Temporal Representation Learning on Monocular Videos for 3D Human Pose Estimation

Sina Honari[1], Victor Constantin[1], Helge Rhodin[2], Mathieu Salzmann[1], Pascal Fua[1]

[1]CVLab, EPFL, Lausanne, Switzerland

[2]Imager Lab, UBC, Vancouver, Canada

Abstract—In this paper we propose an unsupervised feature extraction method to capture temporal information on monocular videos, where we detect and encode subject of interest in each frame and leverage contrastive self-supervised (CSS) learning to extract rich latent vectors. Instead of simply treating the latent features of nearby frames as positive pairs and those of temporally-distant ones as negative pairs as in other CSS approaches, we explicitly disentangle each latent vector into a time-variant component and a time-invariant one. We then show that applying contrastive loss only to the time-variant features and encouraging a gradual transition on them between nearby and away frames while also reconstructing the input, extract rich temporal features, well-suited for human pose estimation. Our approach reduces error by about 50% compared to the standard CSS strategies, outperforms other unsupervised single-view methods and matches the performance of multi-view techniques. When 2D pose is available, our approach can extract even richer latent features and improve the 3D pose estimation accuracy, outperforming other state-of-the-art weakly supervised methods.

Index Terms—Temporal Feature Extraction, Unsupervised Representation Learning, Contrastive Learning, 3D Human Pose.

✦

1 INTRODUCTION

While supervised body pose estimation is rapidly becoming a mature field, the bottleneck remains the availability of sufficiently large training datasets containing in-the-wild images or depicting all kinds of human motions. While some datasets [1], [2] provide relatively large-scale annotations in controlled environments, the set of motions is limited and does not capture the vast diversity of possible human poses, such as the ones observed in athletics or diving. An effective way to address this is to leverage unsupervised data to learn a low-dimensional representation of pose-related features. Then, it only takes very little annotated data to train a regressor to predict 3D poses from this representation. Key to the success is a good unsupervised learning objective, for which many existing techniques exploit the availability of multi-view footage [3], [4], [5], [6], [7]. However, capturing data with multiple cameras increases complexity because it requires synchronizing the cameras and calibrating them to obtain the camera parameters that several of these models [3], [5], [6], [7] require during training.

In this paper, we introduce an alternative unsupervised representation learning strategy using videos acquired by a *single* RGB camera. To this end, we build on the idea of contrastive self-supervised (CSS) learning [8], [9], [10]. For any given sample, CSS aims to maximize the similarity to a positive sample and the dissimilarity to a negative one. In our context, one way to do so is to treat a video frame close to the one of interest as positive, and a temporally-distant one as negative [11], [12]. Unfortunately, this simple strategy is insufficient because it does not account for the fact that, when someone moves, some features, such as garment appearance, remain constant over time, even as the person's pose changes. As a result, contrastive learning should be applied only to the features changing across time and not to all features. To overcome this, we learn a representation that explicitly separates each latent vector into a *time-variant* component and a *time-invariant* one and apply the contrastive loss only to the *time-variant* component.

Our approach, depicted in Fig. 1, differs from typical CSS strategies in four ways: First, we apply the contrastive formulation only to the time-variant component, as opposed to the entire latent space. Second, we encourage a gradual transition of the time-variant component from temporally close to temporally distant frames. Third, we perform image re-synthesis by mapping the latent vectors back to images. Finally, when observing subjects in free fall such as divers, we incorporate a gravity-encoding prior to help with the detection of the subject. Our experiments demonstrate that these four components help the model learn in an unsupervised manner a rich, low-dimensional representation of human motions from ordinary video sequences acquired using standard monocular RGB cameras. We also show when 2D labels are available, one can use them to extract richer time-variant features, such that they contain more relevant pose information.

We demonstrate the benefits of our approach for pose estimation on five real-world RGB datasets. It outperforms other unsupervised monocular approaches and yields results on par with unsupervised multi-view setups, while being much easier to deploy. In particular, it can exploit both monocular and multi-view videos when available, which a purely multi-view setup cannot, and can handle changing backgrounds. When using 2D labels, our approach outperforms other weakly supervised methods. The disentangled latent representation also lends itself for editing operations, such as pose and appearance transfer.
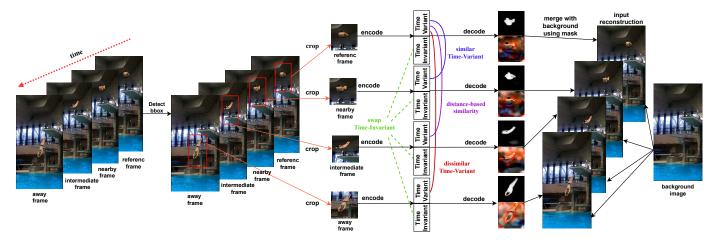
sina.honari@epfl.ch

Fig. 1: Approach. Given a reference frame, we sample a temporally close, far, and intermediate—between close and far —frames. The model first detects the person in all frames using a spatial transformer network. The detected regions are then cropped and encoded into *time-variant* and *time-invariant* components. The time-variant of the reference frame is made similar to that of the nearby frame and dissimilar to that of the away frame, while a distance-dependent similarity or dissimilarity is enforced between the reference and intermediate frames. To promote invariance of the time-invariant components, they are swapped with other frames in the video. The model then decodes each frame into a segmentation mask and an image. The decoded image is then merged with the background using the bounding box location and the decoded mask. No annotations are used (neither for detection, nor for segmentation) and the same network weights are used for all frames. At test time, the model can process a single frame.

2 RELATED WORK

Unsupervised and self-supervised learning [13], [14] techniques have received considerable attention in recent years as a means to decrease the amount of annotated data required to train deep networks. We briefly review some of the methods most closely related to our work.

2.1 Single-View Approaches

Unsupervised learning for general video processing has been extensively researched. For example, LSTMs have been used to learn representations of video sequences [15]. Similarly, optical-flow and color diversity across frames have been leveraged to learn latent features corresponding to motion and appearance [16]. Multiple papers [17], [18], [19], [20] use the temporal order of frames to learn latent features from videos. In particular, the models of [17], [19] classify whether a set of frames have correct ordering or not, whereas those of [18], [20] predict the order of the frames or clips as a classification task. In [21], instead of relying on temporal ordering, a model is trained to predict the pace of a video subsampled at different rates. The aforementioned work also leverages a CSS loss across clips, where positive pairs are clips from the same video and negative pairs are clips from different ones. None of these techniques, however, were designed for pose estimation. For example, predicting temporal order [17], [18], [19], [20] is best suited for activity recognition, for which only the order of the observed event truly matters. Furthermore, the contrastive learning strategy used in [21] aims to extract features that are invariant across video frames. In contrast to these works, to predict a continuously changing body pose, we need to capture variations between the frames.

Several works have nonetheless focused on extracting time-varying features, yet in different contexts. Harley *et al.* [22] use RGB and depth from n time-steps to map 2.5D features to a 3D latent space, and then use the camera pose at time-step $n + 1$ to predict the 2D latent features for that time-step. While [22] takes depth as input in addition to RGB, other works [10], [11], [12] leverage CSS, with [12] and [11] using it for only single-view RGB videos. However, they do no use any decoder to reconstruct the images from the latent features, thus opening the door to losing valuable information about the observed context. In addition, these approaches do not split the latent space, which is equivalent to only learning the time-variant component. We will demonstrate that for the pose estimation task our encoding/decoding approach with latent split that applies the contrastive loss only to the time-variant component yields richer features.

Unsupervised learning methods have also been developed for 2D keypoint estimation [23], [24], [25], [26], [27], [28], [29]. Some of these [24], [25], [26], [28], [29] learn separate features for pose and appearance, as we do. However, being designed for 2D keypoint estimation, they either explicitly detect 2D keypoints or generate heatmaps, and extending them to 3D pose estimation remains unadressed. Furthermore, whereas we learn to detect our subject without supervision, these methods rely on ground-truth bounding box locations.

2.2 Multi-View Approaches

The difficulty of performing unsupervised learning using a single camera has led to the development of several multi-view approaches [5], [7], [12], [30], [31], [32]. While some of them are semi-supervised [6], [33], here, we focus only on fully unsupervised learning, where no pose labels are used when training the feature extraction model.

Among the unsupervised methods, the strategy of [7], [32] is closest in spirit to ours. It uses a multi-camera setup to learn a low-dimensional latent vector that disentangles pose from appearance. An effective mapping from this latent space to 3D poses can then be trained using a small amount of annotated data. In contrast to this approach, we aim to learn a latent representation from videos acquired by a *single* camera. To do so, we leverage contrastive learning and propose multiple innovations compared to the standard CSS approaches.
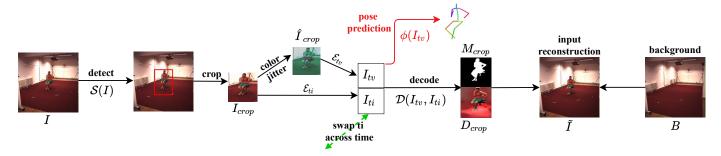
Fig. 2: Model architecture. Using spatial transformer network (STN) $\mathcal{S}$, a bounding box is detected and cropped on an input image $\mathbf{I}$. The cropped image, $\mathbf{I}_{\mathrm{crop}}$, is then passed to an encoder $\mathcal{E}_{ti}$ that predicts time-invariant $\mathbf{I}_{ti}$ component. This component is swapped across time with other frames to enforce disentanglement. A color jittered variant of the crop $\hat{\mathbf{I}}_{\mathrm{crop}}$ is passed to another encoder $\mathcal{E}_{tv}$ to predict time-variant $\mathbf{I}_{tv}$ features. Both encoders share parameters, except in the output layer. The combined features are then passed to a decoder $\mathcal{D}$ that outputs a mask $\mathbf{M}_{\mathrm{crop}}$ and RGB crop $\mathbf{D}_{\mathrm{crop}}$. These two outputs are then put to the image resolution using inverse STN operation and then merged with background $\mathbf{B}$ to output the reconstruct image $\tilde{\mathbf{I}}$. This constitutes the self-supervised model, which is then frozen. In the second phase (shown in red), predicted $\mathbf{I}_{tv}$ is passed to a shallow network Φ to predict the 3D pose.

2.3 Disentanglement

The goal of disentanglement is to find independent and interpretable factors of variation [34], [35], [36], [37], [38], [39] that makes the data easier to process or understand. For example, [40] encodes facial images into disentangled components such as illumination, head yaw, and pitch angles. On the other hand, [39] disentangles encoded image features from camera's view point to a canonical representation that can be more easily processed by the model for 3D pose estimation. In [38] it has been used to separate appearance and motion in a generative model. One closely related approach is DrNet [41], which disentangles each frame into pose and content. It uses a discriminator that distinguishes pose component extracted from the same video from the ones extracted from different videos and encourages the encoder to produce latent representations that do not carry any information about whether they come from the same or different videos. This, however, fails in cases where the two videos feature different kinds of activities—for example, running in one and sitting on the floor in the other. In this scenario, the pose still indicates if the frames come from the same video or not, which precludes disentanglement.

Another closely related approach is [6] that use a multi-view setup to disentangle pose from appearance, by rotating the pose from one view to another while swapping the appearance of a subject with other frames of the same subject in the dataset. Contrary to these approaches, we split features on monocular videos into time-variant and time-invariant components, by using contrastive learning to ensure time-variant contains only temporal features that change across time, while also making sure it does not capture identity information, such as appearance, through input jittering to time-variant and swapping of time-invariant across time.

3 METHOD

Our goal is to extract rich latent features, without any supervision, from single-camera videos featuring people. We aim to capture information only about the foreground subjects and to leverage the similarity or dissimilarity of frames given their distance in time. A key observation is that global object appearance tends to remain constant over time while pose usually changes from frame to frame, which is something our latent vectors should reflect.

Encoding the entire image is expected to yield inferior performance because the resulting latent variables would also capture irrelevant background information. To address this, we guide the model to detect the subject of interest in an unsupervised manner, using image reconstruction. We then only encode features of the foreground subject, to which we apply the contrastive loss.

We train an encoder that splits latent features into *time-variant* and *time-invariant* components. The latter capture features that remain consistent over time, such as a person's clothing and appearance, while the former models the time-changing elements in each frame, such as body pose. We then apply the contrastive loss only to the time-variant elements. Without this split, the latent components learned using CSS would struggle to reliably differentiate temporally-distant frames, because they would still share information, such as identity. Finally, we take time-invariant and time-variant components to re-synthesize the input. We show in the experiments that decoding images help the model learn much richer features, as the model has to capture all relevant information on the subject to reconstruct the input. Fig. 1 summarizes our approach.

