

DECA: Deep viewpoint-Equivariant human pose estimation using Capsule Autoencoders

Nicola Garau¹, Niccolò Bisagno¹, Piotr Bródka¹, and Nicola Conci¹

{[nicola.garau](mailto:nicola.garau@unitn.it), [niccolo.bisagno](mailto:niccolo.bisagno@unitn.it)}@unitn.it, piotrbrodka95@gmail.com, nicola.conci@unitn.it

¹University of Trento, Via Sommarive, 9, 38123 Povo, Trento TN




Figure 1: **[Better seen in color]**. Overview of the proposed solution. Two different views of the same subject are shown for each image: (a) 2D/3D Human Pose Estimation (HPE) and (b) 3D Human Mesh Recovery (HMR) methods achieve good accuracy on the front-view (second row). Changing the viewpoint turns into performance degradation (first row). Our method (c) promotes viewpoint equivariance, showing good results in both the RGB and depth domains.

Abstract

Human Pose Estimation (HPE) aims at retrieving the 3D position of human joints from images or videos. We show that current 3D HPE methods suffer a lack of viewpoint equivariance, namely they tend to fail or perform poorly when dealing with viewpoints unseen at training time. Deep learning methods often rely on either scale-invariant, translation-invariant, or rotation-invariant operations, such as max-pooling. However, the adoption of such procedures does not necessarily improve viewpoint generalization, rather leading to more data-dependent methods. To tackle this issue, we propose a novel capsule autoencoder network with fast Variational Bayes capsule routing, named DECA. By modeling each joint as a capsule

entity, combined with the routing algorithm, our approach can preserve the joints' hierarchical and geometrical structure in the feature space, independently from the viewpoint. By achieving viewpoint equivariance, we drastically reduce the network data dependency at training time, resulting in an improved ability to generalize for unseen viewpoints. In the experimental validation, we outperform other methods on depth images from both seen and unseen viewpoints, both top-view, and front-view. In the RGB domain, the same network gives state-of-the-art results on the challenging viewpoint transfer task, also establishing a new framework for top-view HPE. The code can be found at <https://github.com/mmlab-cv/DECA>.

1. Introduction

Human pose estimation is key for many applications, such as action recognition, animation, gaming, to name a few [16, 29, 28]. State of the art methods [2, 32] that rely on RGB images can correctly localize human joints (e.g. torso, elbows, knees) in images, also in presence of occlusions. However, they tend to fail when dealing with challenging scenarios. The top-view perspective, in particular, turns out to be a difficult task; on the one hand, it causes the largest amount of joints occlusions, and on the other hand, it suffers the scarcity of suitable training data, as shown in Fig. 1.

When presented with unseen viewpoints, humans display a remarkable ability to estimate human poses, even in the presence of occlusions and unconventional joints configurations. This is not always true in computer vision. In fact, available methods are trained in relatively constrained settings [15], with a limited variability between different viewpoints. Limited data, especially from the top-viewpoint, along with limited capabilities of modeling the hierarchical and geometrical structure of the human pose, results in poor generalization capabilities.

This generalization problem, known as the *viewpoint problem*, depends on how the network activations vary with the change of the viewpoint, usually after a transformation (translation, scaling, rotation, shearing). Convolutional Neural Networks (CNNs) scalar activations are not suitable to effectively manage these viewpoint transformations, thus needing to rely on max-pooling and aggressive data augmentation [4, 9, 22, 36]. By doing so, CNNs aim at achieving *viewpoint invariance*, defined as

$$f(Tx) = f(x) \quad (1)$$

According to this formulation, applying a viewpoint transformation T on the input image x , *does not change* the outcome of the network activations.

However, a more desirable property would be to capture and retain the transformation T applied to the input image x , thus obtaining a network that is aware of the different transformations applied to the input. Being able to model network activations that change in a structured way according to the input viewpoint transformations is also called *viewpoint equivariance* and it is defined as:

$$f(Tx) = Tf(x). \quad (2)$$

This is achieved by introducing *capsules*: groups of neurons that explicitly encode the intrinsic viewpoint-invariant relationship existing between different parts of the same object. Capsule networks (CapsNets) can learn part-whole relationships between so-called *entities* across different viewpoints [12, 26, 13], similarly to how our visual cortex system operates, according to the recognition-by-components

theory [1]. Unlike traditional CNNs, which usually retain viewpoint invariance, capsule networks can explicitly model and jointly preserve a viewpoint transformation T through the network activations, achieving *viewpoint equivariance* (Eq. 2).

Developing viewpoint-equivariant methods for 3D HPE networks leads to multiple advantages: (i) the learned model is more robust, interpretable, and suitable for real-world applications, (ii) the viewpoint is treated as a learnable parameter, allowing to disentangle the 3D data of the skeleton from each specific view, (iii) the same annotated data can be used to train a network for different viewpoints, thus less training data is required.

In this work, we address the problem of viewpoint-equivariant human pose estimation from single depth or RGB images. Our contribution is summarised as follows:

- We present a **novel Deep viewpoint-Equivariant Capsule Autoencoder architecture (DECA)** which jointly addresses multiple tasks, such as 3D and 2D human pose estimation.
- We show how our network works with limited training data, no data augmentation, and across different input domains (RGB and depth images).
- We show how the feature space organization, defined by routing the input information to build capsule entities, improves when the tasks are jointly addressed.
- We evaluate our method on the ITOP [9] dataset for the depth domain and on the PanopTOP31K [5] dataset for the RGB domain. We establish a new baseline for the viewpoint transfer task and in the RGB domain.

2. Related work

In recent years, human pose estimation has been a subject of multiple studies, particularly for real-time 2D HPE [2], 3D HPE [32] and human mesh recovery (HMR) approaches [19, 18]. In this work, we focus on HPE from single views, using either RGB [2, 10] or depth images [9, 22, 36].

