen OpenPose+PolishNetU     | 80.8 | 89.6 | 44.9 | 73.6 | 81.6 | 79.1 | 88.6  | 89.3 | 85.3 | 38.9 | 32.8  | 36.7 | 63.4 |
| Frozen OpenPose                | 0.0  | 0.0  | 0.0  | 0.0  | 0.0  | 0.0  | 0.0   | 0.0  | 0.0  | 0.0  | 0.0   | 0.0  | 0.0  |
| Re-trained CPN+PolishNetU      | 92.4 | 95.3 | 95.2 | 94.9 | 95.1 | 94.9 | 94.9  | 95.6 | 95.6 | 95.1 | 95.0  | 95.3 | 95.4 |
| Re-trained CPN                 | 96.4 | 98.5 | 98.6 | 98.5 | 98.4 | 97.9 | 97.7  | 98.6 | 98.6 | 98.4 | 98.4  | 98.5 | 98.4 |
| Frozen CPN+PolishNetU          | 48.4 | 94.8 | 86.9 | 52.1 | 86.1 | 98.1 | 100.0 | 97.3 | 77.4 | 79.0 | 100.0 | 94.4 | 85.7 |
| Frozen CPN                     | 0.0  | 0.0  | 0.0  | 0.5  | 0.0  | 0.5  | 0.0   | 0.0  | 0.0  | 0.0  | 0.2   | 0.0  | 12.8 |



**Fig. 6** We provide examples of the performance of the explored architectures for estimating pose from input colorized pressure maps. Here, we show the raw pressure data (first row), colorized pressure data and ground truth keypoints (second row), frozen OpenPose predictions on colorized images (third row), PolishNetU with frozen OpenPose predictions on polished images (fourth row), and PolishNetU with re-trained OpenPose estimations on polished images (fifth row). We

removed the predictions that were too far from the actual ground-truth for better presentation. The addition of PolishNetU to the frozen OpenPose causes a boost in pose estimation performance, while re-trained OpenPose with PolishNetU makes slightly more detailed predictions. Furthermore, polished images inherit a better visual fidelity of human body limbs

**Fig. 7** Illustration of examples from PmatData where PolishNetU has reconstructed weak body parts. Specifically note the arms, knees, and the head



**Table 3** The AP5 (%) - MPJPE (mm) values are presented for two datasets over all possible solutions

| Pose Estimator | Model                          | PolishNetU  | Dataset     |             |
|----------------|--------------------------------|-------------|-------------|-------------|
|                |                                | Fine-tuning | PmatData    | HRL-ROS     |
| OpenPose [3]   | Frozen OpenPose                | –           | 0.0 – 429.2 | 0.0 – 279.5 |
|                | Frozen OpenPose+PolishNetU     | –           | 78.1 – 29.7 | 77.4 – 36.6 |
|                | Re-trained OpenPose            | –           | 99.8 – 15.1 | 98.4 – 16.3 |
|                | Re-trained OpenPose+PolishNetU | No          | 99.6 – 15.3 | 98.6 – 16.2 |
|                | Re-trained OpenPose+PolishNetU | Yes         | 99.9 – 14.8 | 98.7 – 16.1 |
| CPN [7]        | Frozen CPN                     | –           | 0.5 – 557.0 | 0.0 – 208.9 |
|                | Frozen CPN+PolishNetU          | –           | 71.8 – 78.1 | 85.6 – 39.2 |
|                | Re-trained CPN                 | –           | 99.4 – 20.6 | 98.0 – 14.4 |
|                | Re-trained CPN+PolishNetU      | No          | 95.9 – 28.8 | 95.9 – 18.4 |
|                | Re-trained CPN+PolishNetU      | Yes         | 99.6 – 18.4 | 96.4 – 16.7 |

average 5% threshold is equal to approximately 1.2 pixels or 32 mm considering the size of the input pressure images. Moreover, we also provide another evaluation metric used in pose estimation studies called mean-per-joint-position error (MPJPE) of the predictions [4, 9, 36]. MPJPE is measured by averaging the body joint prediction errors in mm, calculated in euclidean space.