3.1 Model Architecture and Training

We tackle real-world applications where the foreground subject is small and occupies only a fraction of the image. To this end, we use a neural network architecture with an attention mechanism to separate foreground from background. As shown in Fig. 2, given an input image $\mathbf{I}$, a spatial transformer network (STN) [42] $\mathcal{S}$ extracts four parameters that define the subject's bounding box, two non-homogeneous scales s^x, s^y and its center coordinates u^x, u^y, yielding $\mathcal{S}(\mathbf{I}) = (s^x, s^y, u^x, u^y)$. These parameters are used to crop an image patch $\mathbf{I}_{\mathrm{crop}}$. A color-jittered version of the crop $\hat{\mathbf{I}}_{\mathrm{crop}}$ is passed to the time-variant encoder $\mathcal{E}_{tv}$ such that $\mathcal{E}_{tv}(\hat{\mathbf{I}}_{\mathrm{crop}}) = \mathbf{I}_{tv}$, with $\mathbf{I}_{tv}$ the time-variant latent component. A crop of another frame $\mathbf{I}'_{\mathrm{crop}}$ of the same video is passed to a time-invariant encoder $\mathcal{E}_{ti}$ such that $\mathcal{E}_{ti}(\mathbf{I}'_{\mathrm{crop}}) = \mathbf{I}'_{ti}$, where $\mathbf{I}'_{ti}$ is the time-invariant component of the latent vector for image $\mathbf{I}'$. Both encoders share parameters except at the output layer. A decoder $\mathcal{D}$ then takes $\mathbf{I}_{tv}$ and $\mathbf{I}'_{ti}$ as inputs and decodes them into a mask $\mathbf{M}_{\mathrm{crop}}$ and an RGB crop $\mathbf{D}_{\mathrm{crop}}$. The mask and RGB crop are then used to re-synthesize a full-resolution image by using the

inverse transformation of the STN and merging the result with a background image $\mathbf{B}$. This operation can be written as

$$\tilde{\mathbf{I}} = \mathbf{M} \times \mathbf{D} + (1 - \mathbf{M}) \times \mathbf{B} , \qquad (1)$$

where $\tilde{\mathbf{I}}$ is the reconstructed image. Fig. 2 depicts this process. Note that no labels are used here, neither for detection, nor for segmentation. We train the networks $\mathcal{S}$, $\mathcal{E}$, and $\mathcal{D}$ such that they jointly accomplish three goals:

- Encoding information that relates only to the foreground object (Section 3.1.1);
- Disentangling the time-variant and time-invariant components (Section 3.1.3);
- Encoding time-variant components such that they are more similar for poses that are close in time than for temporally distant ones (Section 3.1.2).

3.1.1 Reconstruction Loss

To compare the reconstructed image from Eq. (1) with the original one, we use the L2 loss on the image pixels as well as on the features extracted from layers 1 to 3 of the ResNet18 [43] model. This can be written as

$$\mathcal{L}_{\text{reconst}} = \lambda \|\mathbf{I} - \tilde{\mathbf{I}}\|_2^2 + \rho \sum_{l=1}^{3} \|\text{Res}_l(\mathbf{I}) - \text{Res}_l(\tilde{\mathbf{I}})\|_2^2 . \quad (2)$$

As discussed above, we use a background image when re-synthizing the input image, hence this reconstruction loss encourages our latent representation to focus on the foreground object.

3.1.2 Contrastive Loss

To encourage time-variant components that are close in time to be close to each other and those that are far in time not to be, we rely on a contrastive loss. However, we diverge from standard CSS formulations [8], [10], [11], [12] in the two following aspects: 1) Instead of applying the contrastive loss to the entire latent space, we enforce it only on the time-variant features, as only part of the latent features should evolve over time. 2) We enforce a gradual transition of the latent representation from nearby frames to away frames. Below, we discuss this in more detail, starting from the standard CSS formulation.

Given a reference input $\mathbf{I}$, a positive sample $\mathbf{I}^+$, and $n-1$ negative samples $\mathbf{I}_k^-$ for $1 \leq k < n$, standard CSS approaches [8], [10], [11] minimize the loss

$$\mathcal{L}_{\text{CSS}} = -\log \frac{\exp(\text{sim}(\mathcal{F}(\mathbf{I}), \mathcal{F}(\mathbf{I}^+))/\tau)}{\sum_{k=1}^{n} \exp(\text{sim}(\mathcal{F}(\mathbf{I}), \mathcal{F}(\mathbf{I}_k^{-/+}))/\tau)} , \quad (3)$$

where $\mathbf{I}_k^{-/+}$ indicates a set of one positive and $n-1$ negative examples. $\mathcal{F}(\mathbf{I})$ represents the features encoded by the network $\mathcal{F}$ given frame $\mathbf{I}$ as input and $\text{sim}(\mathbf{m}, \mathbf{n}) = (\mathbf{m}^T \mathbf{n})/(\|\mathbf{m}\|\|\mathbf{n}\|)$. In [12], the contrastive loss is defined instead as a triplet loss

$$\mathcal{L}_{\text{CSS}} = \left[\|\mathcal{F}(\mathbf{I}) - \mathcal{F}(\mathbf{I}^+)\|^2 + \beta \right] < \|\mathcal{F}(\mathbf{I}) - \mathcal{F}(\mathbf{I}^-)\|^2 , \quad (4)$$

with $\beta > 0$, meaning that the squared L2 distance between a positive pair should be smaller than that between a negative one by at least β. In any event, neither definitions of $\mathcal{L}_{\text{CSS}}$ above define the loss relative to the temporal distance of the frames. This is what we change by introducing a distance-based similarity loss.

Distance-based Similarity Loss (DSL). Our goal is to compare every pair of samples directly and to define a loss that enforces

similarity or dissimilarity between any two samples depending on their temporal proximity.

Let $\mathbf{I}_{tv}^m$ and $\mathbf{I}_{tv}^n$ be the time-variant component of the vectors extracted from two images captured at times t_m and t_n. If $|t_m - t_n|$ is small, we would like $\text{sim}(\mathbf{m}, \mathbf{n})$, where $\text{sim}()$ is the similarity function used in Eq. (3), to be close to one. By contrast, if $|t_m - t_n|$ is larger than a threshold distance $d_{\max}$, we would like $\text{sim}(\mathbf{m}, \mathbf{n})$ to be zero. To this end, we introduce the loss function

$$\mathcal{L}_{\text{DSL}}(\mathbf{I}_{tv}^m, \mathbf{I}_{tv}^n) = \qquad (5)$$
$$\begin{cases} (1 - \frac{2d}{d_{\max}}) \, |1 - \text{sim}(\mathbf{I}_{tv}^m, \mathbf{I}_{tv}^n)|, & \text{if } 0 \leq d \leq d_{\max}/2 , \\ \min(\frac{2d}{d_{\max}} - 1, 1) \, |\text{sim}(\mathbf{I}_{tv}^m, \mathbf{I}_{tv}^n)|, & \text{otherwise,} \end{cases}$$

where $d = |t_m - t_n|$ is the temporal distance between the two frames. As shown in Fig. 3, $\mathcal{L}_{\text{DSL}}$ is piecewise linear in terms of $\text{sim}(\mathbf{I}_{tv}^m, \mathbf{I}_{tv}^n)$. It favors similarity for $0 \leq d \leq d_{\max}/2$, and dissimilarity for $d > d_{\max}/2$, with similarity decreasing and dissimilarity increasing as d increases.
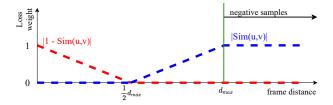

Fig. 3: Distance-based Similarity Loss (DSL). The red line shows the **weight**, $(1 - \frac{2d}{d_{\max}})$, applied to $|\text{sim}(\mathbf{m}, \mathbf{n}) - 1|$ when two frames are close (similarity measure) and the blue line shows the **weight**, $\min(\frac{2d}{d_{\max}} - 1, 1)$, applied to $|\text{sim}(\mathbf{m}, \mathbf{n})|$ when the frames are away (dissimilarity measure). Both weights depend on the temporal distance between the frames. As the frame distance increases, the similarity weight is reduced, eventually going to zero at $d_{\max}/2$, and the dissimilarity weight increases, eventually reaching one at $d_{\max}$, the frame from which negative samples are taken.

In practice, to learn time-variant components in our framework, for any given reference input frame $\mathbf{I}^r$, we sample three other frames, a nearby one $\mathbf{I}^n$ (positive sample), a far away one $\mathbf{I}^a$ (negative sample taken at least at $d_{\max}$), and an intermediate one, $\mathbf{I}^{in}$ (acquired after $\mathbf{I}^n$ and before $\mathbf{I}^a$). We then compute

$$\mathcal{L}_{\text{DSL}}^{\text{all}} = \mathcal{L}_{\text{DSL}}(\mathbf{I}_{tv}^r, \mathbf{I}_{tv}^n) + \mathcal{L}_{\text{DSL}}(\mathbf{I}_{tv}^r, \mathbf{I}_{tv}^a) + \mathcal{L}_{\text{DSL}}(\mathbf{I}_{tv}^r, \mathbf{I}_{tv}^{in}) . \quad (6)$$

The above loss allows us to compare the reference frame with similar, dissimilar, and intermediate frames directly.

3.1.3 Disentanglement

To force the network to disentangle the latent into time-variant and time-invariant components, we apply three things: 1) we swap the time-invariant ones of different frames of the same video during training. This prevents the time-invariant components from encoding information about the subject's pose, 2) we use a color-jittered variant of the cropped image $\hat{\mathbf{I}}_{\text{crop}}$ to predict the time-variant features $\mathbf{I}_{tv}$. This forces $\mathcal{E}_{tv}$ to focus on the pose and not the texture when predicting $\mathbf{I}_{tv}$, and 3) we use contrastive loss to make temporally distant time-invariant dissimilar. Note that points 1 and 3 enforce the time-invariant component to capture the identity information, as we will later show in Section 4.4.3. Also note that reconstruction is needed to make points 1 and 2 impactful, as otherwise the model does not encode relevant information.
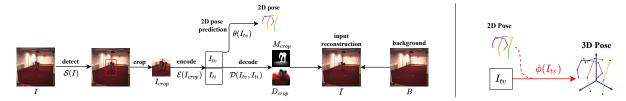
Fig. 4: Using Weak Labels. (left)When 2D labels are available in the self-supervised phase, we apply 2D regression on time-variant features in addition to image reconstruction. (right) In the 2nd phase, $\mathbf{I}_{tv}$ features and optionally 2D pose can be used to predict the 3D pose.

3.2 Gravity as a Supervisory Signal

The work reported here was initially motivated by a project that involves modeling competitive divers. Because the divers are in free fall, we can leverage Newtonian mechanics as follows. The diver's center of gravity (CG) is subject to gravitational acceleration. Ignoring the air resistance that remains small until the diver reaches the water, $\mathbf{p}(t)$, the CG location at time t, can be written as $\mathbf{p}(t) = 0.5\mathbf{g}t^2 + \mathbf{v}_0 t + \mathbf{p}_0$, where $\mathbf{g}$ is the gravitational vector, $\mathbf{v}_0$ the CG initial speed, and $\mathbf{p}_0$ its initial location. We know neither $\mathbf{v}_0$ and $\mathbf{p}_0$ nor when the jump started. Nevertheless, we can leverage the fact that the acceleration, that is, the second derivative of $\mathbf{p}$ is constant and equal to $\mathbf{g}$. To this end, we only need to consider four *equidistant* times t_1, t_2, t_3, t_4. We can then approximate the acceleration at t_2 and t_3 using finite differences and write that they should be equal. This yields

$$2\mathbf{p}(t_2) - \mathbf{p}(t_1) - \mathbf{p}(t_3) = 2\mathbf{p}(t_3) - \mathbf{p}(t_2) - \mathbf{p}(t_4)$$
$$\Rightarrow 3\mathbf{p}(t_2) - 3\mathbf{p}(t_3) + \mathbf{p}(t_4) - \mathbf{p}(t_1) = 0 . \quad (7)$$

In practice, two difficulties arise in imposing this simple constraint. First, it is not obvious where someone's CG is given that it depends on their posture and on the weight of each body part. Second, we only have access to 2D projections of $\mathbf{p}(t)$ in the image plane and not the 3D locations.

Recall from Section 3.1 that our model first applies an STN $\mathcal{S}$, which, for each frame, outputs a bounding box around the subject parameterized by two scales and two translations, s^x, s^y, u^x, u^y. We address the first problem by approximating the projection of $\mathbf{p}(t)$ as the center of the bounding box. To address the second problem, we consider that the diver's CG travels in a plane. If it were parallel to the image plane, the value of u^y would vary linearly with the CG height and the relationship of Eq. (7) would hold exactly, implying that

$$3u_2^y - 3u_3^y + u_4^y - u_1^y = 0 , \quad (8)$$

where u_t^y denotes the value of u^y at time t. Even when the diver does not travel in a plane parallel to the image one, the frame-to-frame change in distance to the camera of the CG typically remains small and perspective effects are negligible with respect to the vertical motion. Eq. (8) is therefore a good approximation, which we will show in Section S.3 of the supplementary material. We therefore impose the constraint of Eq. (8) as a soft constraint by adding a new loss term into our training procedure.

As shown in the supplementary material, displacement along the x axis occurs is much smaller than that along the y axis, which we therefore focus on. Hence, we take the new loss term to be

$$\mathcal{L}_{\text{const-acc}} = \|(u_1^y + 3\,u_3^y) - (u_4^y + 3\,u_2^y)\|^2 . \quad (9)$$

This definition does not differentiate between moving up or down, whereas, in practice, we know that the diver goes down. In our implementation, this corresponds to u_t^y decreasing. Hence, we also introduce an *order loss* $\mathcal{L}_{\text{order}}$ that we define as

$$\mathcal{L}_{\text{order}} = \sum_{t=1}^{3} \max\left(0, \tau - \left(u_t^y - u_{t+1}^y\right)\right) , \quad (10)$$

where τ is a threshold to satisfy a minimum distance between frames t and $t + 1$.