Viewpoint-invariant HPE from RGB images. 3D HPE usually leverages on additional cues, such as 2D predictions [32, 34, 30], multiple images [38], pre-trained models [17] and pose dictionaries [27]. Other recent works aim at end-to-end, learning-based 3D HPE [25, 31, 21]. In the RGB domain, common HPE datasets such as Human3.6M [14], provide images from multiple views, like front-view or side-view, while the top-view component is generally missing. It is then evident that the lack of suitable multi-view (top-view in particular) data implies that state-of-the-art methods [2, 32, 19, 18] necessarily perform poorly when presented with an unseen viewpoint at test time, as shown in Fig. 1(a).

Viewpoint-invariant HPE from depth images. Viewpoint invariant HPE methods have been developed using depth images [9, 22, 36] from top-view and side-view, using datasets like the K2HPD Body Pose Dataset [35] and the ITOP dataset [9]. To take advantage of the 3D information encoded in 2D depth images, one recent research trend is to resort to 3D deep learning. The paid efforts can be generally categorized into 3D CNN-based and point-set-based families. To enhance the 3D properties of depth data and compute more significant features, current methods rely on 3D CNNs [9, 22] or 2D CNNs with dense features [36].

3D CNN-based methods [9, 22] perform a voxelization operation on pixels to transform them into 3D objects. To process the 3D data, each network performs costly 3D convolutions on the input data. These operations are responsible for the high computational burden and the difficulty to properly tune a high number of parameters in 3D CNNs. In the domain of 2D CNNs, Xiong *et al.* [36] capture the 3D structure by computing dense features in an ensemble way, thus avoiding computationally intensive CNN layers, but they still rely on a backbone pre-trained network to extract 2D features. Still, the above-mentioned approaches usually achieve weak viewpoint-invariance but fail to model viewpoint-equivariance. Moreover, we argue that the 3D geometry of the data should be interpreted by the network without relying on the voxelization embedding, or a 2D pre-trained feature extraction network.

Capsule networks for HPE. Capsule networks have shown the ability to model the geometric nature of training data thanks to the network structure and features [26, 13, 20]. Sabour *et al.*, introduce a routing algorithm for vector capsules, called *routing-by-agreement* as a better max-pooling substitute. Hinton *et al.* [13] further improve accuracy through a more complex matrix capsule structure and an Expectation-Maximization routing (*EM-routing*) for capsules. Unfortunately, the EM-routing and the 4×4 pose matrix embedded in the capsule contribute to increasing the training time, when compared to both CNNs and vector CapsNets. Kosiorek *et al.* [20] introduce for the first time an unsupervised capsule-based autoencoder. Ribeiro *et al.* in [24] build upon the EM-routing version of capsule by proposing for the first time a Variational Bayes capsule routing (VB routing) fitting a mixture of transforming Gaussians. They present state-of-the-art results using $\sim 50\%$ fewer capsules, achieving both performance gain and network complexity reduction. However, all the mentioned works only consider small datasets, such as MNIST, small-NORB, and CIFAR-10 for benchmarking.

In the RGB domain, Ramírez [23] tackles the problem of RGB HPE using dynamic vector capsule networks [26] to solve the 3D HPE problem in an end-to-end fashion. However, their work only exploits lateral viewpoints from the Human3.6M dataset and only considering RGB data.

In this work, we use matrix capsules [13], along with a different capsule routing algorithm and a new encoding-decoding pipeline with GELU activations. We argue that matrix capsules are better suited than vector capsules for the 3D HPE task, as the 4×4 pose matrix used for the routing can capture 3D geometry better than a dynamic vector structure.

3. Method

We now analyze the proposed autoencoder, DECA, starting with the capsule encoder and the multi-task decoders. DECA can be trained end-to-end, without any pre-training or data augmentation, and it works in real-time in the inference phase. An overview of the proposed architecture is shown in Fig. 2.

3.1. Capsule encoder

The encoding module of the network (light blue in Fig. 2) is divided in: (i) an input pre-processor I , (ii) a CNN encoder E and (iii) four layers of Matrix Capsules with Variational Bayes Routing [24].

(i) I is a layer which normalizes the different type of data (RGB images, depth images, top-view, side-view, free-view) in the interval $[0, 1]$.

(ii) The normalised input is then forwarded to a CNN encoder E , built using four convolutional layers with inputs $[N_{ch}, 64, 128, 256]$, instance normalisation and GELU activations [11], as shown in Eq. 3. N_{ch} is the number of channels, which may vary depending on the input.

$$\text{GELU}(x) \approx 0.5x(1 + \tanh \left[\sqrt{\frac{2}{\pi}}(x + 0.044715x^3) \right]) \quad (3)$$

(iii) The output of the CNN encoder E feeds our capsule layers. It has been shown in previous works [26, 13, 20] that capsules provide a superior understanding of the viewpoint and the relationship between parts and parent objects, thus aiming at true viewpoint equivariance. Given the multiple degrees of freedom of each joint, we adopt the matrix capsules model [13] instead of vector capsules [26], enriching the description of single joints as hierarchically linked capsule entities. We deploy the novel capsule routing based on Variational Bayes (VB) [24], which is proven to speed up the training of our matrix capsules layers, at the same time improving performances. The last iteration of the VB routing is also called *ClassRouting* and it is used to route the highest-level information to the last layer of capsules before the feature space \mathcal{F} .

