

Tracking People by Predicting 3D Appearance, Location and Pose

Jathushan Rajasegaran, Georgios Pavlakos, Angjoo Kanazawa, Jitendra Malik
UC Berkeley

Abstract

We present an approach for tracking people in monocular videos by predicting their future 3D representations. To achieve this, we first lift people to 3D from a single frame in a robust manner. This lifting includes information about the 3D pose of the person, their location in the 3D space, and the 3D appearance. As we track a person, we collect 3D observations over time in a tracklet representation. Given the 3D nature of our observations, we build temporal models for each one of the previous attributes. We use these models to predict the future state of the tracklet, including 3D appearance, 3D location, and 3D pose. For a future frame, we compute the similarity between the predicted state of a tracklet and the single frame observations in a probabilistic manner. Association is solved with simple Hungarian matching, and the matches are used to update the respective tracklets. We evaluate our approach on various benchmarks and report state-of-the-art results. Code and models are available at: <https://brjathu.github.io/PHALP>.

1. Introduction

When we watch a video, we can segment out individual people, cars, or other objects and track them over time. The corresponding task in computer vision has been studied for several decades now, with a fundamental choice being whether to do the tracking in 2D in the image plane, or of 3D objects in the world. The former seems simpler because it obviates the need for inferring 3D, but if we do take the step of back-projecting from the image to the world, other aspects such as dealing with occlusion become easier. In the 3D world the tracked object doesn't disappear, and even young infants are aware of its persistence behind the occluder. In our recent work [35], we presented experimental evidence that performance is better with 3D representations. In this paper, we will take this as granted, and proceed to develop a system in the 3D setting of the problem. While our approach broadly applies to any object category where parameterized 3D models are available and can be inferred from images, we will limit ourselves in this paper to study-



Figure 1. **Tracking people by predicting and matching in 3D.** The top row shows our tracking results at three different frames. The results are visualized by a colored head-mask for unique identities. The second and third rows show renderings of the 3D states of the two people in their associated tracklets. The bottom row shows the bottom-up detections in each image frame which, after being lifted to 3D, will be matched with the 3D predictions of each tracklet in the corresponding frame. Note how in the middle frame of second row, the 3D representation of the person persists even though they are occluded in the image. Readers are encouraged to watch the videos at the [project website](#).

ing people, the most important case in practice.

Once we have accepted the philosophy that we are tracking 3D objects in a 3D world, but from 2D images as raw data, it is natural to adopt the vocabulary from control theory and estimation theory going back to the 1960s. We are interested in the “state” of objects in 3D, but all we have access to are “observations” which are RGB pixels in 2D. In an online setting, we observe a person across multiple time frames, and keep recursively updating our estimate of the person’s state — their appearance, location in the world, and pose (configuration of joint angles). Since we have a dynamic model (a “tracklet”), we can also predict states at future times. When the next image frame comes in, we detect the people in it, lift them to 3D, and in that setting solve the association problem between these bottom-up detections and the top-down predictions of the different tracklets for this frame. Once the observations have been associated with the tracklets, the state of each person is re-

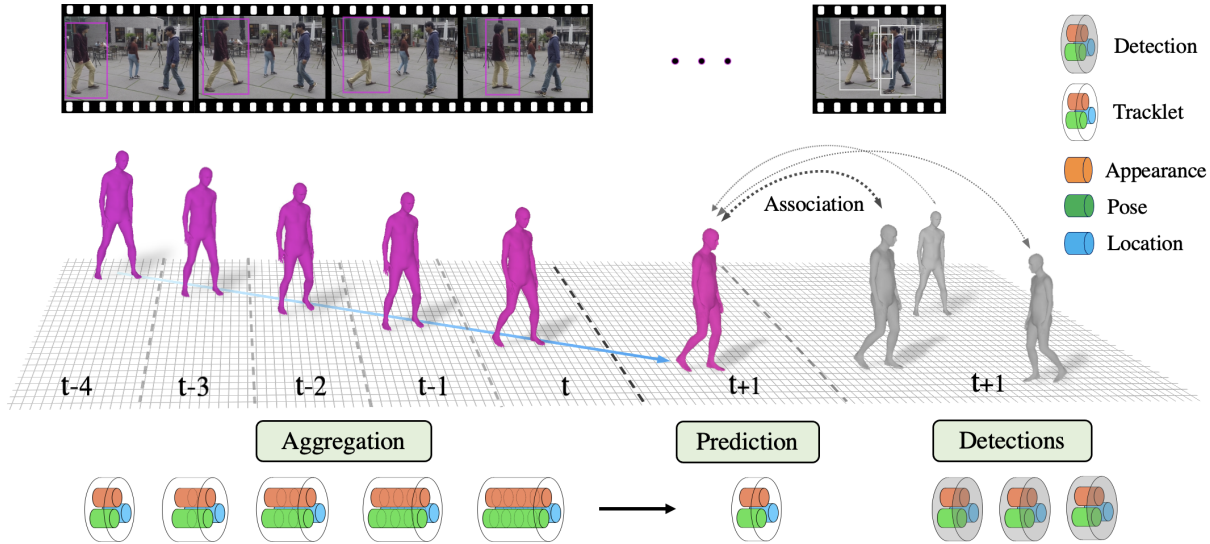


Figure 2. **PHALP: Predicting Human Appearance, Location and Pose for Tracking:** We perform tracking of humans in 3D from monocular video. For every input bounding box, we estimate a 3D representation based on the 3D appearance, 3D pose and 3D location of the person. During tracking, these are integrated to form corresponding tracklet-based representations. We perform tracking by predicting the future representation of each person and using it to solve for association given the detected bounding boxes of a future frame.

estimated and the process continues. Fig. 1 shows this process on a real video. Note that during a period of occlusion of a tracklet, while no new observations are coming in, the state of the person keeps evolving following their dynamics. It is not the case that “Out of sight, out of mind”!

In an abstract form, the procedure sketched in the previous paragraph is basically the same as that followed in multiple computer vision papers from the 1980s and 1990s. The difference is that in 2022 we can actually make it work thanks to the advances brought about by deep learning and big data, that enable consistent and reliable lifting of people to 3D. For this initial lifting, we rely on the HMAR model [35]. This is applied on every detected bounding box of the input video and provides us with their initial, single frame, observations for 3D pose, appearance as well as location of the person in the 3D space.