In free-fall, the detector could still make the bounding box bigger to compensate for the aforementioned losses. We therefore prevent the scales s_x and s_y of the bounding boxes from varying abruptly from frame to frame by defining a *scale loss* $\mathcal{L}_{\text{sc}}$ as

$$\mathcal{L}_{\text{sc}} = \sum_{t=1}^{3} \left\| (s_t^x, s_t^y) - (s_{t+1}^x, s_{t+1}^y) \right\|^2 . \quad (11)$$

Putting this all together, the tracking loss is then equal to

$$\mathcal{L}_{\text{track}} = \mathcal{L}_{\text{const-acc}} + \mathcal{L}_{\text{order}} + \mathcal{L}_{\text{sc}} . \quad (12)$$

3.3 Leveraging Weak Labels

When 2D pose labels are available, we can use them to learn richer latent features. In particular, as shown in Fig. 4, we apply a multi-layer-perceptron network Θ only on the time-variant feature $\mathbf{I}_{tv}$ and regress the 2D pose $\mathbf{q}_{2D}^k$, yielding

$$\mathcal{L}_{\text{pose}}^{2D} = \frac{1}{N} \sum_{k=1}^{N} \|\Theta(\mathbf{I}_{tv}^k) - \mathbf{q}_{2D}^k\|^2 . \quad (13)$$

3.4 Training and Inference

To train our network, we minimize the total training loss

$$\mathcal{L}_{\text{total}} = \mathcal{L}_{\text{reconst}} + \alpha\mathcal{L}_{\text{DSL}}^{\text{all}} + \gamma\mathcal{L}_{\text{track}} + \eta\mathcal{L}_{\text{pose}}^{2D} , \quad (14)$$

where $\mathcal{L}_{\text{reconst}}$ of Eq. (2) favors a good image reconstruction, $\mathcal{L}_{\text{DSL}}^{\text{all}}$ of Eq. (6) applies a contrastive loss, and $\mathcal{L}_{\text{track}}$ of Eq. (12) is used when the subject is in free fall. $\mathcal{L}_{\text{pose}}^{2D}$ is applied only when 2D pose is available. We carry out experiment with and without using 2D poses . During the minimization, we permute randomly the time-invariant components of the latent vectors to promote disentanglement as described in Section 3.1.3. This part of the training is unsupervised or weakly-supervised (depending on the usage of 2D labels).

Once the self-supervised model is trained, its parameters are frozen and we use a small number of images with corresponding ground-truth poses for supervised training of a regressor Φ designed to associate a pose vector to the time-variant latent vector $\mathbf{I}_{tv}$ that the encoder $\mathcal{E}_{tv}$ has extracted from image $\mathbf{I}_{\text{crop}}$. To this end, we adjust the weights of Φ by minimizing

$$\mathcal{L}_{\text{pose}}^{3D} = \frac{1}{N} \sum_{k=1}^{N} \|\Phi(\mathbf{I}_{tv}^k, \hat{\mathbf{q}}_{2D}^k) - \mathbf{q}_{3D}^k\|^2 , \quad (15)$$

where $\mathbf{I}_{tv}^k$ is the time-variant vector associated to the k^{th} image and $\mathbf{q}_{3D}^k$ the corresponding ground-truth 3D pose, in camera coordinates and pelvis centered, represented by a concatenation of the Cartesian coordinates of body joints. $\hat{\mathbf{q}}_{2D}^k$ is the predicted 2D pose by the model, used only when 2D is leveraged in the initial self-supervised learning phase of the model.

4 RESULTS

In this section, we present the datasets and baselines models used in our experiments, followed by quantitative results, ablation studies and analyses of the proposed approach.

4.1 Datasets.

We evaluate our model on the following five datasets:

4.1.1 Human3.6M [1] (H36M)

As in [7], we use subjects 1, 5, 6, 7, and 8 to train the self-supervised models and sub-sample every 5 frames of the training set, which yields 308,760 frames. We test our model on two different test sets:

- Following [7], we use the PoseTrack2018 challenge [44] setup to evaluate the accuracy of 3D human pose estimation. It involves 35,832 training images with 3D pose labels and 19,312 test samples.
- When we use weak 2D labels, we report results on subjects 9 and 11, sub-sampled every 10 frames. As this test-set is used more often, it helps comparison with other weakly supervised methods.

4.1.2 MPI-INF-3DHP [2] (MPI)

On this dataset, we consider two settings:

- **MV-data:** The self-supervised models are trained using subjects S1 to S8 captured using multiple cameras. Since the background of the test-videos is different from that of the training ones, we train the self-supervised models on images with random backgrounds to help with generalization. To this end, we exploit the foreground masks from the dataset to put the subject in front of random images from the Internet. Note that these masks are used only to synthesize new training images but not by the networks themselves.
- **All-data:** Since self-supervised models do not require labels, we can leverage all the data to train them, whether single or multi-view. This corresponds to the transductive learning setting, and the self-supervised models are trained using S1 to S8 (multi-view) and TS1 to TS4 (single-view).

The video sequences are subsampled every 5 frames, yielding 198,096 training frames for MV-data and 202,905 for all-data. In both cases, we use subjects S1, S2, S3, S4, S5, and S8 to train the pose model, and S6 and S7 to form the validation set. For testing, we use TS1 to TS4. Because there are many pose labels, we use at most 2% of the pose data for training, that is, 17,128 images. The test set comprises 2207 frames.

4.1.3 Diving

We acquired video sequences featuring one male and one female competitive diver. The dataset features dives from heights of 3, 5, 7.5, and 10 meters, most captured using multiple cameras and some using a single camera. We use both subjects to train the self-supervised models. As in the MPI case, we consider two settings:

- **MV-data:** Only multi-view data containing 16,469 frames.
- **All-data:** Both single- and multi-view data, yielding 19,464 training frames.

The 3D pose is obtained by first annotating 2D pose in each view and then using triangulation on all views. This makes the 3D pose labels prone to noise, due to both self-occlusion in some views, and also mis-annotation by labelers, which contribute to the difficulty of 3D pose estimation on this dataset.

To train the pose model, we apply two settings:

- **Diver-split:** We use the female subject for training—splitting the videos into training and validation sets—and the male subject for testing. This setup yields up to 2,419 frames for supervised training, 2,288 frames for validation, and 2,112 frames for testing. This is the default setting.
- **Dive-split:** Due to the strong bias of training on only one diver, we also split the dataset by dives, using unique and distinct dives in each split. While diver-split can better show generalization capability, it does not allow for proper evaluation due to the lack of identity diversity in the training set. The dive-split setup instead allows us to analyze how well the model can generalize to new dives, while not suffering from identity-related limitations. This setup yields 2,789 training, 1,278 validation, and 2,724 testing frames.

4.1.4 Skiing

To evaluate our model on a dataset with moving cameras we use a Ski dataset [45] that features 6 professional skiers, each going down a slalom course twice. The data is captured with two cameras. We use Subjects 1 to 5 for training and Subject 6 to evaluate the models. This yields 3,124 images for training and 636 images for testing for both self-supervised and pose models.

4.1.5 3DPW

We use this dataset for evaluation on in-the-wild images. Since our approach is designed for single subjects we only use single-subject videos, yielding 12,265 training and 14,803 testing frames.

4.2 Baselines

We compare our model to other unsupervised feature extraction methods, which are

- **CSS [10], [11]:** Single-view CSS loss using Eq. (3), without detection, image reconstruction, splitting the latent space into $\mathbf{I}_{tv}$ and $\mathbf{I}_{ti}$ components, or the DSL loss of Eq. (5).
- **MV-CSS [12]:** Multi-view CSS loss using the triplet loss of Eq. (4), also without detection, reconstruction, splitting, or DSL loss.
- **DrNet [41]:** Single-view approach using both reconstruction and latent splitting. It uses a discriminator for latent splitting.

- **NSD [7]:** Multi-view approach that performs both reconstruction and latent vector splitting.

We also evaluate variants of our own approach:

- **DSL:**. This variant uses the contrastive loss of Eq. (5) without detection, reconstruction, or decoding.
- **DSL-STN:**. This variant adds detection to the DSL model.
- **DSL-STN-Dec:**. This variant adds decoding to the DSL-STN model.
- **DSL-STN-Dec-Split (Ours):** This is our full approach that adds latent splitting to DSL-Dec-STN. It uses detection together with image-reconstruction and latent splitting. It applies Eq. (5) for the contrastive loss.
- **AE-STN-Dec:**. This variant uses detection together with an auto-encoding image-reconstruction, but without contrastive training or latent splitting.
- **CSS-STN-Dec-Split:** This variant uses detection together with image-reconstruction and latent splitting, similar to Ours. However, it uses Eq. (3) instead of Eq. (5) for the contrastive loss, which does not leverage the DSL formulation.

For a fair comparison, we followed the training protocol of [7] and used the same detector, encoder, and decoder for all models including the baselines. In other words, unless explicitly specified, all these models are trained using the same architecture and model capacity, so that only the training losses differentiate them. For MV-CSS, we follow the recommendation of [12] to sample one positive sample from another view at the same timestamp and one negative sample from a distant frame in the same video. For CSS and CSS-STN-Dec-Split, we use one positive sample from a nearby frame and two negative samples from away frames from the same video. For DSL-STN, we replace one of the away frames with an intermediate frame, sampled between nearby and away frames. We follow the procedures of [41] and [7] to train respectively DrNet and NSD.

4.3 Comparison against Unsupervised Methods

We compare our results to other unsupervised feature-extraction models on five datasets and if not otherwise specified report the N-MPJPE as a function of the percentage of labeled pose samples being used.

4.3.1 H36M Dataset.

In Table 1 we compare against both multi- and single-view approaches. We consistently outperform the other single-view methods and come close to NSD, even though it uses multiple views. Note that the contrastive approaches without decoding, whether single- or multi-view perform poorly on this task, which confirms that decoding helps to learn richer features. This further highlights the importance of using decoding together with contrastive learning and detection, especially when the subject occupies only a small portion of the image and small details should be captured for accurate pose estimation.

As discussed in Section 2, some 2D keypoint estimation approaches [24], [27], [28], [29] separate features for pose and appearance, as we do. For comparison purposes, we therefore also evaluate our extracted time-variant features for 2D keypoint estimation, using the same setup as in Table 1 of [29], that is, we train a single linear layer from the time-variant features to the 2D keypoints. As in [29], we use subjects 1,5,6,7,8 for training and

	Model	Decodes?	Percentage of Labeled Pose Data				
			0.3% (100)	1% (500)	14% (5K)	50% (17K)	all (35K)
MV	MV-CSS [12]	✗	216.15	206.08	203.24	202.65	202.52
	NSD [7]	✓	140.92	115.37	94.48	91.01	88.89
SV	CSS [10], [11]	✗	230.77	215.96	197.22	195.70	194.69
	DrNet [41]	✓	158.81	126.18	108.21	106.41	100.38
	Ours	✓	149.32	122.44	100.23	97.38	95.39

TABLE 1: Comparison on the PoseTrack H36M test-set. The top rows characterize the unsupervised multi-view (MV) models and the bottom rows the single-view (SV) ones. The last 5 columns show N-MPJPE results when using different percentage of labeled 3D pose samples. The value in parenthesis are the corresponding number of images with labels.

9,11 for testing. As can be seen in Table 2, our model outperforms the other approaches, while it does not use any ground-truth bounding box locations for object detection.

Model	%-MSE Error
Thewlis et al. [27]	7.51
Zhang et al. [24]	4.14
Lorenz et al. [28]	2.79
Jakab et al. [29]	2.73
Ours	**2.53**

TABLE 2: Comparison against state-of-the-art 2D keypoint estimation models on H36M dataset. %-MSE normalized by image size is reported. All models predict 32 keypoints on 6 actions of wait, pose, greet, direct, discuss, and walk.

4.3.2 MPI Dataset.

In this dataset as the S1 to S8 videos are multi-view and the TS1 to TS4 ones are single-view, we can only evaluate NSD in the MV-data setting. We report results in Table 3, where we compare with the best-performing models of Table 1. Our model again performs closely to NSD in this setting with the advantage that, because our approach is single-view, we can also train on the TS1 to TS4 videos, thus further improving performance. In a practical setting, this means that we could exploit additional videos as they become available, whether or not they are multi-view, something NSD cannot do.

	Model	Percentage of Labeled Pose Data				
		.02% (171)	.1% (856)	.2% (1712)	5% (4K)	2%(17K)
MV-data	NSD [7]	**256.34**	**241.12**	**239.81**	**237.56**	**235.61**
	Ours	261.99	245.09	241.34	239.48	236.85
All-data	DrNet [41]	253.47	243.87	233.09	228.09	218.18
	Ours	**235.62**	**219.72**	**211.26**	**202.60**	**192.11**

TABLE 3: Comparison on MPI-INF-3DHP. N-MPJPE results on pose prediction as a function of the quantity of annotated data. The top rows show results in the MV-data setting, where the self-supervised models have access only to multi-view data. The bottom rows show results in the All-data setting where self-supervised models have access to both multi-view and single-view data. Note that, in this scenario, our results are now better than those of NSD because it can also exploit single-view data.

4.3.3 Diving Dataset

We compare with NSD in the MV-data setting and with DrNet in the All-data setting. As can be seen in Table 4, our single-view approach outperforms DrNet, performs on par with multi-view NSD in the MV-data setting and better in the All-data setting. In other words, because our approach can benefit from both multi- and single-view data, it can leverage more of the available images.