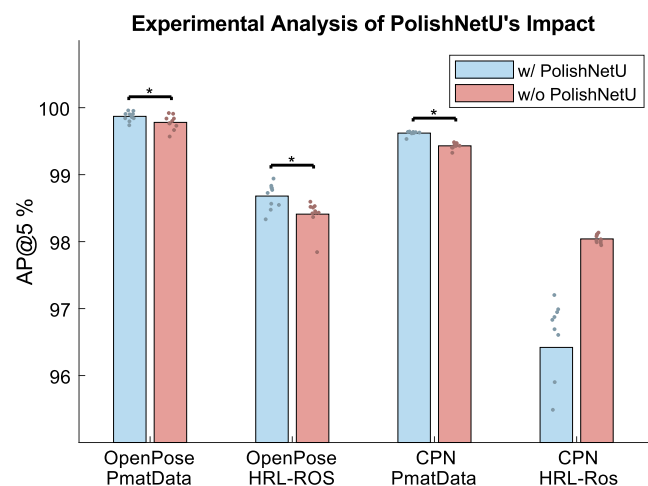
We analyze the aforementioned solutions by comparing them based on AP, plotted against normalized distances (defined by threshold × torso length) for different body parts in Figs. 4 and 5 for PmatData and HRL-ROS datasets. We omitted the frozen pose estimators without PolishNetU since they showed poor detection rates on both datasets. In contrast, the highest performance in both figures belongs to the models in which the pose estimation network is re-trained, where the combination of PolishNetU and OpenPose is slightly better than the rest, especially for PmatData dataset. Finally, solutions utilizing PolishNetU and frozen pose recognition networks show a weaker performance at 5% threshold compared to the others, but they also reach above 98% detection rate after the 10% threshold (50mm) for all of the body parts. It is also shown that the wrists, ankles, and the head are the most challenging body parts, where models containing OpenPose perform better than CPN. A more in-depth comparison of the AP5 for our selected models are presented in Tables 1 and 2

**Table 4** The number of parameters used in the models and their backbone feature extractors are presented

| Model      | Parameters | BackBone  |
|------------|------------|-----------|
| OpenPose   | 52.3M      | Resnet50  |
| CPN        | 46.0M      | VGG16     |
| PolishNetU | 13.6M      | UNet      |
| PolishNet  | 436K       | HourGlass |

for PmatData and HRL-ROS, respectively. We refrained from comparing AP10 of our models since in most cases they achieved near perfect accuracy, making the comparisons uninformative. It can be seen that the frozen pose estimators alone are not able to correctly identify poses without PolishNetU.

Several examples depicting the performance of PolishNetU are presented in Fig. 6, comparing 3 of our explored models. Frozen OpenPose is rarely able to predict the correct pose, thus being unreliable. It also miss-identifies the left and right sides of the human body since it was trained on real images that were mostly captured when facing the front of human subjects, as opposed to pressure data that records the image from the backside. As presented in the fourth row, PolishNetU with the frozen OpenPose pipeline has accurately identified the poses for vague input pressure maps while only miss-identifying very blurry areas such as



**Fig. 8** Statistical analysis of 10 repetitions of our experiments on re-trained pose estimators with and without PolishNetU (fine-tuned) is illustrated, showing that significant improvement is made by PolishNetU in 3 out of the 4 scenarios

wrists. Finally, the re-trained OpenPose with PolishNetU, illustrated in the fifth row, has made the best pose estimations among other methods, able to identify invisible limbs correctly.

Since PolishNetU was trained to synthesize images compatible with the image space by which OpenPose was trained, the polished outputs inherit less noise and show a higher resemblance to common standing human poses from behind. See Fig. 7. Notice that PolishNetU has reconstructed and connected the limbs and weak pressure areas that are not clearly visible in the colorized pressure maps. We have highlighted some of these reconstructed regions in Fig. 7. Moreover, in some instances (e.g. last column of Fig. 7), PolishNetU has even attempted to interestingly synthesize *oufits* for the subjects in order to make the output images look more natural and consistent with the input image space of the pose estimator.

Finally, we perform an ablation study on the explored models in Table 3. We see that the frozen pose estimators without PolishNetU achieve very poor estimation results, where the addition of PolishNetU boosts their performance by a significant amount. Moreover, we observe that re-training the pose estimator block has a higher impact on the performance compared to the previous approach, reducing the MPJPE on both datasets to less than half. Finally, while comparing OpenPose with CPN, we observe that OpenPose outperforms CPN on both datasets based on the AP5 metric, which can be because of reasons such as the number of parameters or their features extractor backbones (see Table 4). In contrast, CPN performs better for the HRL-ROS dataset based on the MPJPE criteria.