As we link individual detections into tracklets, these representations are aggregated across each tracklet, allowing us to form temporal models, *i.e.*, functions for the aggregation and prediction of each representation separately (see right side of Fig. 2). More specifically, for appearance, we use the canonical UV map of the SMPL model to aggregate appearance, and employ its most recent prediction as person’s appearance. For pose, we aggregate information using a modification of the HMMR model [17], where through its “movie strip” representation, we can produce 3D pose predictions. Finally, for 3D location, we use linear regression to predict the future location of the person.

This modeling enables us to develop our tracking system, **PHALP** (Predicting Human Appearance, Location and Pose for tracking), which aggregates information over

time, uses it to predict future states, and then associates the predictions with the detections. First, we predict the 3D location, 3D pose and 3D appearance for each tracklet for a short period of time. For a future frame, these predictions need to be associated with the detected people of the frame. To measure similarity, we adopt a probabilistic interpretation and compute the posterior probabilities of every detection belonging to each one of the tracklets, based on the three basic attributes. With the appropriate similarity metric, association is then easily resolved by means of the Hungarian algorithm. The newly linked detections can now update the temporal model of the corresponding tracklets for 3D appearance, 3D location and 3D pose in an online manner and we continue the procedure by rolling-out further prediction steps. The final output is an identity label for each detected bounding box in the video. Notably, this approach can also be applied on videos with shot changes, *e.g.*, movies [12], with minor modifications. Effectively, we modify our similarity to include only appearance and 3D pose information for these transitions, since they (unlike location) are not affected by the shot boundary.

2. Related work

Tracking. Object tracking is studied in various settings such as single object tracking, multi-object tracking for humans, and multi-object tracking for vehicles etc. The tracking literature is vast and we refer readers to [6, 8, 44] for a comprehensive summary. In general, tracking can be applied on any category, however, in this section we discuss the methods that focus on tracking people. These approaches often work in a tracking-by-detection setting,

where 2D location, 2D keypoint features [11, 37, 40] and 2D appearance [32, 42, 43] are used to associate detections over time. Quality of the detection plays a key role in this case and many works jointly learn or fine-tune their own detection models [4, 32]. In this work, we are interested in the effectiveness of 3D representations for tracking and thus assume that detected bounding boxes are provided, which we associate through our representations. On the other hand, tracking-by-regression [4, 10] predicts future locations using the knowledge of the past detections. While this alleviates the requirement for good quality detections, most of the works regress in the image plane. The projection from 3D world to the image plane makes it hard to make this prediction, therefore these methods need to learn non-linear motion models [2, 4, 47]. Compared to these methods, PHALP predicts short-term location in 3D coordinates, by simple linear regression. Additionally, we also predict appearance and pose features for better association.

Finally, there are methods that incorporate 3D information in tracking [30], however most of these approaches assume multiple input cameras [27, 46] or 3D point cloud observation from LiDAR [39]. In this paper we focus on the setting where the input is a monocular video. Recent works track occluded people based on object permanence [19, 38]. These methods rely on depth estimation network to get 3D location [19] or need to learn permanence using sequence of annotated data [38]. However, by placing full human mesh in 3D space and predicting their location, pose and appearance, object permanence is already built into our system.

Monocular 3D human reconstruction. Although there is a long history of monocular 3D human reconstruction methods, *e.g.*, [5, 13], here we focus on more recent works. Many of the relevant approaches rely on the SMPL model [28], which offers a low-dimensional parameterization of the human body. HMR [16] has been one of the most notable ones, using a neural network to regress the SMPL body parameters from a single image. Follow-up works have improved the robustness [22, 25] and added additional features like estimation of camera parameters [23], or probabilistic estimation of pose [26]. Recently, Rajasegaran *et al.* [35] introduced HMAR, by adding an appearance head to HMR. Other works have focused on extending HMR to the temporal dimension, *e.g.*, HMMR [17] and VIBE [21]. In this work, we make use of a modification of HMAR [35] as the main feature backbone, while employing a model similar to HMMR [17] for temporal pose prediction, but instead, using a transformer to aggregate pose information over time [34]. Regarding human motion prediction, Kanazawa *et al.* [17], regress future poses from the temporal pose representation of HMMR, the “movie-strip”. Zhang *et al.* [45] extend this to PHD, employing autoregressive prediction of human motion. Aksan *et al.* [1] also regress future motion in an autoregressive manner, using a transformer.

3. Method

Tracking humans using 3D representations has significant advantages, including that 3D appearance is invariant to pose variations and 3D provides the ability to have amodal completion for humans during partial occlusion. Our tracking algorithm accumulates these 3D representations over time, to achieve better association with the detections. PHALP has three main stages: 1) estimating 3D representations for each human detection, 2) aggregating representations over time and predicting their future state, 3) associating tracklets with detections using predicted representations in a probabilistic framework. We explain each stage in the next sections.

3.1. Single-frame processing

The input to our system is a set of person detections along with their masks, estimated by conventional detection networks, like Mask-RCNN [14]. Each detection is processed by our backbone that computes the basic representations for pose, appearance and location on a single-frame basis. For this feature extraction we use a modification of the HMAR model [35]. HMAR returns a feature representation for the 3D pose \mathbf{p} , for appearance \mathbf{a} , while it can recover an estimate for the 3D location \mathbf{l} of the person.

The HMAR model takes as input the pixels in the bounding box corresponding to a detected person. This means that in a crowded, multi-person scenario, the input can contain pixels from more than one person, potentially confusing the network. To deal with this problem, we modify HMAR to take as additional input, the pixel level mask of the person of interest (obtained from Mask R-CNN [14]) and re-train HMAR. Obviously, we cannot expect this step to be perfect, since there can be inaccuracies in the bounding box or mask estimates. However, we observed that the model is more robust in the case of close person-person interactions, which are common in natural videos.