As discussed above, on this dataset, we use the tracking loss presented in Section 3.2 for pre-training purposes for all models, including the baselines. Otherwise, it would take longer for them to detect the subject. The tracking loss reduces the required number of training iterations from 50K to 5K, and thus the training time, saving 38 training hours on a Tesla V100 GPU. Table 4

Fig. 5: Swapping time-invariant and time-variant components. In each column, the time-invariant component is taken from the image on the top, and the time-variant component is taken from the image on the bottom and concatenated to generate the image shown in bottom, using the decoder.

features results for an additional variant of our model, Ours (no $\mathcal{L}_{track}$), which is trained from scratch without gravity pre-training using the same total number iterations as our full method with pre-training. The pose estimation error increases by 21%.

Model		Percentage of Labeled Pose Data				
		5% (120)	10% (241)	20% (483)	50% (1209)	100%(2419)
MV-data	NSD [7]	270.42	238.23	221.25	200.28	195.13
	Ours	279.82	242.78	224.45	205.89	197.40
All-data	DrNet [41]	320.07	290.93	271.87	258.14	248.31
	Ours (no $\mathcal{L}_{track}$)	296.22	268.88	255.02	236.49	225.77
	Ours	**265.09**	**232.66**	**216.14**	**197.82**	**186.45**

TABLE 4: Comparison on the Diving dataset (Diver-split). N-MPJPE results as a function of the quantity of annotated data used to train the pose predictor. Top two rows show models trained only on multi-view data. The bottom rows show models trained using both single-view and multi-view data.

4.3.4 Ski Dataset

To check the performance of our model given a changing background, we evaluate our model on a ski dataset [45] where the camera follows skiers down the slope. On this dataset, we obtain an N-MPJPE error of 113.20 compared to 180.31 for the single-view approach of DrNet.

In Fig. S6 we present qualitative results, with input images, the corresponding estimated backgrounds, and the background-subtracted images. Note that, because of noise, background subtraction does not yield images containing only the subject of interest. In rows 4 to 6 we show different thresholdings of the background-subtracted inputs. Note that thresholding by itself cannot eliminate all the background noise and the model needs to learn to focus on the subject and ignore the remaining background. We further show samples of detected bounding boxes, reconstructed images, decoded masks, and 3D pose estimates.

4.4 Analysis

4.4.1 Ablation Study

To analyze the different components of our model, we perform an ablation study on H36M by removing different model components and report our results in Table 5.

In the first two rows (DSL, DSL-STN), we only apply the DSL contrastive loss, without latent splitting and reconstruction. Compared to our full approach in the last row of the table, the accuracy is degraded by a factor of almost 2. Since these two

models do not reconstruct the image, they cannot learn the time-invariant components and all the latent components are time-variant. This being the case, the basic contrastive approach does not suffice.

In the third row, we add reconstruction of the input to the DSL-STN model but without splitting the latent variables. The error of DSL-STN-Dec in the last column improves by 36% compared to DSL-STN. Reconstructing the image helps the model to both detect the subject of interest and to capture more details.

In the final row, we report again the results of the full model, which adds latent splitting to the DSL-STN-Dec model of the third row. It improves pose estimation by 25% compared to DSL-STN-Dec, showing splitting the latent space is beneficial in learning better time-variant features, as the model can apply the contrastive loss only to this component. Compared to the basic DSL approach, we improve by 50%. In other words, detection, reconstruction and latent vector splitting all help deliver the best possible accuracy.

We also compare against AE-STN-Dec, which only applies detection and decoding, without any contrastive loss or latent splitting. The error in the last column increases by about 20% compared to our full approach. Finally, in the fifth row, we change only the DSL loss of Eq. (5) to the CSS loss of Eq. (3) and still observe an increase in error by 4% in the last column. This gap is even more noticable in previous columns of the table, that is, when using fewer annotations. We observe an error increase of about 10% for the columns corresponding to 0.3%, 1%, and 14% of the data, and of about 7% for the column corresponding to 50%.

In short, all of the proposed modules have a role and disabling any of them reduces performance.

4.4.2 Validity of our Distance-Based Similarity Loss

To verify whether the DSL loss proposed in Section 3.1.2 makes sense, we measure how much 3D pose differs over time in videos. Fig. 6 shows the average MPJPE distance over the training set of pelvis-centered poses as a function of frames temporal distances. For different datasets, the pose distance is different because the actions are different. For example, in the Diving dataset, pose distances change more quickly and reach a higher level than in the other datasets, mostly because the divers rotate during their dives. On average, while poses are similar in nearby frames they get more dissimilar as the temporal distance increases before plateauing. This trend matches the assumption of how frames change over time that underpins our formulation of the DSL loss of Eq. (5).
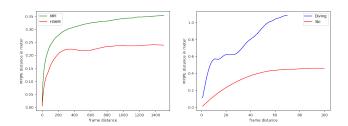

Fig. 6: Average MPJPE distance of pelvis-centered poses at different frame distances.

4.4.3 Disentanglement

For us disentanglement is not a goal by itself but rather a means to an end. However, to verify the disentanglement of the time-variant (TV) and time-invariant (TI) components, we applied KMeans

Model	Decodes?	uses STN?	contrastive loss	latent split?	Percentage of Labeled Pose Data				
					0.3% (100)	1% (500)	14% (5K)	50% (17K)	all (35K)
DSL	✗	✗	DSL	✗	227.24	220.64	198.34	193.08	191.21
DSL-STN	✗	✓	DSL	✗	236.32	211.85	202.57	199.16	198.22
DSL-STN-Dec	✓	✓	DSL	✗	185.75	159.34	135.93	129.62	127.52
AE-STN-Dec	✓	✓	None	✗	187. 25	161.53	130.21	117.044	114.73
CSS-STN-Dec-Split	✓	✓	CSS	✓	163.61	137.52	110.61	104.23	99.32
DSL-STN-Dec-Split (Ours)	✓	✓	DSL	✓	**149.32**	**122.44**	**100.23**	**97.38**	**95.39**

TABLE 5: Ablation Study. The last row shows our full model with all features applied. In the first two rows we apply the DSL contrastive loss without decoding and latent split. In the third row, we keep all the features except we do not split the latent components. In the fourth row, we apply only detection and reconstruction without any application of contrastive loss. In the fifth row, we only change the contrastive loss of our full model from DSL to CSS. In all settings, we observe a performance drop, however, the models without decoding suffer the most.

clustering separately on the TV and TI features of our model on the H36M test-set and compared them to the subject classes. We obtained 52.3% and 95.9% classification accuracy for TV and TI, showing that TI captures identity while TV does not. We also trained a pose model on the TI features of H36M and obtained an N-MPJPE of 201.3 for the 100%-data case, compared to 95.39 for TV (presented in Table 1), showing that TV is almost twice better than TI for pose estimation. Hence, disentanglement is indeed happening.

1st-phase	2nd-phase	Train-set	MPJPE	N-MPJPE	P-MPJPE
Resynthesis (w STN)	I_{tv}	H36M	10.17	10.09	7.38
Resynthesis (w/o STN)	I_{tv}	H36M	10.03	9.93	7.49
Resynthesis (w/o STN) + 2D	I_{tv}	H36M	7.45	7.41	5.73
Resynthesis (w/o STN) + 2D	I_{tv} + 2D	H36M	7.33	7.29	5.61
Resynthesis (w/o STN) + 2D *	I_{tv}	MPI	7.46	7.42	5.72
Resynthesis (w/o STN) + 2D *	I_{tv} + 2D	MPI	7.39	7.35	5.68
Resynthesis (w/o STN) + 2D *	2D	H36M	6.42	6.37	5.01
Resynthesis (w/o STN) + 2D *	I_{tv}	H36M	6.26	6.23	4.82
Resynthesis (w/o STN) + 2D *	I_{tv} + 2D	H36M	6.01	5.97	4.60
Resynthesis (w/o STN) + 2D *	I_{tv} + 2D alpha	H36M	**5.69**	**5.65**	**4.36**

TABLE 6: Comparison on H36M test-set (in cm). First two rows apply image re-synthesis with and without detection in the first phase of training, where in the latter the cropped image is passed to the model. In the 3rd and 4th rows 2D is used together with image re-synthesis in the first phase of training. The bottom rows with ∗ show models where the ResNet50 backbone is replaced with a Resnet 101. The 5th and 6th rows show results where the model is trained on MPI in the first phase and then tested on H36M in the 2nd phase. When 2D is used in the 2nd phase, it is the prediction of the model itself. In the last row, however, we pass I_{tv} together with 2D pose predicted by Alphapose to 3D pose estimation model in the 2nd phase. All models use a ResNet MLP in the 2nd phase.

1st-phase	2nd-phase	Train-set	MPJPE	N-MPJPE	P-MPJPE
Resynthesis	I_{tv} (2-hid MLP)	MPI	18.06	17.49	12.31
Resynthesis	I_{tv} (ReNet MLP)	MPI	16.80	16.37	11.15
Resynthesis + 2D *	I_{tv} (ReNet MLP)	H36M	11.50	11.40	7.54
Resynthesis + 2D *	I_{tv} + 2D (ReNet MLP)	H36M	11.39	11.29	7.49
Resynthesis + 2D *	I_{tv} (ReNet MLP)	MPI	10.21	10.15	6.54
Resynthesis + 2D *	I_{tv} + 2D (ReNet MLP)	MPI	**9.30**	**9.24**	**5.95**

TABLE 7: Comparison on MPI-INF-3DHP test-set (in cm). First two rows apply only image re-synthesis, where the first row uses 2-hid MLP for 3D pose regression. The last rows with ∗ show models where the ResNet50 backbone is replaced with a Resnet 101. The 3rd and 4th rows show results where the model is trained on H36M in the first phase and then tested on MPI in the 2nd phase.

dataset	1st-phase	2nd-phase	MPJPE	N-MPJPE	P-MPJPE
Diver-split	Resynthesis	I_{tv}	16.4	15.7	10.7
Diver-split	Resynthesis + 2D	I_{tv}	15.2	14.1	9.43
Dive-split	Resynthesis	I_{tv}	9.19	8.75	5.81
Dive-split	Resynthesis + 2D	I_{tv}	**8.53**	**8.26**	**5.60**

TABLE 8: Comparison on the Diving dataset (in cm). The first two rows compare the results in the diver-split setting, whereas the last two rows evaluate the dive-split setup. All models use a ResNet50 backbone and use a 6-layer MLP as pose model. Only the time-variant features are used for pose estimation to compare each pair of rows with a similar input.

	Model	MPJPE	N-MPJPE	P-MPJPE
Unsup	Rhodin [32]	-	115.0	-
	Kundu et al. [46]	**99.2**	-	-
	Ours	100.3	**99.3**	**74.9**
Weakly-Sup	AIGN [47]	-	-	79.0
	MCSS [33]	94.25	92.60	72.48
	Chen et al. [48]	-	-	68.0
	HMR [49]	-	-	66.5
	RepNet [50]	89.9	-	65.1
	Wang et al. [51]	86.4	-	62.8
	Kolotouros et al. [52]	-	-	62.0
	Iqbal et al. [53]	67.4	64.5	54.5
	CanonPose [54]	74.3	-	53.0
	Kundu et al. [46]	62.4	-	-
	Ours	**60.1**	**59.7**	**46.0**

TABLE 9: Comparison with models on H36M test-set (in mm). Unsupervised methods do not use any sort of labels in the feature extraction phase. Weakly-supervised methods use either 2D labels or a pre-trained 2D model.

	Model	MPJPE	N-MPJPE	P-MPJPE
Weakly-Sup	Kolotouros et al. [52]	124.8	-	80.4
	HMR [49]	169.5	-	113.2
	Iqbal et al. [53]	113.8	102.2	-
	CanonPose [54]	104.0	-	70.3
	Chen et al. [48]	-	-	71.1
	RepNet [50]	97.8	-	-
	Kundu et al. [46]	93.9	-	-
	Ours	**93.0**	**92.4**	**59.5**

TABLE 10: Comparison with models on MPI-INF-3DHP test-set (in mm). Weakly-supervised methods use either 2D labels or a pre-trained 2D model.
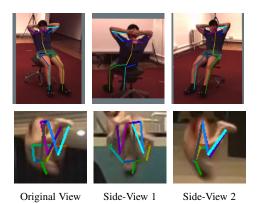

Original View   Side-View 1   Side-View 2

Fig. 7: Side-view visualization. For each triplet, the first image shows the prediction of the model on the input view, while side-views 1 and 2 show projection of the 3D pose to other views.

In Fig. 5, we show results obtained by merging time-variant and time-invariant components from different sources to reconstruct new images. We observe that the model correctly splits information into time-variant and time-invariant components, which it can then use to synthesize a new image.