The addition of PolishNetU (fine-tuned) to re-trained pose estimators achieves the best results by a small, yet statistically significant margin. We use a non-parametric t-test on 10 different repetitions of our experiments with random initialization to compare the retrained pose estimators to the retrained pose estimators with PolishNetU (fine-tuned), and achieve  $p < 0.05$ , showing that the improvements caused by PolishNetU (fine-tuned) are statistically significant. We illustrate this experimental analysis in Fig. 8, showing the positive impact of PolishNetU in 3 out of the 4 cases. Furthermore, Fig. 8

highlights the effectiveness of OpenPose compared to CPN, achieving better performance on both datasets.

#### 4.5 Discussion and comparison

Although the field of in-bed pose estimation has attracted a considerable amount of recent works, most of the prior works on pressure datasets do not use a unified evaluation method and datasets, making comparisons challenging. We evaluate the solutions based on 2D MPJPE and AP5, which corresponds to 1.2 pixels or a 32 mm threshold for correct prediction. We show the effectiveness of PolishNetU by demonstrating the pose estimation performance of pre-existing pose estimators combined with it. Specifically, we show that PolishNetU with a frozen OpenPose achieves near-perfect pose estimation with AP5 values of 99.9% and 98.7%, and MPJPE of 14.8 and 16.1 mm, while PolishNetU with a frozen OpenPose performs with AP5 of 78.1% and 85.6%, and MPJPE of 29.7 and 39.2 mm on PmatData and HRL-ROS datasets respectively.

Table 5 compares the results of other works with our best configuration, which is the combination of PolishNetU and the re-trained OpenPose, achieving an AP5 of 99.9%. Moreover, we show that on the PmatData dataset, we are able to obtain PCK@25 of 100%, outperforming previous works [14]. In another study [9], a kinematic-based convolutional neural network was used for 3D joint prediction on HRL-ROS dataset, achieving MPJPE of 73.5 mm. Although their model has the advantage of 3D joint prediction, several of our explored solutions containing a re-trained pose estimator or the PolishNetU are able to achieve a more accurate prediction in 2D space. In a more recent study [10], the same authors were able to obtain MPJPE of 111.8 mm on a different dataset by training a model called PressureNet on synthetic data and testing it on real images. In another study [4], in-bed pressure data were collected from 6 subjects, reporting an MPJPE of 68 mm using a deep fully convolutional pose estimation model. Lastly, [5] used a camera-based approach on 3 subjects. They reported an average accuracy, at 15 pixels threshold, of 80.5% and 91.6% using a frozen OpenPose and a combination of Kalman filter with a pose estimation network, respectively.

**Table 5** A comparison of our best results (PolishNetU+Re-trained OpenPose) with other works

| Ref  | Dataset  | Evaluation  | AP5  | MPJPE (mm)        | PCK25 (%) |
|------|----------|-------------|------|-------------------|-----------|
| [9]  | HRL-ROS  | Leave 1 Out | N/A  | 73.5 <sup>a</sup> | N/A       |
| [14] | PmatData | Leave 1 Out | N/A  | N/A               | 95.8      |
| Ours | HRL-ROS  | Leave 4 Out | 98.0 | 16.1              | 100       |
| Ours | PmatData | Leave 4 Out | 99.9 | 14.8              | 100       |

<sup>a</sup>Denotes the performance of 3D pose estimation

Although in some cases our explored methods and datasets are different, we were able to obtain higher performances on a much more larger and complex data space.

## 5 Conclusions and future work

In-bed pressure data can provide valuable information for the estimation of a user's pose, which is of high value for clinical and smart home monitoring. However, pressure-based pose estimation deals with a number of challenges, including the lack of large annotated datasets and proper fine-tuned frameworks. Additionally, pressure data impose some inherent limitations such as weak pressure areas caused by supported or raised body parts. In this paper, we explored several end-to-end models for performing pose estimation with in-bed pressure maps, including direct use of off-the-shelf models and re-training them for our purpose. As a part of our analysis, we exploited the novel idea of learning a domain adaptation fully convolutional network, PolishNetU, which generates images as robust representations that work well for common pre-trained pose estimation models, in this case, OpenPose and CPN. This method utilized a compound objective function which integrates the pose identification loss, reconstructing lost body parts caused by weak pressure points, and a pixel loss penalizing large deviations from the original pressure maps. The explored pipelines showed effective performance on highly unclear pressure data. Our evaluation results demonstrated that while re-training the pre-existing pose estimation models have the most impact on performance, if they are kept frozen, PolishNetU can boost the performance significantly as well. Given the performance of PolishNetU with two different pose estimators and on two datasets, we believe this model can be used as a pre-trained block prior to other pose estimation networks for identifying pose from pressure data.