In our CapsNet, we employ four layers: a *primary capsules* layer encapsulates the output features of E into 16-dimensional capsules, two *convolutional capsules* layers refine the capsule features, and a final *class capsules* layer

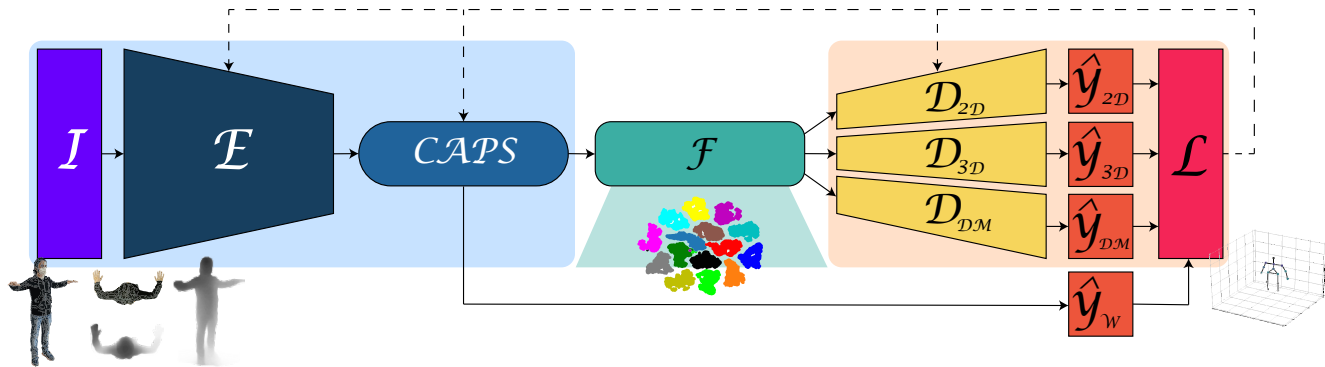


Figure 2: [Better seen in color]. Overview of the proposed architecture. In light blue, the encoding module (Input, CNN encoder, Capsule layers), in green the interpretable feature space with capsule entities, in light orange the decoding module (fully connected decoders with multiple tasks and self-balancing loss).

encodes the output into a J -dimensional features in the latent space \mathcal{F} , where J is the number of joints, also called *entities*.

Given each lower-level capsule i and the corresponding higher-level capsule j , we define M_i as the proposed lower level pose matrix and $W_{ij} \in \mathbb{R}^{4 \times 4}$ as a trainable viewpoint-equivariant transformation matrix such that:

$$V_{j|i} = M_i W_{ij} \quad (4)$$

where $V_{j|i}$ is the vote coming from lower capsules i for higher capsules j . The voting procedure takes place inside the VB routing and it allows each lower capsule i to route its information to a higher capsule j of its choice, thus allowing to build the hierarchical structure typical of CapsNets.

To promote the viewpoint equivariance in Eq. 2, we introduce an inverse matrix \hat{y}_W in the *class capsules*, which aims at satisfying the Inverse Graphics constraint:

$$\hat{y}_W W_{ij} = I \quad (5)$$

meaning that the learned inverse matrix \hat{y}_W effectively acts as an approximated inverse of the rendering operation, as it is commonly found in computer graphics [12].

At the output of the encoder, each *entity* corresponding to each joint of the skeleton is defined by a flattened vector of 16 elements, or, in other words, a 4×4 matrix, which is sufficient to grasp the complete pose (translation + rotation) of each joint.

An overview of the capsule encoder is shown in Algorithm 1. In the algorithm, s_{3D} , s_{2D} , s_{DM} , s_W are weights used for the self-balancing of the loss, w_c are the convolutional layer weights, a are the activations of each Capsule layer, and $\{\cdot\}$ represents parameters used only when in the RGB domain.

Algorithm 1: Capsule encoder

```

CapsuleEncoder ( $x$ )
  inputs :  $x = x_0 \dots x_{BS}$ ,  $BS$  = batch size of
            RGB or depth images
  outputs:  $\mathcal{F} = J$  16-dimensional entities;
             $\hat{y}_W$  = trainable Inverse Graphics
            matrix
   $s_{3D}, s_{2D}, \{s_{DM}\}, s_W \leftarrow 1$ ;
   $w_c \leftarrow xavier_{uniform}() \quad \forall c \in ConvLayers$ ;
  foreach  $i \in ConvLayers$  do
     $x \leftarrow Conv2d_i(x)$ ;
     $x \leftarrow InstanceNorm2d_i(x)$ ;
     $x \leftarrow GELU(x)$ ;
   $a, x \leftarrow PrimaryCapsules(x)$ ;
  foreach  $j \in ConvCapsuleLayers$  do
     $a, x \leftarrow ConvCapsules_j(a, x)$ ;
     $a, x \leftarrow VBRouting_j(a, x)$ ;
   $a, x, \hat{y}_W \leftarrow ClassCapsules(a, x)$ ;
   $a, x \leftarrow ClassRouting(a, x)$ ;
   $\mathcal{F} \leftarrow entities(x)$ ;
  return  $\mathcal{F}, \hat{y}_W$ ;

```

3.2. Multi-task decoders

Starting from the 16-dimensional entities in the capsule feature space \mathcal{F} , we design a decoding module (light orange block in Fig. 2) that allows us to simultaneously retrieve multiple predictions for different tasks from the same feature space \mathcal{F} . Each decoder D_τ in the decoding module is configured as an independent fully connected block, with 0.5 Dropout and GELU activations [11]. We employ no weight sharing or layer sharing across the decoders to enforce the multi-task loss, as explained in section 3.3.

We define different tasks (τ) with different objectives:

- $3D$: minimise the distance between ground truth and predicted 3D joints in 3D space \hat{y}_{3D} ;
- $2D$: as above, but without relying on 3D joints predictions, and rather predicting 2D joints \hat{y}_{2D} as seen from the current viewpoint in camera frame coordinates;
- \mathcal{DM} : reconstruct the depth map $\hat{y}_{\mathcal{DM}}$ of the input RGB image. It is used only in the RGB domain;
- \mathcal{W} Inverse Graphics loss : learn the inverse graphics matrix $\hat{y}_{\mathcal{W}}$ to promote the *de-rendering* of input pixels into isolated capsule entities, as explained in Sec. 3.1, Eq. 5.