3.2. 3D tracklet prediction

The 3D estimates provide a rich and expressive representation for each detection. However, this is the result of single-frame processing. During tracking, as we expand each tracklet, we have access to more information that is representative of the state of the tracked person. To properly leverage this information, our tracking algorithm builds a tracklet representation after each step of the online processing, which allows us to predict the future states for each tracklet. In this section we describe how we build this tracklet representation, and more importantly, how we use it to *predict* the future state of each tracklet.

Appearance: The appearance pathway is used to integrate appearance information for each person over multiple frames. The single frame appearance representation

for the person i at time step t , \mathbf{A}_t^i , is taken from the HMAR model by combining the UV image of that person $\mathbf{T}_t^i \in \mathcal{R}^{3 \times 256 \times 256}$ and the corresponding visibility map $\mathbf{V}_t^i \in \mathcal{R}^{1 \times 256 \times 256}$ at time step t :

$$\mathbf{A}_t^i = [\mathbf{T}_t^i, \mathbf{V}_t^i] \in \mathcal{R}^{4 \times 256 \times 256} \quad (1)$$

Note that the visibility mask $\mathbf{V}_t^i \in [0, 1]$ indicates whether a pixel in the UV image is visible or not, based on the MaskRCNN mask. Now, if we assume that we have established the identity of this person in neighboring frames, we can integrate the partial appearance coming from the independent frames to an overall tracklet appearance for the person. Using the set of single frame appearance representations $\mathcal{A}^i = \{\mathbf{A}_t^i, \mathbf{A}_{t-1}^i, \mathbf{A}_{t-2}^i, \dots\}$, after every new detection we create a single per-tracklet appearance representation:

$$\hat{\mathbf{A}}_t^i = \Phi_A(\hat{\mathbf{A}}_{t-1}^i, \mathbf{A}_t^i) = (1 - \alpha) * \hat{\mathbf{A}}_{t-1}^i + \alpha * \mathbf{A}_t^i, \quad (2)$$

where $\alpha = \begin{cases} \alpha_0, & \text{if } \hat{\mathbf{V}}_{t-1}^i = 1 \text{ and } \mathbf{V}_t^i = 1 \\ 1, & \text{if } \hat{\mathbf{V}}_{t-1}^i = 0 \text{ and } \mathbf{V}_t^i = 1 \\ 0, & \text{if } \hat{\mathbf{V}}_{t-1}^i = 1 \text{ and } \mathbf{V}_t^i = 0. \end{cases}$

Here, Φ_A is the appearance aggregation function, which takes a weighted sum of the appearance representations from the previous tracklet and the new detection. Note that, at the start of the tracklet we simply assign the initial single-frame appearance to the tracklet ($\hat{\mathbf{A}}_0^i = \mathbf{A}_0^i$). With this definition of Φ_A , we can aggregate appearance information over time, while allowing the representation to change slowly to account for slight appearance changes of the person during a video. Moreover, the UV image provides appearance of each point on the body surface independently of body pose and viewpoint which enables summation on the pixel space, without any learnable components. Figure 3 shows how the UV image of the person is aggregated over time and used for association of new detections.

For prediction, we make the realistic assumption that human appearance will not change rapidly over time. Then, the appearance of the tracklet $\hat{\mathbf{A}}_t^i$ functions as a reasonable prediction for the future appearance of the person. Therefore, we use $\hat{\mathbf{A}}_t^i$ as the prediction for appearance and use it to measure similarity against detections in future frames.

Location: Lifting humans from pixels into the 3D space allows us to place them in a 3D coordinate frame. Let us assume that a person i at time t has an estimated 3D location $\mathbf{l}_t^i \in \mathcal{R}^3$. Although, we can get an estimate for the location of the person in the global camera frame, this tends to be noisy, particularly along the z -axis. To avoid any instabilities when it comes to predicting future location, instead of performing our prediction on the Euclidean $(X, Y, Z)^T$ space, we express our locations in an equivalent $\mathbf{l}_t^i = (x, y, n)^T$ space where (x, y) is the location of the root of the person in the pixel space and n is *nearness*, defined

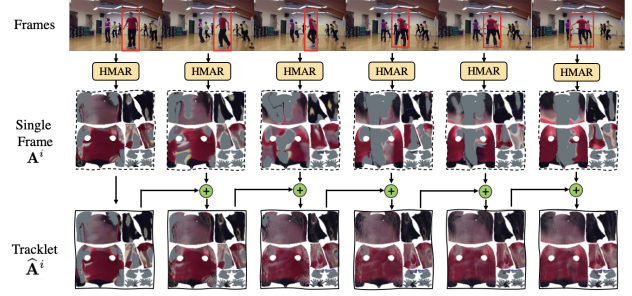


Figure 3. **Prediction of appearance:** The single frame appearance \mathbf{A}^i is aggregated over time for the prediction of the tracklet appearance $\hat{\mathbf{A}}^i$. At the start, only the front side of the *person* is visible, however as the person moves their visibility changes, making only their back side visible. For the single frame appearance, the visible regions change corresponding to the visibility of the person in the frame. However, for the tracklet, the appearance is accumulated over time. Even if the front side is not visible in the last frame, the tracklet can predict these regions using its past.

as log inverse depth $n = \log(1/z)$. Nearness is a natural parameterization of depth in vision settings, *e.g.*, [24], because of the $1/z$ scaling of perspective projection. In our case it corresponds to the scale of the human figures that we estimate directly from images. We independently linearly regress the location predictions for x, y and n . This is related to the Constant Velocity Assumption (CVA) used in past tracking literature [44], but there is a subtlety here because constant velocity in 3D need not give rise to constant velocity in 2D (a person would appear to speed up in the image as they approaches the camera). But local linearization is always a reasonable approximation to make, which is what we do.

Let us assume that a tracklet has a set of past locations $\mathcal{L}^i = \{\mathbf{l}_t^i, \mathbf{l}_{t-1}^i, \mathbf{l}_{t-2}^i, \dots\}$. Then, the prediction of the location for time step $t + 1$ is given by:

$$\hat{\mathbf{l}}_{t+1}^i = (\hat{x}_{t+1}^i, \hat{y}_{t+1}^i, \hat{n}_{t+1}^i)^T, \quad (3)$$

where $\hat{x}_{t+1}^i = \Phi_L(\{x_t^i, x_{t-1}^i, x_{t-2}^i, \dots, x_{t-w}^i\}, t + 1)$.