4.5 Leveraging Weak Labels

We evaluate how leveraging 2D poses, when available as described in Section 3.3, can help improve the learned I_{tv} features. Tables 6 shows the impact of using 2D poses in the feature extraction phase

Model	2nd-phase	MPJPE	N-MPJPE	P-MPJPE
ImageNet + 2D	$\mathbf{I}_{tv}$ + 2D	18.38	18.07	11.44
Resynthesis + 2D (Ours)	$\mathbf{I}_{tv}$ + 2D	13.35	13.29	9.20
ImageNet + 2D *	$\mathbf{I}_{tv}$ + 2D	16.47	16.29	11.02
Resynthesis + 2D * (Ours)	$\mathbf{I}_{tv}$ + 2D	**10.42**	**10.38**	**7.47**

TABLE 11: Results on 3DPW (in cm). We compare the impact of our pre-training approach to the ImageNet pre-trained models on in-the-wild images. The first two rows use ResNet50 as backbone, whereas the bottom two rows, indicated by *, use ResNet101.

on H36M dataset. In particular, comparing 2nd and 3rd rows of the table show using 2D pose can improve MPJPE by about 2.5 cm. In the 2nd row we provide cropped images to the network. Comparing the first two rows show the models perform almost similarly with and without detection, indicating the efficiency of our unsupervised detection in the first row.

In the bottom rows, we replace the ResNet50 backbone with a ResNet101, which further improves the accuracy. This indicates a more sophisticated architecture can yield improved performance. In the final row we take 2D poses from pre-trained Alphapose model [55], which is fine-tuned on H36M, and pass it together with $\mathbf{I}_{tv}$ to 3D pose estimation network. Rows 7th to 9th also show that adding 2D in the 2nd phase to $\mathbf{I}_{tv}$ yields the best results rather than using them individually.

Table 7 similarly shows the impact of 2D on MPI dataset. In the first two rows, where the model only applies re-synthesis, we compare the impact of pose-regressor architecture. The 2nd row shows that using a better ResNet MLP architecture, take from [56], can further boost the performance. The 3rd and 4th rows of Table 7 as well as 5th and 6th rows of Table 6, show the impact of training on one dataset and testing on another. As expected, the results are about 1 to 2 cm worse than when training and testing on the same dataset. Nevertheless, the results are still decent on the target dataset, which indicates our color augmentation helps the model not overfit on the background or the person.

Table 8 shows our results on the diving dataset in the diver-split and dive-split cases. Similarly to the previous datasets, in each setup we obtain better features when using 2D labels. The gap between the two different splits indicates that the model suffers from diver diversity in the diver-split case. Hence, diversity in identities is essential to help the model generalize better.

Comparison with the state-of-the-art models. In Table 9 we report results on the test set of H36M dataset and compare against unsupervised models (without any sort of supervision in the feature extraction phase) and also weakly supervised methods that leverage either 2D labels or pre-trained 2D networks. As can be observed our approach compares favorably with respect to these methods. The model of [46] leverages some in-house private dataset in addition to an off-the-shelf body-part segmentation, making it not completely unsupervised. In Table 10 we compare methods on MPI-INF-3DHP dataset. Similar to H36M our approach obtains better result in the weakly-supervised setup.

Comparison with supervised pre-training. Similar to the recent unsupervised learning methods [57], [58], [59], we compare the effectiveness of the features extraction using a ResNet50 backbone with our approach to those obtained with a supervised ImageNet pre-trained model. The results are provided in Figure 8. For a fair comparison between our variant that leverages 2D and the ImageNet pre-trained model, we also fine-tune the latter on 2D labels. Ours without 2D improves upon the ImageNet baseline by at least 4.5 cm, while Ours+2D also performs better than the ImageNet+2D variant by 1 to 2 cm in MPJPE. The results clearly

show that the features extracted by Ours outperform the ones extracted by a supervised ImageNet model.

To further evaluate the quality of weight transfer, we take the best models in Figure 8 and use them to initialize networks for end-to-end training on 3D labels. This experiment compares the two models on the quality of the transferred parameters as model initialization. The results depicted in Figure 9 clearly show the advantage of our pre-training.

3DPW dataset. To evaluate the impact of our feature extraction network on in-the-wild images, we use the 3DPW dataset. We take our best models pre-trained on the H36M dataset in the weakly supervised setup, and compare the quality of the extracted features to the supervised ImageNet variants. For a fair comparison with our Resynthesis+2D models, we also fine-tune the ImageNet models on the H36M 2D labels; without this (as also shown in Figure 8) the gap with our approach was even larger. The result are provided in Table 11. Comparing every pair of rows clearly shows that the features extracted by our approach are much better suited for pose estimation than those obtained with ImageNet pre-trained models.

5 LIMITATIONS

While we do not need the cameras to be static, as shown by our experiments on the Ski dataset, we still need a way to obtain a background image. Furthermore, our current formulation can only handle a single subject of interest. We leave the extension to multiple subjects for future work. One approach would be to add depth estimation to the self-supervised model, as proposed in [7], and use the resulting depth to reason about occlusions when merging the decoded crops of different subjects to reconstruct the input image.

6 CONCLUSION

In this work, we have presented an unsupervised feature extraction model for monocular RGB videos. It improves upon the standard CSS approaches by applying contrastive learning only to the time-variant components of the latent representations and enforcing latent similarity given the distance between the frames. When combined with detection and an additional image decoding, this allows our model to extract richer features, well suited for capturing transient dynamics in videos, such as human motion. We have evidenced the benefits of our approach for pose estimation on five benchmark datasets; our approach was shown to outperform other single-view self-supervised learning strategies and to match the performance of multi-view ones. Finally, we show we can leverage 2D poses to enhance the quality of extracted features in the time-variant component, making 3D pose estimation more accurate.

REFERENCES

[1] C. Ionescu, I. Papava, V. Olaru, and C. Sminchisescu, "Human3.6M: Large Scale Datasets and Predictive Methods for 3D Human Sensing in Natural Environments," *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 2014.
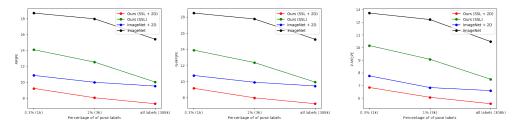
Fig. 8: Comparison of extracted feature by Ours versus supervised ImageNet on H36M test-set using a ResNet50 architecture. The three plots from left to right show results on MPJPE, NMPJPE, and PMPJPE (in cm) for different percentage of labeled 3D data. Ours SSL and SSL+2D are variants trained without and with 2D labels. For ImageNet variants we take the features of the layer before the classification one. In ImageNet+2D we also fine-tune this network on 2D labels for a fair comparison with Ours (SSL+2D).

Fig. 9: Comparison of the quality of the learned parameters for model (ResNet50 architecture) initialization by Ours versus the supervised ImageNet on H36M test-set. The three plots show results on MPJPE, NMPJPE, and PMPJPE (in cm) for different percentage of 3D labels.

[2] D. Mehta, H. Rhodin, D. Casas, P. Fua, O. Sotnychenko, W. Xu, and C. Theobalt, "Monocular 3D Human Pose Estimation in the Wild Using Improved CNN Supervision," in *International Conference on 3D Vision*, 2017.

[3] G. Pavlakos, X. Zhou, K. D. G. Konstantinos, and D. Kostas, "Harvesting Multiple Views for Marker-Less 3D Human Pose Annotations," in *Conference on Computer Vision and Pattern Recognition*, 2017.

[4] H.-Y. Tung, H.-W. Tung, E. Yumer, and K. Fragkiadaki, "Self-Supervised Learning of Motion Capture," in *Advances in Neural Information Processing Systems*, 2017, pp. 5242–5252.

[5] T. Zhou, M. Brown, N. Snavely, and D. Lowe, "Unsupervised Learning of Depth and Ego-Motion from Video," in *Conference on Computer Vision and Pattern Recognition*, 2017, pp. 1851–1858.

[6] H. Rhodin, J. Spoerri, I. Katircioglu, V. Constantin, F. Meyer, E. Moeller, M. Salzmann, and P. Fua, "Learning Monocular 3D Human Pose Estimation from Multi-View Images," in *Conference on Computer Vision and Pattern Recognition*, 2018.

[7] H. Rhodin, V. Constantin, I. Katircioglu, M. Salzmann, and P. Fua, "Neural Scene Decomposition for Human Motion Capture," in *Conference on Computer Vision and Pattern Recognition*, 2019.

[8] A. v. d. Oord, Y. Li, and O. Vinyals, "Representation Learning with Contrastive Predictive Coding," in *arXiv Preprint*, 2018.

[9] A. Hyvarinen and H. Morioka, "Unsupervised Feature Extraction by Time-Contrastive Learning and Nonlinear ICA," in *Advances in Neural Information Processing Systems*, 2016, pp. 3765–3773.

[10] T. Chen, S. Kornblith, M. Norouzi, and G. Hinton, "A Simple Framework for Contrastive Learning of Visual Representations," in *International Conference on Machine Learning*, 2020.

[11] A. Anand, E. Racah, S. Ozair, Y. Bengio, M.-A. Côté, and R. D. Hjelm, "Unsupervised State Representation Learning in Atari," in *Advances in Neural Information Processing Systems*, 2019, pp. 8769–8782.

[12] P. Sermanet, C. Lynch, Y. Chebotar, J. Hsu, E. Jang, S. Schaal, S. Levine, and G. Brain, "Time-Contrastive Networks: Self-Supervised Learning from Video," in *International Conference on Robotics and Automation*, 2018, pp. 1134–1141.

[13] Y. Bengio, A. Courville, and P. Vincent, "Representation Learning: A Review and New Perspectives," *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 2013.

[14] C. Doersch and A. Zisserman, "Multi-Task Self-Supervised Visual Learning," in *International Conference on Computer Vision*, October 2017.

[15] S. N., E. Mansimov, and R. Salakhudinov, "Unsupervised Learning of Video Representations Using LSTMs," in *International Conference on Machine Learning*, 2015, pp. 843–852.

[16] J. Wang, J. Jiao, L. Bao, S. He, Y. Liu, and W. Liu, "Self-Supervised Spatio-Temporal Representation Learning for Videos by Predicting Motion and Appearance Statistics," in *Conference on Computer Vision and Pattern Recognition*, 2019, pp. 4006–4015.

[17] I. Misra, C. L. Zitnick, and M. Hebert, "Shuffle and Learn: Unsupervised Learning Using Temporal Order Verification," in *European Conference on Computer Vision*, 2016, pp. 527–544.

[18] H.-Y. Lee, J.-B. Huang, M. Singh, and M.-H. Yang, "Unsupervised Representation Learning by Sorting Sequences," in *Conference on Computer Vision and Pattern Recognition*, 2017, pp. 667–676.

[19] B. Fernando, H. Bilen, E. Gavves, and S. Gould, "Self-Supervised Video Representation Learning with Odd-One-Out Networks," in *Conference on Computer Vision and Pattern Recognition*, 2017, pp. 3636–3645.

[20] D. Xu, J. Xiao, Z. Zhao, J. Shao, D. Xie, and Y. Zhuang, "Self-Supervised Spatiotemporal Learning via Video Clip Order Prediction," in *Conference on Computer Vision and Pattern Recognition*, 2019, pp. 10 334–10 343.

[21] J. Wang, J. Jiao, and Y. Liu, "Self-Supervised Video Representation Learning by Pace Prediction," in *European Conference on Computer Vision*, 2020.

[22] A. W. Harley, S. K. Lakshmikanth, F. Li, X. Zhou, H.-Y. F. Tung, and K. Fragkiadaki, "Learning from Unlabelled Videos Using Contrastive Predictive Neural 3D Mapping," in *International Conference on Learning Representations*, 2020.

[23] J. Thewlis, H. Bilen, and A. Vedaldi, "Unsupervised Learning of Object Landmarks by Factorized Spatial Embeddings," in *International Conference on Computer Vision*, 2017.

[24] Y. Zhang, Y. Guo, Y. Jin, Y. Luo, Z. He, and H. Lee, "Unsupervised Discovery of Object Landmarks as Structural Representations," in *Conference on Computer Vision and Pattern Recognition*, 2018, pp. 2694–2703.

[25] O. Wiles, A. Koepke, and A. Zisserman, "Self-Supervised Learning of a Facial Attribute Embedding from Video," in *British Machine Vision Conference*, 2018.

[26] T. Jakab, A. Gupta, H. Bilen, and A. Vedaldi, "Unsupervised Learning of Object Landmarks through Conditional Image Generation," in *Advances in Neural Information Processing Systems*, 2018.

[27] J. Thewlis, S. Albanie, H. Bilen, and A. Vedaldi, "Unsupervised Learning of Landmarks by Descriptor Vector Exchange," in *International Conference on Computer Vision*, 2019, pp. 6361–6371.

[28] D. Lorenz, L. Bereska, T. Milbich, and B. Ommer, "Unsupervised Part-Based Disentangling of Object Shape and Appearance," in *Conference on Computer Vision and Pattern Recognition*, 2019, pp. 10 955–10 964.

[29] T. Jakab, A. Gupta, H. Bilen, and A. Vedaldi, "Self-Supervised Learning of Interpretable Keypoints from Unlabelled Videos," in *Conference on Computer Vision and Pattern Recognition*, 2020, pp. 8787–8797.

[30] J. R. A. Moniz, C. Beckham, S. Rajotte, S. Honari, and C. Pal, "Unsupervised Depth Estimation, 3D Face Rotation and Replacement," in *Advances in Neural Information Processing Systems*, 2018, pp. 9736–9746.

[31] Y. Tian, D. Krishnan, and P. Isola, "Contrastive Multiview Coding," in *European Conference on Computer Vision*, 2020.

[32] H. Rhodin, M. Salzmann, and P. Fua, "Unsupervised Geometry-Aware Representation for 3D Human Pose Estimation," in *European Conference on Computer Vision*, 2018.