For each task $\tau = 3D, 2D, \mathcal{DM}$, a decoder D_τ takes as input the feature space \mathcal{F} and it outputs the predictions $\hat{Y} = [\hat{y}_{3D}, \hat{y}_{2D}, \{\hat{y}_{\mathcal{DM}}\}]$ to the loss function. For \mathcal{W} , the $\hat{y}_{\mathcal{W}}$ matrix is forwarded to the loss function directly from the encoder.

An overview of the capsule decoders is shown in Algorithm 2.

Algorithm 2: Capsule decoders

```

CapsuleDecoders ( $x$ )
  inputs :  $\mathcal{F} = J$  16-dimensional entities
  outputs:  $\hat{Y} = [\hat{y}_{3D}, \hat{y}_{2D}, \{\hat{y}_{\mathcal{DM}}\}]$ 
   $x \leftarrow \mathcal{F}$ ;
  foreach  $i \in Y$  do
     $x \leftarrow \text{Dropout}_{0,5}(x)$ ;
     $x \leftarrow \text{Linear}(x)$ ;
     $\hat{y}_i \leftarrow \text{GELU}(x)$ ;
  return  $\hat{Y} = [\hat{y}_{3D}, \hat{y}_{2D}, \{\hat{y}_{\mathcal{DM}}\}]$ ;

```

3.3. Self-balancing multi-task loss

Tasks are associated to the different input domains, as follows:

	$3D$	$2D$	\mathcal{DM}	\mathcal{W}
Depth	✓	✓		✓
RGB	✓	✓	✓	✓

Each task is assigned a loss \mathcal{L}_τ , defined as:

- $\mathcal{L}_{2D}, \mathcal{L}_{3D}$: Mean Square Error (MSE) loss for the $3D$ and $2D$ joints prediction tasks.
- $\mathcal{L}_{\mathcal{DM}}$: masked L1 loss for the depth estimation task \mathcal{DM} , in the RGB domain, where *mask* is a function that applies the L1 loss only on pixels over a certain depth threshold, to promote the depth estimation over non-background areas.

- $\mathcal{L}_{\mathcal{W}}$: inverse graphics loss \mathcal{W} , which role is to enforce invertibility for the capsule weight matrices. The notation $\|\cdot\|_F$ defines the Frobenius norm of a matrix.

$$\mathcal{L}_\tau = \begin{cases} \mathcal{L}_{2D,3D} = \frac{1}{BS} \sum_{i=0}^{BS} (y_i - \hat{y}_i)^2 \\ \mathcal{L}_{\mathcal{DM}} = \frac{\sum_{i=0}^{BS} [mask |y_i - \hat{y}_i| + |y_i - \hat{y}_i|]}{2 * BS} \\ \mathcal{L}_{\mathcal{W}} = \|\hat{y}_{\mathcal{W}} W_{ij}\|_F \end{cases} \quad (6)$$

Considering \mathcal{T} as the set of the employed tasks τ , the overall balanced loss for all the tasks is expressed as:

$$\mathcal{L} = \sum_{\tau \in \mathcal{T}} (s_\tau + e^{-s_\tau} \mathcal{L}_\tau) \quad (7)$$

where $s_\tau = [s_{3D}, s_{2D}, s_{\mathcal{DM}}, s_{\mathcal{W}}]$ are the trainable weights associated with each loss in \mathcal{T} , initialised to 1 in algorithm 1, and \mathcal{L}_τ is each loss of the enabled decoders, as defined in Eq. 6.

4. Experiments

4.1. Datasets

ITOP Dataset of depth images. The ITOP dataset [9] contains depth images from top and front view. The training split and the test split consist of 40k and 10k images, respectively. The depth images display 15 videos of 20 actors in a constrained setting. The dataset is recorded using two Axus Xtion Pro cameras. The 3D skeleton model consists of 15 joints.

PanopTOP31K dataset of depth and RGB images. The PanopTOP dataset [?] consists of 31k top-view and 31k front view images coming from video sequences of 24 different actors, available both in the RGB and depth domain, for a total of 68k images. The ground truth 3D skeleton consists of 19 joints.

4.2. Evaluation metrics

Following the works of [9, 22, 36], we choose the mean average precision (mAP) as the evaluation metric for the depth domain. It is defined as the percentage of all predicted joints which fall in an interval smaller than 0.10 meters. In the RGB domain, we use the Mean Per Joint Position Error (MPJPE) in millimeters as in many HPE works [2, 32, 23].

4.3. Implementation details

Our network is trained in an end-to-end fashion using Pytorch Lightning. Input images are normalized in the interval $[0, 1]$ with a resolution of 256x256 pixels for depth

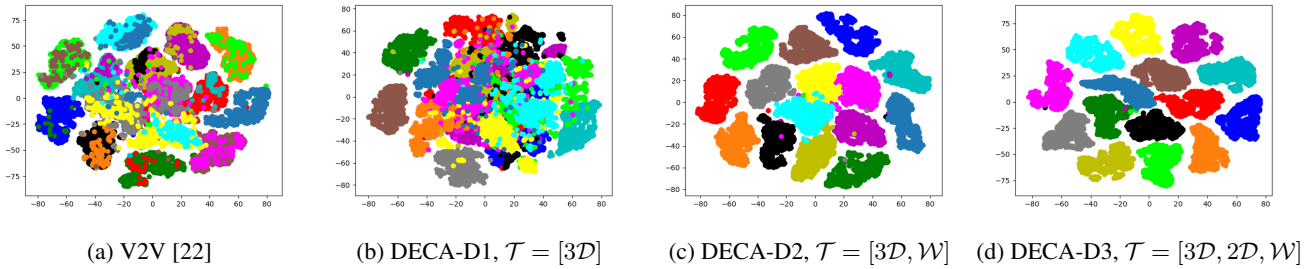


Figure 3: 2D representation on the 16-dimensional latent space obtained using t-SNE [33]. Each dot corresponds to an entity E_{jt} representing a joint jt of the skeleton from the test set of ITOP [9]. V2V network [22] relies on CNNs, thus is not able to cluster together samples corresponding to the same entity (a). When trained to satisfy only the 3D prediction constraint our DECA-D1 network performs slightly better than V2V (b). The 15 clusters, corresponding to the 15 joints of the skeleton model, are clearly distinguishable in DECA-D2 (c) and DECA-D3 (d), with (d) displaying better cluster separation and fewer outliers.

images and 256x256 pixels for RGB ones. We do not perform any augmentations on the input datasets. The batch size is set to 128 for ITOP and 128 for PanopTOP31K. We initialize the weights with the Xavier initialization [6]. The learning rate is set to $1e^{-5}$, the weight decay is set to 0, and Adam is the optimizer of choice. We train our network for 20 epochs on the ITOP dataset and 15 epochs on PanopTOP31K.