Here, Φ_L is the location aggregation function and we use a simple linear regression for prediction in our tracking algorithm. We predict \hat{y}_{t+1}^i and \hat{n}_{t+1}^i in a similar fashion. Φ_L takes the last w observations to fit a line by least squares and regress the future location for x, y and n independently. From the theory of linear regression, the prediction interval for x at a time step t' is given by the equation below:

$$\delta_x(t') = t_{(1-\alpha/2)} \times \sqrt{MSE \times \left(1 + \frac{1}{w} + \frac{(t' - \bar{t})^2}{\sum (t - \bar{t})^2}\right)}. \quad (4)$$

Here, $t_{(1-\alpha/2)}$ is the Student's t distribution with confidence α and degree of freedom $w - 2$. MSE is the mean squared error on the predicted locations and \bar{t} is the mean of the time stamps for the previous observations. We similarly compute prediction intervals Δ_y, Δ_n for y, n respectively.




Figure 4. **Prediction of Pose:** We use a modified version of HMMR [17] with transformer backbone. Having transformer as the backbone gives us the flexibility to have missing people in the tracklet (by masking the attention [34]), while still allowing us to predictions of future poses. Finally, the transformer gives us a movie-strip representation and that is used to regress future poses.

Pose: For the pose pathway, we integrates pose information across the tracklet and predicts poses for the near future. To do this, we follow the HMMR architecture [17]. Effectively, we learn a function Φ_P that takes as input a series of pose embeddings of a person $\mathcal{P}^i = \{\mathbf{p}_t^i, \mathbf{p}_{t-1}^i, \mathbf{p}_{t-2}^i, \dots\}$ and computes a temporal pose embedding $\hat{\mathbf{p}}_t$. We train this temporal pose aggregation function Φ_P to smooth the pose $\hat{\mathbf{p}}_t^i$ at frame t , and regress future pose representations $\{\hat{\mathbf{p}}_{t+1}^i, \hat{\mathbf{p}}_{t+2}^i, \dots, \hat{\mathbf{p}}_{t+c}^i\}$ (for up to $c = 12$ frames in the future). We use a transformer to compute Φ_P [34]. This choice allows for additional flexibility, since there are frames where the identity cannot be detected (*e.g.*, due to occlusions). Transformer can handle this scenario gracefully, by not attending to these frames.

3.3. Tracking with predicted 3D representations

Given the bounding boxes and their single-frame 3D representations, our tracking algorithm associates identities across frames in an online manner. At every frame, we make future predictions for each tracklet and we measure the similarity with the detected single-frame representation. More specifically, let us assume that we have a tracklet T_i , which has been tracked for a sequence of frames and has information for appearance, pose and location. The tracklet predicts its appearance $\hat{\mathbf{A}}$, pose $\hat{\mathbf{p}}$ and location $\hat{\mathbf{l}}$ for the next frame, and we need to measure a similarity score between these predictions of the tracklet T_i and a detection D_j to make an association. Our tracklet representation has three different attributes (appearance, location and pose), so, directly combining their similarities/distances would not be ideal, since, each attribute has different characteristics. Instead, we investigate the conditional distributions of inliers and outliers of the attributes. Figure 5 presents the corresponding probability distributions for the PoseTrack dataset [3]. The characteristics of these distributions motivate our design decisions for our further modeling.

Assume that tracklet T_i has an appearance representation $\hat{\mathbf{A}}_t^i$. On the detection side, the detection D_j has a single-frame appearance representation \mathbf{A}_t^j . Both of these representations are in the UV pixel space, therefore we first

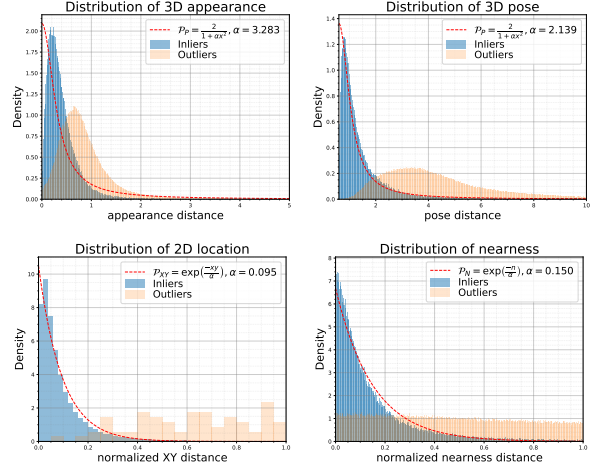


Figure 5. **Conditional distributions of the attribute distances:** We plot the data for the distances between the tracklet prediction and the single frame detection using annotated data from PoseTrack [3]. The curves show the distributions of correct matches (inliers) and incorrect matches (outliers). For 2D location and nearness the distances are normalized by the prediction interval.

encode them into an embedding space using the HMMR appearance-encoder network. This gives us an appearance embedding $\hat{\mathbf{a}}_t^i$ and \mathbf{a}_t^j for the prediction of the tracklet T_i and detection D_j , respectively. We are interested in estimating the posterior probability of the event where the detection D_j belongs to the tracklet T_i , given some distance measure of the appearance feature (Δ_a). Assuming that the appearance distance is $\Delta_a = \|\hat{\mathbf{a}}_t^i - \mathbf{a}_t^j\|_2^2$, then the posterior probability is proportional to the conditional probability of the appearance distances, given correct assignments based on Bayes rule. We model this conditional probability as a Cauchy distribution, based on the observations from the inlier distribution of appearance distances (see also Fig 5):

$$\mathcal{P}_A(D_j \in T_i | d_a = \Delta_a) \propto \frac{1}{1 + \beta_a \Delta_a} \quad (5)$$

The distribution has one scaling hyper-parameter β_a .