[33] R. Mitra, N. B. Gundavarapu, A. Sharma, and A. Jain, "Multiview-Consistent Semi-Supervised Learning for 3D Human Pose Estimation," in *Conference on Computer Vision and Pattern Recognition*, 2020, pp. 6907–6916.

[34] C. Kim, F. Li, and J. M. Rehg, "Multi-Object Tracking with Neural Gating Using Bilinear LSTM," in *European Conference on Computer Vision*, 2018.

[35] R. Chen, X. Li, R. G. B, and F. Duvenaud, "Isolating Sources of Disentanglement in Variational Autoencoders," in *Advances in Neural Information Processing Systems*, 2018.

[36] T. Karras, S. Laine, and T. Aila, "A Style-Based Generator Architecture for Generative Adversarial Networks," in *Conference on Computer Vision and Pattern Recognition*, 2019.

[37] X. Chen, Y. Duan, R. Houthooft, J. Schulman, I. Sutskever, and P. Abbeel, "Infogan: Interpretable Representation Learning by Information Maximizing Generative Adversarial Nets," in *Advances in Neural Information Processing Systems*, 2016.

[38] Y. Wang, P. Bilinski, F. Bremond, and A. Dantcheva, "G3AN: Disentangling Appearance and Motion for Video Generation," in *Conference on Computer Vision and Pattern Recognition*, 2020.

[39] E. Remelli, S. Han, S. Honari, P. Fua, and R. Wang, "Lightweight Multi-View 3D Pose Estimation through Camera-Disentangled Representation," in *Conference on Computer Vision and Pattern Recognition*, 2020.

[40] D. Worrall, S. Garbin, D. Turmukhambetov, and G. Brostow, "Interpretable Transformations with Encoder-Decoder Networks," in *International Conference on Computer Vision*, 2017.

[41] E. L. Denton *et al.*, "Unsupervised Learning of Disentangled Representations from Video," in *Advances in Neural Information Processing Systems*, 2017, pp. 4414–4423.

[42] M. Jaderberg, K. Simonyan, A. Zisserman, and K. Kavukcuoglu, "Spatial Transformer Networks," in *Advances in Neural Information Processing Systems*, 2015, pp. 2017–2025.

[43] K. He, X. Zhang, S. Ren, and J. Sun, "Deep Residual Learning for Image Recognition," in *Conference on Computer Vision and Pattern Recognition*, 2016, pp. 770–778.

[44] "Posetrack Dataset and Benchmark," 2018.

[45] M. Ostrek, H. Rhodin, P. Fua, E. Mueller, and J. Spoerri, "Are Existing Monocular Computer Vision-Based 3D Motion Capture Approaches Ready for Deployment? a Methodological Study on the Example of Alpine Skiing," *Sensors*, vol. 19, no. 19, p. 4323, 2019.

[46] J. N. Kundu, S. Seth, V. Jampani, M. Rakesh, R. V. Babu, and A. Chakraborty, "Self-Supervised 3D Human Pose Estimation via Part Guided Novel Image Synthesis," in *Conference on Computer Vision and Pattern Recognition*, 2020.

[47] H.-Y. Tung, A. Harley, W. Seto, and K. Fragkiadaki, "Adversarial Inverse Graphics Networks: Learning 2D-To-3D Lifting and Image-To-Image Translation from Unpaired Supervision," in *International Conference on Computer Vision*, 2017.

[48] C. Chen, A. Tyagi, A. Agrawal, D. Drover, R. MV, S. Stojanov, and J. M. Rehg, "Unsupervised 3D Pose Estimation with Geometric Self-Supervision," in *Conference on Computer Vision and Pattern Recognition*, 2019.

[49] A. Kanazawa, M. J. Black, D. W. Jacobs, and J. Malik, "End-To-End Recovery of Human Shape and Pose," in *Conference on Computer Vision and Pattern Recognition*, 2018.

[50] B. Wandt and B. Rosenhahn, "Repnet: Weakly Supervised Training of an Adversarial Reprojection Network for 3D Human Pose Estimation," in *Conference on Computer Vision and Pattern Recognition*, 2019.

[51] G. Wang, W. Li, M. Aertsen, J. Deprest, S. Ourselin, and T. Vercauteren, "Aleatoric Uncertainty Estimation with Test-Time Augmentation for Medical Image Segmentation with Convolutional Neural Networks," *Neurocomputing*, vol. 338, pp. 34–45, 2019.

[52] N. Kolotouros, G. Pavlakos, M. J. Black, and K. Daniilidis, "Learning to Reconstruct 3D Human Pose and Shape via Model-Fitting in the Loop," in *Conference on Computer Vision and Pattern Recognition*, 2019.

[53] U. Iqbal, P. Molchanov, and J. Kautz, "Weakly-Supervised 3D Human Pose Learning via Multi-View Images in the Wild," in *Conference on Computer Vision and Pattern Recognition*, 2020.

[54] B. Wandt, M. Rudolph, P. Zell, H. Rhodin, and B. Rosenhahn, "Canonpose: Self-Supervised Monocular 3D Human Pose Estimation in the Wild," in *Conference on Computer Vision and Pattern Recognition*, 2021.

[55] H. Fang, S. Xie, Y. Tai, and C. Lu, "RMPE: Regional Multi-Person Pose Estimation," in *International Conference on Computer Vision*, 2017.

[56] J. Martinez, R. Hossain, J. Romero, and J. Little, "A Simple Yet Effective Baseline for 3D Human Pose Estimation," in *International Conference on Computer Vision*, 2017.

[57] K. He, H. Fan, Y. Wu, S. Xie, and R. Girshick, "Momentum Contrast for Unsupervised Visual Representation Learning," in *Conference on Computer Vision and Pattern Recognition*, 2020, pp. 9729–9738.

[58] H. Bao, L. Dong, S. Piao, and F. Wei, "BEiT: BERT Pre-Training of Image Transformers," in *International Conference on Learning Representations*, 2022.

[59] K. He, X. Chen, S. Xie, Y. Li, P. Dollár, and R. Girshick, "Masked Autoencoders Are Scalable Vision Learners," in *Conference on Computer Vision and Pattern Recognition*, 2022, pp. 16 000–16 009.

[60] D. P. Kingma and J. Ba, "Adam: A Method for Stochastic Optimisation," in *International Conference on Learning Representations*, 2015.

[61] H. Li, Z. Xu, G. Taylor, C. Studer, and T. Goldstein, "Visualizing the Loss Landscape of Neural Nets," in *Advances in Neural Information Processing Systems*, 2018, pp. 6389–6399.

[62] D. Bieler, S. Gunel, P. Fua, and H. Rhodin, "Gravity as a Reference for Estimating a Person's Height from Video," in *International Conference on Computer Vision*, 2019.

[63] C. Clauser, J. McConville, and J. Young, "Weight, Volume, and Center of Mass Segments of the Human Body," *Journal of Occupational and Environmental Medicine*, vol. 13, no. 5, p. 270, 1971.

Sina Honari is a post-doctoral reseacher in computer vision laboratory at EPFL, where he works on 3D human pose estimation, multi-view geometry, road-obstacle detection, and animal neural decoding. He received his PhD from University of Montreal and Mila lab, and then was a post-doctoral researcher at Polytechnique Montreal. His research interests are computer vision and deep learning, in particular self-supervised and semi-supervised methods.

Victor Constantin received his M.Sc. degree in Computer Science from EPFL in 2016. He was working as a Research Engineer in Computer Vision Laboratory, EPFL. His research interests include deep learning, 3D pose estimation and garment draping. He is now a machine learning engineer at Uplift Labs.

Helge Rhodin is an Assistant Professor at the University of British Columbia. His research interests range from computer graphics and augmented reality, over 3D computer vision, to self-supervised machine learning. Helge received the BSc and MSc degree in Computer science from Saarland University. He graduated with a PhD in 2016 for is work at the Max-Planck Institute for Informatics and was a postdoctoral researcher and lecturer at EPFL.

Mathieu Salzmann is a Senior Researcher at EPFL and an Artificial Intelligence Engineer at ClearSpace. Previously, he was a Senior Researcher and Research Leader in NICTA's computer vision research group, a Research Assistant Professor at TTI-Chicago, and a postdoctoral fellow at ICSI and EECS at UC Berkeley. He obtained his PhD in 2009 from EPFL. His research interests lie at the intersection of machine learning and computer vision.

Pascal Fua received an engineering degree from Ecole Polytechnique, Paris, in 1984 and a Ph.D. in Computer Science from the University of Orsay in 1989. He joined EPFL in 1996 as a Professor in the School of Computer and Communication Science. He is the head of the Computer Vision Lab. He has (co)authored over 300 publications in refereed journals and conferences and received several ERC grants. He is an IEEE Fellow and has been an Associate Editor of IEEE Transactions for Pattern Analysis and Machine Intelligence.

Temporal Representation Learning on Monocular Videos for 3D Human Pose Estimation — Supplementary Information

In Section S.1, we present training details and model architecture. In Section S.2, we provide additional examples on disentanglement of pose and appearance, as well as qualitative results on pose estimation. Finally, we present details on gravity sensitivity analysis in Section S.3.

S.1 TRAINING PROCEDURE

We use images of resolution 500×500 in the H36 and MPI datasets, 500×1000 in the Diving dataset, and 640×360 in the Ski dataset. We train all models using the Adam optimizer [60] and a fixed learning rate of 1e-4. In theory, we could train our full model by simply minimizing the total loss $\mathcal{L}_{\text{total}}$ of Eq.(14). However, because the model has to learn first where to detect the object, it takes a long time for the network to find the correct locations for the bounding boxes and convergence is slow. To speed it up, we introduce priors for the bounding box locations.

For **H36M, MPI, and Ski**, we apply a Gaussian prior on the location and size of the bounding boxes. We set their average detected size over the mini-batch to be about 50% of the input image resolution. On H36M we also set a prior that their average locations be in the middle of the frame. The latter was not needed on the MPI and Ski datasets. Note that these priors are applied to the average prediction of the mini-batch and not individual ones, which would otherwise imply a strong bias. We use this prior for all models including the baselines, except when spatial transformer networks are not used, as in the ablation study.

For **Diving**, which depicts free-fall motion, we use the gravity loss introduced in Section 3.2 in addition to the other objectives. It requires sampling frames that are temporally equidistant. Since this sampling is different from the one required for the DSL terms and the mini-batch size is limited, we developed a two-stage training strategy. First, the network is pre-trained with $\mathcal{L}_{\text{DSL}}^{\text{all}}$ of Eq. (14) with only gravity and reconstruction enabled, which helps localize the divers even though they occupy only a small fraction of image frame. Then, we switch on the DSL loss of $\mathcal{L}_{\text{DSL}}^{\text{all}}$ of Eq. (14) with its sampling to learn a more precise human model for which the gravity loss is no longer required. We use the gravity prior as pre-training for all models including the baselines because, without it, it takes them longer to detect the subject. The only exception is Ours (no $\mathcal{L}_{track}$), which acts as an ablation study on the gravity prior. We use $\mathcal{L}_{track}$ only on this dataset and we set γ to 1.

Background Estimation. As our model requires a background image, we apply the following procedure to estimate the background in each dataset. On H36M and MPI we use the per-pixel median, where the median is taken over all video frames. On the Diving dataset, we use the bottom half of the first frame and the top half of the last frame because the diver goes from the top of the frame to the bottom in the video. In the Ski dataset, the cameras rotate. We therefore compute a homography between consecutive frames using SIFT correspondences and generate a panorama in which we again take the median pixel values, as shown in Fig. S5.

Hyper-parameters. On the H36M, MPI, and Diving datasets we set λ and ρ to 2. On the Ski dataset, to help the model better capture the foreground, we reduce the RGB coefficient λ to 0.05 while keeping ρ at 2. This encourages the model to focus more on higher-level ImageNet features corresponding to structures in the image and hence to encode the skier better. We set α to 0.01 on the all datasets. We used the following temporal distance for sampling in models that exploit a contrastive loss: In H36M, a maximum of 10 temporal frame distance for positive samples and a minimum of 200 for negative samples. In MPI, 10 for positive and 450 for negative samples. In Diving and Ski, poses change more quickly within a few frames, so in Ski we use 5 for positive and 40 for negative samples whereas in Diving we use 3 for positive and 20 for negative samples. These hyper-parameters were selected based on how fast the poses change frame-to-frame and how the choice of these hyper-parameters impacts the quality of image reconstruction in the self-supervised models, which is our validation set model-selection criterion in the self-supervised model. We used the same sampling procedure for all models that use a contrastive loss.