For future work we aim to propose a modified objective function making use of pose priors as a constraint, preventing the model from outputting unlikely poses. Moreover, we aim to investigate the use of generative adversarial networks and integrate a discriminator in our model, which we anticipate, may enhance the reconstruction of weak pressure areas.

**Acknowledgements** The Titan XP GPU used for this research was donated by the NVIDIA Corporation.

## Declarations

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

## References

- Achilles F, Ichim AE, Coskun H, Tombari F, Noachtar S, Navab N (2016) Patient mocap: human pose estimation under blanket occlusion for hospital monitoring applications. In: International conference on medical image computing and computer-assisted intervention, pp 491–499
- Andriluka M, Pishchulin L, Gehler P, 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), pp 3686–3693
- Cao Z, Simon T, Wei SE, Sheikh Y (2017) Realtime multi-person 2d pose estimation using part affinity fields. In: IEEE conference on computer vision and pattern recognition (CVPR), pp 1302–1310
- Casas L, Navab N, Demirci S (2019) Patient 3d body pose estimation from pressure imaging. *Int J Comput Assist Radiol Surg* 14(3):517–524
- Chen K, Gabriel P, Alasfour A, Gong C, Doyle WK, Devinsky O, Friedman D, Dugan P, Melloni L, Thesen T et al (2018a) Patient-specific pose estimation in clinical environments. *IEEE J Transl Eng Health Med* 1
- Chen Y, Shen C, Wei XS, Liu L, Yang J (2017) Adversarial posenet: A structure-aware convolutional network for human pose estimation. In: Proceedings of the IEEE international conference on computer vision (ICCV), pp 1212–1221
- Chen Y, Wang Z, Peng Y, Zhang Z, Yu G, Sun J (2018b) Cascaded pyramid network for multi-person pose estimation. In: IEEE conference on computer vision and pattern recognition (CVPR), pp 7103–7112
- Chu X, Yang W, Ouyang W, Ma C, Yuille AL, Wang X (2017) Multi-context attention for human pose estimation. In: IEEE conference on computer vision and pattern recognition (CVPR), pp 5669–5678
- Clever HM, Kapusta A, Park D, Erickson Z, Chitalia Y, Kemp CC (2018) 3d human pose estimation on a configurable bed from a pressure image. In: IEEE/RSJ international conference on intelligent robots and systems (IROS), pp 54–61
- Clever HM, Erickson Z, Kapusta A, Turk G, Liu K, Kemp CC (2020) Bodies at rest: 3d human pose and shape estimation from a pressure image using synthetic data. In: IEEE conference on computer vision and pattern recognition (CVPR), pp 6215–6224
- Cunha JPS, Choupina HMP, Rocha AP, Fernandes JM, Achilles F, Loesch AM, Vollmar C, Hartl E, Noachtar S (2016) Neurokinect: a novel low-cost 3d video-eeeg system for epileptic seizure motion quantification. *PLoS One* 11(1):e0145669
- Davoodnia V, Etemad A (2019) Identity and posture recognition in smart beds with deep multitask learning. In: IEEE international conference on systems, man and cybernetics (SMC), pp 3054–3059
- Davoodnia V, Slinowsky M, Etemad A (2020) Deep multitask learning for pervasive bmi estimation and identity recognition in smart beds. *J Ambient Intell Humaniz Comput* :1–15
- Davoodnia V, Ghorbani S, Etemad A (2021) In-bed pressure-based pose estimation using image space representation learning. In: IEEE international conference on acoustics, speech and signal processing (ICASSP)
- Goldberger AL, Amaral LA, Glass L, Hausdorff JM, Ivanov PC, Mark RG, Mietus JE, Moody GB, Peng CK, Stanley HE (2000) Physiobank, physiotoolkit, and physionet: components of a new research resource for complex physiologic signals. *Circulation* 101(23):e215–e220
- Goldberger AL, Amaral LAN, Glass L, Hausdorff JM, Ivanov PC, Mark RG, Mietus JE, Moody GB, Peng CK, Stanley HE (2000)