Similarly, for pose, we use Cauchy distribution to model the conditional probability of inlier distances. We measure pose distance $\Delta_p = \|\hat{\mathbf{p}}_t^i - \mathbf{p}_t^j\|_2^2$ between the predicted pose representation $\hat{\mathbf{p}}_t^i$ from the tracklet T_i and the pose representation \mathbf{p}_t^j of detection D_j . The posterior probability that the detection belongs to the track, given the pose distance is:

$$\mathcal{P}_P(D_j \in T_i | d_p = \Delta_p) \propto \frac{1}{1 + \beta_p \Delta_p} \quad (6)$$

Here, $\Delta_p = \|\hat{\mathbf{p}}_t^i - \mathbf{p}_t^j\|_2^2$ and β_p is the scaling factor.

For location, let us assume the tracklet T_i has predicted a location $\hat{\mathbf{l}}_t^i = (\hat{x}_t^i, \hat{y}_t^i, \hat{n}_t^i)^T$ with a set of prediction intervals $\{\delta_x, \delta_y, \delta_n\}$, and the detection D_j is at a 3D location $\mathbf{l}_t^j = (x_t^j, y_t^j, n_t^j)^T$. We treat the 3D coordinates x, y and

the nearness term n coordinates independently, and compute the posterior probabilities of the detection belongs to the tracklet given the location distance. We model the conditional probability distribution as an exponential distribution, based on the findings from the empirical data. The Fig 5 shows the distribution of 2D distance and nearness distance, scaled by the confidence interval, of inliers approximately follow the exponential distribution:

$$\mathcal{P}_{XY}(D_j \in T_i | d_{xy} = \Delta_{xy}) \propto \frac{1}{\beta_{xy}} \exp\left(\frac{-\Delta_{xy}}{\beta_{xy}\delta_{xy}}\right). \quad (7)$$

Here, β_{xy} is a scaling parameter for the exponential distribution, Δ_{xy} is the 2D pixel distance between the predicted tracklet and the detection and $\delta_{xy} = \sqrt{\delta_x^2 + \delta_y^2}$ is the prediction interval for the 2D location prediction. The expression for the posterior probability for nearness \mathcal{P}_N is similar:

$$\mathcal{P}_N(D_j \in T_i | d_n = \Delta_n) \propto \frac{1}{\beta_n} \exp\left(\frac{-\Delta_n}{\beta_n\delta_n}\right). \quad (8)$$

Here, β_n is the scaling parameter for the exponential distribution, δ_n is the confidence interval for the nearness prediction, and Δ_n is the L_1 distance between the nearness of the tracklet prediction and the detection.

Given the conditional probabilities of the detection belonging to a tracklet conditioned on the individual cues (appearance, location, pose), we can compute the overall conditional probability of the detection D_j belonging to the tracklet T_i , given all the cues (by making the simplifying assumption that they are independent):

$$\mathcal{P}(D_j \in T_i | \Delta_a, \Delta_p, \Delta_{xy}, \Delta_n) \propto \mathcal{P}_A \mathcal{P}_P \mathcal{P}_{XY} \mathcal{P}_N. \quad (9)$$

This allow us to estimate how probable an association is based on various attribute distances. Finally, we map the similarity measures (probability values up to a scale), to cost values, for solving association. The cost function between the detection representations and a predicted representations of the tracklet is defined as:

$$\begin{aligned} \Phi_{C(D_j, T_i)} &= -\log(\mathcal{P}(D_j \in T_i)) \\ &= -\log(\mathcal{P}_A) - \log(\mathcal{P}_P) - \log(\mathcal{P}_{XY}) - \log(\mathcal{P}_N), \end{aligned} \quad (10)$$

where the second equality is up to an additive constant. Once the cost between all the tracklets and the detection is computed, the Hungarian algorithm solves the association.

Estimating the parameters of the cost function: The cost function Φ_C has 4 parameters ($\beta_a, \beta_p, \beta_{xy}$ and β_n). Additionally, the Hungarian algorithm has one parameter β_{th} to decide whether the tracklet is not a match to the detection. Therefore, overall we have five parameters for the association part of our method. We treat this as an empirical risk minimization problem and optimize the β values based on a loss function. We initialize $\beta_a, \beta_p, \beta_{xy}$ and β_n with the values from the estimated density functions and use frame level

Algorithm 1 Tracking Algorithm

```

1: procedure PHALP TRACKING
2: Require: All active tracklets  $\mathcal{T}$ , all detections and
   their single frame 3D representations at time  $t$ ,  $\mathcal{D}$  and
   maximum age of a tracklet  $t_{max}$ .
3:   for  $T_j \in \mathcal{T}$  do
4:     # Predict all attributes for the next frame.
5:      $\hat{\mathbf{A}}_t^j \leftarrow \Phi_A(\{\mathbf{A}_{t-1}^j, \mathbf{A}_{t-2}^j, \dots\})$ 
6:      $\hat{\mathbf{P}}_t^j \leftarrow \Phi_P(\{\mathbf{P}_{t-1}^j, \mathbf{P}_{t-2}^j, \dots\})$ 
7:      $\hat{\mathbf{I}}_t^j \leftarrow \Phi_L(\{\mathbf{I}_{t-1}^j, \mathbf{I}_{t-2}^j, \dots\})$ 
8:     # Compute cost between tracklets and detections.
9:      $\mathbf{C}_{i,j} \leftarrow \Phi_C(D_i, T_j)$  for all  $D_i \in \mathcal{D}$  and  $T_j \in \mathcal{T}$ 
10:    # Hungarian to assign detections to tracklets.
11:     $\mathcal{M}, \mathcal{T}_u, \mathcal{D}_u \leftarrow \text{Assignment}(\mathbf{C})$ 
12:    # Update the matched tracklets.
13:     $\mathcal{T} \leftarrow \{T_j(D_i), \forall (i, j) \in \mathcal{M}\}$ 
14:    # Increase the age of unmatched tracklets.
15:     $\mathcal{T} \leftarrow \{T_j(\text{age})+ = 1, \forall (j) \in \mathcal{T}_u\}$ 
16:    # Make new tracklets with unmatched detections.
17:     $\mathcal{T} \leftarrow \{T_j(D_i), \forall (i) \in \mathcal{D}_u, j = |\mathcal{T} + 1|\}$ 
18:    Kill the tracklets with age  $\geq t_{max}$ .
19: return Tracklets  $\mathcal{T}$ 

```

association error as a loss function for the optimization. We use the Nelder–Mead algorithm [33] for this optimization. Finally, the optimized β values are used for the cost function across all datasets. The sketch of the tracking algorithm is shown in Algorithm 1.