4.4. Feature space entities and ablation study

We report experiments on the top-view of the ITOP dataset [9] to validate the 3D representation provided by our network and to show how the multi-tasks decoder influences the overall performances.

To do so, we deploy 4 configurations, 3 on depth data and 1 on RGB data, with different sets of tasks \mathcal{T} of our method:

- DECA-D1, with $\mathcal{T} = [3D]$
- DECA-D2, with $\mathcal{T} = [3D, \mathcal{W}]$
- DECA-D3, with $\mathcal{T} = [3D, 2D, \mathcal{W}]$
- DECA-R4, with $\mathcal{T} = [3D, 2D, \mathcal{DM}, \mathcal{W}]$

where the letter D or R indicates the depth or RGB domains, and the number defines how many tasks are assigned to the network. Since we are evaluating the performances on the 3D HPE, the $\tau = [3D]$ is used for all the different configurations.

Loss effectiveness analysis. The results are reported in the last 3 columns of Table 1. As shown in the Table, increasing the number of tasks in \mathcal{T} generally leads to an increase in the network’s performances. DECA-D1 already achieves similar results to the state-of-the-art, thanks to the CapsNets’ capability to interpret the geometrical nature of

the input data. When the inverse graphics loss \mathcal{W} is employed (DECA-D2 and DECA-D3), the enforced invertibility of the weights matrix leads to an immediate gain in performances. In DECA-D3, the introduction of the $2D$ loss leads to an additional improvement in terms of accuracy. Hence, we argue that the network performances improve when more tasks are given because we achieve a better representation of the entities in the latent space.

Latent space analysis. To analyze the latent space, we use the features of the test set extracted after the capsule modules. Each feature $f \in \mathcal{F}$ is linearised to obtain a vector of length L_{feat} . At this stage, each entity E_{jt} corresponding to each joint jt is defined by dividing each feature vector by the number of joints, resulting in vectors of length $\frac{L_{feat}}{\#of\ joints}$. For visualisation purposes, we use t -SNE [33] to project the entities on a 2-dimensional space. The results are displayed in Fig. 3. We compare our latent space against the publicly available version of the V2V [22] encoder/decoder structure. We show how our DECA network can better cluster and separate each entity E_{jt} with respect to V2V. Our solution provides a better organization of the latent space, with bigger inter-class margins and fewer outliers. The latent space organization improves drastically when we employ the $\tau = \mathcal{W}$ task (DECA-D2), thus enforcing the inverse graphics constraint. In DECA-D3 we add the $\tau = 2D$ task. The resulting organization of the latent space improves, thus further establishing a correlation between the growing number of tasks and the improvement in performances.

4.5. Comparison with state-of-the-art methods

Depth data: ITOP dataset. We compare our DECA against common state-of-the-art method for human pose estimation on depth images [28, 37, 3, 9, 8, 22, 36]. The results are reported in Tab. 1. Our DECA outperforms exist-

Body part	ITOP front-view								ITOP top-view									
	RF[28]	RTW[37]	IEF[3]	VI [9]	REN9x6x6[8]	V2V[22]	A2J[36]	DECA-D3	RF[28]	RTW[37]	IEF[3]	VI [9]	REN9x6x6[8]	V2V[22]	A2J[36]	DECA-D1	DECA-D2	DECA-D3
Head	63.80	97.80	96.20	98.10	98.70	98.29	98.54	93.87	95.40	98.40	83.80	98.10	98.20	98.40	98.38	94.41	95.31	95.37
Neck	86.40	95.80	85.20	97.50	99.40	99.07	99.20	97.90	98.50	82.20	50.00	97.60	98.90	98.91	98.91	98.86	99.16	98.68
Shoulders	83.30	94.10	77.20	96.50	96.10	97.18	96.23	95.22	89.00	91.80	67.30	96.10	96.60	96.87	96.26	96.12	97.51	96.57
Elbows	73.20	77.90	45.40	73.30	74.70	80.42	78.92	84.53	57.40	80.10	40.20	86.20	74.40	79.16	75.88	76.86	81.67	84.07
Hands	51.30	70.50	30.90	68.70	55.20	67.26	68.35	56.49	49.10	76.90	39.00	85.50	50.70	62.44	59.35	44.41	45.97	54.33
Torso	65.00	93.80	84.70	85.60	98.70	98.73	98.52	99.04	80.50	68.20	30.50	72.90	98.10	97.78	97.82	99.46	99.70	99.46
Hip	50.80	80.30	83.50	72.00	91.80	93.23	90.85	97.42	20.00	55.70	38.90	61.20	85.50	86.91	86.88	97.84	97.87	97.42
Knees	65.70	68.80	81.80	69.00	89.00	91.80	90.75	94.56	2.60	53.90	54.00	51.60	70.00	83.28	79.66	88.01	88.19	90.84
Feet	61.30	68.40	80.90	60.80	81.10	87.60	86.91	92.04	0.00	28.70	62.40	51.50	41.60	69.62	58.34	79.30	83.53	81.88
Upper Body	-	-	-	84.00	-	-	-	83.03	-	-	-	91.40	-	-	-	78.51	80.60	83.00
Lower Body	-	-	-	67.30	-	-	-	95.30	-	-	-	54.70	-	-	89.96	91.27	91.39	-
Mean	65.80	80.50	71.00	77.40	84.90	88.74	88.00	88.75	47.40	68.20	51.20	75.50	75.50	83.44	80.5	83.85	85.58	86.92

Table 1: Comparison with the state-of the art for ITOP front-view and top-view (metric: 0.1m mAP).

ing methods on the front-view task, improving the accuracy by a wide margin on the more challenging top viewpoint. In general, we consistently perform better than other methods on most of the joints and the average. The gain of our method is particularly large when dealing with the lower body, which is often occluded in the top-view.