- PhysioBank, PhysioToolkit, and PhysioNet: components of a new research resource for complex physiologic signals. *Circulation* 101(23):e215–e220. <https://doi.org/10.1161/01.CIR.101.23.e215>, <http://circ.ahajournals.org/content/101/23/e215>
17. Grimm R, Sukkau J, Hornegger J, Greiner G (2011) Automatic patient pose estimation using pressure sensing mattresses. In: *Bildverarbeitung für die Medizin*. Springer, pp 409–413
  18. He K, Cao X, Shi Y, Nie D, Gao Y, Shen D (2018) Pelvic organ segmentation using distinctive curve guided fully convolutional networks. *IEEE Trans Med Imaging* 38(2):585–595
  19. Hu H, Gu J, Zhang Z, Dai J, Wei Y (2018) Relation networks for object detection. In: *IEEE conference on computer vision and pattern recognition (CVPR)*, pp 3588–3597
  20. Imabuchi T, Prima ODA, Ito H (2018) Automated assessment of nonverbal behavior of the patient during conversation with the healthcare worker using a remote camera. *eTELEMED* 2018. p 29
  21. Insafutdinov E, Pishchulin L, Andres B, Andriluka M, Schiele B (2016) Deepercut: A deeper, stronger, and faster multi-person pose estimation model. In: *European conference on computer vision (ECCV)*. Springer, pp 34–50
  22. Insafutdinov E, Andriluka M, Pishchulin L, Tang S, Levinkov E, Andres B, Schiele B (2017) Arttrack: articulated multi-person tracking in the wild. In: *IEEE conference on computer vision and pattern recognition (CVPR)*, pp 1293–1301
  23. Javaid AQ, Gupta R, Mihalidis A, Etemad SA (2017) Balance-based time-frequency features for discrimination of young and elderly subjects using unsupervised methods. In: *IEEE EMBS international conference on biomedical & health informatics (BHI)*, pp 453–456
  24. Ke L, Chang MC, Qi H, Lyu S (2018) Multi-scale structure-aware network for human pose estimation. In: *Proceedings of the european conference on computer vision (ECCV)*, pp 713–728
  25. Koppula HS, Saxena A (2016) Anticipating human activities using object affordances for reactive robotic response. *IEEE Trans Pattern Anal Mach Intell* 38(1):14–29
  26. Lee CH, Kim DK, Kim SY, Rhee CS, Won TB (2015) Changes in site of obstruction in obstructive sleep apnea patients according to sleep position: a dise study. *Laryngoscope* 125(1):248–254
  27. Lee J, Kwon H, Seo J, Shin S, Koo JH, Pang C, Son S, Kim JH, Jang YH, Kim DE et al (2015) Conductive fiber-based ultrasensitive textile pressure sensor for wearable electronics. *Adv Mater* 27(15):2433–2439
  28. Lin F, Song C, Xu X, Cavuoto L, Xu W (2017) Patient handling activity recognition through pressure-map manifold learning using a footwear sensor. *Smart Health* 1:77–92
  29. Liu JJ, Huang MC, Xu W, Sarrafzadeh M (2014) Bodypart localization for pressure ulcer prevention. In: *Annual international conference of the IEEE engineering in medicine and biology society (EMBC)*, pp 766–769
  30. Liu S, Ostadabbas S (2017) A vision-based system for in-bed posture tracking. In: *Proceedings of the IEEE international conference on computer vision workshops*, pp 1373–1382
  31. Liu S, Ostadabbas S (2019) Seeing under the cover: A physics guided learning approach for in-bed pose estimation. In: *International conference on medical image computing and computer-assisted intervention*. Springer, pp 236–245
  32. Liu S, Yin Y, Ostadabbas S (2019) In-bed pose estimation: Deep learning with shallow dataset. *IEEE J Transl Eng Health Med* 7:1–12
  33. McCabe SJ, Xue Y (2010) Evaluation of sleep position as a potential cause of carpal tunnel syndrome: preferred sleep position on the side is associated with age and gender. *Hand* 5(4):361–363
  34. McCabe SJ, Gupta A, Tate DE, Myers J (2011) Preferred sleep position on the side is associated with carpal tunnel syndrome. *Hand* 6(2):132–137
  35. Ostadabbas S, Pouyan MB, Nourani M, Kehtarnavaz N (2014) In-bed posture classification and limb identification. In: *IEEE biomedical circuits and systems conference*, pp 133–136
  36. Rhodin H, Salzmann M, Fua P (2018) Unsupervised geometry-aware representation for 3d human pose estimation. In: *Proceedings of the european conference on computer vision (ECCV)*, pp 750–767
  37. Ruvalcaba-Cardenas AD, Scoleri T, Day G (2018) Object classification using deep learning on extremely low-resolution time-of-flight data. In: *IEEE digital image computing: techniques and applications (DICTA)*, pp 1–7
  38. Singh V, Ma K, Tamersoy B, Chang YJ, Wimmer A, O'Donnell T, Chen T (2017) Darwin: Deformable patient avatar representation with deep image network. In: *International conference on medical image computing and computer-assisted intervention*, pp 497–504
  39. Tang W, Yu P, Wu Y (2018) Deeply learned compositional models for human pose estimation. In: *European conference on computer vision (ECCV)*, pp 190–206
  40. Tulyakov S, Liu MY, Yang X, Kautz J (2018) Mocogan: Decomposing motion and content for video generation. In: *Proceedings of the IEEE conference on computer vision and pattern recognition*, pp 1526–1535
  41. Wang CW, Hunter A (2010) Robust pose recognition of the obscured human body. *Int J Comput Vis (ICCV)* 90(3):313–330
  42. Wei SE, Ramakrishna V, Kanade T, Sheikh Y (2016) Convolutional pose machines. In: *IEEE conference on computer vision and pattern recognition*, pp 4724–4732
  43. Woo KY, Sears K, Almost J, Wilson R, Whitehead M, VanDenKerkhof EG (2017) Exploration of pressure ulcer and related skin problems across the spectrum of health care settings in ontario using administrative data. *Int Wound J* 14(1):24–30
  44. Xiao Y, Zijie Z (2020) Infrared image extraction algorithm based on adaptive growth immune field. *Neural Process Lett* 51(3):2575–2587
  45. Yang W, Li S, Ouyang W, Li H, Wang X (2017) Learning feature pyramids for human pose estimation. In: *IEEE international conference on computer vision (ICCV)*, pp 1290–1299
  46. Zhao ZQ, Zheng P, Xu St, Wu X (2019) Object detection with deep learning: A review. *IEEE Trans Neural Netw Learn Syst* 30(11):3212–3232
  47. Zhou P, Ni B, Geng C, Hu J, Xu Y (2018) Scale-transferrable object detection. In: *IEEE conference on computer vision and pattern recognition (CVPR)*, pp 528–537