3.4. Extension to shot changes

Our framework can be easily extended to also handle shot changes, which are common in edited media (movies, TV shows, sports, etc). We use a shot detector [15] to reliably identify frames that indicate shot changes. Knowing the shot boundary, during tracking, we update the distance metric accordingly. More specifically, since appearance and 3D pose are invariant to the viewpoint, we keep these factors in the distance computation, while we drop the location distance from the distance metric, because of the change in the camera location. Then, the association is computed based on this updated metric. We demonstrate this utility on the AVA dataset [12] and present results in Section 4.

4. Experiments

In this section, we present the experimental evaluation of our approach. We report results on three datasets: PoseTrack [3], MuPoTS [31] and AVA [12], which capture a diverse set of sequences, including sports, casual interactions and movies. Our method operates on detections and masks coming from an off-the-shelf Mask-RCNN network [14], and returns the identity label for each one of them. Therefore, the metrics we use to report results also focus on

Method	IDs↓	PoseTrack	
		MOTA↑	IDF1↑
w/o 3D appearance	632	58.4	74.9
w/o 3D pose	558	58.9	76.2
w/o location	948	57.3	71.6
w/o nearness	622	58.5	74.8
w/o pose, appearance	645	57.5	74.7
w/o mask	573	58.1	75.3
Full system	541	58.9	76.4

Table 1. **Ablation of the main components of PHALP on PoseTrack [3].** Removing each tracking cue (3D appearance, 3D pose or 3D location) leads to degradation in the performance.

identity tracking at the level of the bounding box. More specifically, we report results using Identity switches (IDs), Multi-Object Tracking Accuracy (MOTA) [18], ID F1 score (IDF1) [36] and HOTA [29]. In all cases, we adopt the protocols of Rajasegaran *et al.* [35] for evaluation.

First, we ablate the main components of our approach. Specifically, we investigate the effect of each one of the tracking cues we employ, *i.e.*, 3D appearance, 3D location and 3D pose, and how they affect the overall tracking pipeline. For this comparison, we report results on the PoseTrack dataset [3]. The full results are presented in Table 1. As we can see, removing each one of the main cues leads to degradation in the performance of the system, where 3D location seems to have the largest effect on the performance, followed by appearance and 3D pose. While the effect of 3D pose is not significant as other cues, computing the 3D pose is essential because a) it allows us to get a good location estimate, and b) it guides the appearance calculation. Moreover, this ablation also highlights the importance of having the nearness term in the cost function, a feature that is not available to purely 2D tracking methods. Additionally, we evaluate our system when only the 3D location is included in the cost function. This has a negative effect on the tracking performance, validating the importance of the other cues. Finally, we evaluate our method without masks, during the appearance aggregation and HMR stage. Our experiment shows that access to masks improves performance.

Next, we evaluate our approach in comparison with the state-of-the-art methods. The results are presented in Table 2. We report results on PoseTrack [3], MuPoTS [31] and AVA [12]. Our method outperforms the previous baselines, as well as the state-of-the-art approach of Rajasegaran [35]. The gains are significant across all metrics. Our method also outperforms the other approaches in the HOTA metric.

Finally, we provide qualitative results of our method in Fig 6. These results indicate that our method performs reliably even in very hard occlusion cases, while it can recover the correct identity over multiple successive occlusions. Please note the robustness of our method in com-

plex motion sequences, shot changes and long trajectories. For visualization we use a colored head-mask to represent a unique track.

More design choices: PHALP has a modular design, where each components can be easily replaced. We demonstrate this by altering some of our original choices. For example, we replace MaskRCNN [14] with PointRend [20] at the detection stage. However, we did not observe any significant improvement on tracking metrics with this change. Moreover, we replace the inferred 3D location from HMAR with the equivalent location inferred by PARE [22]. This change improves overall tracking performance by 5% in terms of ID switches. Finally, we also investigate the option of postponing the final decision about the identity of a detection. We refer to this as LMC (Lookback Merge Check), which can connect tracklets that have been over-segmented. With LMC, for every new tracklet, we regress backwards (*i.e.*, past location) and check if there is any older tracklet in that location with similar appearance to the new tracklet. If such a tracklet exists, then we merge these two tracklets. This lookback merge and check (LMC) can help to achieve roughly 20 less ID switches in the PoseTrack dataset. For more implementation details and for detailed results for all these experiments, please see the SupMat.

5. Discussion

We presented PHALP, an approach for monocular people tracking, by predicting appearance, location and pose in 3D. Our method relies on a powerful backbone for 3D human mesh recovery, modeling on the tracklet level for collecting information across a tracklet’s trajectory, and eventually predicting the future states of each tracklet. One of the main benefits of PHALP is that the association aspect requires tuning of only five parameters, which makes it very friendly for training on small-scale tracking datasets, where annotating the identity of every person in a video can be expensive. Our approach can be naturally extended to make use of more attributes, *e.g.*, a face embedding, which could be useful for cases with close-ups, like movies.

The main assumptions for PHALP are that we have access to a good object detector for the initial bounding box/mask detection, and a strong HMAR network for single-frame lifting of people to 3D. If the performance of these components is not satisfactory, it can also affect tracking quality. Regarding societal impact, tracking systems can be used to monitor patients or help with some treatments [7]. On the other hand, tracking systems have often been used for human surveillance. We do not condone such use. Instead, we believe that a tracking system will be valuable for studying social-human interactions.

Acknowledgements: This work was supported by ONR MURI (N00014-14-1-0671), the DARPA Machine Common Sense program, as well as BAIR and BDD sponsors.