S.1.1 Model Architecture

To implement our model, we use a ResNet18 [43] as detector that takes as input an image downsampled by factor 4 in both dimensions and returns $\mathcal{S}(\mathbf{I}) = (s^x, s^y, u^x, u^y)$. This crops a patch $\mathbf{I}_{\text{crop}}$ that is then resized to 128×128 using the affine-grid PyTorch function. Unless otherwise specified, we use a ResNet50 [43] model as the encoder for both $\mathcal{E}_{ti}$ and $\mathcal{E}_{tv}$. They share parameters, except in the output layer, where $\mathbf{I}_{ti}$ has a resolution of 129 and $\mathbf{I}_{tv}$ a resolution of 600 in H36M and Ski datasets and 900 in MPI and Dive datasets. In the decoder $\mathcal{D}$, $\mathbf{I}_{tv}$ features are first passed to a fully-connected layer with ReLU and dropout with an output size of $32{,}768$, which is then reshaped to 128 feature maps of resolution 16×16. The $\mathbf{I}_{ti}$ features are passed to another fully-connected layer with ReLU and dropout to produce an output of shape 128×16^2. The two feature maps are then concatenated into a single one of size 256×16^2. This feature map is then passed to three upsampling layers, each with one bilinear upsampling with factor 2 and three convolutional layers with kernel size 3×3. This upsamples the features maps from 256×16^2, in order, to 128×32^2, 64×64^2, and finally to 32×128^2. A final convolutional layer it then applied to make the resolution of these feature maps 4×128^2, where the first three channels are the decoded RGB image crop $\mathbf{D}_{\text{crop}}$ and the last channel is the foreground subject mask crop $\mathbf{M}_{\text{crop}}$. Finally, the inverse STN function is used to map the resolution of $\mathbf{D}_{\text{crop}}$ and $\mathbf{M}_{\text{crop}}$ to that of the original image and reconstructing the final image $\hat{\mathbf{I}}$ according to Eq.(1). Figure S1 depicts architecture details.

Pose Model. For the 3D pose estimation results presented in Sections 4.3 and 4.4, following [7] we use a two-layer MLP for Φ. It features two hidden layers of size 2048 with 50% dropout applied to each hidden layer. The output layer dimension is $3 \times K$,
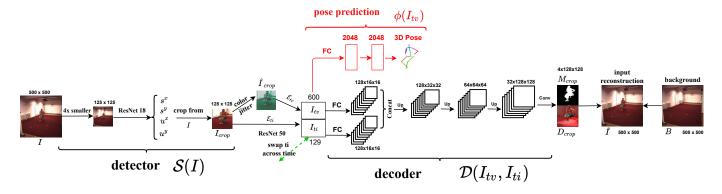
Fig. S1: Model architecture details. This image depicts the implementation details described in Section S.1.1.

Fig. S2: Swapping time-invariant and time-variant components. Top two rows show two source images S1 and S2. In each column, in the third row the time-invariant component is taken from S2, and the time-variant component is taken from S1 and concatenated to generate the image shown, using the decoder. In the last row, the sources are swapped. As can be observed the model manages to correctly swap time-invariant and time-variant components. However, due to low-dimensional feature representation of the latent components, not all details are maintained. For examples, the clothes details are lost.

where K is the number of joints, which is equal to 17 in the H36M and Ski datasets, 24 for the 3DPW dataset, and to 13 in the Diving and MPI datasets.

For experiments in Section 4.5, unless otherwise specified, we use a multi-layer perceptron (MLP) ResNet architecture, taken from [56] for Φ, which helps yielding more accurate 3D pose estimation. The ResNet consists of three residual blocks, each containing a two-layer MLP with ReLU non-linearity of dimension 1024. The three blocks is preceded by a linear layer that maps from time-variant resolution to 1024, and is followed by a linear layer that outputs the 3D pose.

When training the pose network, the self-supervised parameters are frozen, in order to evaluate the quality of extracted features in the self-supervised phase. The only except is the experiment in Figure 8, where the impact of pre-training for model initialization is evaluated. For all models that split the latent features, we use only the time-variant component for pose estimation. For the models that do not—CSS, MV-CSS, AE-STN-Dec—we use all the latent features.

Weak Labels. When 2D pose label is used, we use the same architecture as our 3D MLP ResNet (without sharing parameters), with the difference that it outputs the 2D pose. The input to this model is the time-variant features. We set η to 10 in all experiments that use 2D.

Considering the upgraded encoder backbone in Section 4.5, we replace the ResNet50 with a ResNet101 module, taken from a pre-trained Alphapose model [55] on CrowdPose [61] dataset. Note that this module is only part of the Alphapose model, hence we do not use its 2D predictions. The input image resolution to this model is 256×320 and it outputs a feature of dimentionality
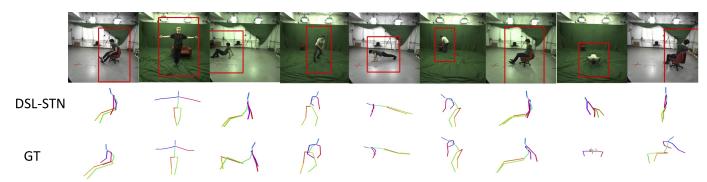
Fig. S3: Qualitative results on MPI Dataset. Samples are taken from the test set. In the top row, the red rectangle shows the bounding box detected by our model for image cropping. The two bottom rows, respectively show the pose estimation by our model and the ground truth pose. The last two columns show examples with high error.
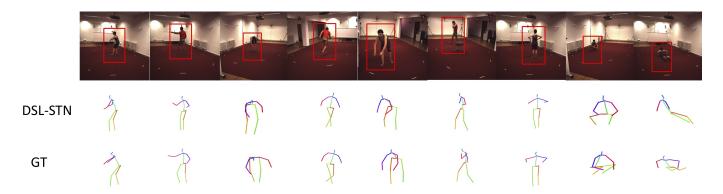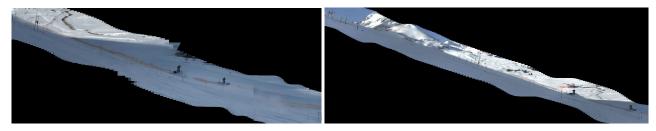

Fig. S4: Qualitative results on H36M Dataset. Samples are taken from the test set. In the top row, the red rectangle shows the bounding box detected by our model for image cropping. The two bottom rows, respectively show the pose estimation by our model and the ground truth pose. The last two columns show examples with high error.

Fig. S5: Ski dataset background samples. Two samples of estimated ski backgrounds are shown.

$2048 \times 10 \times 8$. To map it to our low-dimensional time-variant (tv) and time-invariant (ti) features, we use the same MLP residual block similar to the pose model described above, by changing its input to take ResNet101 features and its output to tv and ti dimensionality.

3D Pose Results. We report the results using the normalized mean per joint position error (N-MPJPE), expressed in mm, in Sections 4.3 and 4.4. The results in Section 4.5 are reported with mean per joint position error (MPJPE), the normalized N-MPJPE, and the procrustes-aligned P-MPJPE. To test the performance of

the models given only small sets of labeled data, we use different fractions of the training set. To this end, when the dataset is sub-sampled, we sub-sample every K frames of the training videos. Hence, sub-sampling is applied uniformly across dataset subjects and actions.

S.1.2 Computational Complexity.

A forward pass of the self-supervised network, including detector, encoder and decoder, for a mini-batch of size 32 takes 0.002 seconds on a a Tesla V100 GPU. We use the same backbone

and model architecture for DrNet and NSD models. Thus, their computational complexity is almost the same as our model. By contrast, CSS and MV-CSS do not use the detection and decoding components. This makes training these models faster. However, as we showed in Table 1, these models are worse than our approach by a factor of around 2 and have a big gap compared to other models that reconstruct the image.

Considering the results in Section 4.3, on H36M, the self-supervised models are trained for 250K iterations, which takes 4.5 days on a Tesla V100 GPU. The models converge at about 100K iterations, and further iterations yield negligible improvement. The pose models are trained for 50K iterations, which takes 4 hours on the same GPU. On the MPI dataset, the self-supervised models are trained for 550K iterations and the pose models for 50K iterations. On the Diving dataset, we trained self-supervised models for 500K iterations (taking about 7.25 days on a Tesla V100 GPU). Due to GPU limitations we trained on mini-batches of size 16 for this dataset. We used mini-batched of size 32 on all other datasets. The pose models were trained for 40K iterations, taking about 9 hours on the same GPU. On the Ski dataset, the self-supervised models are trained for 400K (taking about 4 days on Tesla V100 GPU) and the pose models for 100K iterations. These same settings were applied to all models trained on a given dataset.

Considering the results in Section 4.5, on the H36M and MPI datasets, the self-supervised models are trained for 200K and the pose models for 100K iterations. On the Diving dataset, the self-supervised models are trained for 300K and the pose models for 200K iterations.

S.2 QUALITATIVE RESULTS

S.2.1 Disentanglement.

In Figures S2 we take time-invariant component from an image and time-variant component from a different image to reconstruct an output image. As can be observed, the model learns to separate these two components.

S.2.2 Prediction Samples.

In Figures S3 and S4 we show pose estimation results on MPI and H36M test sets. In Figures S7, S8, S9, we show samples on the Diving dataset and visualize detected bounding box, segmentation mask, decoded image, as well as final pose predictions by different models. We also visualize side-view results on H36M and Diving datasets in Figures S10 and S11 to show the quality of the learned depth.

Following [62], we measure the center of gravity (CG) of the body by using the ground truth 3D joints, where we first get the CG location of each body-part (such as forearm, upper-arm, thigh), as provided in [63], then we measure the CG of the entire body as the weighted average of CG of each body-part multiplied in its weight ratio (as provided in [63]), yielding $\text{CG}_{\text{body}} = \sum_j^J w_j \text{CG}_j$, where w_j is the weight ratio of body part j w.r.t to the entire body weight and CG_j is the 3D location of the CG of body part j. We then project the estimated CG to the image using camera projection matrix and also visualize our center of bounding box (bbox), which estimates the CG. While the estimated CG and center of bbox have a gap, our estimated center of bbox can guide the model to correctly localize the person.

S.3 ACCELERATION SENSITIVITY ANALYSIS

An object's location $\mathbf{p}(t)$ at time t falling down in 3D space using gravity is formulated by $\mathbf{p}(t) = 0.5\mathbf{g}t^2 + \mathbf{v}_0 t + \mathbf{p}_0$, where $\mathbf{g}$, $\mathbf{v}_0$, and $\mathbf{p}_0$ respectively indicating gravity, initial speed, and initial location. We can write the projection of the center of gravity CG in the image plane in projective coordinates as

$$\mathbf{p}_{\text{proj}} = [p_{\text{proj}}^x, p_{\text{proj}}^y, p_{\text{proj}}^z] = \mathbf{K}[\mathbf{R}|\mathbf{T}](0.5\mathbf{g}t^2 + \mathbf{v}_0 t + \mathbf{p}_0), \tag{16}$$

where $\mathbf{g}$, $\mathbf{v}_0$, and $\mathbf{p}_0$ are represented in homogeneous coordinates, each being a column vector of 4×1. this yields

$$\begin{bmatrix} p_{\text{proj}}^x \\ p_{\text{proj}}^y \\ p_{\text{proj}}^z \end{bmatrix} = \begin{bmatrix} \mathbf{K}_1^T[\mathbf{R}|\mathbf{T}](0.5\mathbf{g}t^2 + \mathbf{v}_0 t + \mathbf{p}_0) \\ \mathbf{K}_2^T[\mathbf{R}|\mathbf{T}](0.5\mathbf{g}t^2 + \mathbf{v}_0 t + \mathbf{p}_0) \\ \mathbf{K}_3^T[\mathbf{R}|\mathbf{T}](0.5\mathbf{g}t^2 + \mathbf{v}_0 t + \mathbf{p}_0) \end{bmatrix} \tag{17}$$

where $\mathbf{K}_i^T$ is the i-th row of matrix $\mathbf{K}$. The observed x and y locations are then equal to

$$p^x = p_{\text{proj}}^x / p_{\text{proj}}^z, \quad p^y = p_{\text{proj}}^y / p_{\text{proj}}^z .$$

When the image plane is parallel to the plane in which the subject travels or in orthographic projection, the depth p_{proj}^z remains fixed or is ignored in the course of dive, so we can treat is as a constant c. In this case, we then get

$$\begin{bmatrix} p^x \\ p^y \end{bmatrix} = \frac{1}{c} \begin{bmatrix} \mathbf{K}_1^T[\mathbf{R}|\mathbf{T}](0.5\mathbf{g}t^2 + \mathbf{v}_0 t + \mathbf{p}_0) \\ \mathbf{K}_2^T[\mathbf{R}|\mathbf{T}](0.5\mathbf{g}t^2 + \mathbf{v}_0 t + \mathbf{p}_0) \end{bmatrix} \tag{18}$$

Getting the second order derivative yields

$$\begin{bmatrix} a^x \\ a^y \end{bmatrix} = \frac{1}{c} \begin{bmatrix} \mathbf{K}_1^T[\mathbf{R}|\mathbf{T}]\mathbf{g} \\ \mathbf{K}_2^T[\mathbf{R}|\mathbf{T}]\mathbf{g} \end{bmatrix} \tag{19}$$

with a^x and a^y being acceleration in x and y directions. This makes acceleration a constant value dependent on intrinsic $\mathbf{K}$, rotation $\mathbf{R}$, translation $\mathbf{T}$, and gravity $\mathbf{g}$, hence not a function of time t. In this case, our assumption of constant acceleration strictly holds.