Depth data: Viewpoint-equivariant ITOP. We test DECA on the viewpoint transfer task, meaning training on one viewpoint, either top-view or front-view, and testing on the other one, unseen at training time. The comparison against available state-of-the-art methods [28, 37, 3, 9] are reported in Tab. 2. We consistently outperform other methods by a wide margin, thus making a step forward toward viewpoint equivariance. While other methods provide only the best subset of viewpoint transfer results (Tab. 2), omitting entirely the train on top and test on front scenario, we provide results for all the joints and all the viewpoint transfer combinations in Tab. 3. Our DECA achieves better results than the top-most of the other methods on many different joints (e.g. shoulders, lower body). In Tab 3, training DECA on top-view or front-view achieves comparable lower body accuracy. This means that when the network is trained on top view, where the lower body is mostly occluded, it can retrieve the occluded joints from previously unseen front views, and vice versa. This shows how our network has learned the viewpoint as a parameter, and it is thus able to generalize in a similar fashion in all the viewpoint transfer combinations.

RGB data: Viewpoint-equivariant PanopTOP31K. To the best of our knowledge, we are the first to tackle the problem of viewpoint transfer between top-view and front-view in the RGB domain. We report results with training and testing on both seen and unseen viewpoints in Tab. 4. The chosen metric is the mean per-joint projection error (MPJPE). We report results with and without the Procrustes alignment [7] of the predicted poses. It is interesting to notice how DECA can reduce the gap between the same viewpoint results and the results of the viewpoint transfer tasks. In the case of viewpoint transfer, we train on viewpoint A, validate on the same viewpoint A and test on viewpoint B.

Body part	ITOP				
	Train on front, test on top				
	RF [28]	RTW [37]	IEF [3]	VI [9]	DECA-D3
Head	48.10	1.50	47.90	55.60	46.27
Neck	5.90	8.10	39.00	40.90	73.14
Torso	4.70	3.90	41.90	35.00	85.94
Upper Body	19.70	2.20	23.90	29.40	45.00
Full Body	10.80	2.00	17.40	20.40	51.85

Table 2: Comparison with the state-of the art for the ITOP viewpoint transfer task (metric: 0.1m mAP). Training on front-view, validating on front-view, testing on top-view (top-view data is unseen in validation).

Body part	DECA-D3	
	Train on front, test on top	Train on top, test on front
Head	46.27	18.51
Neck	73.14	44.77
Shoulders	69.02	25.18
Elbows	43.87	16.23
Hands	9.41	2.19
Torso	85.94	68.63
Hip	72.15	64.75
Knees	49.31	68.15
Feet	42.46	46.12
Upper Body	45.00	18.81
Lower Body	59.11	60.95
Mean	51.85	38.48

Table 3: DECA-D3 complete results for the ITOP viewpoint transfer tasks (metric: 0.1m mAP). Test data is unseen during validation for both the cases.

4.6. Qualitative results

In Fig. 4 we show some qualitative results from DECA-R4 configuration on RGB data. We deploy our network training and testing on all the possible viewpoint combinations. The network takes as input either the top-view RGB (Fig. 4a) image or the front view (Fig. 4b) one. When trained and tested on the same viewpoint (Fig. 4d, 4e), the network produces similar outputs, thus confirming its abil-

DECA-R4									
Body part	Train on front, test on front		Train on top, test on top		Train on front, test on top		Train on top, test on front		
	No Procrustes	Procrustes	No Procrustes	Procrustes	No Procrustes	Procrustes	No Procrustes	Procrustes	
Neck	4.02	2.37	4.55	2.51	16.02	4.16	8.21	5.06	
Nose	5.66	3.75	6.98	3.89	16.83	7.67	10.72	6.76	
Body Center	0.56	4.63	1.23	3.63	1.01	31.20	0.83	11.59	
Shoulders	4.56	2.76	5.14	3.07	17.43	5.33	8.51	5.35	
Elbows	9.82	7.14	9.64	7.51	29.70	18.52	23.20	15.47	
Hands	13.88	10.82	14.02	12.34	47.01	38.29	36.78	28.25	
Hips	18.75	4.87	2.71	3.89	5.10	30.07	3.64	10.88	
Knees	9.54	5.14	7.59	4.84	52.98	28.65	20.11	9.28	
Feet	11.53	5.08	9.83	5.10	69.18	28.75	26.36	11.07	
Eyes	6.19	4.00	7.44	3.79	19.33	11.00	11.40	7.45	
Ears	5.50	3.73	7.15	3.74	23.56	13.00	11.22	7.16	
Upper Body	6.93	5.21	7.66	5.46	23.69	16.56	15.54	11.60	
Lower Body	7.65	5.03	6.71	4.61	42.42	29.16	16.71	10.41	
Mean	7.16	5.15	7.36	5.19	29.60	20.54	15.91	11.22	

Table 4: DECA-R4 results on the PanopTOP31K RGB dataset, with and without the Procrustes transformation [7] (metric: MPJPE). Tasks: (i) 3D pose estimation from the front and top viewpoints (ii) viewpoint transfer for both front and top views. Test data is unseen during validation for both the viewpoint transfer tasks.