Method	Posetrack				MuPoTS				AVA	
	IDs↓	MOTA↑	IDF1↑	HOTA↑	IDs↓	MOTA↑	IDF1↑	HOTA↑	IDs↓	IDF1↑
Trackformer [32]	1263	33.7	64.0	46.7	43	24.9	62.7	53.2	716	40.9
Tracktor [4]	702	42.4	65.2	38.5	53	51.5	70.9	50.3	289	46.8
AlphaPose [9]	2220	36.9	66.9	37.6	117	37.8	67.6	41.8	939	41.9
FlowPose [41]	1047	15.4	64.2	38.0	49	21.4	67.1	43.0	452	52.9
T3DP [35]	655	55.8	73.4	50.6	38	62.1	79.1	59.2	240	61.3
PHALP	541	58.9	76.4	52.9	22	66.2	81.4	59.4	227	62.7

Table 2. **Comparison with state-of-the-art tracking methods.** We compare our method, PHALP, with various tracking methods in three different datasets. Our approach outperforms the other baselines across all datasets and metrics.




Figure 6. **Qualitative Results:** We show the tracking performance of PHALP in various datasets (frame number is shown at the top left corner). While all the tracks have full 3D reconstructions, we only use the **head-masks** to visualize the tracks since it is easier to spot errors. The first three rows are from the PoseTrack dataset [3]. These results show that even during successive occlusions PHALP is able to track the identity of the correct person. Note that, in the first row, although the two persons with the **green head-mask** and the **purple head-mask** have similar appearance, our method can track each one of them successfully. In the second row, the **player** is going through multiple occlusions, yet recovered correctly. The third row shows the robustness of our linearization approximation for 3D location prediction, even when the motions of the players are very complex. In the MuPoTS dataset [31] (4th row), our method can handle very close interactions between people. This is due to the fact that, our modification of HMAR recovers meshes conditioned on the detected mask. We also show results (5th row) on the AVA dataset [12]. After the 3rd frame, there is a shot change in the video, and the person with the **yellow head-mask** is tracked successfully across the shots. Finally, we show qualitative results on a MOT17 sequence. The person with the **blue head-mask** is tracked for the whole sequence while they are going through multiple occlusions for a long time. More results at the [PHALP website](#).

References

- [1] Emre Aksan, Peng Cao, Manuel Kaufmann, and Otmar Hilliges. A spatio-temporal transformer for 3D human motion prediction. In *3DV*, 2020. 3
- [2] Alexandre Alahi, Kratharth Goel, Vignesh Ramanathan, Alexandre Robicquet, Li Fei-Fei, and Silvio Savarese. Social LSTM: Human trajectory prediction in crowded spaces. In *CVPR*, 2016. 3
- [3] Mykhaylo Andriluka, Umar Iqbal, Eldar Insafutdinov, Leonid Pishchulin, Anton Milan, Juergen Gall, and Bernt Schiele. PoseTrack: A benchmark for human pose estimation and tracking. In *CVPR*, 2018. 5, 6, 7, 8
- [4] Philipp Bergmann, Tim Meinhardt, and Laura Leal-Taixe. Tracking without bells and whistles. In *ICCV*, 2019. 3, 8
- [5] Christoph Bregler and Jitendra Malik. Tracking people with twists and exponential maps. In *CVPR*, 1998. 3
- [6] Gioele Ciaparrone, Francisco Luque Sánchez, Siham Tabik, Luigi Troiano, Roberto Tagliaferri, and Francisco Herrera. Deep learning in video multi-object tracking: A survey. *Neurocomputing*, 381:61–88, 2020. 2
- [7] R James Cotton. Posepipe: Open-source human pose estimation pipeline for clinical research. *arXiv preprint arXiv:2203.08792*, 2022. 7
- [8] Patrick Dendorfer, Aljosa Osep, Anton Milan, Konrad Schindler, Daniel Cremers, Ian Reid, Stefan Roth, and Laura Leal-Taixé. MOTChallenge: A benchmark for single-camera multiple target tracking. *IJCV*, 129(4):845–881, 2021. 2
- [9] Hao-Shu Fang, Shuqin Xie, Yu-Wing Tai, and Cewu Lu. RMPE: Regional multi-person pose estimation. In *ICCV*, 2017. 8
- [10] Christoph Feichtenhofer, Axel Pinz, and Andrew Zisserman. Detect to track and track to detect. In *ICCV*, 2017. 3
- [11] Rohit Girdhar, Georgia Gkioxari, Lorenzo Torresani, Manohar Paluri, and Du Tran. Detect-and-track: Efficient pose estimation in videos. In *CVPR*, 2018. 3
- [12] Chunhui Gu, Chen Sun, David A Ross, Carl Vondrick, Caroline Pantofaru, Yeqing Li, Sudheendra Vijayanarasimhan, George Toderici, Susanna Ricco, Rahul Sukthankar, Cordelia Schmid, and Jitendra Malik. AVA: A video dataset of spatio-temporally localized atomic visual actions. In *CVPR*, 2018. 2, 6, 7, 8
- [13] Peng Guan, Alexander Weiss, Alexandru O Balan, and Michael J Black. Estimating human shape and pose from a single image. In *ICCV*, 2009. 3
- [14] Kaiming He, Georgia Gkioxari, Piotr Dollár, and Ross Girshick. Mask R-CNN. In *ICCV*, 2017. 3, 6, 7
- [15] Qingqiu Huang, Yu Xiong, Anyi Rao, Jiase Wang, and Dahua Lin. MovieNet: A holistic dataset for movie understanding. In *ECCV*, 2020. 6
- [16] Angjoo Kanazawa, Michael J Black, David W Jacobs, and Jitendra Malik. End-to-end recovery of human shape and pose. In *CVPR*, 2018. 3
- [17] Angjoo Kanazawa, Jason Y Zhang, Panna Felsen, and Jitendra Malik. Learning 3D human dynamics from video. In *CVPR*, 2019. 2, 3, 5
- [18] Rangachar Kasturi, Dmitry Goldgof, Padmanabhan Soundararajan, Vasant Manohar, John Garofolo, Rachel Bowers, Matthew Boonstra, Valentina Korzhova, and Jing Zhang. Framework for performance evaluation of face, text, and vehicle detection and tracking in video: Data, metrics, and protocol. *PAMI*, 2008. 7
- [19] Tarasha Khurana, Achal Dave, and Deva Ramanan. Detecting invisible people. In *ICCV*, 2021. 3
- [20] Alexander Kirillov, Yuxin Wu, Kaiming He, and Ross Girshick. PointRender: Image segmentation as rendering. In *CVPR*, 2020. 7
- [21] Muhammed Kocabas, Nikos Athanasiou, and Michael J Black. VIBE: Video inference for human body pose and shape estimation. In *CVPR*, 2020. 3
- [22] Muhammed Kocabas, Chun-Hao P Huang, Otmar Hilliges, and Michael J Black. PARE: Part attention regressor for 3D human body estimation. In *ICCV*, 2021. 3, 7
- [23] Muhammed Kocabas, Chun-Hao P Huang, Joachim Tesch, Lea Muller, Otmar Hilliges, and Michael J Black. SPEC: Seeing people in the wild with an estimated camera. In *ICCV*, 2021. 3
- [24] Jan J Koenderink. Optic flow. *Vision research*, 26(1):161–179, 1986. 4
- [25] Nikos Kolotouros, Georgios Pavlakos, Michael J Black, and Kostas Daniilidis. Learning to reconstruct 3D human pose and shape via model-fitting in the loop. In *ICCV*, 2019. 3
- [26] Nikos Kolotouros, Georgios Pavlakos, Dinesh Jayaraman, and Kostas Daniilidis. Probabilistic modeling for human mesh recovery. In *ICCV*, 2021. 3
- [27] Oh-Hun Kwon, Julian Tanke, and Juergen Gall. Recursive bayesian filtering for multiple human pose tracking from multiple cameras. In *ACCV*, 2020. 3
- [28] Matthew Loper, Naureen Mahmood, Javier Romero, Gerard Pons-Moll, and Michael J Black. SMPL: A skinned multi-person linear model. *ACM Transactions on Graphics (TOG)*, 34(6):1–16, 2015. 3
- [29] Jonathon Luiten, Aljosa Osep, Patrick Dendorfer, Philip Torr, Andreas Geiger, Laura Leal-Taixé, and Bastian Leibe. HOTA: A higher order metric for evaluating multi-object tracking. *IJCV*, 2021. 7
- [30] Dushyant Mehta, Oleksandr Sotnychenko, Franziska Mueller, Weipeng Xu, Mohamed Elgharib, Pascal Fua, Hans-Peter Seidel, Helge Rhodin, Gerard Pons-Moll, and Christian Theobalt. XNect: Real-time multi-person 3D motion capture with a single RGB camera. *ACM Transactions on Graphics (TOG)*, 39(4):82–1, 2020. 3
- [31] Dushyant Mehta, Oleksandr Sotnychenko, Franziska Mueller, Weipeng Xu, Srinath Sridhar, Gerard Pons-Moll, and Christian Theobalt. Single-shot multi-person 3D pose estimation from monocular RGB. In *3DV*, 2018. 6, 7, 8
- [32] Tim Meinhardt, Alexander Kirillov, Laura Leal-Taixe, and Christoph Feichtenhofer. TrackFormer: Multi-object tracking with transformers. In *CVPR*, 2022. 3, 8
- [33] John A Nelder and Roger Mead. A simplex method for function minimization. *The computer journal*, 7(4):308–313, 1965. 6