On the other hand, if the subject travels on a plane not parallel to the image plane, the depth of observed location changes over time, making it a function of t. In this case for y location we have (p^x is obtained similarly)

$$p^y = \frac{p_{\text{proj}}^y}{p_{\text{proj}}^z} = \frac{\mathbf{K}_2^T[\mathbf{R}|\mathbf{T}](0.5\mathbf{g}t^2 + \mathbf{v}_0 t + \mathbf{p}_0)}{\mathbf{K}_3^T[\mathbf{R}|\mathbf{T}](0.5\mathbf{g}t^2 + \mathbf{v}_0 t + \mathbf{p}_0)} \tag{20}$$

Getting derivative w.r.t t we have

$$v^y = \frac{\partial p^y}{\partial t} = \tag{21}$$

$$\frac{\mathbf{K}_2^T[\mathbf{R}|\mathbf{T}](\mathbf{g}t + \mathbf{v}_0)(p_{\text{proj}}^z) - \mathbf{K}_3^T[\mathbf{R}|\mathbf{T}](\mathbf{g}t + \mathbf{v}_0)(p_{\text{proj}}^y)}{(p_{\text{proj}}^z)^2} =$$

$$\frac{(p_{\text{proj}}^z \mathbf{K}_2^T - p_{\text{proj}}^y \mathbf{K}_3^T)([\mathbf{R}|\mathbf{T}](\mathbf{g}t + \mathbf{v}_0))}{(p_{\text{proj}}^z)^2}$$

Getting second derivative w.r.t t we obtain

$$a^y = \frac{\partial v^y}{\partial t} = \tag{22}$$

$$\frac{(p_{\text{proj}}^z \mathbf{K}_2^T - p_{\text{proj}}^y \mathbf{K}_3^T)([\mathbf{R}|\mathbf{T}]\mathbf{g})(p_{\text{proj}}^z)^2}{(p_{\text{proj}}^z)^4} -$$

$$\frac{2p_{\text{proj}}^z(\frac{\partial}{\partial t} p_{\text{proj}}^z)(p_{\text{proj}}^z \mathbf{K}_2^T - p_{\text{proj}}^y \mathbf{K}_3^T)([\mathbf{R}|\mathbf{T}](\mathbf{g}t + \mathbf{v}_0))}{(p_{\text{proj}}^z)^4} =$$

$$\frac{p_{\text{proj}}^z(p_{\text{proj}}^z \mathbf{K}_2^T - p_{\text{proj}}^y \mathbf{K}_3^T)[\mathbf{R}|\mathbf{T}](p_{\text{proj}}^z \mathbf{g} - (2(\frac{\partial}{\partial t} p_{\text{proj}}^z)(\mathbf{g}t + \mathbf{v}_0)))}{(p_{\text{proj}}^z)^4}$$
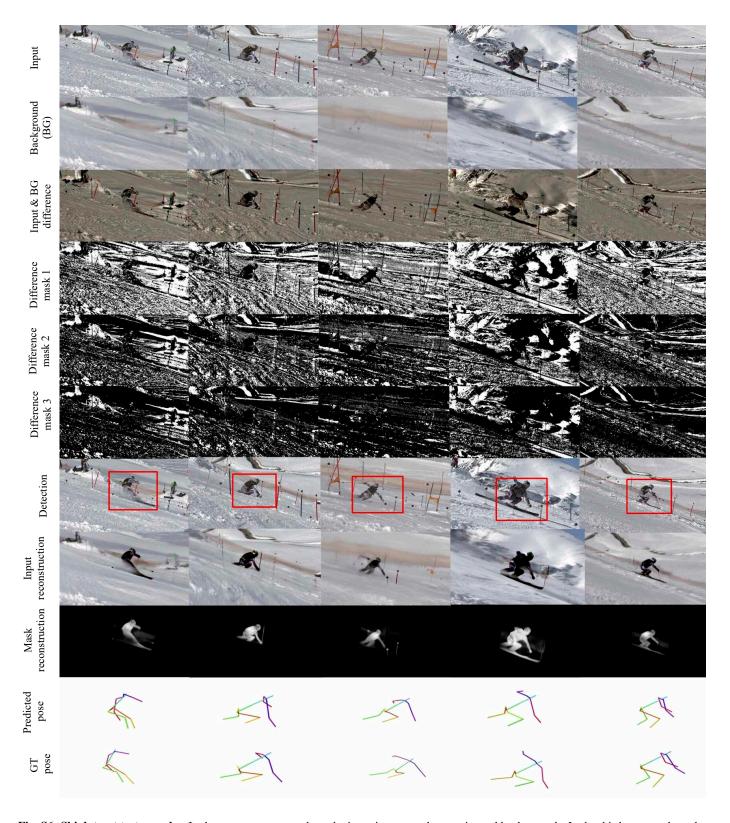
Fig. S6: Ski dataset test samples. In the top two rows we show the input images and our estimated backgrounds. In the third row we show the background subtracted input image. As can be observed the background is still present in the image. We also apply thresholding to this image and create masks by using three different thresholds. The resulted images are shown in rows 4 to 6. The thresholding still does not remove the background and even increasing the threshold does not help. In the next four rows we show the output of our model, which are detected subject bounding box, the reconstructed image, the predicted mask, and finally the estimated pose.
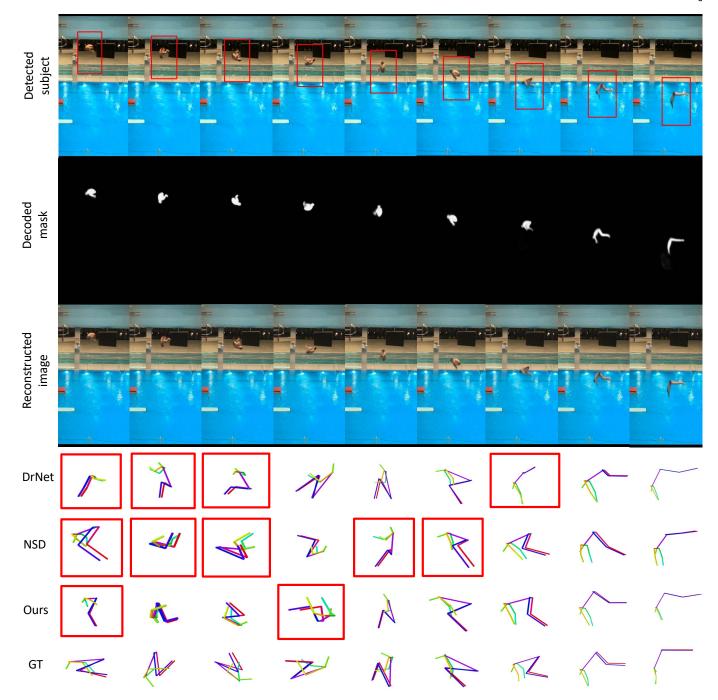
Fig. S7: Qualitative results on Diving Dataset. Samples are taken from the test set. In the top row, the red rectangle shows the bounding box detected by our model for image cropping. The second and third rows, show respectively the decoded mask and image. Bottom rows, show the pose estimation by different models, where the red rectangle indicates pose estimation with a high error. GT in the bottom row indicates the ground truth labels.

As can be seen in the last line of Eq.(22), both the nominator and denominator depend on p_{proj}^z, which is the depth. It changes over time and is therefore a function of t. This violates our assumption that the acceleration is constant in the image plane when the subject moves in a plane not parallel to the image plane. However, p_{proj}^z only changes slowly for points far away from the camera because perspective projection gets increasingly close to orthographic projection as one moves away from the camera.
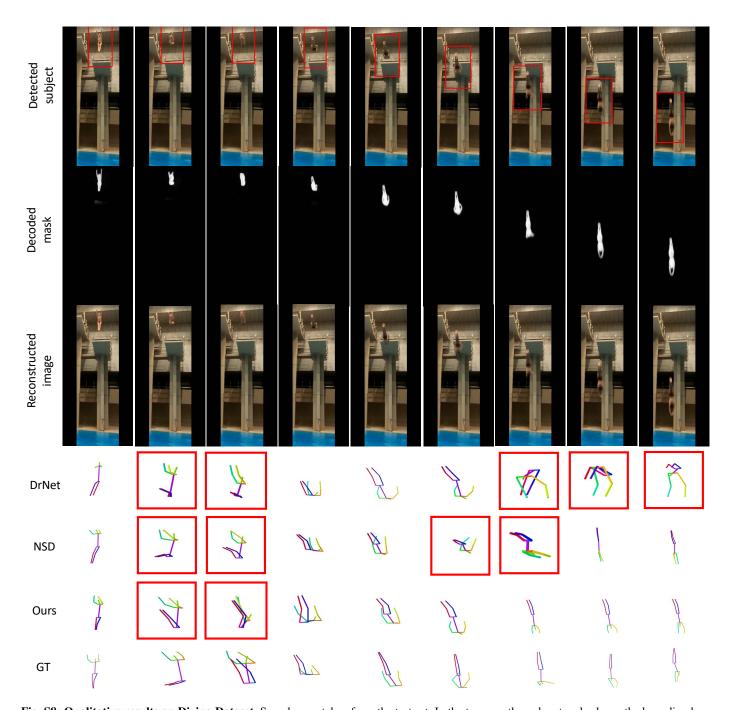
Fig. S8: Qualitative results on Diving Dataset. Samples are taken from the test set. In the top row, the red rectangle shows the bounding box detected by our model for image cropping. The second and third rows, show respectively the decoded mask and image. Bottom rows, show the pose estimation by different models, where the red rectangle indicates pose estimation with a high error. GT in the bottom row indicates the ground truth labels.
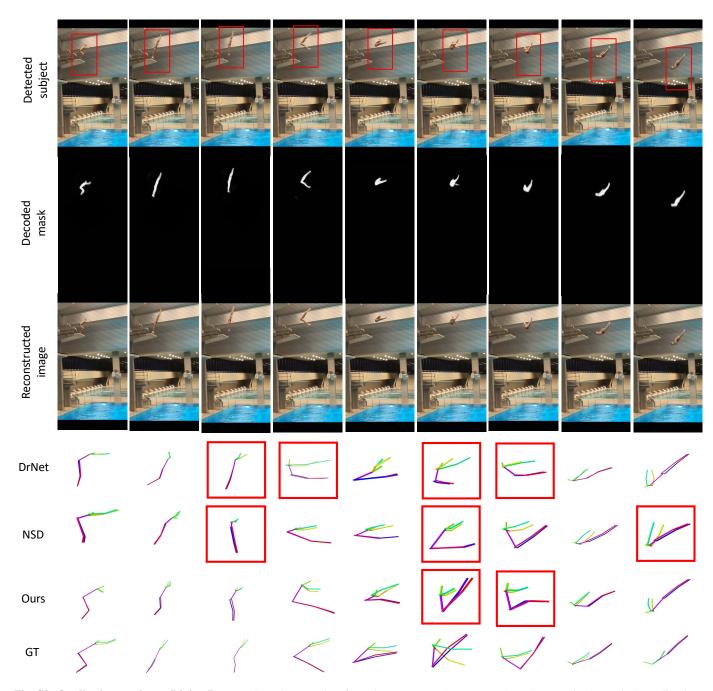
Fig. S9: Qualitative results on Diving Dataset. Samples are taken from the test set. In the top row, the red rectangle shows the bounding box detected by our model for image cropping. The second and third rows, show respectively the decoded mask and image. Bottom rows, show the pose estimation by different models, where the red rectangle indicates pose estimation with a high error. GT in the bottom row indicates the ground truth labels.
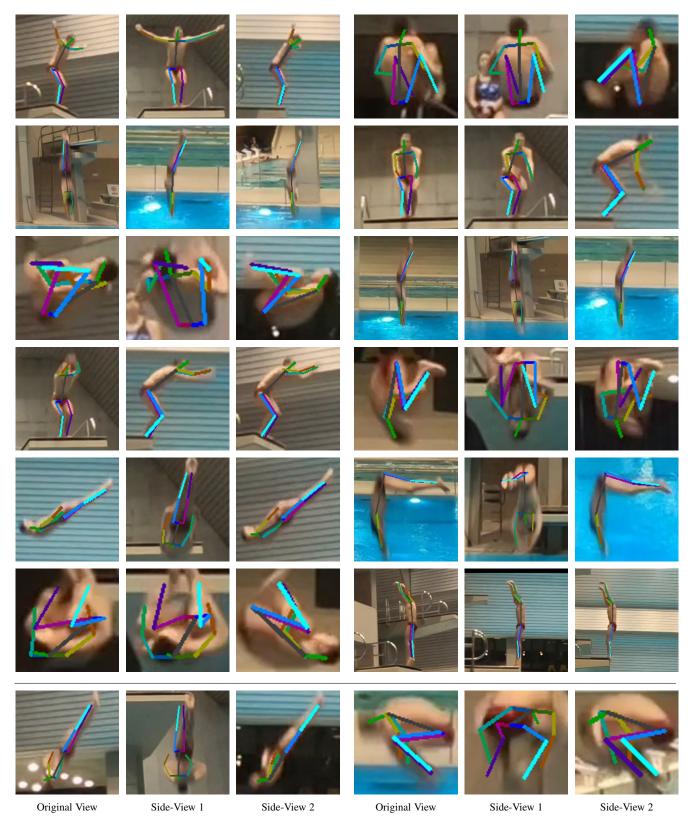
| Original View | Side-View 1 | Side-View 2 | Original View | Side-View 1 | Side-View 2 |

Fig. S10: Side-view visualization on diving dataset. For each triplet, the first image shows the prediction of the model on the input view, while side-views 1 and 2 show projection of the 3D pose to other views. The last row shows erroneous cases where the model's prediction in the original view is plausible while it does not correspond to correct poses in other views.

| Original View | Side-View 1 | Side-View 2 | Original View | Side-View 1 | Side-View 2 |

Fig. S11: Side-view visualization on H36M dataset. For each triplet, the first image shows the prediction of the model on the input view, while side-views 1 and 2 show projection of the 3D pose to other views. The last row shows erroneous cases where the model's prediction in the original view is plausible while it does not correspond to correct poses in other views.