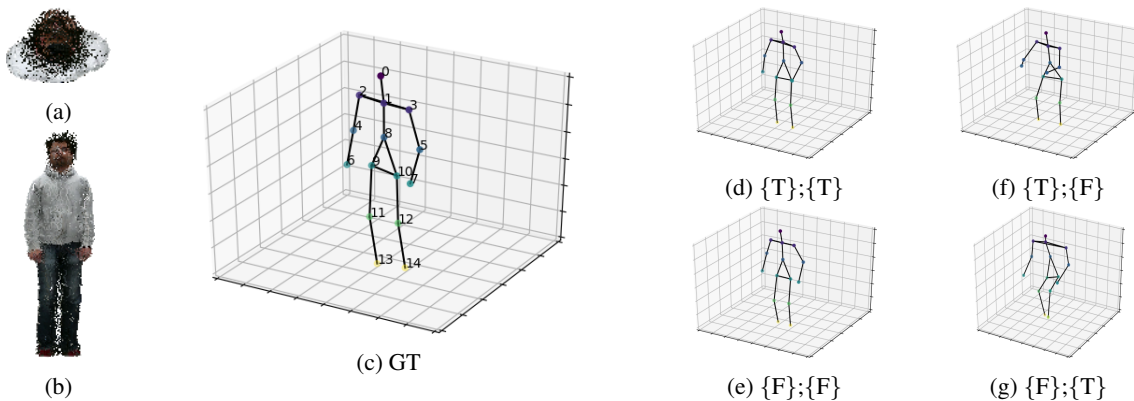


Figure 4: DECA-R4 qualitative results on the PanopTOP31K dataset. On the left (a, b) the types of input accepted by DECA (top-view or front-view). DECA can also accept inputs in the depth domain. In the center (c), the corresponding 3D ground truth. On the right, the possible combinations of training/testing experiments. **T** stands for **top** and **F** stands for **front**. As an example, in (f), $\{T\};\{F\}$ means that DECA has been trained exclusively on **top** data and tested on previously unseen (not even at validation time) **front** data.

ity to deal with the challenging top-view scenario. When training on the top view and testing on the front one (Fig. 4f), the network can accurately retrieve the positions of the lower body joints. DECA can retrieve parts of the body mostly occluded at training time, thus displaying its generalization capabilities. When training on the front view and testing on the top one (Fig. 4g), the network can retrieve the positions of the upper body joints, which are visible in both images but from different perspectives, proving that DECA can internally model the viewpoint.

5. Conclusions

We presented DECA, a deep viewpoint-equivariant method for human pose estimation on single RGB/depth images using capsule autoencoders. We show how CapsNets are better suited to deal with the 3D nature of raw data and how they allow taking a step forward to viewpoint equivariance. We have shown how our method can effectively generalize and achieve state-of-the-art results in both RGB and depth domains, as well as in the viewpoint transfer task. In future work, we aim at improving hands pose estimation and employing matrix capsules on bigger RGB datasets.