- [34] Georgios Pavlakos, Jitendra Malik, and Angjoo Kanazawa. Human mesh recovery from multiple shots. In *CVPR*, 2022. 3, 5
- [35] Jathushan Rajasegaran, Georgios Pavlakos, Angjoo Kanazawa, and Jitendra Malik. Tracking people with 3D representations. In *NeurIPS*, 2021. 1, 2, 3, 7, 8
- [36] Ergys Ristani, Francesco Solera, Roger Zou, Rita Cucchiara, and Carlo Tomasi. Performance measures and a data set for multi-target, multi-camera tracking. In *ECCV*, 2016. 7
- [37] Michael Snower, Asim Kadav, Farley Lai, and Hans Peter Graf. 15 keypoints is all you need. In *CVPR*, 2020. 3
- [38] Pavel Tokmakov, Jie Li, Wolfram Burgard, and Adrien Gaidon. Learning to track with object permanence. In *ICCV*, 2021. 3
- [39] Xinshuo Weng, Yongxin Wang, Yunze Man, and Kris M Kitani. GNN3DMOT: Graph neural network for 3D multi-object tracking with 2D-3D multi-feature learning. In *CVPR*, 2020. 3
- [40] Bin Xiao, Haiping Wu, and Yichen Wei. Simple baselines for human pose estimation and tracking. In *ECCV*, 2018. 3
- [41] Yuliang Xiu, Jiefeng Li, Haoyu Wang, Yinghong Fang, and Cewu Lu. Pose Flow: Efficient online pose tracking. In *BMVC*, 2018. 8
- [42] Jiarui Xu, Yue Cao, Zheng Zhang, and Han Hu. Spatial-temporal relation networks for multi-object tracking. In *ICCV*, 2019. 3
- [43] Yihong Xu, Aljosa Osep, Yutong Ban, Radu Horaud, Laura Leal-Taixé, and Xavier Alameda-Pineda. How to train your deep multi-object tracker. In *CVPR*, 2020. 3
- [44] Alper Yilmaz, Omar Javed, and Mubarak Shah. Object tracking: A survey. *Acm computing surveys (CSUR)*, 38(4):13–es, 2006. 2, 4
- [45] Jason Y Zhang, Panna Felsen, Angjoo Kanazawa, and Jitendra Malik. Predicting 3D human dynamics from video. In *ICCV*, 2019. 3
- [46] Yuxiang Zhang, Liang An, Tao Yu, Xiu Li, Kun Li, and Yebin Liu. 4D association graph for realtime multi-person motion capture using multiple video cameras. In *CVPR*, 2020. 3
- [47] Yang Zhang, Hao Sheng, Yubin Wu, Shuai Wang, Weifeng Lyu, Wei Ke, and Zhang Xiong. Long-term tracking with deep tracklet association. *IEEE Transactions on Image Processing*, 29:6694–6706, 2020. 3