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



**Vaddad Davoodnia** is a Ph.D. student and a member of the Ambient Intelligence and Interactive Machines (AIIM) lab at the Department of Electrical and Computer Engineering, Queen's University, Canada. He is also a member of the Ingenuity Labs Research Institute. He received his B.Sc. degree in electrical and computer engineering from Sharif University, Iran. Prior to joining Queen's, he was a part-time researcher with the Brain

Engineering Center, Institute for Research in Fundamental Sciences (IPM) for three years, resulting in publications in the fields of machine learning and neuroscience. His main research interests are machine learning, computer vision, and affective computing using wearable technologies.



**Saeed Ghorbani** is a Ph.D. student in Electrical Engineering and Computer Science at York University, Toronto, advised by Dr. Niko Troje. He is also a VISTA and CVR trainee. His research interests are machine learning, computer vision, computer graphics, and computer animation. His current research focuses on leveraging novel deep probabilistic models for realistic character animation.



**Ali Etemad** is an Assistant Professor, as well as a Mitchell Professor in AI for Human Sensing & Understanding at the Department of Electrical and Computer Engineering, Queen's University, Canada, where he leads the Ambient Intelligence and Interactive Machines (AIIM) lab. He is also a faculty member at Ingenuity Labs Research Institute. He received his M.A.Sc. and Ph.D. degrees in Electrical and Computer Engineering from Carleton University,

Ottawa, Canada, in 2009 and 2014, respectively. His main area of research is machine learning and deep learning focused on human-centered applications with wearables, smart devices, and smart environments. Prior to joining Queen's, he held several industrial positions as lead scientist. He has co-authored over 80 articles and is a co-inventor of several patents. He has been a Senior PC, PC, and an active reviewer for several venues. He has received a number of awards including the Supervisor of the Year Award, the Instructor of the Year Award, and Best Paper Awards at conferences such as CGI, among others. Dr. Etemad's lab and research program have been funded by the Natural Sciences and Engineering Research Council (NSERC) of Canada, Ontario Centers of Excellence (OCE), Canadian Foundation for Innovation (CFI), Mitacs, and other organizations, as well as the private sector.