References

- [1] Irving Biederman. Recognition-by-components: a theory of human image understanding. *Psychological review*, 94(2):115, 1987.
- [2] Z. Cao, T. Simon, S. Wei, and Y. Sheikh. Realtime multi-person 2d pose estimation using part affinity fields. In *2017 IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 1302–1310, 2017.
- [3] Joao Carreira, Pulkit Agrawal, Katerina Fragkiadaki, and Jitendra Malik. Human pose estimation with iterative error feedback. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 4733–4742, 2016.
- [4] Taco S Cohen, Mario Geiger, Jonas Köhler, and Max Welling. Spherical cnns. *arXiv preprint arXiv:1801.10130*, 2018.
- [5] Nicola Garau, Giulia Martinelli, Piotr Brødka, Niccoló Bisagno, and Nicola Conci. Panoptop: a framework for generating viewpoint-invariant human pose estimation datasets. In *Proceedings of the IEEE/CVF International Conference on Computer Vision Workshops*, 2021.
- [6] Xavier Glorot and Yoshua Bengio. Understanding the difficulty of training deep feedforward neural networks. volume 9 of *Proceedings of Machine Learning Research*, pages 249–256, Chia Laguna Resort, Sardinia, Italy, 13–15 May 2010. JMLR Workshop and Conference Proceedings.
- [7] Colin Goodall. Procrustes methods in the statistical analysis of shape. *Journal of the Royal Statistical Society: Series B (Methodological)*, 53(2):285–321, 1991.
- [8] Hengkai Guo, Guijin Wang, Xinghao Chen, and Cairong Zhang. Towards good practices for deep 3d hand pose estimation. *arXiv preprint arXiv:1707.07248*, 2017.
- [9] Albert Haque, Boya Peng, Zelun Luo, Alexandre Alahi, Serena Yeung, and Li Fei-Fei. Towards viewpoint invariant 3d human pose estimation. In *European Conference on Computer Vision*, pages 160–177. Springer, 2016.
- [10] Kaiming He, Georgia Gkioxari, Piotr Dollár, and Ross Girshick. Mask r-cnn. In *Proceedings of the IEEE international conference on computer vision*, pages 2961–2969, 2017.
- [11] Dan Hendrycks and Kevin Gimpel. Gaussian error linear units (gelus). *arXiv preprint arXiv:1606.08415*, 2016.
- [12] Geoffrey E Hinton, Alex Krizhevsky, and Sida D Wang. Transforming auto-encoders. In *International conference on artificial neural networks*, pages 44–51. Springer, 2011.
- [13] Geoffrey E Hinton, Sara Sabour, and Nicholas Frosst. Matrix capsules with EM routing. In *International Conference on Learning Representations*, 2018.
- [14] Catalin Ionescu, Dragos Papava, Vlad Olaru, and Cristian Sminchisescu. Human3.6m: Large scale datasets and predictive methods for 3d human sensing in natural environments. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 36(7):1325–1339, jul 2014.
- [15] Hanbyul Joo, Tomas Simon, Xulong Li, Hao Liu, Lei Tan, Lin Gui, Sean Banerjee, Timothy Scott Godisart, Bart Nabbe, Iain Matthews, Takeo Kanade, Shohei Nubuhara, and Yaser Sheikh. Panoptic studio: A massively multiview system for social interaction capture. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 2017.
- [16] M Esat Kalfaoglu, Sinan Kalkan, and A Aydin Alatan. Late temporal modeling in 3d cnn architectures with bert for action recognition. In *European Conference on Computer Vision*, pages 731–747. Springer, 2020.
- [17] Isinsu Katircioglu, Bugra Tekin, Mathieu Salzmann, Vincent Lepetit, and Pascal Fua. Learning Latent Representations of 3D Human Pose with Deep Neural Networks. *International Journal of Computer Vision*, 126(12):1326–1341, 2018.
- [18] Muhammed Kocabas, Nikos Athanasiou, and Michael J. Black. Vibe: Video inference for human body pose and shape estimation. In *The IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, June 2020.
- [19] Nikos Kolotouros, Georgios Pavlakos, Michael J Black, and Kostas Daniilidis. Learning to reconstruct 3d human pose and shape via model-fitting in the loop. In *Proceedings of the IEEE International Conference on Computer Vision*, pages 2252–2261, 2019.
- [20] Adam Roman Kosiorek, Sara Sabour, Yee Whye Teh, and Geoffrey Hinton. Stacked capsule autoencoders. 2019.
- [21] J. Liu, H. Ding, A. Shahroudy, L. Duan, X. Jiang, G. Wang, and A. C. Kot. Feature boosting network for 3d pose estimation. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 42(2):494–501, 2020.
- [22] Gyeongsik Moon, Ju Yong Chang, and Kyoung Mu Lee. V2v-posenet: Voxel-to-voxel prediction network for accurate 3d hand and human pose estimation from a single depth map. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 5079–5088, 2018.
- [23] Iván Ramírez, Alfredo Cuesta-Infante, Emanuele Schiavi, and Juan José Pantrigo. Bayesian capsule networks for 3d human pose estimation from single 2d images. *Neurocomputing*, 379:64 – 73, 2020.
- [24] Fabio Ribeiro, Georgios Leontidis, and Stefanos Kollias. Capsule routing via variational bayes. *Proceedings of the AAAI Conference on Artificial Intelligence*, 34:3749–3756, 04 2020.
- [25] Grégory Rogez, Philippe Weinzaepfel, and Cordelia Schmid. LCR-Net++: Multi-person 2D and 3D Pose Detection in Natural Images. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 2019.
- [26] Sara Sabour, Nicholas Frosst, and Geoffrey E. Hinton. Dynamic routing between capsules. In *Proceedings of the 31st International Conference on Neural Information Processing Systems, NIPS’17*, page 3859–3869, Red Hook, NY, USA, 2017. Curran Associates Inc.
- [27] Marta Sanzari, Valsamis Ntouskos, and Fiora Pirri. Bayesian image based 3d pose estimation. In Bastian Leibe, Jiri Matas, Nicu Sebe, and Max Welling, editors, *Computer Vision – ECCV 2016*, pages 566–582, Cham, 2016. Springer International Publishing.
- [28] Jamie Shotton, Andrew Fitzgibbon, Mat Cook, Toby Sharp, Mark Finocchio, Richard Moore, Alex Kipman, and Andrew Blake. Real-time human pose recognition in parts from single depth images. In *CVPR 2011*, pages 1297–1304. Ieee, 2011.

- [29] Sebastian Starke, He Zhang, Taku Komura, and Jun Saito. Neural state machine for character-scene interactions. *ACM Trans. Graph.*, 38(6):209–1, 2019.
- [30] Bugra Tekin, Pablo Márquez-Neila, Mathieu Salzmann, and Pascal Fua. Learning to fuse 2d and 3d image cues for monocular body pose estimation. In *Proceedings of the IEEE International Conference on Computer Vision*, pages 3941–3950, 2017.
- [31] Yan Tian, Wei Hu, Hangsen Jiang, and Jiachen Wu. Densely connected attentional pyramid residual network for human pose estimation. *Neurocomputing*, 347:13 – 23, 2019.
- [32] Denis Tome, Chris Russell, and Lourdes Agapito. Lifting from the deep: Convolutional 3d pose estimation from a single image. 07 2017.
- [33] Laurens Van der Maaten and Geoffrey Hinton. Visualizing data using t-sne. *Journal of machine learning research*, 9(11), 2008.
- [34] K. Wang, L. Lin, C. Jiang, C. Qian, and P. Wei. 3d human pose machines with self-supervised learning. *IEEE Transactions on Pattern Analysis & Machine Intelligence*, 42(05):1069–1082, may 2020.
- [35] Keze Wang, Shengfu Zhai, Hui Cheng, Xiaodan Liang, and Liang Lin. Human pose estimation from depth images via inference embedded multi-task learning. In *Proceedings of the 24th ACM international conference on Multimedia*, pages 1227–1236, 2016.
- [36] Fu Xiong, Boshen Zhang, Yang Xiao, Zhiguo Cao, Taidong Yu, Joey Tianyi Zhou, and Junsong Yuan. A2j: Anchor-to-joint regression network for 3d articulated pose estimation from a single depth image. In *Proceedings of the IEEE International Conference on Computer Vision*, pages 793–802, 2019.
- [37] Ho Yub Jung, Soochahn Lee, Yong Seok Heo, and Il Dong Yun. Random tree walk toward instantaneous 3d human pose estimation. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 2467–2474, 2015.
- [38] X. Zhou, M. Zhu, S. Leonardos, K. G. Derpanis, and K. Daniilidis. Sparseness meets deepness: 3d human pose estimation from monocular video. In *2016 IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 4966–4975, 2016.